6. Martingales

- For casino gamblers, a **martingale** is a betting strategy where (at even odds) the stake doubled each time the player loses. Players follow this strategy because, since they will eventually win, they argue they are guaranteed to make money!
- A stochastic process $\{Z_n, n \geq 1\}$ is a **martingale** if $\mathbb{E}\big[\,|Z_n|\,\big] < \infty$ and $\mathbb{E}\big[Z_{n+1} \,|\, Z_1, \ldots, Z_n\big] = Z_n$.
- Think of Z_{n+1} as being a gambler's earnings after $n+1$ games. If the game if **fair**, then $\mathbb{E}\big[Z_{n+1} \,|\, Z_n\big] = Z_n$. This is true whatever stake the gambler places-the stake at game $n+1$ can depend on $Z_1, \ldots, Z_n$.
- We will show that no strategy can guarantee success at a fair game. This is a generalization of Wald's equation.

- Since martingales can have rather general dependence (the only constraint is an conditional expectations), they are a powerful tool for dependent stochastic processes.
- Identifying an embedded martingale can lead to elegant solutions.

Examples

(i) Suppose $X_1, X_2, \ldots$ are iid, mean μ. Then $Z_n = \sum_{i=1}^{n}(X_i - \mu)$ is a Martingale. Why?

(ii) **Product martingale**. If $X_1, X_2, \ldots$ are iid with mean 1, then $Z_n = \prod_{i=1}^{n} X_i$ is a martingale. Why?

(iii) **Martingale differences**. Let $X_1, X_2, \ldots$ be arbitrary dependent random variables with $\mathbb{E}\big[\,|X_i|\,\big] < \infty$ and $\mathbb{E}\big[X_k \,|\, X_1, \ldots, X_{k-1}\big] = 0$. Then $Z_n = \sum_{i=1}^{n} X_i$ is a martingale. Why?

- Correspondingly, for any martingale $\{Z_n\}$ we can construct **martingale differences** $X_k = Z_k - Z_{k-1}$
- In particular, for any sequence $X_1, X_2, \ldots$, $Z_n = \sum_{i=1}^{n} \big\{X_i - \mathbb{E}\big[X_i \,|\, X_1, \ldots X_{i-1}\big]\big\}$ is a martingale. Why?

(iv) Let $\mathbb{E}\big[\,|X|\,\big] < \infty$ and take $Y_1, Y_2, \ldots$ to be arbitrary random variables. Set $Z_n = \mathbb{E}\big[X \,|\, Y_1, \ldots, Y_n\big]$. Then Z_n is a Martingale. Why?

(v) **Continuous-time martingales**. $Z(t)$ is a martingale in continuous time if $\mathbb{E}\big[\,|Z(t)|\,\big] < \infty$ and $\mathbb{E}\big[Z(t) \,|\, Z(u), 0 \leq u \leq s\big] = Z(s)$ for $s < t$. We will not study the continuous time case thoroughly, but similar results apply.

- If $N(t)$ is a rate λ Poisson counting process, $Z(t) = N(t) - \lambda t$ is a martingale.

- If $N(t)$ is a general renewal process, then $N(t) - m(t)$ is not a Martingale. Why? Find a related process which is a Martingale.

- Identify a Martingale corresponding to a continuous time birth-death process, $X(t)$, with rates λ_n and μ_n.

- Recall that a non-negative integer-valued random variable N is a **stopping time** for $\{Z_n, n \geq 1\}$ if $\{N = n\}$ is determined by $Z_1, \ldots, Z_n$.

- Here, we do not require $\mathbb{E}[N] < \infty$.

- More generally, if we allow the possibility that $\mathbb{P}\big[N{=}\infty\big] > 0$, N is a **random time**.

- The **stopped process** $\big\{\overline{Z}_n, n \geq 1\big\}$ is given by

$$\overline{Z}_n = \begin{cases} Z_n & \text{if } n \leq N \\ Z_N & \text{if } n > N \end{cases}$$

- $\big\{\overline{Z}_n\big\}$ inherits the martingale property from $\{Z_n\}$. Why?

Solution continued

Theorem (**Martingale Stopping Theorem**)

If N is a stopping time for a martingale $\{Z_n\}$ then $\mathbb{E}\big[Z_N\big] = \mathbb{E}\big[Z_1\big]$, provided one of the following conditions is satisfied:

(i) $\overline{Z}_n$ is uniformly bounded (i.e., there exist a and b with $a < Z_n < b$ whenever $n \leq N$).

(ii) N is bounded.

(iii) $\mathbb{E}[N] < \infty$ and, for some $M < \infty$,
$\mathbb{E}\big[\,|Z_{n+1} - Z_n|\,\big|\,Z_1, \ldots, Z_n\big] < M$.

Proof

Example: What does the stopping theorem imply about the "martingale" betting strategy (keep doubling the stakes until you win, then eventually you are guaranteed to make a profit).

Example: A gambler plays a fair game at even odds, each play results in winning/losing \$1 with probability 1/2. The gambler starts with $\$\,a$ and stops when he goes broke or reaches $\$\,b$. Find the chance of reaching $\$\,b$.

Example: Expected time to see a given pattern.

A monkey hits random letters on a keyboard. What is the expected number of hits until typing ABRACADABRA.

Solution

Convergence of Martingales

- A useful property of martingales is that, if their expected absolute value is uniformly bounded, they converge with probability 1.
- To develop these ideas, we first study some inequalities.

Definition If $\{Z_n, n \geq 1\}$ is a stochastic process with $\mathbb{E}\big[\,|Z_n|\,\big] < \infty$ then it is a

submartingale if $\mathbb{E}\big[Z_{n+1}\,|\,Z_1,\ldots,Z_n\big] \geq Z_n$

supermartingale if $\mathbb{E}\big[Z_{n+1}\,|\,Z_1,\ldots,Z_n\big] \leq Z_n$

- Most casino games are super martingales, as far as the player is concerned, i.e., subfair. Allegedly, there are systems to make the player's winnings at blackjack a submartingale, i.e., superfair.
- Note that the definition implies
  $\mathbb{E}\big[Z_{n+1}\big] \geq \mathbb{E}\big[Z_n\big]$ for a submartingale.
  $\mathbb{E}\big[Z_{n+1}\big] \leq \mathbb{E}\big[Z_n\big]$ for a supermartingale.

Example: If f is a convex function and $\{Z_n\}$ is a martingale, then $\{f(Z_n)\}$ is a submartingale.

Proof. We need to use **Jensen's inequality**: If f is a convex function, then $\mathbb{E}[f(X)] \geq f(\mathbb{E}[X])$.

Stopping for Sub(Super) Martingales

If N is a stopping time for $\{Z_n\}$ satisfying any of the conditions for the martingale stopping theorem,

$$\left.\begin{array}{ll} \mathbb{E}\big[Z_N\big] \geq \mathbb{E}\big[Z_1\big] & \text{for a submartingale} \\ \mathbb{E}\big[Z_N\big] \leq \mathbb{E}\big[Z_1\big] & \text{for a supermartingale} \end{array}\right\} (1)$$

If furthermore, N is bounded, say $N \leq n$, then

$$\left.\begin{array}{ll} \mathbb{E}\big[Z_n\big] \geq \mathbb{E}\big[Z_N\big] \geq \mathbb{E}\big[Z_1\big] & \text{submartingale} \\ \mathbb{E}\big[Z_n\big] \leq \mathbb{E}\big[Z_N\big] \leq \mathbb{E}\big[Z_1\big] & \text{supermartingale} \end{array}\right\} (2)$$

Proof. For the submartingale case, how does (1) follow from the martingale stopping theorem? How does (2) follow from (1)?

Proof continued

Kolmogorov's submartingale inequality

If $\{Z_n\}$ is a non-negative submartingale, then $\mathbb{P}\big[\max(Z_1, \ldots, Z_n) \geq a\big] \leq \frac{\mathbb{E}[Z_n]}{a}$ for $a > 0$.

- Note that Jensen's inequality then gives, for any martingale $\{Z_n\}$,
  $\mathbb{P}\big[\max(\,|Z_1|\,, \ldots, \,|Z_n|\,) \geq a\big] \leq \mathbb{E}\big[\,|Z_n|\,\big]/a$
  and
  $\mathbb{P}\big[\max(\,|Z_1|\,, \ldots, \,|Z_n|\,) \geq a\big] \leq \mathbb{E}\big[Z_n^2\big]/a^2$.

Proof of the submartingale inequality

Example: An urn initially contains one white and one black ball. At each stage a ball is drawn, and is then replaced in the urn along with another ball of the same color. Let Z_n be the fraction of white balls in the urn after the n^{th} iteration.

(a) Show that $\{Z_n\}$ is a martingale.

(b) Show that the probability that the fraction of white balls is ever as large as $3/4$ is at most $2/3$.

Martingale Convergence Theorem

If $\{Z_n, n \geq 1\}$ is a martingale and $\mathbb{E}\big[\,|Z_n|\,\big] \leq M$ then, with probability 1, $\lim_{n\to\infty} Z_n$ exists and is finite.

- Note: write $Z_\infty = \lim_{n\to\infty} Z_n$. This limit is generally random, but Z_∞ may sometimes be a constant. This theorem asserts that for (almost) every outcome s in the sample space $\mathbb{S}$, $\lim_{n\to\infty} Z_n(s) = Z_\infty(s)$.
- Note that the theorem applies to <u>any</u> non-negative martingale, since then $\mathbb{E}\big[\,|Z_n|\,\big] = \mathbb{E}[Z_n] = \mathbb{E}[Z_1]$.

Proof

<u>Proof continued</u>

Example: Let $\{Z_n, n \geq 1\}$ be a sequence of random variables such that $Z_1 = 1$ and, given $Z_1, \ldots, Z_{n-1}$, the distribution of Z_n is conditionally Poisson with mean Z_{n-1} for $n > 1$. What happens to Z_n as n gets large?

Example: Let X_n be the population size of the n^{th} generation of a branching process, with each individual having, on average, m offspring. Describe the behavior of X_n for large n.

Martingales to Analyze Random Walks

- The general random walk, $\{S_n, n \geq 0\}$, is defined by $S_0 = 0$ and $S_n = \sum_{i=1}^{n} X_i$ for $n > 0$, where $X_1, X_2, \ldots$ are iid.

- A random walk can be considered as a generalization of a renewal process, where we drop the requirement that $X_i \geq 0$.

- The most obvious martingale is $S_n - n\mu$ where $\mu = \mathbb{E}[X_1]$. Here, μ is called the **drift**.

- Another useful martingale is $\exp\{\theta S_n\}$ where θ solves $\mathbb{E}[e^{\theta X_1}] = 1$. This equation has one solution at $\theta = 0$, and it usually has exactly one other solution, with $\theta > 0$, if $\mathbb{E}[X_1] < 0$. Why?

Example: Let $N = \min\{n : S_n \geq A \text{ or } S_n \leq -B\}$. Use martingale arguments to find (approximately) $\mathbb{P}[S_N \geq A]$ and $\mathbb{E}[N]$.

- Note: this models a general situation where we accumulate rewards, and at some point we quit and declare failure (if $S_N \leq -B$), or quit having achieved our goal (if $S_N \geq A$). An example is sequential analysis of clinical trails.

Solution

Solution continued