3. Martingales I.

Let us start with a sequence $\{X_i\}$ of independent random variables with $E[X_i] = 0$ and $E[X_i^2] = 1$. We saw earlier that for a sequence $\{a_j\}$ of constants

$$S = \sum_{i=1}^{\infty} a_i X_i$$

will converge with probability 1 and in mean square provided

$$\sum_j a_j^2 < \infty$$

We shall now see that actually $a_j(X_1, X_2, \ldots, X_{j-1})$ can be a function of $X_1, \ldots, X_{j-1}$. If they satisfy

$$\sum_j E[a_j(X_1, \ldots, X_{j-1})^2] < \infty$$

then the series

$$S = \sum_{j=1}^{\infty} a_j(X_1, X_2, \ldots, X_{j-1}) X_j$$

will converge. Let us show for the present convergence of

$$S_n = \sum_{j=1}^{n} a_j(X_1, X_2, \ldots, X_{j-1}) X_j$$

to S in $L_2(P)$. This requires us to show that

$$\lim_{\substack{n\to\infty \\ m\to\infty}} E[|S_n - S_m|^2] = 0.$$

A straight forward computation of

$$E\big[[\sum_{j=m+1}^{n} a_j(X_1, \ldots, X_{j-1})\, X_j]^2\big]$$

shows that the non-diagonal terms are 0. If $i \neq j$, either X_i or X_j sticks out and makes the expectation 0. On the other hand if $i = j$

$$E[a_i^2(X_1, \ldots, X_{k-1}) X_i^2] = E[a_i(X_1, \ldots, X_{i-1})^2]$$

resulting for $n > m$,

$$E[|S_n - S_m|^2] = \sum_{j=m+1}^{n} E[a_j(X_1, \ldots, X_{i-1})^2]$$

proving the convergence of S_n in $L_2(P)$.

Actually one does not need even independence of $\{X_j\}$. If

$$E[X_j|X_1,\ldots,X_{j-1}]=0$$

and

$$E[X_j^2|X_1,\ldots,X_{j-1}]=\sigma_j^2(X_1,\ldots,X_{k-1})$$

then $E[S_n]=0$ and

$$E[|S_n-S_m|^2]=\sum_{j=m+1}^{n}E[a_j^2(X_1,\ldots,X_{j-1})^2\,\sigma_j^2(X_1,\ldots,X_{j-1})]$$

and the convergence of

$$\sum_{j=1}^{\infty}E[a_j^2(X_1,\ldots,X_{j-1})^2\,\sigma_j^2(X_1,\ldots,X_{j-1})]$$

implies the existence of the limit S of S_n in $L_2(P)$.

All of this leads to the following definition. We have a probability space $(\Omega,\mathcal{F},P)$ and a family of sub σ-fields $\mathcal{F}_i$ with $\mathcal{F}_{i-1}\subset\mathcal{F}_i$ for every i. A sequence $X_i(\omega)$ of random variables is called a (square integrable) **martingale** with respect to $(\Omega,\{\mathcal{F}_i\},P)$ if

1) X_i is $\mathcal{F}_i$ measurable for every $i\ge 0$.

2) For every $i\ge 1$, $E[X_i-X_{i-1}|\mathcal{F}_{i-1}]=0$

3) $E[X_i^2]<\infty$ for $i\ge 0$.

If we dente by $\sigma_i^2(\omega)=E[(X_i-X_{i-1})^2|\mathcal{F}_{i-1}]$ then $E[(X_i-X_{i-1})^2]=E[\sigma_i^2(\omega)]<\infty$. If X_i is a (square integrable) martingale then $\xi_i=X_i-X_{i-1}$ is called a (square integrable) martingale difference. It has the properties

1) ξ_i is $\mathcal{F}_i$ measurable for every $i\ge 1$.

2) For every $i\ge 1$, $E[\xi_i]|\mathcal{F}_{i-1}]=0$

3) $E[\xi_i^2]<\infty$ for $i\ge 1$.

X_0 is often 0 or a constant. It does not have to be. It is $\mathcal{F}_0$ measurable and for $n\ge 1$,

$$X_n=X_0+\sum_{i=1}^{n}\xi_i$$

It is now an easy calculation to conclude that

$$E[X_n^2]=E[X_0]^2+\sum_{j=1}^{n}E[\xi_j^2]$$

If we dente by $\sigma_i^2(\omega) = E[\xi_i^2|\mathcal{F}_{i-1}]$ then $E[\xi_i^2] = E[\sigma_i^2(\omega)] < \infty$. If we take a sequence of bounded functions $a_i(\omega)$ that are $\mathcal{F}_{i-1}$ measurable, then we can define

$$Y_n = \sum_{i=1}^{n} a_i(\omega)\xi_i$$

It is easy to see that Y_n is again a martingale and $E[(Y_i - Y_{i-1})^2|\mathcal{F}_{i-1}] = a_i^2(\omega)\sigma_i^2(\omega)$. One just needs to note that if $\{\xi_i\}$ is a sequence of martingale differences and a_i is $\mathcal{F}_{i-1}$ measurable, then $\eta_i = a_i\xi_i$ is again a martingale difference and $E[\eta_i^2|\mathcal{F}_{i-1}] = a_i^2(\omega)\sigma_i^2(\omega)$. $\{Y_n\}$ is called a martingale transform of X_n and one writes

$$(Y_{n+1} - Y_n) = a_n(X_{n+1} - X_n)$$

or

$$\nabla Y_n = a_n \nabla X_n$$

Notice that in the definition of the increment ∇X_n and ∇Y_n the increments stick out. This is important because otherwise the cross terms in the calculation of the expectation of the square of the sum will not equal 0. Martingale transforms generates new martingale from a given martingale. There are obvious identities.

$$\nabla Y_n = a_n \nabla X_n,\ \nabla Z_n = b_n \nabla Y_n \Rightarrow \nabla Z_n = a_n b_n \nabla X_n$$

leading to the inversion rule

$$\nabla Y_n = a_n \nabla X_n \Rightarrow \nabla X_n = b_n \nabla Y_n$$

with $b_n = [a_n]^{-1}$.

One should think of martingale differences ξ_i as possible returns on a standard bet of one unit at the i-th round of a game. There is some randomness. The game is "fair" if the expected return is 0 no matter what the past history through $i-1$ rounds. $\mathcal{F}_{i-i}$ represents historical information through $i-1$ rounds is. One should think of a_i as the leverage, with negative values representing short positions or betting in the opposite direction. The dependence on ω is the strategy that can only be based on information available through the previous round. No matter what the strategy is for the leverage, the game is always "fair". You can not find a winning (or losing) strategy in a "fair" game.

Stopping Times. A special set of choices for $a_i(\omega)$ is a decision about when to stop. $\tau(\omega)$ is a non-negative integer valued function on Ω and $a_i(\omega)$ is defined by

$$a_i(\omega) = \begin{cases} 1 \text{ for } i < \tau(\omega) \\ 0 \text{ for } i \geq \tau(\omega) \end{cases}$$

The condition that $a_i(\omega)$ is $\mathcal{F}_{i-1}$ measurable leads to, for $i \geq 1$,

$$\{\omega : \tau(\omega) \leq i\} \in \mathcal{F}_i$$

You can not decide to "stop" after looking into the "future". In particular quitting while ahead can not be a winning strategy.

A paradox. What if in a game of even odds you double your bet every time you lose. If there is any chance of winning at all, sooner or later you will win a round. You will exit with exactly \$1 in winnings. This is indeed a stopping time. But this requires ability to play as many games as it takes. Even in a fair game you can have a long run of bad luck and you can not really afford it. You wont live that long or before that you will exceed you credit limit. If you put a limit on the number of games to be played, say N, then with probability 2^{-N} you will have a big loss of $\$(2^N - 1)$ and with probability $1 - 2^{-N}$ a gain of \$1. Technically, in order to show that the game remains "fair" , i.e $E[X_\tau] = E[X_0]$, one has to assume that stopping times τ is bounded i.e. $\tau(\omega) \leq T$ for some non-random T. There is maximum number of rounds before which one has to stop.

With this provision a proof can be made. Let $a_i(\omega) = 1$ if $\tau(\omega) \geq i$ and 0 otherwise. The martingale transform

$$\nabla Y_n = a_n \nabla X_n$$

yields

$$Y_n = X_\tau, \quad \text{if} \quad n \geq \tau$$

In particular $E[X_\tau - X_0] = E[Y_T - Y_0] = 0$. Actually

$$Y_n = X_{\tau \wedge n}$$

and we have proved

Theorem. Let X_n be a martingale and $\tau(\omega)$ a stopping time. Then $Y_n = X_{\tau \wedge n}$ is again a martingale.

Remark. If τ is unbounded, then to show that $E[X_\tau - X_0] = 0$ one has to let $n \to \infty$ in $E[X_{\tau \wedge n} - X_0] = 0$. This requires uniform integrability assumptions on X_n. In particular if $\sup_{0 \leq n \leq \tau} |X_n| \leq C$, for some constant C, there is no difficulty.

There is a natural σ-field $\mathcal{F}_\tau$ associated with a stopping time. Intuitively this is all the information one has gathered up to the stopping time. For example if a martingale X_n is stopped when it reaches a level a or above , i.e $\tau = \{\inf n : X_n \geq a\}$, then the set

$$A = \{\omega : X_n \text{ goes below level } b \text{ before going above level } \ a\}$$

will be in $\mathcal{F}_\tau$ but not $\{\omega : X_n \geq a\}$ unless we know for sure that $\tau \geq a$. Technically

Definition. A set $A \in \mathcal{F}_\tau$ if for every k

$$A \cap \{\tau \leq k\} \in \mathcal{F}_k$$

Equivalently one can demand that for every k

$$A \cap \{\tau = k\} \in \mathcal{F}_k$$

The following facts are easy to verify. $\mathcal{F}_\tau$ is a σ-field for any stopping time τ. If $\tau = \ell$ is a constant then $\mathcal{F}_\tau = \mathcal{F}_\ell$. If $\tau_1 \le \tau_2$ then $\mathcal{F}_{\tau_1} \subset \mathcal{F}_{\tau_2}$. If τ_1, τ_2 are stopping times so are $\tau_1 \wedge \tau_2, \tau_1 \vee \tau_2$. If τ is a stopping time and $f(t)$ is an increasing function of t satisfying $f(t) \ge t$, then $f(\tau)$ is a stopping time.

Theorem (Doob's stopping theorem). If $\tau_1 \le \tau_2 \le T$ are two bounded stopping times then

$$E[X_{\tau_2}|\mathcal{F}_{\tau_1}] = X_{\tau_1}$$

It is enough to prove that for any stopping time bounded by T

$$E[X_T|\mathcal{F}_\tau] = X_\tau$$

Let $A \in \mathcal{F}_\tau$. We need to show

$$\int_A X_T dP = \int_A X_\tau dP$$

It suffices to show that for each k

$$\int_{A\cap\{\tau=k\}} X_T dP = \int_{A\cap\{\tau=k\}} X_\tau dP = \int_{A\cap\{\tau=k\}} X_k dP$$

we can then sum over k. But $A \cap \{\tau = k\} \in \mathcal{F}_k$ and for $B \in \mathcal{F}_k$, that

$$\int_B X_T dP = \int_B X_k dP$$

follows from $E[X_T|\mathcal{F}_k] = X_k$.

Remark. Although we have worked with square integrable martingales the definition of a martingale and Doob's stopping theorem only needs the existence of the mean. Just integrability of $|X_n|$ is all that is needed for the definition to make sense. Of course any calculation that involves the second moment needs the variances to be finite.