Stochastic Processes and Stochastic Calculus - 4
Martingales

Maurizio Pratelli

Università degli Studi di Pisa

San Miniato - 13 September 2016

Overview

A fair game

A closed box contains 2 billiard balls, 1 billiard black and 1 billiard white (no difference except for the colors). You extract blindly one ball.

- if the ball extracted is white, I give you 1 €
- if the ball is black, you give me x €.

If $x = 2$, is this game fair, i.e. both players are treated equally? NO.

Which value(s) of x should we put so the game becomes fair?

A small variant: there are 3 black balls and 2 white ones, and again

- if the ball extracted is white, I give you 1 €
- if the ball is black, you give me x €.

Which value(s) of x should we put so this game becomes fair?

A criterion of fairness

There are 3 black balls and 2 white ones:

- if the ball extracted is white, I give you 1 €
- if the ball is black, you give me x €.

A criterion could be that on average we win or lose the same money:

$$P(\text{I win})x€ = P(\text{You win})1€, \quad \text{and} \quad P(\text{I lose})1€ = P(\text{You lose})x€.$$

In this game, the amount of money I get corresponds to the money you lose, and viceversa (zero-sum game), so the two conditions reduce to

$$P(\text{I win})x€ = P(\text{I lose})1€$$

which is a condition on a single player. We can also introduce the (random) gain

$$X = I_{\{\text{I win}\}}x€ - I_{\{\text{I lose}\}}1€$$

Fair game

We say the game is fair if

$$E[X] = 0, \quad \text{i.e.} \quad P(\text{I win})x€ - P(\text{I lose})1€ = 0$$

Let us compute x in the example.

$$P(\text{You lose}) = P(\text{I win}) = P(\text{black is extracted}) = \frac{3}{5}$$

$$P(\text{You win}) = P(\text{I lose}) = P(\text{white is extracted}) = \frac{2}{5}$$

Then the condition $E[X] = 0$ becomes

$$\frac{3}{5}x - \frac{2}{5}1 = 0, \quad \Rightarrow \quad x = \frac{2}{3}.$$

The condition $E[X] = 0$ is just a criterion.

Not everyone is willing to play “fair” games according to this criterion.
Consider the following example
With probability

- 10^{-6} you win 10^9 €
- 10^{-4} you lose 10^6 €
- $1 - 10^{-4} - 10^{-6}$ you lose 1 €

This game even better than fair, since the expected gain (in €) is

$$E[X] = 10^{-6} \cdot 10^9 - 10^{-4} \cdot 10^6 - \left(1 - 10^{-4} - 10^{-6}\right) 1 \sim 899 > 0$$

Would you play?

The problem seems to be that we choose to measure gain and losses in a linear way, but to model reality it could be better to introduce some non-linear function.

Linearity is very useful to generalize this notion to sequences of games.

Fair games and information

We want to generalize our criterion

$$\text{fair} \quad \Leftrightarrow \quad E[X] = 0$$

to situations where we play sequences of "fair" games (e.g. at a casino).

We want to take into the picture also the information we get about our games.

Information is essential:

- we could think a game is fair, but in reality it could be tricked
- we could get better information than our competitors and beat them.

We know that any σ-algebras $\mathcal{B}$ encode possible information that we may get. Therefore, we could say that a game is still fair given the information $\mathcal{B}$, if the gain X satisfies

$$E[X|\mathcal{B}](\omega) = 0, \quad \text{for a.e. } \omega \in \Omega.$$

Recall the rule

$$E[E[X|\mathcal{B}]] = E[X] \quad \Rightarrow \quad E[X] = 0.$$

Notice that if X is $\mathcal{B}$-measurable (i.e. the outcome is known), then the game is fair only if $E[X|\mathcal{B}] = X = 0$: "no risk, no gain".

Sequences of fair games

Assume a player is given a sequence of $N \geq 1$ games, with (uncertain) gains

$$X_1, X_2, \ldots, X_N$$

and a sequence of increasing σ-algebras (filtrations)

$$\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \ldots \subseteq \mathcal{F}_N$$

representing the information the player get as time goes on.

Moreover, for every $n \in \{1, \ldots, N\}$, $\mathcal{F}_n$ contains the knowledge of the outcome of X_n, i.e.

$$X_n \text{ is } \mathcal{F}_n\text{-measurable}$$

Then we could say that the N games are fair if, for every $n \in \{1, \ldots, N\}$,

$$E\left[X_n | \mathcal{F}_{n-1}\right] = 0, \quad \text{a.e. on } \Omega.$$

Martingales

Instead of working with the gain/losses of each game, we can also study how the total capital evolves, i.e. define

$$M_n := M_0 + X_1 + \ldots + X_n$$

when M_0 is the initial capital at our disposal (constant or $\mathcal{F}_0$-measurable).

$X_n = M_n - M_{n-1} \Rightarrow$ the condition $E\left[X_n | \mathcal{F}_{n-1}\right] = 0$ becomes

$$E\left[M_n - M_{n-1} | \mathcal{F}_{n-1}\right] = 0 \quad \text{or} \quad E\left[M_n | \mathcal{F}_{n-1}\right] = M_{n-1}.$$

Notice that this implies

$$E\left[M_n | \mathcal{F}_{n-2}\right] = E\left[E\left[M_n | \mathcal{F}_{n-1}\right] | \mathcal{F}_{n-2}\right] = E\left[M_{n-1} | \mathcal{F}_{n-2}\right] = M_{n-2}$$

and more generally, for $k \leq n$,

$$E\left[M_n | \mathcal{F}_k\right] = M_k.$$

Martingales (general definitions)

Definition (adapted process)

A process $(Y_t)_{t\in\mathcal{T}}$ is adapted to a filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$ if, for every $t \in \mathcal{T}$

$$Y_t \text{ is } \mathcal{F}_t\text{-measurable.}$$

Definition (martingale)

A process $(M_t)_{t\in\mathcal{T}}$ is a martingale with respect to a filtration $(\mathcal{F}_t)_{t\in\mathcal{T}}$ if

- it is adapted to $(\mathcal{F}_t)_{t\in\mathcal{T}}$
- for every $t \in \mathcal{T}$, $E[|M_t|] < \infty$
- for every $s, t \geq 0$ with $s \leq t$,

$$E[M_t|\mathcal{F}_s] = M_s.$$

A useful fact is that $E[M_t] = E[E[M_t|\mathcal{F}_0]] = E[M_0]$ is constant.

We call martingale differences the random variables $M_t - M_s$, for $s \leq t$.

Examples - Sum of independent variables

Let $X_1, \ldots, X_N$ be independent (real) random variables with

$$E[X_1] = E[X_2] = \ldots = E[X_N] = 0.$$

Let $\mathcal{F}_n = \sigma(X_1, \ldots, X_n)$ be the natural filtration. Then $M_n := X_1 + \ldots + X_n$ is a martingale. Indeed, by independence, for $k \leq n$,

$$E[M_n | \mathcal{F}_k] = E[M_k + (X_{k+1} + \ldots + X_n) | \mathcal{F}_k] = M_k + E[X_{k+1} + \ldots + X_n] = M_k.$$

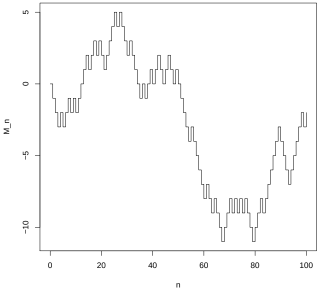

(iid uniform jumps on $\{-1, 1\}$)

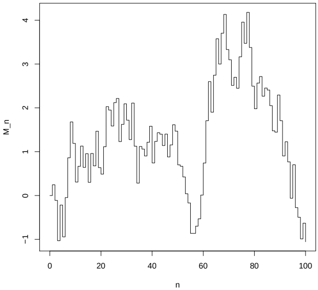

(iid uniform jumps on $[-1, 1]$)

Examples - Poisson process (continuous time)

Recall the definition of the Poisson process N_t, $t \geq 0$. Since $E[N_t] = t$, N is not a martingale. It turns out that

$$M_t := N_t - t$$

is a martingale (with respect to the natural filtration of N).

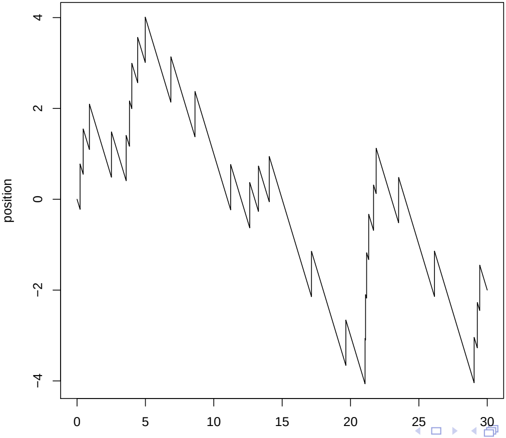

Examples - Collecting information about a random variable

Another way to build a martingale, is to “approximate” a given random variable X by means of conditional expectations w.r.t. $\mathcal{F}_t$, i.e.

$$M_t = E\left[X|\mathcal{F}_t\right].$$

The martingale property follows again from

$$E\left[E\left[X|\mathcal{F}_t\right]|\mathcal{F}_s\right] = E\left[X|\mathcal{F}_s\right].$$

Betting strategies

Let us go back to the situation of N fair games, with gains $X_1, \ldots, X_N$.

Problem

Can we set up betting strategies so that we have surely non-negative (or even positive) earnings after playing all games (free lunch)?

It turns out the answer is NO (more about this on the afternoon lecture), unless you break some rules (e.g. insider trading).

We define a betting strategy $(C_n)_{n=1}^N$ as the decision to bet a (positive or negative) amount of money C_n for the game $n \Rightarrow$ earning after playing is

$$C_n \cdot X_n$$

The strategy $(C_n)_{n=1}^N$ is not deterministic, but can only depend on the information that you have just before playing the game n, i.e.

$$C_n \text{ is } \mathcal{F}_{n-1}\text{-measurable.}$$

The total capital at time n (i.e. after we played game n) is

$$M_0 + \sum_{k=1}^{n} C_n \cdot X_n.$$

Martingale transforms

Let us give formal definitions.

Definition (predictable process)

A process $(Y_n)_{n=1}^N$ is predictable with respect to a filtration $(\mathcal{F}_n)_{n=0}^N$ if, for every $n = 1, \ldots, N$

$$Y_n \text{ is } \mathcal{F}_{n-1}\text{-measurable.}$$

Hence, betting strategies must be predictable.

Definition (martingale transform)

Given a martingale $(M_n)_{n=1}^N$ and a predictable process $(C_n)_{n=1}^N$, we define the martingale transform $(C \cdot M)_n$ as the process

$$(C \cdot M)_n := M_0 + \sum_{k=1}^{n} C_k \cdot (M_k - M_{k-1}).$$

Theorem (Martingale transforms are martingales)

Let M be a martingale and C be a bounded predictable process. Then

$$(C \cdot M)_{n=0}^{N}$$

is a martingale. In particular,

$$E\left[(C \cdot M)_N\right] = E\left[M_0\right].$$

The boundedness assumption

$$\sup_{n,\omega} |C_n|(\omega) \leq c < \infty$$

is natural: could/would you bet unlimited amounts of money?

In words: If the games are fair and you play fair, there is no free lunch.

Proof of Theorem

We have to check the three conditions that make $(C \cdot M)_{n=0}^{N}$ a martingale

1) $(C \cdot M)_n = M_0 + \sum_{k=1}^{n} C_k(M_k - M_{k-1})$ is $\mathcal{F}_n$-measurable: it is a function only of

$$C_k\text{'s with } k \leq n \Rightarrow \mathcal{F}_{n-1} \text{ measurable}$$

$$M_k\text{'s with } k \leq n \Rightarrow \mathcal{F}_{n} \text{ measurable}$$

2) $|(C \cdot M)_n| \leq |M_0| + \sum_{k=1}^{n} |C_k| \left(|M_k| + |M_{k-1}|\right) \leq 2c \sum_{k=0}^{n} |M_k|$, where

$$\sup_{n,\omega} |C_n| (\omega) \leq c < \infty,$$

hence $E\left[|(C \cdot M)_n|\right] < \infty$.

3) For $n \geq 1$, we have

$$E[(C\cdot M)_n|\mathcal{F}_{n-1}] = E\left[M_0 + \sum_{k=1}^{n} C_k(M_k - M_{k-1})\middle|\mathcal{F}_{n-1}\right]$$

$$(E[\cdot|\mathcal{F}_{n-1}] \text{ is linear}) = M_0 + \sum_{k=1}^{n} E[C_k(M_k - M_{k-1})|\mathcal{F}_{n-1}]$$

$$(M_k\ C_k \text{ are } \mathcal{F}_{n-1}\text{-meas.}) = M_0 + \sum_{k=1}^{n-1} C_k(M_k - M_{k-1}) + E[C_n(M_n - M_{n-1})|\mathcal{F}_{n-1}]$$

$$(C_n \text{ is } \mathcal{F}_{n-1}\text{-meas.}) = M_0 + \sum_{k=1}^{n-1} C_k(M_k - M_{k-1}) + C_n E[(M_n - M_{n-1})|\mathcal{F}_{n-1}]$$

$$(M \text{ is mart.}) = M_0 + \sum_{k=1}^{n-1} C_k(M_k - M_{k-1}) = (C\cdot M)_{n-1}.$$

A natural betting strategy could be as follows:

keep playing until some condition chosen in advance is realized, then leave immediately the game.

Example Fix in advance some $\lambda > 0$. Then keep playing until your capital

$$M_0 + M_1 + \ldots + M_n$$

gets larger than λ. As soon as this happens, quit the game.

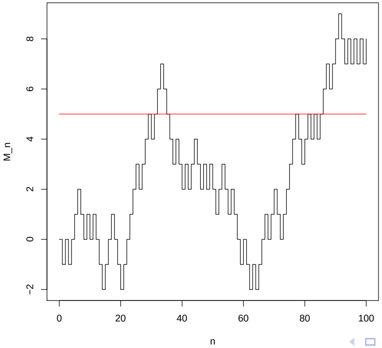

In order to be a predictable strategy, we must be able to choose whether to leave or stay before we play the game!

Example Leave the game once you reach the maximal gain that you can reach

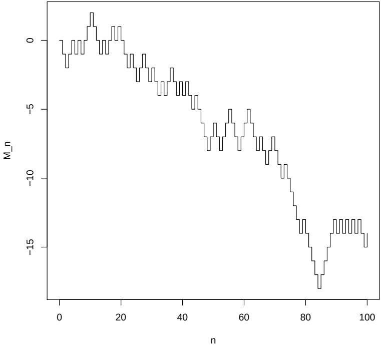

How could you tell at $n \sim 17$ that the overall maximum was already attained?

Let us denote τ the random time at which we have to leave the games.

Definition (Stopping time)

A stopping time $\tau : \Omega \to \{0, 1, \dots, N\}$ is a random variable such that

$$\{\tau > n-1\} \in \mathcal{F}_{n-1}, \quad \text{for } n = 1, \dots, N.$$

If τ is a stopping time, the betting strategy "play up to game τ and then leave", in formulas

$$C_n = I_{\{\tau > n-1\}}, \quad n \in \{1, \dots, N\}$$

is predictable.

We have that

$$M_\tau(\omega) := M_{\tau(\omega)}(\omega)$$

is the final capital if we follow the strategy to play up to (included) the stopping time τ.

We have no gain in expectation:

$$E[M_\tau] = E[M_0]$$

Back to the example

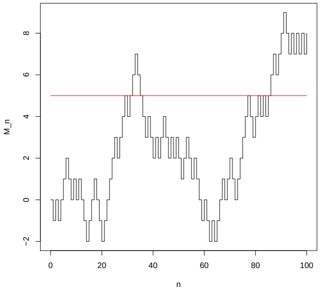

Fix in advance some $\lambda > 0$. Then keep playing until your capital

$$M_n = M_0 + X_1 + \ldots + X_n$$

gets larger than λ. As soon as this happens, quit the game.

We have

$$\{\tau > n - 1\} = \{M_0 < \lambda, M_1 < \lambda, \ldots, M_{n-1} < \lambda\} \in \mathcal{F}_{n-1}$$

Let $A = \{\text{there is } k \in \{0, 1, \ldots, N\} \text{ such that } M_n \geq \lambda\}$. Then

$$M_\tau \geq \lambda I_A + M_N I_{A^c} \quad \Rightarrow \quad E[M_0] = E[M_\tau] \geq \lambda P(A) + E[M_N I_{A^c}]$$

$$P(\{\text{there is } k \in \{0, 1, \ldots, N\} \text{ such that } M_n \geq \lambda\}) \leq \frac{1}{\lambda} E[|M_N| + |M_0|]$$

Predictable quadratic variation

We know that the variance of a random variable X is a good index of uncertainty/variability.

Given a fair game with gain X, we have

$$E[X] = 0, \quad \text{hence} \quad \mathrm{Var}(X) = E\left[X^2\right].$$

Problem

Given a martingale $(M_n)_{n=0}^N$, what can we say on the variance of M_n?

Intuitively, the variance increases as n increases. We can be more precise.

For any martingale difference $X_n := M_n - M_{n-1}$, we compute a conditional variance

$$\mathrm{Var}(X|\mathcal{F}_{n-1}) = E\left[(X_n - E[X_n|\mathcal{F}_{n-1}])^2 |\mathcal{F}_{n-1}\right] = E\left[X_n^2|\mathcal{F}_{n-1}\right] \geq 0,$$

which measures the variability of the gain X_n given the information $\mathcal{F}_{n-1}$.

Given a martingale M, introduce the predictable process, $\langle M \rangle_0 = 0$ and, for $n \geq 1$,

$$\langle M \rangle_n := \sum_{k=1}^{n} \mathrm{Var}(M_k - M_{k-1} | \mathcal{F}_{k-1}) = \sum_{k=1}^{n} E\left[(M_k - M_{k-1})^2 | \mathcal{F}_{k-1}\right]$$

Notice that $n \mapsto \langle M \rangle_n (\omega)$ is increasing.
Beware than $\langle \cdot \rangle$ is quadratic (like the variance), e.g. $\langle \lambda M \rangle_n = \lambda^2 \langle M \rangle_n$

Theorem

Let $M = (M_n)_{n=0}^N$ be a martingale with $E\left[M_N^2\right] < \infty$. Then, the process

$$n \mapsto M_n^2 - \langle M \rangle_n$$

is a martingale. Moreover,

$$\mathrm{Var}(M_n) = \mathrm{Var}(M_0) + E\left[\langle M \rangle_n\right]$$

We check only the third condition. For $n \geq 1$, write $X_n = M_n - M_{n-1}$,

$$
\begin{aligned}
E\left[M_n^2 - \langle M \rangle_n \mid \mathcal{F}_{n-1}\right] &= E\left[(M_{n-1} + X_n)^2 - E\left[X_n^2 | \mathcal{F}_{n-1}\right] - \langle M \rangle_{n-1} \mid \mathcal{F}_{n-1}\right] \\
&= E\left[M_{n-1}^2 + 2M_{n-1}X_n + X_n^2 - E\left[X_n^2 | \mathcal{F}_{n-1}\right] - \langle M \rangle_{n-1} \mid \mathcal{F}_{n-1}\right] \\
&= M_{n-1}^2 - \langle M \rangle_{n-1} - E\left[X_n^2 | \mathcal{F}_{n-1}\right] + E\left[2M_{n-1}X_n + X_n^2 \mid \mathcal{F}_{n-1}\right] \\
&= M_{n-1}^2 - \langle M \rangle_{n-1} + E\left[2M_{n-1}X_n \mid \mathcal{F}_{n-1}\right] \\
&= M_{n-1}^2 - \langle M \rangle_{n-1} + 2M_{n-1}E\left[X_n \mid \mathcal{F}_{n-1}\right] \\
&= M_{n-1}^2 - \langle M \rangle_{n-1}
\end{aligned}
$$

The expectation of a martingale is constant, hence

$$
E\left[M_0^2\right] = E\left[M_0^2 - \langle M \rangle_0\right] = E\left[M_n^2 - \langle M \rangle_n\right] = E\left[M_n^2\right] - E\left[\langle M \rangle_n\right]
$$

and since $E\left[M_0\right] = E\left[M_n\right]$,

$$
\mathrm{Var}(M_0) = E\left[M_0^2\right] - E\left[M_0\right]^2 = E\left[M_n^2\right] - E\left[M_n\right]^2 - E\left[\langle M \rangle_n\right] = \mathrm{Var}(M_n) - E\left[\langle M \rangle_n\right]
$$