Lecture 26

David Aldous

28 October 2015

For a real-valued RV X and a σ-field $\mathcal{F}$, we can define the conditional expectation $\mathbb{E}(X|\mathcal{F})$.

- The **gambling interpretation** of $\mathbb{E}(X|\mathcal{F})$ is as the fair stake Z to pay today in order to receive X tomorrow, when $\mathcal{F}$ is the known information,
- The **abstract math definition** of $\mathbb{E}(X|\mathcal{F})$ is as the $\mathcal{F}$-measurable RV Z such that

$$\mathbb{E}[Z 1_A] = \mathbb{E}[X 1_A] \text{ for all } A \text{ in } \mathcal{F}.$$

- In the case $\mathcal{F} = \sigma(Y)$ we have $\mathbb{E}(X|\mathcal{F}) = \mathbb{E}(X|Y)$ as defined before.

Analogous to rules in algebra/calculus, there are many rules for manipulating conditional expectations; will develop as we go.

Rules for manipulating conditional expectation

All RVs assumed integrable.

1. $\mathbb{E}(X \pm Y|\mathcal{F}) = \mathbb{E}(X|\mathcal{F}) \pm \mathbb{E}(Y|\mathcal{F})$

2. If X is $\mathcal{F}$-measurable then $\mathbb{E}(X|\mathcal{F}) = X$.

3. If X is independent of $\mathcal{F}$ then $\mathbb{E}(X|\mathcal{F}) = \mathbb{E}X$.

4. If W is $\mathcal{F}$-measurable then $\mathbb{E}(WX|\mathcal{F}) = W\mathbb{E}(X|\mathcal{F})$.

5. If $\mathcal{F} \subseteq \mathcal{G}$ then $\mathbb{E}[\mathbb{E}(X|\mathcal{G})\ |\mathcal{F}] = \mathbb{E}(X|\mathcal{F})$. In particular $\mathbb{E}[\mathbb{E}(X|\mathcal{G})] = \mathbb{E}X$.

Most material in this lecture is in [BZ] chapter 3, different notation.

A martingale is a process $(X_0, X_1, X_2, \ldots)$ such that

- $\mathbb{E}(X_{t+1}|\mathcal{F}_t) = X_t$ for each $t \geq 0$.

If no filtration is specified then we take the natural filtration $\mathcal{F}_t = \sigma(X_0, \ldots, X_t)$. A martingale represents your successive fortune in a fair game. In particular

$$\mathbb{E}X_t = \mathbb{E}X_0, \ \ t = 1, 2, 3, \ldots.$$

In more advanced Probability, one studies limit theorems and inequalities for martingales, which can be used to prove results about other stochastic processes. In this course, our theory will involve **gambling strategies** and **stopping times** and the **optional sampling theorem**. We then see how to use the theory to do calculations.

Take a process (M_t), and view it as results of 1 unit bets at each time, or the value of 1 unit of stock at the end of successive days. A **gambling strategy** is a decision, each time, of how many units to bet, or how many units of stock to hold.

Thinking of the "stock" case, suppose you can buy/sell only at end of trading on day t. So the number α_t of shares you hold during day t is chosen by you at end of day $t-1$ based on information known then. So

$$X_t = \text{your fortune at end of day } t$$

is given by

$$X_t - X_{t-1} = \alpha_t (M_t - M_{t-1}).$$

Theorem

If (M_t) is a martingale and α_t is $\mathcal{F}_{t-1}$-measurable for each t then (X_t) is a martingale.

[outline on board]. The conceptual point is that, with a "fair game", there is no "system" (varying the amounts you bet each time) which makes the game favorable to you.

A **stopping time** τ is a RV taking values in $\{0, 1, 2, \ldots; \infty\}$ such that

$$\{\tau = t\} \in \mathcal{F}_t \text{ for each } 0 \leq t < \infty.$$

In words: your decision when to stop depends on past and present information only – you cannot see the future.

Most stopping times we use are defined as "the first time" something happens. Note that "the last time" is usually **not** a stopping time.

Given a process $(X_t, 0 \leq t < \infty)$ and a stopping time τ, the "stopped process" is defined as

$$X_t^* = X_{\min(t,\tau)}, \quad t = 0, 1, 2, \ldots.$$

[board: notationally more convenient to stop the process changing than to stop time.] Mathematically, this is just a simple gambling strategy, so

If (X_t) is a martingale then the stopped process (X_t^*) is a martingale.

$$X_t^* = X_{\min(t,\tau)}, \quad t = 0, 1, 2, \ldots.$$

If (X_t) is a martingale then the stopped process (X_t^*) is a martingale.

- -

Now suppose the stopping time τ is such that, for some constant t_0,

$$\mathbb{P}(\tau \le t_0) = 1.$$

Then

$$X_\tau = X_\tau^* = X_{t_0}^*; \quad X_0 = X_0^*$$

and because X^* is a martingale we have $\mathbb{E}X_{t_0}^* = \mathbb{E}X_0^*$, that is

$$\mathbb{E}X_\tau = \mathbb{E}X_0.$$

This is a special case of a general theorem.

Theorem (Optional Sampling Theorem)

If (X_t) is a martingale and τ is a stopping time, then (under extra technical conditions)

$$\mathbb{E}X_\tau = \mathbb{E}X_0.$$

Advanced Probability courses give different versions of the "extra technical conditions" – see [BZ] Theorem 3.1 for one version of these conditions. In the examples I will give, it is not hard to show the conditions hold.

The "double when you lose" strategy shows that some extra condition is necessary. [board]

Conceptual point: The Optional Sampling Theorem and the previous "gambling systems" theorem constitute an informal "conservation of fairness" principle: the overall results of any "system" based on fair games is like a single fair bet. Even in models not explicitly involving gambling, one can do calculations by inventing hypothetical gambling strategies and using this principle.

Example: patterns in coin-tossing or dice-throwing.

Throw a die until we see a specified sequence, say 5 2 5 of outcomes. This requires τ throws. Calculate $\mathbb{E}\tau$.

[board – outline below]
Consider "strategy 17":
- bet 1 that throw 17 will be "5";
- if win (now have 6 units) bet 6 units that throw 18 will be "2";
- if win (now have 36 units) be 36 units that throw 19 will be "5"
- if win then the game stops.

Now consider analogous "strategy n" for each $1 \leq n \leq \tau$. The overall gain from all these strategies works out as $216 + 6 - \tau$. But by the "conservation of fairness" principle the expected gain must be zero. So $\mathbb{E}\tau = 216 + 6 = 222$.