m3f22l13tex

Lecture 13 4.11.2016

Optional Stopping Theorem (continued).

The OST is important in many areas, such as sequential analysis in statistics. We turn in the next section to related ideas specific to the gambling/financial context.

Write $X_n^T := X_{n \wedge T}$ for the sequence (X_n) *stopped* at time T.

Proposition. (i) If (X_n) is adapted and T is a stopping-time, the stopped sequence $(X_{n \wedge T})$ is adapted.
(ii) If (X_n) is a martingale [supermartingale] and T is a stopping time, (X_n^T) is a martingale [supermartingale].

Proof. If $\phi_j := I\{j \le T\}$,

$$X_{T \wedge n} = X_0 + \sum_1^n \phi_j (X_j - X_{j-1}).$$

Since $\{j \le T\}$ is the complement of $\{T < j\} = \{T \le j-1\} \in \mathcal{F}_{j-1}$, $\phi_j = I\{j \le T\} \in \mathcal{F}_{j-1}$, so (ϕ_n) is previsible. So (X_n^T) is adapted.

If (X_n) is a martingale, so is (X_n^T) as it is the martingale transform of (X_n) by (ϕ_n). Since by previsibility of (ϕ_n)

$$E[X_{T \wedge n} | \mathcal{F}_{n-1}] = X_0 + \sum_1^{n-1} \phi_j (X_j - X_{j-1}) + \phi_n (E[X_n | \mathcal{F}_{n-1}] - X_{n-1}), \qquad \text{i.e.}$$

$$E[X_{T \wedge n} | \mathcal{F}_{n-1}] - X_{T \wedge (n-1)} = \phi_n (E[X_n | \mathcal{F}_{n-1}] - X_{n-1}),$$

$\phi_n \ge 0$ shows that if (X_n) is a supermg [submg], so is $(X_{T \wedge n})$. //

§7. The Snell Envelope and Optimal Stopping.

Definition. If $Z = (Z_n)_{n=0}^N$ is a sequence adapted to a filtration $(\mathcal{F}_n)$, the sequence $U = (U_n)_{n=0}^N$ defined by

$$\begin{cases} U_N := Z_N, \\ U_n := \max(Z_n, E[U_{n+1} | \mathcal{F}_n]) \qquad (n \le N-1) \end{cases}$$

is called the *Snell envelope* of Z (J. L. Snell in 1952; [N] Ch. 6). U is adapted, i.e. $U_n \in \mathcal{F}_n$ for all n. For, Z is adapted, so $Z_n \in \mathcal{F}_n$. Also $E[U_{n+1} | \mathcal{F}_n] \in \mathcal{F}_n$

(definition of conditional expectation). Combining, $U_n \in \mathcal{F}_n$, as required.

The Snell envelope (see IV.8 L20) is exactly the tool needed in pricing American options. It is the *least supermg majorant* (also called the *réduite* or *reduced function* – crucial in the mathematics of gambling):

Theorem. The Snell envelope (U_n) of (Z_n) is a supermartingale, and is the smallest supermartingale dominating (Z_n) (that is, with $U_n \geq Z_n$ for all n).

Proof. First, $U_n \geq E[U_{n+1}|\mathcal{F}_n]$, so U is a supermartingale, and $U_n \geq Z_n$, so U dominates Z.

Next, let $T = (T_n)$ be any other supermartingale dominating Z; we must show T dominates U also. First, since $U_N = Z_N$ and T dominates Z, $T_N \geq U_N$. Assume inductively that $T_n \geq U_n$. Then

$$T_{n-1} \geq E[T_n|\mathcal{F}_{n-1}] \qquad \text{(as } T \text{ is a supermartingale)}$$

$$\geq E[U_n|\mathcal{F}_{n-1}] \qquad \text{(by the induction hypothesis)}$$

and

$$T_{n-1} \geq Z_{n-1} \qquad \text{(as } T \text{ dominates } Z\text{)}.$$

Combining,

$$T_{n-1} \geq \max(Z_{n-1}, E[U_n|\mathcal{F}_{n-1}]) = U_{n-1}.$$

By backward induction, $T_n \geq U_n$ for all n, as required. //

Note. It is no accident that we are using induction here *backwards in time*. We will use the same method – also known as *dynamic programming (DP)* – in Ch. IV below when we come to pricing American options.

Proposition. $T_0 := \min\{n \geq 0 : U_n = Z_n\}$ is a stopping time, and the stopped sequence $(U_n^{T_0})$ is a martingale.

We omit the proof (not hard, but fiddly – for details, see e.g. L13, 2014). Because U is a supermartingale, we knew that stopping it would give a supermartingale, by the Proposition of §6. The point is that, using the special properties of the Snell envelope, we actually get a *martingale*.

Write $\mathcal{T}_{n,N}$ for the set of stopping times taking values in $\{n, n+1, \cdots, N\}$ (a finite set, as Ω is finite). We next see that the Snell envelope solves the *optimal stopping problem*: it *maximises* the expectation of our final value of

Z – the value when we choose to quit – conditional on our present (publicly available) information. This is the best we can hope to do in practice (without cheating – insider trading, etc.)

Theorem. T_0 solves the optimal stopping problem for Z:

$$U_0 = E[Z_{T_0}|\mathcal{F}_0] = \max\{E[Z_T|\mathcal{F}_0] : T \in \mathcal{T}_{0,N}\}.$$

Proof. As $(U_n^{T_0})$ is a martingale (above),

$$\begin{aligned} U_0 &= U_0^{T_0} \qquad \text{(since } 0 = 0 \wedge T_0\text{)} \\ &= E[U_N^{T_0}|\mathcal{F}_0] \qquad \text{(by the martingale property)} \\ &= E[U_{T_0}|\mathcal{F}_0] \qquad \text{(since } T_0 = T_0 \wedge N\text{)} \\ &= E[Z_{T_0}|\mathcal{F}_0] \qquad \text{(since } U_{T_0} = Z_{T_0}\text{),} \end{aligned}$$

proving the first statement. Now for any stopping time $T \in \mathcal{T}_{0,N}$, since U is a supermartingale (above), so is the stopped process (U_n^T) (§6). So

$$\begin{aligned} U_0 &= U_0^T \qquad (0 = 0 \wedge T\text{, as above}) \\ &\geq E[U_N^T|\mathcal{F}_0] \qquad ((U_n^T) \text{ a supermartingale}) \\ &= E[U_T|\mathcal{F}_0] \qquad (T = T \wedge N) \\ &\geq E[Z_T|\mathcal{F}_0] \qquad ((U_n) \text{ dominates } (Z_n)), \end{aligned}$$

and this completes the proof. //

The same argument, starting at time n rather than time 0, gives an apparently more general version:

Theorem. If $T_n := \min\{j \geq n : U_j = Z_j\}$,

$$U_n = E[Z_{T_n}|\mathcal{F}_n] = \sup\{E[Z_T|\mathcal{F}_n] : T \in \mathcal{T}_{n,N}\}.$$

To recapitulate: as we are attempting to maximise our payoff by stopping $Z = (Z_n)$ at the most advantageous time, the Theorem shows that T_n gives the best stopping-time that is realistic: it maximises our *expected payoff* given only information *currently available* (it is easy, but irrelevant, to maximise things with hindsight!). We thus call T_0 (or T_n, starting from time n) the *optimal* stopping time for the problem.

§8. Doob Decomposition.
Theorem. Let $X = (X_n)$ be an adapted process with each $X_n \in L_1$. Then X has an (essentially unique) Doob decomposition

$$X = X_0 + M + A : \qquad X_n = X_0 + M_n + A_n \qquad \forall n \qquad (D)$$

with M a martingale null at zero, A a previsible process null at zero. If also X is a submartingale ('increasing on average'), A is increasing: $A_n \leq A_{n+1}$ for all n, a.s.

The proof in discrete time is quite easy (see L13, 2014). It is hard in continuous time – but more important there (see Ch. V: quadratic variation (QV) and the Itô integral). This illustrates the contrasts between the theories of stochastic processes in discrete and continuous time.

§9. Examples.
1. *Simple random walk.* Recall the simple random walk: $S_n := \sum_1^n X_k$, where the X_n are independent tosses of a fair coin, taking values ± 1 with equal probability 1/2. We decide to bet until our net gain is first $+1$, then quit – at time T, a stopping time. This has been analysed in detail; see e.g. [GS] GRIMMETT, G. R. & STIRZAKER, D.: *Probability and random processes*, OUP, 3rd ed., 2001 [2nd ed. 1992, 1st ed. 1982], §5.2:
(i) $T < \infty$ a.s.: the gambler will certainly achieve a net gain of $+1$ eventually;
(ii) $E[T] = +\infty$: the mean waiting-time till this happens is infinity. So:
(iii) No bound can be imposed on the gambler's maximum net loss before his net gain first becomes $+1$.

At first sight, this looks like a foolproof way to make money out of nothing: just bet till you get ahead (which happens eventually, by (i)), then quit. But as a gambling strategy, this is hopelessly impractical: because of (ii), you need unlimited time, and because of (iii), you need unlimited capital!

Notice that the Optional Stopping Theorem fails here: we start at zero, so $S_0 = 0$, $E[S_0] = 0$; but $S_T = 1$, so $E[S_T] = 1$. This shows two things:
(a) The Optional Stopping Theorem does indeed need conditions, as the conclusion may fail otherwise [none of the conditions (i) - (iii) in the OST are satisfied in the example above],
(b) Any practical gambling (or trading) strategy needs to have some integrability or boundedness restrictions to eliminate such theoretically possible but practically ridiculous cases.
2. *The doubling strategy.* Similarly for the doubling strategy (§3).