spl17.tex
Lecture 17. 14.11.2010

III. Stochastic processes; Martingales; Brownian motion

1. Filtrations; Finite-dimensional Distributions

We take a stochastic basis (II.16) (Ω, $\{\mathcal{F}_t,\}$, $\mathcal{F}$, P) (or *filtered probability space*), which following Meyer we assume satisfies the *usual conditions* (conditions habituelles):
a. completeness: each $\mathcal{F}_t$ contains all P-null sets of $\mathcal{F}$;
b. the filtration is right-continuous, i.e. $\mathcal{F}_t = \mathcal{F}_{t+} := \cap_{s>t}\mathcal{F}_s$.
A *stochastic process* $X = (X(t))_{t\geq 0}$ is a family of random variables defined on a stochastic basis. Call X *adapted* if $X(t) \in \mathcal{F}_t$ (i.e. $X(t)$ is $\mathcal{F}_t$-measurable) for each t: thus $X(t)$ is known when $\mathcal{F}_t$ is known, at time t.

If $\{t_1, \cdots, t_n\}$ is a finite set of time points in $[0,\infty)$, $(X(t_1), \ldots, X(t_n))$ is a random n-vector, with a distribution, $\mu(t_1, \ldots, t_n)$ say. The class of all such distributions as $\{t_1, \ldots, t_n\}$ ranges over all finite subsets of $[0,\infty)$ is called the class of all *finite-dimensional distributions* of X. These satisfy certain obvious consistency conditions:
DK1. deletion of one point t_i can be obtained by 'integrating out the unwanted variable', as usual when passing from joint to marginal distributions;
DK2. permutation of the times t_i permutes the arguments of the measure $\mu(t_1, \ldots, t_n)$ on $\mathbf{R}^n$ in the same way.
Conversely, a collection of finite-dimensional distributions satisfying these two consistency conditions arises from a stochastic process in this way (this is the content of the *Daniell-Kolmogorov theorem*). This classical result (due to P.J. Daniell in 1918 and A.N. Kolmogorov in 1933) is the basic existence theorem for stochastic processes. For the proof, see e.g. [K].

Important though it is as a general existence result, however, the Daniell-Kolmogorov theorem does not take us very far. It gives a stochastic process X as a random function on $[0,\infty)$, i.e. a random variable on $\mathbf{R}^{[0,\infty)}$. This is a vast and unwieldy space; we shall usually be able to confine attention to much smaller and more manageable spaces, of functions satisfying regularity conditions. The most important of these is continuity: we want to be able to realize $X = (X(t,\omega))_{t\geq 0}$ as a random continuous function, i.e. a member of $C[0,\infty)$; such a process X is called path-continuous (since the map $t \to X(t,\omega)$ is called the sample path, or simply path, given by ω) – or more briefly, continuous. This is possible for the extremely important case of Brownian motion, for example, and its relatives. Sometimes we need to allow our random function $X(t, \omega)$ to have jumps. It is then customary, and convenient, to require $X(t)$ to be right-continuous with left limits (RCLL), or càdlàg (*continu à droite, limite à gauche*) – i.e. to have X in the space $D[0, \infty)$ of all such functions (the Skorohod space). This is the case, for instance, for the Poisson process and its relatives (see below).

General results on realisability – whether or not it is possible to realize, or obtain, a process so as to have its paths in a particular function space – are known; see for example the Kolmogorov-Ĉentsov theorem. For our purposes, however, it is usually better to construct the processes we need directly on the function space on which they naturally live.

Given a stochastic process X, it is sometimes possible to improve the regularity of its paths without changing its distribution (that is, without changing its finite-dimensional distributions). For background on such results (separability, measurability, versions, regularization etc.) see e.g. [D].

There are several ways to define 'sameness' of two processes X and Y. We say
(i) X and Y have the *same finite-dimensional distributions* if, for any integer n and $\{t_1, \cdots, t_n\}$ a finite set of time points in $[0, \infty)$, the random vectors $(X(t_1), \ldots, X(t_n))$ and $(Y(t_1), \ldots, Y(t_n))$ have the same distribution;
(ii) Y is a *modification* of X if, for every $t \geq 0$, we have $P(X_t = Y_t) = 1$;
(iii) X and Y are *indistinguishable* if almost all their sample paths agree:

$$P\left[X_t = Y_t; \forall 0 \leq t < \infty\right] = 1.$$

Indistinguishable processes are modifications of each other; the converse is not true in general, but is true for processes with right-continuous paths. This will cover the processes we encounter in this course.

A process is called *progressively measurable* if the map $(t, \omega) \mapsto X_t(\omega)$ is measurable, for each $t \geq 0$. Progressive measurability holds for adapted processes with right-continuous (or left-continuous) paths – and so always in the generality in which we work.

A random variable $\tau : \Omega \to [0, \infty]$ is a *stopping time* if $\{\tau \leq t\} \in \mathcal{F}_t$ for all $t \geq 0$. If $\{\tau < t\} \in \mathcal{F}_t$ for all t, τ is called an *optional time*. For right-continuous filtrations (as here, under the usual conditions) the concepts of stopping and optional times are equivalent.

For a set $A \subset \mathbf{R}^d$ and a stochastic process X, we can define the *hitting time* of A for X as

$$\tau_A := \inf\{t > 0 : X_t \in A\}.$$

For our usual situation (RCLL processes and Borel sets) hitting times are stopping times.

We will also need the *stopping time σ-algebra* $\mathcal{F}_\tau$ defined as

$$\mathcal{F}_\tau = \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t.$$

Intuitively, $\mathcal{F}_\tau$ represents the events known at time τ.

The continuous-time theory is technically much harder than the discrete-time theory, for two reasons:
1. questions of path-regularity arise in continuous time but not in discrete time;
2. uncountable operations (such as taking the supremum over an interval) arise in continuous time. But measure theory is constructed using countable operations: uncountable operations risk losing measurability.
This is why discrete and continuous time are often treated separately.

2. 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. Mgs have many connections with harmonic functions in probabilistic potential theory. Supermgs correspond to superharmonic functions, submgs 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 τ taking values in $\{0, 1, 2, \ldots; +\infty\}$ is a *stopping time* if

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

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

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

Call a stopping time τ *bounded* if there is a constant K such that $P(\tau \leq K) = 1$. (Since $\tau(\omega) \leq K$ 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. $B = [c, \infty)$):

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

Now $\{\tau \leq n\} = \bigcup_{k \leq n} \{X_k \in B\} \in \mathcal{F}_n$ 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

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

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

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

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

Proof. Clearly $\Omega, \emptyset$ are in $\mathcal{F}_\tau$. Also for $A \in \mathcal{F}_\tau$ we find

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

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

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

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

One can show similarly that for σ, τ stopping times with $\sigma \leq \tau$, $\mathcal{F}_\sigma \subseteq \mathcal{F}_\tau$. 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 now give 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 $\sigma \leq \tau$. 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 $X^\tau = (X^\tau_n)$ for the sequence $X = (X_n)$ *stopped* at time τ, where we define $X^\tau_n(\omega) := X_{\tau(\omega)\wedge n}(\omega)$. 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.

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: $A_n \leq A_{n+1}$ 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. //