Lecture 19 · 22.6.2021 · 14:15–16:00 via Zoom

Optimal stopping problems (II)

In the last lecture we proved the basic result of optimal stopping in discrete time and finite horizon N:

Theorem. *Let $(Y_n)_{n\in[\![N]\!]}$ be an adapted process such that $\mathbb{E}[|Y_n|]<\infty$ for all $n\in[\![N]\!]$, recall then $[\![N]\!]=\{1,\ldots,N\}$. Define $(Z_n)_{n\in[\![N]\!]}$ as*

$$Z_N=Y_N, \qquad Z_n=\sup(\mathbb{E}[Z_{n+1}|\mathscr{F}_n],Y_n), \qquad n\in[\![N-1]\!],$$

and let

$$T^*:=\inf\{k\in[\![N]\!]:Y_k=Z_k\}.$$

Then the process $Z_\bullet$ is the Snell envelope of $Y_\bullet$, the process $(Z_n^{T^*})_{n\in[\![N]\!]}$ is a martingale and

$$J_N:=\sup_{T\in\mathscr{T}_N}\mathbb{E}[Y_T]=\mathbb{E}[Y_{T^*}]=\mathbb{E}[Z_{T^*}]=\mathbb{E}[Z_1]$$

where $\mathscr{T}_N$ is the set of all stopping time taking values in $[\![N]\!]$. The stopping time T^* is optimal.

Remark. The process $(Z_n)_{n\in[\![N]\!]}$ is also the average gain for playing optimally from time n on, that is

$$Z_n=\sup_{T\in\mathscr{T}_N:T\geqslant n}\mathbb{E}[Y_T|\mathscr{F}_n].$$

One can actually introduce this process in this way and then proceed to prove all the properties we have shown, i.e. Snell's envelope property and also the representation for optimal stopping times.

Markovian problems

In many optimal stopping problems the following assumptions are satisfied:

Assumption. There exists an adapted process $(X_n)_{n\geqslant 0}$ taking values in a measure space $(E,\mathscr{E})$ for which, for any $n\geqslant 0$ we have

a) For any bounded measurable function $f:E\to\mathbb{R}$

$$\mathbb{E}[f(X_{n+1})|\mathscr{F}_n]=\mathbb{E}[f(X_{n+1})|X_n]=(P_{n+1}f)(X_n)$$

for some probability kernel on E, that is a function $P_{n+1}:E\times\mathscr{E}\to[0,1]$ such that it is a probability measure in the second variable and for any $A\in\mathscr{E}$ the function $x\in E\mapsto P_{n+1}(x,A)$ has to be ($\mathscr{E}$-)measurable. This is called the Markov property for the process $(X_n)_{n\geqslant 0}$ wrt. $(\mathscr{F}_n)_{n\geqslant 0}$. Much more on this later on in the course.

b) For any $n\in[\![N]\!]$ the gain Y_n can be expressed as a function of X_n, that is $Y_n=\varphi_n(X_n)$ for some measurable function $\varphi_n:E\to\mathbb{R}$.

When this assumption is satisfied we say that the optimal stopping problem is *Markovian*.

Remark. Recall that the notation $P_n f: E \to \mathbb{R}$ means the function defined via the integral

$$(P_n f)(x) = \int_E f(y) P_n(x, \mathrm{d}y).$$

When the problem is Markovian the Snell envelope and the optimal stopping time T^* has a very simple and concrete representation which is very important for real-world applications.

Let's find it. First note that

$$Z_N = Y_N = \varphi_N(X_N),$$

moreover

$$Z_{N-1} = \sup(\mathbb{E}[Z_N|\mathscr{F}_{N-1}], Y_{N-1}) = \sup(\mathbb{E}[\varphi_N(X_N)|\mathscr{F}_{N-1}], \varphi_{N-1}(X_{N-1}))$$

and using the Markovian structure of the problem we have now

$$Z_{N-1} = \sup((P_N \varphi_N)(X_{N-1}), \varphi_{N-1}(X_{N-1})) = v_{N-1}(X_{N-1})$$

where $v_{N-1}(x) := \sup((P_N \varphi_N)(x), \varphi_{N-1}(x))$ for any $x \in E$. So we see a structure emerging, in particular by recursion Z_n will be a function of X_n which we call $v_n: E \to \mathbb{R}$, namely we set $(v_n: E \to \mathbb{R})_{n \in [\![N]\!]}$ satisfying

$$v_N = \varphi_N, \qquad v_n = \sup(P_{n+1} v_{n+1}, \varphi_n). \tag{1}$$

And with this definition you can check very easily that indeed we have

$$Z_n = v_n(X_n)$$

for all $n \in [\![N]\!]$. Moreover

$$J_N = \mathbb{E}[Z_1] = \mathbb{E}[v_1(X_1)]$$

and the optimal stopping time T^* is given by the first time $k \in [\![N]\!]$ when $Y_k = Z_k$, namely when $\varphi_k(X_k) = v_k(X_k)$. We can restate this by defining the stopping region at time $k \in [\![N]\!]$ to be

$$\mathscr{S}_k := \{x \in E: \varphi_k(x) = v_k(x)\} \in \mathscr{E}$$

and then

$$T^* = \inf\{k \in [\![N]\!]: X_k \in \mathscr{S}_k\}.$$

The Markovian assumption allow to transfer the optimal stopping problem from the abstract measure space $(\Omega, \mathscr{F})$ (which in particular depends on the horizon N) to the "concrete" measure space $(E, \mathscr{E})$ (which in particular is independent of N).

The Moser problem

Let's discuss a specific very simple example called the Moser problem. This is a Markovian optimal stopping problem in discrete time and horizon N where $(X_n)_{n \in [\![N]\!]}$ is just a family of i.i.d. random variables with values in $E = [0, 1]$ with the Borel σ-algebra $\mathscr{E} = \mathscr{B}([0, 1])$ and where we take $\varphi_n: [0, 1] \to \mathbb{R}$ to be the identity: $\varphi_n(x) = x$, namely we have $Y_n = X_n$, at every step n we gain the value of the process X_n. The natural filtration in this problem is then $\mathscr{F}_n = \sigma(X_k: k \in \{1, \ldots, n\})$ the filtration generated by the process $X_\bullet$.

One can check that due to independence $(X_n)_{n\in[\![N]\!]}$ is indeed Markovian wrt. this filtration, i.e.:

$$\mathbb{E}[f(X_{n+1})|\mathscr{F}_n] = \mathbb{E}[f(X_{n+1})|X_n,\ldots,X_1] = \mathbb{E}[f(X_{n+1})] = \mathbb{E}[f(X_{n+1})|X_n] = (P_{n+1}f)(X_n)$$

where the kernel $P_{n+1}(x,\mathrm{d}y)$ is idependent of n and of $x\in E$ (the first variable) and given by

$$P_{n+1}(x,\mathrm{d}y) = \mu(\mathrm{d}y)$$

where $\mu = \mathrm{Law}(X_1)$ is the law of the r.v. X_1. We assume that we know the law μ.

We can know use the theory we constructed above to solve this optimal stopping problem. Since this problem has a Markovian structure we know that to compute the Snell's envelope we can just find the family of functions $v_n\colon[0,1]\to\mathbb{R}$ defined above with eq. (1):

When $n=N$ we have

$$v_N(x) = \varphi_N(x) = x$$

and then for any $n<N$ we compute:

$$v_n(x) = \max\left((P_{n+1}v_{n+1})(x), \varphi_n(x)\right)$$

we have $\varphi_n(x) = x$ and also

$$(P_{n+1}v_{n+1})(x) = \int_{[0,1]} v_{n+1}(y)P_{n+1}(x,\mathrm{d}y) = \int_{[0,1]} v_{n+1}(y)\,\mu(\mathrm{d}y)$$

$$v_n(x) = \max\left(x, \int_{[0,1]} v_{n+1}(y)\,\mu(\mathrm{d}y)\right).$$

Note that

$$J_N = \mathbb{E}[Z_1] = \mathbb{E}[v_1(X_1)] = \int_{[0,1]} v_1(y)\,\mu(\mathrm{d}y).$$

To simplify this representation we let

$$\rho_n := \int_{[0,1]} v_n(y)\,\mu(\mathrm{d}y),$$

then note that $v_n(x) = \max(x, \rho_{n+1})$ and therefore that

$$\rho_n = \int_{[0,1]} \max(x, \rho_{n+1})\,\mu(\mathrm{d}y)$$

so all the information on the solution of this optimal stopping problem is contained in sequence of numbers $(\rho_n)_{n\in[\![N]\!]}$ which then give the functions $v_n(x) = \max(x, \rho_{n+1})$ and therefore the Snell's envelope $Z_n = \max(X_n, \rho_{n+1})$ and the optimal stopping time

$$T^* = \inf\{k\in[\![N-1]\!]\colon X_k \geqslant \rho_{k+1}\} \wedge N$$

that is, once we have computed the numbers $(\rho_k)_{k\in[\![N]\!]}$ we stop the first time k we see a number X_k bigger than ρ_{k+1}, or we stop at N if this never happens before.

In particular

$$\rho_N = \int_{[0,1]} v_N(y)\,\mu(\mathrm{d}y) = \int_{[0,1]} y\,\mu(\mathrm{d}y) = \mathbb{E}[X_1].$$

Note that $\rho_n \geqslant \rho_{n+1} \geqslant \cdots \geqslant \rho_N = \mathbb{E}[X_1]$.

Exercise 1. Compute the levels $(\rho_n)_{n\in[\![N]\!]}$ for $N = 10$ and $\mu = \text{Uniform}([0,1])$ and determine J_{10}.

Remark. Recall that $\int_E f(y)\,\mu(\mathrm{d}y)$ is a notation for the integral of the measurable function f: $E \to \mathbb{R}$ with respect to the measure μ, alternative notations are:

$$\int_E f(y)\,\mu(\mathrm{d}y) = \int_E f(y)\mathrm{d}\mu(y) = \int_E f\,\mathrm{d}\mu$$

Markov chains (Markov processes in discrete time)

Markov chains are an important class of stochastic processes in discrete time which are very useful in applications and which have a nice and complex theory.

To introduce the subject, let us look at a specific example of random recurrences.

Random recurrences

Take a sequence of i.i.d random variables $(U_n)_{n\geqslant 1}$ with values in $[0,1]$ and uniformly distributed there, and let $(\phi_n: E\times[0,1]\to E)_{n\geqslant 1}$ be a family of measurable functions where $(E,\mathscr{E})$ is a given measure space. We let also ν to be a probability measure on $(E,\mathscr{E})$. Then define a new stochastic process $(X_n)_{n\geqslant 0}$ as follows: X_0 is independent of $(U_n)_{n\geqslant 1}$ and with law ν, i.e. $\nu = \text{Law}_{\mathbb{P}}(X_0)$ and then let regursively for all $n \geqslant 1$

$$X_n := \phi_n(X_{n-1}, U_n).$$

This is like in a board game: at time $n-1$ we are in position X_{n-1}, we throw the dice U_n and the rule of the game ϕ_n tells us where to be in the next round, i.e. our new position X_n.

This is a random recurrence. Many phenomena in science and in applications can be modeled with such stochastic processes. Two extreme situations:

- Take $\phi_n(x,y) = \varphi_n(y)$ to be independent of the variable $x\in E$, then $(X_n)_{n\geqslant 1}$ is just a family of independent random variables $X_n = \varphi_n(U_n)$, moreover if φ_n do not depend on n they are also indentically distributed.

- Take $\phi_n(x,y) = f_n(x)$ to be independent of its second variable $y\in[0,1]$, then the random sequence $(U_n)_n$ is not influencing anymore the process $(X_n)_{n\geqslant 0}$ which become a deterministic recursion, i.e. $X_n = f_n(X_{n-1})$, possibly with a random initial condition X_0. This represents a phenomenon which evolves according a deterministic rule from state to state.

Random recurrences interpolate naturally between these two extreme situations: we have a mechanism to evolve the current state into a future state, but this mechanism is a combination of deterministic and random behaviours.

In this problem we have a natural filtration, which is the one for example generated by the process $(X_n)_{n\geqslant 0}$ itself, i.e. $\mathscr{F}_n = \sigma(X_0,\ldots,X_n)$. There is also another filtration which is the one generated by the $(U_n)_{n\geqslant 1}$ togetgher with X_0, i.e. $\mathscr{G}_n = \sigma(X_0, U_1,\ldots,U_n)$.

As an exercise, I leave you to prove the following lemma.

Lemma. *Let $(X_n)_{n\geqslant 0}$ be the random recurrence just defined. For any measurable bounded function $f: E\to\mathbb{R}$ we have*

$$\mathbb{E}[f(X_{n+1})|\mathscr{G}_n] = \mathbb{E}[f(X_{n+1})|X_n]$$

and a similar property for the filtration $(\mathscr{F}_n)_{n\geqslant 0}$. *This property is called the Markov property of the process* $(X_n)_{n\geqslant 0}$ *wrt. the filtration* $(\mathscr{G}_n)_{n\geqslant 0}$ *(or* $(\mathscr{F}_n)_{n\geqslant 0}$*).*

Remark. The Markov property intuitively says that the “future” of the process depends on its “past” only via the “present”. Not all stochastic processes are markov processes.
