Lecture 18 · 18.6.2021 · 10:15–12:00 via Zoom

Optimal stopping problems

Optimal stopping problems involve the given of a stochastic process $(Y_n)_{n\geqslant 1}$ which we assume adapted to given filtration $\mathscr{F}_\bullet = (\mathscr{F}_n)_{n\geqslant 0}$ representing our knowledge of the system. The goal is to optimize the average value of Y stopped at some random time $T:\Omega\to\mathbb{N}$. In applications is natural to take this random time to be a stopping time and therefore what we want to study is the quantity $\mathbb{E}[Y_T]$ over all admissible stopping times T in our specific problem.

We interpret this situation as a game: we imagine that $Y_n(\omega)$ is the gain which we obtain if we decide to stop at time n in the situation ω. We then try to find a stopping strategy which maximizes the average gain. Stopping times are the natural class of allowed stopping strategies since they do not involve knowledge from the future, i.e. $\{T=n\}\in\mathscr{F}_n$, stopping at time n can be decided only using the information contained in $\mathscr{F}_n$.

We will only consider problems in *finite horizon*, meaning only involving a finite set of time indexed $n\in[\![N]\!]:=\{1,\ldots,N\}$ with $N<\infty$. We let $\mathscr{T}_N$ to be the set of all stopping times T bounded by N, i.e. $\mathscr{T}_N=\{T:\Omega\to[\![N]\!]:T$ is a stopping time for the filtration $\mathscr{F}_\bullet\}$.

Our problem is to study to the optimal average gain J_N with horizon N:

$$J_N = \sup_{T\in\mathscr{T}_N}\mathbb{E}[Y_T].$$

We say that $T^*\in\mathscr{T}_N$ is an optimal stopping time iff $J_N=\mathbb{E}[Y_{T^*}]$. It does not have to be unique of course.

Notation. $\inf_N A = (\inf A)\wedge N$ for all $A\subseteq\mathbb{N}$.

Question: How we compute an optimal stopping time?

As with many optimization problems, an efficient solution here goes thru the determination of a suitable *value function* $(Z_n)_{n\in[\![N]\!]}$ associated to the choices still available at time $n\in[\![N]\!]$.

The value function represents the average gain conditional on the information gained up to time n, namely conditional on $\mathscr{F}_n$. Let's see what we know about it:

a) $(Z_n)_n$ must be an adapted process, i.e. Z_n must be measurable wrt. $\mathscr{F}_n$.

b) Provided $n<N$, $Z_n\geqslant Y_n$: we can always stop at time n, i.e. not play further, and get Y_n.

c) Provided $n<N$, $Z_n\geqslant\mathbb{E}[Z_{n+1}|\mathscr{F}_n]$: my current position has a value which is no less than what I will get in average if I continue one more step, given I already know $\mathscr{F}_n$.

At every step $n<N$ I have indeed two possibilities: stop there or continue one more step. At the final time N however I do not have these two choices: I need to stop right away, so if I'm at time N I gain Y_N. This gives us an interesting information: $Z_N=Y_N$. Morever since I have only two possibilies in each other step, the value function must satisfy the following backward recursion:

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

This equation defines uniquely the process $(Z_n)_{n\in[\![N]\!]}$, which by construction is a supermartingale which bounds Y from above.

We will show that $Z_\bullet$ is the *Snell's envelope* of $Y_\bullet$, that is the smallest supermartingale $Q_\bullet$ which bounds $Y_\bullet$ from above, i.e. $Q_n \geqslant Y_n$ for all $n \in [\![N]\!]$.

> **Theorem.** *Let $(Y_n)_{n\in[\![N]\!]}$ be an adapted process such that $\mathbb{E}[|Y_n|] < \infty$ for all $n \in [\![N]\!]$. Define $(Z_n)_{n\in[\![N]\!]}$ as in (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*
>
> $$\mathbb{E}[Z_1] = \mathbb{E}[Z_{T^*}] = \mathbb{E}[Y_{T^*}] = J_N.$$
>
> *The stopping time T^* is optimal.*

Proof. Let us prove that $Z_\bullet$ is the Snell envelope of $Y_\bullet$. Assume $Q_\bullet$ is a supermartingale and such that $Q_n \geqslant Y_n$ for all $n \in [\![N]\!]$. Now let's perform a backward induction on n to prove that $Q_n \geqslant Z_n$ for all $n \in [\![N]\!]$. Base case: for $n = N$ we have $Q_N \geqslant Y_N = Z_N$. Induction step: assume that $Q_{n+1} \geqslant Z_{n+1}$ then by the supermartingale property

$$Q_n \geqslant \mathbb{E}[Q_{n+1}|\mathscr{F}_n] \geqslant \mathbb{E}[Z_{n+1}|\mathscr{F}_n]$$

and since $Q_n \geqslant Y_n$ we have also that $Q_n \geqslant \sup(Y_n, \mathbb{E}[Z_{n+1}|\mathscr{F}_n]) = Z_n$. So the induction is complete and we proved that Q is above Z. Therefore Z is the smallest supermartingale above Y.

Now we want to prove that the process $(Z_n^{T^*})_{n\in[\![N]\!]}$ is a martingale. On the event $\{T^* > n\}$ we have that $Z_n = \mathbb{E}[Z_{n+1}|\mathscr{F}_n]$ from the definition of Z. Recall that $Z_n^{T^*} = Z_{n\wedge T^*}$, then

$$\mathbb{E}[Z_{n+1}^{T^*}|\mathscr{F}_n] = \mathbb{E}[Z_{n+1}^{T^*}\mathbb{1}_{n<T^*}|\mathscr{F}_n] + \mathbb{E}[Z_{n+1}^{T^*}\mathbb{1}_{n\geqslant T^*}|\mathscr{F}_n]$$

$$= \mathbb{E}[Z_{n+1}|\mathscr{F}_n]\mathbb{1}_{n<T^*} + \mathbb{E}[Z_{T^*}\mathbb{1}_{n\geqslant T^*}|\mathscr{F}_n]$$

$$= Z_n\mathbb{1}_{n<T^*} + Z_{T^*}\mathbb{1}_{n\geqslant T^*} = Z_n^{T^*}$$

which proves the martingale property.

Let now consider the two stopping times $n \wedge T^*$ and T^*, we have $n \wedge T^* \leqslant T^*$ and therefore by optional stopping of the martingale $M_\bullet^{T^*}$ we have

$$\mathbb{E}[Z_{T^*}|\mathscr{F}_{n\wedge T^*}] = \mathbb{E}[Z_{T^*}^{T^*}|\mathscr{F}_{n\wedge T^*}] = Z_{n\wedge T^*}^{T^*} = Z_{n\wedge T^*}$$

and taking expectation of this equation we have, for any $T \in \mathscr{T}_N$

$$\mathbb{E}[Y_T] \leqslant_{(1)} \mathbb{E}[Z_T] \leqslant_{(2)} \mathbb{E}[Z_1] = \mathbb{E}[Z_1^{T^*}] =_{(3)} \mathbb{E}[Z_N^{T^*}] = \mathbb{E}[Z_{T^*}] =_{(4)} \mathbb{E}[Y_{T^*}]$$

where (1) follows from the fact that Z bounds Y from above, (2) follow from the supermatingale property of Z and optional stopping, (3) follows from the martingale property of Z^{T^*} and finally (4) follows from the definition of T^*.

This chain on inequalities prove that T^* is an optimal stopping time and thereafore that

$$J_N = \mathbb{E}[Y_{T^*}] = \mathbb{E}[Z_1].$$

□

Corollary. *The stopping time T^* is the smallest optimal stopping time, i.e. if S is another optimal stopping time then $T^* \leqslant S$ almost surely.*

Proof. We prove it by contradiction, assume then that $\mathbb{P}(T^* > S) > 0$. For $\omega \in \Omega$ where $T^*(\omega) > S(\omega)$ we have $Y_S(\omega) < Z_S(\omega)$ by the definition of T^*, everywhere else we have in general only $Y_S(\omega) \leqslant Z_S(\omega)$. Given that $\{T^* > S\}$ has positive probability this implies that $\mathbb{E}[Y_S] < \mathbb{E}[Z_S]$. The supermatingale property of Z and optional stopping give then

$$\mathbb{E}[Y_S] < \mathbb{E}[Z_S] \leqslant \mathbb{E}[Z_1] = J_N$$

which is in contradition with the fact that S is optimal and therefore $J_N = \mathbb{E}[Y_S]$. We conclude that $\mathbb{P}(T^* > S) = 0$ which is what we wanted to prove. □
