ma414l5.tex
Lecture 5. 9.2.2012

5. Martingales: discrete time. We refer for a fuller account to [W]. The classic exposition is Ch. VII in Doob's book [D] of 1953.

Definition. A process $X = (X_n)$ in discrete time is called a *martingale* (mg) relative to $(\{\mathcal{F}_n\}, P)$ if
(i) X is adapted (to $\{\mathcal{F}_n\}$);
(ii) $E|X_n| < \infty$ for all n;
(iii) $[X_n|\mathcal{F}_{n-1}] = X_{n-1}$ P-a.s.
X is a *supermartingale* (supermg) if in place of (iii)

$$E[X_n|\mathcal{F}_{n-1}] \leq X_{n-1} \qquad P - a.s. \qquad (n \geq 1);$$

X is a *submartingale* (submg) if in place of (iii)

$$E[X_n|\mathcal{F}_{n-1}] \geq X_{n-1} \qquad P - a.s. \qquad (n \geq 1).$$

Martingales have a useful interpretation in terms of dynamic games: a mg is 'constant on average', and models a fair game; a supermg is 'decreasing on average', and models an unfavourable game; a submg is 'increasing on average', and models a favourable game.
Note. 1. Martingales have many connections with harmonic functions in probabilistic potential theory. The terminology in the inequalities above comes from this: supermartingales correspond to superharmonic functions, submartingales to subharmonic functions.
2. X is a submg (supermg) iff $-X$ is a supermg (submg); X is a mg if and only if it is both a submg and a supermg.
3. (X_n) is a mg iff $(X_n - X_0)$ is a mg. So w.l.o.g. take $X_0 = 0$ if convenient.
4. If X is a martingale, then for $m < n$ using the iterated conditional expectation and the martingale property repeatedly (all equalities are in the a.s.-sense)

$$E[X_n|\mathcal{F}_m] = E[E(X_n|\mathcal{F}_{n-1})|\mathcal{F}_m] = E[X_{n-1}|\mathcal{F}_m] = \ldots = E[X_m|\mathcal{F}_m] = X_m,$$

and similarly for submgs, supermgs.

The word 'martingale' is taken from an article of harness, to control a horse's head. The word also means a system of gambling which consists in

doubling the stake when losing in order to recoup oneself (1815).
Thackeray: 'You have not played as yet? Do not do so; above all avoid a martingale if you do.'
Examples.
1. Mean zero random walk: $S_n = \sum X_i$, with X_i independent with $E(X_i) = 0$ is a mg (submg: positive mean; supermg: negative mean).
2. Stock prices: $S_n = S_0\zeta_1 \cdots \zeta_n$ with ζ_i independent positive r.vs with finite first moment.
3. Accumulating data about a random variable ([W], pp. 96, 166–167). If $\xi \in L_1(\Omega, \mathcal{F}, \mathcal{P})$, $M_n := E(\xi|\mathcal{F}_n)$ (so M_n represents our best estimate of ξ based on knowledge at time n), then using iterated conditional expectations

$$E[M_n|\mathcal{F}_{n-1}] = E[E(\xi|\mathcal{F}_n)|\mathcal{F}_{n-1}] = E[\xi|\mathcal{F}_{n-1}] = M_{n-1},$$

so (M_n) is a martingale – indeed, a 'nice' mg; see below.
Stopping Times and Optional Stopping

Recall that a random variable τ taking values in $\{0, 1, 2, \ldots; +\infty\}$ is called a *stopping time* if

$$\{\tau \le n\} = \{\omega : \tau(\omega) \le n\} \in \mathcal{F}_n \qquad \forall\, n \le \infty.$$

From $\{\tau = n\} = \{\tau \le n\} \setminus \{\tau \le n-1\}$ and $\{\tau \le n\} = \bigcup_{k \le n}\{\tau = k\}$, we see the equivalent characterization

$$\{\tau = n\} \in \mathcal{F}_n \qquad \forall\, n \le \infty.$$

Call a stopping time τ *bounded* if there is a constant K such that $P(\tau \le K) = 1$. (Since $\tau(\omega) \le K$ for some constant K and all $\omega \in \Omega \setminus N$ with $P(N) = 0$ all identities hold true except on a null set, i.e. a.s.)
Example.

Suppose (X_n) is an adapted process and we are interested in the time of first entry of X into a Borel set B (e.g. $B = [c, \infty)$):

$$\tau = \inf\{n \ge 0 : X_n \in B\}.$$

Now $\{\tau \le n\} = \bigcup_{k \le n}\{X_k \in B\} \in \mathcal{F}_n$ and $\tau = \infty$ if X never enters B. Thus τ is a stopping time. Intuitively, think of τ as a time at which you decide to quit a gambling game: whether or not you quit at time n depends only on the history up to and including time n – NOT the future. Thus stopping times model gambling and other situations where there is no foreknowledge, or prescience of the future; in particular, in the financial context, where there is no insider trading. Furthermore since a gambler cannot cheat the system the expectation of his hypothetical fortune (playing with unit stake) should equal his initial fortune.

Theorem (Doob's Stopping-time Principle). Let τ be a bounded stopping time and $X = (X_n)$ a martingale. Then X_τ is integrable, and

$$E(X_\tau) = E(X_0).$$

Proof. Assume $\tau(\omega) \leq K$ for all ω (K integer), and write

$$X_{\tau(\omega)}(\omega) = \sum_{k=0}^{\infty} X_k(\omega) I(\tau(\omega) = k) = \sum_{k=0}^{K} X_k(\omega) I(\tau(\omega) = k).$$

Then

$$\begin{aligned}
E(X_\tau) &= E[\sum_{k=0}^{K} X_k I(\tau = k)] && \text{(by the decomposition above)} \\
&= \sum_{k=0}^{K} E[X_k I(\tau = k)] && \text{(linearity of } E) \\
&= \textstyle\sum_{k=0}^{K} E[E(X_K|\mathcal{F}_k) I(\tau = k)] && (X \text{ a mg, } \{\tau = k\} \in \mathcal{F}_k\ ) \\
&= \sum_{k=0}^{K} E[X_K I(\tau = k)] && \text{(defn. of conditional expectation)} \\
&= E[X_K \sum_{k=0}^{K} I(\tau = k)] && \text{(linearity of } E) \\
&= E[X_K] && \text{(the indicators sum to 1)} \\
&= E[X_0] && (X \text{ a mg}) \qquad //.
\end{aligned}$$

The stopping time principle holds also true if $X = (X_n)$ is a supermg; then the conclusion is

$$EX_\tau \leq EX_0.$$

Also, alternative conditions such as
(i) $X = (X_n)$ is bounded ($|X_n|(\omega) \leq L$ for some L and all n, ω);
(ii) $E\tau < \infty$ and $(X_n - X_{n-1})$ is bounded; suffice for the proof of the stopping time principle.

The stopping time principle is important in many areas, such as sequential analysis in statistics.

We now wish to create the concept of the σ-algebra of events observable up to a stopping time τ, in analogy to the σ-algebra $\mathcal{F}_n$ which represents the events observable up to time n.

Definition. Let τ be a stopping time. The *stopping time $\sigma-$algebra $\mathcal{F}_\tau$* is defined to be

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \le n\} \in \mathcal{F}_n, \text{ for all } n\}.$$

Proposition. For τ a stopping time, $\mathcal{F}_\tau$ is a $\sigma-$algebra.

Proof. We simply have to check the defining properties. Clearly $\Omega, \emptyset$ are in $\mathcal{F}_\tau$. Also for $A \in \mathcal{F}_\tau$ we find

$$A^c \cap \{\tau \le n\} = \{\tau \le n\} \setminus (A \cap \{\tau \le n\}) \in \mathcal{F}_n,$$

thus $A^c \in \mathcal{F}_\tau$. Finally, for a family $A_i \in \mathcal{F}_\tau,\ i = 1, 2, \dots$ we have

$$\left(\bigcup_{i=1}^{\infty} A_i\right) \cap \{\tau \le n\} = \bigcup_{i=1}^{\infty} (A_i \cap \{\tau \le n\}) \in \mathcal{F}_n,$$

showing $\bigcup_{i=1}^{\infty} A_i \in \mathcal{F}_\tau$. //

One can show similarly that for σ, τ stopping times with $\sigma \le \tau$, $\mathcal{F}_\sigma \subseteq \mathcal{F}_\tau$. Similarly, for any adapted sequence of random variables $X = (X_n)$ and a.s. finite stopping time τ, define

$$X_\tau := \sum_{n=0}^{\infty} X_n I(\tau = n).$$

Then X_τ is $\mathcal{F}_\tau$-measurable.

We are now in position to obtain an important extension of the Stopping-Time Principle.

Theorem (Doob's Optional-Sampling Theorem, OST. Let $X = (X_n)$ be a mg and σ, τ be bounded stopping times with $\sigma \le \tau$. Then

$$E\left[X_\tau | \mathcal{F}_\sigma\right] = X_\sigma, \qquad \text{and so} \qquad E(X_\tau) = E(X_\sigma).$$

Proof. First observe that X_τ and X_σ are integrable (use the sum representation and the fact that τ is bounded by an integer K) and X_σ is $\mathcal{F}_\sigma$-measurable by above. So it only remains to prove that

$$E(I_A X_\tau) = E(I_A X_\sigma) \qquad \forall A \in \mathcal{F}_\sigma.$$

For any such fixed $A \in \mathcal{F}_\sigma$, define ρ by

$$\rho(\omega) = \sigma(\omega) I_A(\omega) + \tau(\omega) I_{A^c}(\omega).$$

Since

$$\{\rho \leq n\} = (A \cap \{\sigma \leq n\}) \cup (A^c \cap \{\tau \leq n\}) \in \mathcal{F}_n$$

ρ is a stopping time, and from $\rho \leq \tau$ we see that ρ is bounded. So the STP implies $E(X_\rho) = E(X_0) = E(X_\tau)$. But

$$E(X_\rho) = E\left(X_\sigma I_A + X_\tau I_{A^c}\right), \qquad E(X_\tau) = E\left(X_\tau I_A + X_\tau I_{A^c}\right).$$

So subtracting yields the result. //

Write $X^\tau = (X_n^\tau)$ for the sequence $X = (X_n)$ *stopped* at time τ, where we define $X_n^\tau(\omega) := X_{\tau(\omega)\wedge n}(\omega)$. One can show
(i) If τ is a stopping time and X is adapted, then so is X^τ.
(ii) If τ is a stopping time and X is a mg (supermg, submg), then so is X^τ.
Examples and Applications.
1. Simple Random Walk. Recall the simple random walk: $S_n := \sum_{k=1}^n X_k$, where the X_n are independent tosses of a fair coin, taking values ± 1 with equal probability 1/2. Suppose we decide to bet until our net gain is first $+1$, then quit. Let τ be the time we quit; τ is a stopping time. The stopping time τ has been analyzed in detail (see e.g. [GS], 5.3, or Ex. 3.4). From this:
(i) $\tau < \infty$ a.s.: the gambler will certainly achieve a net gain of $+1$ eventually;
(ii) $E\tau = +\infty$: the mean waiting-time for this is infinity. Hence also:
(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 until you get ahead (which happens eventually, by (i)), then quit. However, as a gambling strategy, this is hopelessly impractical: because of (ii), you need unlimited time, and because of (iii), you need unlimited capital – neither of which is realistic.

Notice that the Stopping-time Principle fails here: we start at zero, so

$S_0 = 0$, $ES_0 = 0$; but $S_\tau = 1$, so $ES_\tau = 1$. This example shows two things:
1. Conditions are indeed needed here, or the conclusion may fail (none of the conditions in STP or the alternatives given are satisfied in this example).
2. Any practical gambling (or trading) strategy needs to have some integrability or boundedness restrictions to eliminate such theoretically possible but practically ridiculous cases.

The Doubling Strategy.

The strategy of doubling when losing – *the martingale*, according to the *Oxford English Dictionary* (S3.3) – has similar properties. We play until the time τ of our first win. Then τ is a stopping time, and is geometrically distributed with parameter $p = 1/2$. If $\tau = n$, our winnings on the nth play are 2^{n-1} (our previous stake of 1 doubled on each of the previous $n - 1$ losses). Our cumulative losses to date are $1 + 2 + \ldots + 2^{n-2} = 2^{n-1} - 1$ (summing the geometric series), giving us a net gain of 1. The mean time of play is $E(\tau) = 2$ (so doubling strategies accelerate our eventually certain win to give a finite expected waiting time for it). But no bound can be put on the losses one may need to sustain before we win, so again we would need unlimited capital to implement this strategy – which would thus be suicidal in practice.

Theorem (Doob Decomposition). Let $X = (X_n)$ be an adapted process with each $X_n \in \mathcal{L}_\infty$. Then X has an (essentially unique) Doob decomposition

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

with M a martingale null at zero, A a predictable process null at zero. If also X is a submartingale, A is increasing: $A_n \le A_{n+1}$ for all n, a.s.

Proof. If X has a Doob decomposition as above,

$$E[X_n - X_{n-1}|\mathcal{F}_{n-1}] = E[M_n - M_{n-1}|\mathcal{F}_{n-1}] + E[A_n - A_{n-1}|\mathcal{F}_{n-1}].$$

The first term on the right is zero, as M is a martingale. The second is $A_n - A_{n-1}$, since A_n (and A_{n-1}) is $\mathcal{F}_{n-1}$-measurable by predictability. So

$$E[X_n - X_{n-1}|\mathcal{F}_{n-1}] = A_n - A_{n-1},$$

and summation gives

$$A_n = \sum_{k=1}^{n} E[X_k - X_{k-1}|\mathcal{F}_{k-1}], \qquad a.s.$$

So set $A_0 = 0$ and use this formula to *define* (A_n), clearly predictable. We then use the equation in the Theorem to *define* (M_n), then a martingale, giving the Doob decomposition. To see uniqueness, assume two decompositions, i.e. $X_n = X_0 + M_n + A_n = X_0 + \tilde{M}_n + \tilde{A}_n$, then $M_n - \tilde{M}_n = A_n - \tilde{A}_n$. Thus the martingale $M_n - \tilde{M}_n$ is predictable and so must be constant a.s.

If X is a submg, the LHS of the Doob decomposition is ≥ 0, so the RHS is ≥ 0, i.e. (A_n) is increasing. //

Martingale transforms (Burkholder).

If $X = (X_n)$ is a mg [submg, supermg], $C = (C_n)$ is predictable, write

$$(C \bullet X)_n := \sum_1^n C_k(X_k - X_{k-1})$$

($C \bullet N$ is the *martingale* [submg, supermg] *transform* of X by C). Then
(i) if C is bounded and non-negative and X is a submg [supermg], $C \bullet X$ is a submg [supermg] null at 0;
(ii) if C is bounded and X is a mg, $C \bullet X$ is a mg null at 0.
Proof. As C is bounded and X is integrable, $C \bullet X$ is integrable; it is null at 0 (empty sum is 0). As C is predictable, C_n is $\mathcal{F}_{n-1}$-measurable, so

$$E[(C \bullet X)_n - (C \bullet X)_{n-1}|\mathcal{F}_{n-1}] = E[C_n(X_n - X_{n-1}|\mathcal{F}_{n-1}] = C_n E[X_n - X_{n-1}|\mathcal{F}_{n-1}],$$

taking out what is known. This is ≥ 0 in case (i) with $C \geq 0$ and X a submg, and 0 in case (ii) with X a mg. //
Upcrossings.

For a process X and interval $[a, b]$, define stopping times σ_k, τ_k by $\sigma_1 := \min\{n : X_n \leq a\}$, $\tau_1 := \min\{n > \sigma_1 : X_n \geq b\}$, and inductively $\sigma_k := \min\{n > \tau_{k-1} : X_n \leq a\}$, $\tau_k := \min\{n > \sigma_k : X_n \geq b\}$. Call $[\sigma_k, \tau_k]$ an *upcrossing* of $[a, b]$ by X, and write $U_n := U_n([a, b], X)$ for the number of such upcrossings by time n.

Upcrossing Inequality (Doob). If X is a submg,

$$EU_n([a, b], X) \leq E[(X_n - a)^+]/(b - a).$$

Proof. As $(X - a)^+$ is a submg by Q2 (iii) and upcrossings of $[a, b]$ by X correspond to upcrossings of $[0, b - a]$ by $(X - a)^+$, we may (by passing to $(X - a)^+$) take $X \geq 0$, $a = 0$. Write

$$V_n := \sum_{k \geq 1} I(\sigma_k < n \leq \tau_k).$$

Then V is predictable (this comes from the "$<$" above – we know at time $n-1$ whether the kth upcrossing has begun). So $1-V$ is predictable. So by above the transform $(1-V)\bullet X$ is a submg. So

$$E[(1-V)\bullet X)_n] \geq E[(1-V)\bullet X)_0] = 0 : \qquad E[(V\bullet X)_n] \leq E[X_n].$$

Each completed upcrossing contributes at least b to the sum in $(V\bullet X)_n = \sum_1^n V_k(X_k - X_{k-1})$, and the contribution of the last (possibly uncompleted) upcrossing is ≥ 0, so

$$(V\bullet X)_n \geq bU_n.$$

Combining, $bU_n \leq E[(V\bullet X)_n] \leq E[X_n]$: $EU_n \leq E[X_n]/b$. Reverting to the original notation gives the result. //

(Sub-)Martingale Convergence Theorem (Doob). An L_1-bounded submg $X=(X_n)$ (i.e. $E|X_n| \leq K$ for some K and all n) is a.s. convergent.

Proof. For $a<b$ rational, the expected number EU_n of upcrossings of $[a,b]$ up to time n is $\leq (K+|a|)/(b-a) < \infty$, for each n. As U_n increases in n, monotone convergence gives $E[\sup U_n] < \infty$. So $U := \sup U_n < \infty$ a.s. If $X_* := \text{liminf} X_n$, $X^* := \text{limsup} X_n$, $\{X_* < X^*\} = \cup_{a,b}\{X_* < a < b < X^*\}$ ($a<b$ rational). Each such set is null (or U would be infinite). So their union is null, i.e. $X_* = X^*$ a.s.: X is a.s. convergent (its limit X_∞ may be infinite). But $E|X_\infty| = E[\lim(\inf)|X_n|] \leq \text{liminf} E[|X_n|]$ (by Fatou), $\leq K < \infty$. So $|X_\infty| < \infty$ a.s., and $X_n \to X_\infty$ finite, a.s. //

Corollary (Doob). A non-negative supermg X_n is a.s. convergent.

Proof. As X_n is a supermg, EX_n decreases. As $X \geq 0$, $E[X_n] \geq 0$. So $E[|X_n|] = E[X_n]$ is decreasing and bounded below, so (convergent and) bounded: X_n is L_1-bounded. So the submg $-X_n$ is L_1-bounded, so a.s. convergent by Doob's Theorem, so X_n is a.s. convergent.