DISCRETE TIME MARTINGALES

Definition A sequence of random variables $M_n,\ n \geq 0$ on a common probability space $(\Omega, \mathcal{F}, P)$ is said to be a martingale if $M_n \in L^1(P)$ for each n and $\mathbf{E}[M_{n+1}|\mathcal{F}_n] = M_n$ for all $n \geq 0$, where $\mathcal{F}_n = \sigma(M_k : k \leq n)$.

Proposition Let M be a martingale. Let $C_n, n \geq 1$ be a sequence of random variables such that $C_n \in \mathcal{F}_{n-1}$ for each n. Then let $X_n = \sum_{k=1}^{n} C_k(M_k - M_{k-1})$ for $n \geq 1$. Then X is a martingale with $X_0 = 0$ a.s.

Optional sampling theorem If M is a martingale and τ is a (finite a.s.) stopping time for $\mathcal{F}_n$, and $X_n = M_{\tau \wedge n}$, then X is a martingale.

Convergence theorem for martingales Let M be a martingale.

(1) If M is bounded in $L^1(P)$, i.e., $\sup\limits_n \mathbf{E}[|M_n|] < \infty$, then $M_n \stackrel{a.s.}{\to} M$ for some $M \in L^1(P)$.

(2) If M is also uniformly integrable, i.e., $\sup\limits_n \mathbf{E}[|M_n| \mathbf{1}_{|M_n|>A}] \to 0$ as $A \to \infty$, then $M_n \stackrel{L^1, a.s.}{\longrightarrow} M$. In particular, $M_n = \mathbf{E}[M|\mathcal{F}_n]$ for all n and $\mathbf{E}[M] = \mathbf{E}[M_0]$.

(3) If M is bounded in L^p for some $p > 1$, that is $\sup\limits_n \mathbf{E}[|M_n|^p] < \infty$, then $M_n \stackrel{L^p, a.s.}{\longrightarrow} M$. In particular, $M_n = \mathbf{E}[M|\mathcal{F}_n]$ for all n and $\mathbf{E}[M] = \mathbf{E}[M_0]$.

Note that the assumptions and conclusions are both progressively stronger. A natural (and surprisingly useful!) way to construct martingales is to take an arbitrary random variable $M \in L^1(P)$ and any filtration $\mathcal{F}_n$, and then define the sequence $M_n = \mathbf{E}[M|\mathcal{F}_n]$ - this is called a "Doob martingale". The martingale convergence theorem is a converse of sorts to this statement, that any uniformly integrable martingale is a Doob martingale.

For all applications in this course, it is enough to remember that an L^2 bounded martingale converges a.s. and in L^2 (and hence in L^1 too) and therefore $\mathbf{E}[M|\mathcal{F}_n] = M_n$ and $\mathbf{E}[M_n] \to \mathbf{E}[M]$.

A few words about the proofs The first proposition about X being a martingale is a trivial consequence of definitions and properties of conditional expectation (check it!).

The optional sampling theorem is a direct consequence of the proposition by letting $C_n = \mathbf{1}_{\tau \geq n}$ which is $\mathcal{F}_{n-1}$ measurable and $X_n = \sum_{k=1}^{n} C_k(M_k - M_{k-1}) = \sum_{k=1}^{\tau \wedge n}(M_k - M_{k-1}) = M_{\tau \wedge n} - M_0$. By the previous theorem X is a martingale and hence, so is $M_{\tau \wedge n} = X_n + M_0$.

The first statement in the martingale convergence theorem has a truly beautiful proof due to Doob via upcrossing inequalities. We don't give the proof here, but Doob's inequality is presented below. Once the a.s. convergence of an $L^1(P)$ martingale is proved, the other two statements are proved in less spectacular ways.

Doob's upcrossing inequality Let M be a martingale and for any interval $[a, b]$, let $U_N[a, b]$ be the number of upcrossings of $[a, b]$ by the sequence $M_0, M_1, \ldots, M_N$. Then,

$$\mathbf{E}\left[U_N[a, b]\right] \leq \frac{1}{b-a} \mathbf{E}\left[(M_N - a)^-\right] \leq \frac{1}{b-a}\left(|a| + \sup_k \mathbf{E}[|M_k|]\right).$$

For an L^1 bounded martingale, the right hand side does not depend on N at all! Hence for any interval the number of upcrossings is a.s. finite, and as this holds a.s. for all rational intervals simultaneously, it follows that M must converge. Doob's inequality itself is proved by a clever choice of $C_k \in \mathcal{F}_{k-1}$ and applying the proposition above. (Informally, think of a gambler who fixes two numbers $a < b$, and whenever M gets below a, decided that things can only get better now and starts betting a dollar on each game. Once M reaches a level above b, she decides that things can only go downhill now, and stops betting till the next time that M reaches below a, and so on. If M oscillated infinitely many times over $[a, b]$ then the gambler would make profit in the long run, because she makes a profit of $b - a$ for every upcrossing).

Continuous time martingales

Definition A collection of random variables $(M_t)_{t\ge 0}$, on a common filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\ge 0}, P)$ is said to be a martingale if $M_t \in L^1(P)$ for each t and $\mathbf{E}[M_t|\mathcal{F}_s] = M_s$ for all $s < t$, where $\mathcal{F}_s = \sigma(M_u : u \le s)$.

For simplicity we assume that $t \to M_t$ is a.s. continuous (this is sufficient for our purposes). We shall also assume that $\mathcal{F}_\bullet$ is a right continuous and complete filtration. This latter is not just for simplicity, but is essential, but like in case of BM, one can take the natural filtration and make it right continuous and complete it, all without losing the martingale property.

Remark Results about continuous time martingales are proved using the same tools as for discrete time martingales, namely Doob's upcrossing inequalities. Notice that for any deterministic sequence of times $0 = t_0 < t_1 < t_2 < \ldots$, the sequence $X_n = M_{t_n}$ is a discrete time martingale and hence (assuming $\sup_t \mathbf{E}[|M_t|] < \infty$) for any interval $[a, b]$, the expected number of upcrossings of X over $[a, b]$ has a finite expectation (and the bound does not depend on the chosen times $\{t_i\}$). If M had infinitely many upcrossings across $[a, b]$, by the assumed continuity of M, for ϵ small, the martingale $X_n = M_{n\epsilon}$ would have arbitrarily large number of upcrossings across the same interval. Thus M itself must have only finitely many upcrossings across any interval. This means, M_t converges a.s. Along these lines of reasoning, one can prove the following theorems for continuous time martingales.

Optional sampling theorem Let τ be an $\mathcal{F}_\bullet$ stopping time and M an $\mathcal{F}_\bullet$ martingale. Define $\mathcal{G}_t = \mathcal{F}_{\tau\wedge t}$. Then $X_t = M_{\tau\wedge t}$ is a $\mathcal{G}_t$ martingale.

Martingale convergence theorem Let M be an $\mathcal{F}_\bullet$ martingale. Then

(1) If M is bounded in $L^1(P)$, i.e., $\sup_t \mathbf{E}[|M_t|] < \infty$, then $M_t \overset{a.s.}{\to} M$ for some $M \in L^1(P)$.

(2) If M is also uniformly integrable, i.e., $\sup_t \mathbf{E}[|M_t|\mathbf{1}_{|M_t|>A}] \to 0$ as $A \to \infty$, then $M_t \overset{L^1,a.s.}{\longrightarrow} M$. In particular, $M_t = \mathbf{E}[M|\mathcal{F}_t]$ for all t.

(3) If M is bounded in L^p for some $p > 1$, that is $\sup_t \mathbf{E}[|M_t|^p] < \infty$, then $M_t \overset{L^p,a.s.}{\longrightarrow} M$. In particular, $M_t = \mathbf{E}[M|\mathcal{F}_t]$ for all t.

A typical application In our applications we often have a stopping time τ. We use the optional sampling theorem and make up the martingale $M_{\tau\wedge t}$. If this new martingale is uniformly integrable (usually we check that it is bounded in $L^2(P)$) then by the martingale convergence theorem, as $t \to \infty$, we get $M_{\tau\wedge t} \overset{L^1,a.s.}{\longrightarrow} X$ for some random variable X. Since $\tau \wedge t \to \tau$ and M is continuous, it is clear that X must be equal to M_τ. That is, $M_{\tau\wedge t} \overset{L^1,a.s.}{\longrightarrow} M_\tau$ and $M_t = \mathbf{E}[M_\tau|\mathcal{G}_t]$. In particular, setting $t = 0$ in the last equation implies that $\mathbf{E}[M_\tau] = \mathbf{E}[M_0]$.

Example: Gambler's ruin problem Consider standard 1-dimensional Brownian motion B. Clearly B is a martingale (w.r.t the augmented filtration). Fix $-a < 0 < b$ and let $\tau = \inf\{t \;:\; B_t = -a \text{ or } b\}$. Let $p = \mathbf{P}[B_\tau = b]$ and $1 - p = \mathbf{P}[B_\tau = -a]$. Then if $X_t = M_{\tau\wedge t}$, then $|X_t| \le \max\{a, b\}$ and hence X is u.i. Thus the above reasoning goes through, and we get $\mathbf{E}[B_\tau] = \mathbf{E}[B_0] = 0$. Thus $pb - (1 - p)a = 0$, which implies $p = \frac{a}{a+b}$, which is exactly what we found earlier using Strong Markov property.

Example: Gambler's ruin problem continued Now consider the martingale $M_t = B_t^2 - t$. Let a, b, τ be exactly as before. Then, $X_t = M_{\tau\wedge t}$ is again a u.i. martingale and hence $\mathbf{E}[B_\tau^2 - \tau] = 0$. Thus

$$\mathbf{E}[\tau] = \mathbf{E}[B_\tau^2] = pb^2 + (1 - p)a^2 = \frac{ab^2 + ba^2}{a + b} = ab.$$

Thus martingale techniques help us to understand the stopping time distribution itself.

Exercise Check that $M_t = B_t^4 - 6tB_t^2 + 3t^2$ is a martingale. What does this tell us about τ?