

Lecture Notes on Martingales

BY RUSSELL LYONS

Extended from the book by Patrick Billingsley

Suppose we gamble on a sequence of fair games. Let X_n be our fortune after n plays and $\mathcal{F}_n := \sigma(X_1, \dots, X_n)$ or more; e.g., we might use other randomness. Fairness is $X_n = E[X_{n+1} | \mathcal{F}_n]$. Unfavorable gambling has \geq and favorable has \leq . Note $\mathcal{F}_n \subseteq \mathcal{F}_{n+1}$. One way to be favorable is to get handed some money before each fair game, where the amount has conditional expectation ≥ 0 and the future games remain fair.

Definition. Let $X_n \in L^1(\Omega, \mathcal{F}, P)$ and $\mathcal{F}_n \subseteq \mathcal{F}$ be σ -fields. We call $\langle \mathcal{F}_n \rangle$ a **filtration** if $\mathcal{F}_n \subseteq \mathcal{F}_{n+1}$. We call $\langle X_n \rangle$ **adapted** to $\langle \mathcal{F}_n \rangle$ if $X_n \in \mathcal{F}_n$. We call $\langle (X_n, \mathcal{F}_n) \rangle$ a **martingale** (**submartingale**; **supermartingale**) if $\langle X_n \rangle$ is adapted to the filtration $\langle \mathcal{F}_n \rangle$ and

$$\forall n \quad X_n = E[X_{n+1} | \mathcal{F}_n] \text{ a.s. } (\leq ; \geq)$$

We'll later see the reason for the names of the prefixes; it is related to subharmonic functions (convex functions in the case of \mathbb{R}).

We call $\langle X_n \rangle$ a **martingale** if there is some filtration with respect to which it is a martingale. In such a case, $\mathcal{F}_n := \sigma(X_1, \dots, X_n)$ always works. Note that $E[X_m | \mathcal{F}_n] = X_{m \wedge n}$ a.s. and $E[X_n]$ is the same for all n .

Example 35.1. Sums of independent random variables.

Example 35.5. If $Z \in L^1(\Omega, \mathcal{F}, P)$ and $\langle \mathcal{F}_n \rangle$ is a filtration, then $\langle (E[Z | \mathcal{F}_n], \mathcal{F}_n) \rangle$ is a martingale.

Extra Credit Homework: Show that $\langle n\mathbf{1}_{[0,1/n]} ; n \geq 1 \rangle$ is a martingale on $[0, 1]$ with respect to Lebesgue measure.

Homework: 35.2, 6 (correction: $X_0 = E[X_1]$). Also, show that $\langle Z_n \rangle$ is predictable, i.e., for all n , Z_n is \mathcal{F}_{n-1} -measurable.)

Example 35.2. Let ν be a finite, signed measure on (Ω, \mathcal{F}, P) and $\langle \mathcal{F}_n \rangle$ be a filtration. Suppose that $\nu \upharpoonright \mathcal{F}_n \ll P \upharpoonright \mathcal{F}_n$ for all n . Let $X_n := d(\nu \upharpoonright \mathcal{F}_n) / d(P \upharpoonright \mathcal{F}_n)$. Then $\langle (X_n, \mathcal{F}_n) \rangle$ is a martingale:

$$\begin{aligned} E[X_{n+1} | \mathcal{F}_n](P \upharpoonright \mathcal{F}_n) &= (X_{n+1}P) \upharpoonright \mathcal{F}_n = (X_{n+1}P \upharpoonright \mathcal{F}_{n+1}) \upharpoonright \mathcal{F}_n \\ &= (\nu \upharpoonright \mathcal{F}_{n+1}) \upharpoonright \mathcal{F}_n = \nu \upharpoonright \mathcal{F}_n = X_n(P \upharpoonright \mathcal{F}_n). \end{aligned}$$

Example 35.3. Let P be Lebesgue measure on $(0, 1]$ and $\mathcal{F}_n := \sigma(I_k^{(n)} ; 0 \leq k < 2^n)$, where $I_k^{(n)} := (k2^{-n}, (k+1)2^{-n}]$. Since $(A \in \mathcal{F}_n, P(A) = 0) \Rightarrow A = \emptyset$, we have for all ν that $\nu \upharpoonright \mathcal{F}_n \ll P \upharpoonright \mathcal{F}_n$. In this case,

$$X_n = \sum_{k=0}^{2^n-1} \mathbf{1}_{I_k^{(n)}} \nu(I_k^{(n)}) / P(I_k^{(n)}) = \sum_{k=0}^{2^n-1} \mathbf{1}_{I_k^{(n)}} \nu(I_k^{(n)}) 2^n.$$

Example 35.8. If $\langle (X_n, \mathcal{F}_n) \rangle$ is a martingale, then $\langle (|X_n|, \mathcal{F}_n) \rangle$ is a submartingale.

This is a special case of:

Theorem 35.1. Let $\langle X_n \rangle$ be integrable random variables adapted to a filtration $\langle \mathcal{F}_n \rangle$, φ be convex on an interval containing all the ranges of X_n , and $\varphi \circ X_n \in L^1(P)$. Then $\langle (\varphi \circ X_n, \mathcal{F}_n) \rangle$ is a submartingale if either

- (i) $\langle (X_n, \mathcal{F}_n) \rangle$ is a martingale; or
- (ii) $\langle (X_n, \mathcal{F}_n) \rangle$ is a submartingale and φ is increasing.

Proof. In both cases, we have $\varphi(X_n) \leq \varphi(E[X_{n+1} | \mathcal{F}_n])$. By Jensen's inequality, this is $\leq E[\varphi(X_{n+1}) | \mathcal{F}_n]$. ■

We now depart from the book.

Wald's Equation. Let $Z_n \in L^1(P)$ for $n \geq 1$, τ be an \mathbb{N} -valued random variable, and $\mu \in \mathbb{R}$. Suppose that

- (i) for each $n \geq 1$, we have $P[\tau \geq n] > 0 \Rightarrow E[Z_n | \tau \geq n] = \mu$ (\leq ; \geq)

and

- (ii) one of the following holds:

- (a) $\forall n \ Z_n \geq 0$;
- (b) $\sup_{n; P[\tau \geq n] > 0} E[|Z_n| | \tau \geq n] < \infty$ and $E[\tau] < \infty$;
- (c) $E[|\sum_{n=1}^{\tau} Z_n|] < \infty$ and $\lim_{n \rightarrow \infty} E[\sum_{k=1}^n \mathbf{1}_{\{\tau > n\}}] = 0$.

Then

$$E\left[\sum_{n=1}^{\tau} Z_n\right] = \mu E[\tau], \quad (\leq ; \geq)$$

where $0 \cdot \infty := 0$.

Proof. In case (ii)(a), we have by Tonelli's theorem (or the MCT) that

$$\begin{aligned} E\left[\sum_{n=1}^{\tau} Z_n\right] &= E\left[\sum_{n=1}^{\infty} Z_n \mathbf{1}_{[n \leq \tau]}\right] = \sum_{n=1}^{\infty} E[Z_n \mathbf{1}_{[\tau \geq n]}] \\ &= \sum_{n=1}^{\infty} E[Z_n | \tau \geq n] P[\tau \geq n] = \sum_{n=1}^{\infty} \mu P[\tau \geq n] = \mu E[\tau]. \end{aligned} \quad (\text{N1})$$

In case (ii)(b), we have

$$\begin{aligned} E\left[\sum_{n=1}^{\infty} |Z_n| \mathbf{1}_{[n \leq \tau]}\right] &= \sum_{n=1}^{\infty} E[|Z_n| | \tau \geq n] P[\tau \geq n] \\ &\leq E[\tau] \sup_n E[|Z_n| | \tau \geq n] < \infty, \end{aligned}$$

whence we may apply Fubini's theorem as in (N1).

Finally, in case (ii)(c), the LDCT gives

$$\begin{aligned} E\left[\sum_{k=1}^{\tau} Z_k\right] &= \lim_{n \rightarrow \infty} E\left[\sum_{k=1}^{\tau} Z_k \mathbf{1}_{[\tau \leq n]}\right] = \lim_{n \rightarrow \infty} E\left[\sum_{k=1}^n Z_k \mathbf{1}_{[k \leq \tau \leq n]}\right] \\ &= \lim_{n \rightarrow \infty} E\left[\sum_{k=1}^n Z_k (\mathbf{1}_{[\tau \leq k]} - \mathbf{1}_{[\tau > n]})\right] = \lim_{n \rightarrow \infty} \sum_{k=1}^n \mu P[\tau \geq k] = \mu E[\tau]. \end{aligned}$$

(Note the case $\mu = 0$ and $E[\tau] = \infty$.) ■

Definition. Given a filtration $\langle \mathcal{F}_n \rangle$, a random variable $\tau \in \mathbb{N} \cup \{\infty\}$ is called a **stopping time** with respect to $\langle \mathcal{F}_n \rangle$ if $\forall n \in \mathbb{N} [\tau = n] \in \mathcal{F}_n$. Equivalently, $\forall n \in \mathbb{N} [\tau \leq n] \in \mathcal{F}_n$.

Homework: Show that if τ_1 and τ_2 are stopping times with respect to $\langle \mathcal{F}_n \rangle$, then so are $\tau_1 \wedge \tau_2$, $\tau_1 \vee \tau_2$, and $\tau_1 + \tau_2$.

Corollary. Let $\langle (X_n, \mathcal{F}_n) \rangle$ be a martingale (submartingale; supermartingale) and τ an a.s. finite stopping time with respect to $\langle \mathcal{F}_n \rangle$. If

(a) $\exists c \leq 0 \forall n X_{n+1} - X_n \geq c$ a.s. and $E[\tau] < \infty$

or

(b) $E[|X_\tau|] < \infty$ and $\lim_n E[X_n \mathbf{1}_{[\tau > n]}] = 0$,

then

$$E[X_\tau] = E[X_0]. \quad (\geq ; \leq)$$

Proof. In case (a), apply Wald's equation with $Z_n := X_n - X_{n-1} - c$ for $n \geq 1$ and $\mu := -c$. In case (b), apply Wald's equation with $Z_n := X_n - X_{n-1}$ for $n \geq 1$ and $\mu := 0$. Note that $[\tau \geq n] = [\tau \leq n - 1]^c \in \mathcal{F}_{n-1}$. ■

Example. If $\sup \tau < \infty$, then the hypotheses (b) hold. More generally, if $\{X_{\tau \wedge n}; n \geq 0\}$ is uniformly integrable, then the hypotheses (b) hold, because $E[|X_\tau|] = \lim_{n \rightarrow \infty} E[|X_{\tau \wedge n}|]$ and $E[X_n \mathbf{1}_{[\tau > n]}] = E[X_{\tau \wedge n} \mathbf{1}_{[\tau > n]}]$, where the latter random variables are uniformly integrable.

Here are some nice applications.

Example. Let $\langle X_n \rangle$ be simple random walk on \mathbb{Z} starting from $X_0 = 0$. Let $\tau_a := \inf\{n; X_n = a\}$. Then $E[\tau_a] = \infty$ for $a \neq 0$ since the conclusion of the corollary does not hold. But if $\tau := \tau_a \wedge \tau_{-b}$ for $a, b > 0$, then $X_{\tau \wedge n}$ are bounded, so uniformly integrable, whence

$$0 = E[X_\tau] = a P[\tau_a < \tau_{-b}] + b P[\tau_{-b} < \tau_a],$$

giving $P[\tau_a < \tau_{-b}] = b/(a + b)$.

Here, we assumed that $\tau < \infty$ a.s. This is implied by a martingale convergence theorem, Theorem 35.5, but we can also prove it as follows, at the same time as we calculate $E[\tau]$.

Problem 35.2 shows that $\langle (X_n^2 - n, \sigma(X_1, \dots, X_n)) \rangle$ is a martingale. Fix $N \in \mathbb{N}$. Then $\tau \wedge N$ is a bounded stopping time, so we may apply the corollary to get

$$0 = E[X_0^2 - 0] = E[X_{\tau \wedge N}^2 - \tau \wedge N] = E[X_{\tau \wedge N}^2] - E[\tau \wedge N],$$

i.e., $E[\tau \wedge N] = E[X_{\tau \wedge N}^2]$. Now let $N \rightarrow \infty$. The right-hand side is bounded, whence so is the left-hand side; this shows that $E[\tau] < \infty$, whence $\tau < \infty$ a.s. Furthermore, the MCT and the BCT give

$$E[\tau] = E[X_\tau^2] = a^2 \frac{b}{a+b} + b^2 \frac{a}{a+b} = ab.$$

Definition. Let $p(\cdot, \cdot)$ be a transition probability on a state space, S . A function $f: S \rightarrow \mathbb{R}$ is called *harmonic* (*subharmonic*; *superharmonic*) if

$$\forall x \in S \quad f(x) = \sum_{y \in S} p(x, y) f(y). \quad (\leq ; \geq)$$

If $\langle X_n \rangle$ is a Markov chain on S with these transition probabilities and f is harmonic (subharmonic; superharmonic), then $\langle f(X_n) \rangle$ is a martingale (submartingale; supermartingale).

Example. Consider the asymmetric random walk $\langle X_n \rangle$ on \mathbb{Z} with $p(n, n+1) = p$ and $p(n, n-1) = q := 1-p$. Then $f(n) := (q/p)^n$ is harmonic, whence $\langle (q/p)^{X_n} \rangle$ is a martingale. If $\tau := \tau_a \wedge \tau_{-b}$, then this martingale is bounded up to time τ , whence the corollary applies:

$$1 = E[(q/p)^{X_0}] = E[(q/p)^{X_\tau}] = (q/p)^a P[\tau_a < \tau_{-b}] + (q/p)^{-b} P[\tau_{-b} < \tau_a].$$

We can solve this to get another derivation of the probabilities in gambler's ruin when $p \neq q$. (Again, the martingale convergence theorem will apply that $\tau < \infty$ a.s., but it is also a consequence of the following homework.)

Homework: Let $\langle X_n \rangle$ be asymmetric random walk on \mathbb{Z} with $p > q$, $s := \sqrt{q/p}$, and $r := 2\sqrt{pq}$. Show that $\langle s^{X_n} r^{-n} \rangle$ is a martingale. Deduce that if $X_0 = 0$, then $P[\tau_a \geq n] \leq s^{-a} r^n$ for $a \geq 1$.

Example. What is $E[\tau_a]$ when $p > q$ and $a \geq 1$? We use another martingale: $\langle X_n - (p-q)n \rangle$. Fix $N \in \mathbb{N}$ and define $\tau := \tau_a \wedge N$. Then τ is bounded, so $0 = E[X_0 - (p-q)0] = E[X_\tau - (p-q)\tau]$, i.e., $(p-q)E[\tau_a \wedge N] = E[X_{\tau_a \wedge N}]$. By the MCT, $E[\tau_a \wedge N] \rightarrow E[\tau_a]$ as $N \rightarrow \infty$. Using the notation of the homework, we have

$$|E[X_{\tau_a \wedge N}] - a| \leq E[|X_{\tau_a \wedge N} - a|; \tau_a > N] \leq (N+a)P[\tau_a > N] \leq (N+a)s^{-a}r^N \rightarrow 0$$

as $N \rightarrow \infty$, so $E[X_{\tau_a \wedge N}] \rightarrow a$ as $N \rightarrow \infty$. In conclusion,

$$E[\tau_a] = \frac{a}{p-q}.$$

Homework: For $\langle X_n \rangle$ be asymmetric random walk on \mathbb{Z} with $p > q$, calculate $\text{Var}(\tau_a)$ for $a > 0$. *Hint:* Use problem 35.2.

Example. Consider an ordinary deck of 26 red and 26 black cards, well shuffled. The cards are turned over one by one. At any time, and once only, we may bet that the next card is red. What's the best strategy to maximize our chance to be correct?

Let A_k be the event that the k th card is red. If τ is the time we bet, then $[\tau = k] \in \sigma(A_1, \dots, A_{k-1}) =: \mathcal{F}_k$, i.e., τ is a stopping time. At time k , the chance of winning would be $X_k := P[A_k | \mathcal{F}_k]$, so the chance of winning is $E[X_\tau]$. Now

$$E[X_{k+1} | \mathcal{F}_k] = E[P[A_{k+1} | \mathcal{F}_{k+1}] | \mathcal{F}_k] = P[A_{k+1} | \mathcal{F}_k] = X_k,$$

because all cards we have not seen are in random order. Thus, $\langle X_k; 1 \leq k \leq 52 \rangle$ is a martingale, so $E[X_\tau] = E[X_1] = 1/2$. All strategies are equally effective.

Having learned about this martingale, we realize that $\forall k X_k = E[X_{52} | \mathcal{F}_k]$. But $X_{52} = P[A_{52} | \mathcal{F}_{52}] = \mathbf{1}_{A_{52}}$, so $X_k = P[A_{52} | \mathcal{F}_k]$. This shows more directly that any strategy is equivalent to betting on the last card. That's why it does not matter. In fact, we can see this without any use of martingales: betting on the next card has the same chance of winning as betting on the last card, no matter what we have seen, because all cards we have not seen are in random order.

Recall problem 16.9: \mathcal{X} is uniformly integrable iff $\sup\{E[|X|]; X \in \mathcal{X}\} < \infty$ and $\forall \epsilon > 0 \exists \delta > 0 P(A) < \delta \Rightarrow \sup\{E[|X|; A]; X \in \mathcal{X}\} < \epsilon$.

Proposition. If \mathcal{X} is UI and \mathcal{G} is a collection of sub- σ -fields, then $\{E[X | \mathcal{G}]; X \in \mathcal{X}, \mathcal{G} \in \mathcal{G}\}$ is UI.

Example. $\mathcal{X} = \{Z\}$ for some $Z \in L^1(P)$ and \mathcal{G} is a filtration.

Proof. Let $M := \sup\{E[|X|]; X \in \mathcal{X}\}$. Given $\epsilon > 0$, choose α such that

$$\sup\{E[|X|; A]; X \in \mathcal{X}, P(A) \leq M/\alpha\} < \epsilon.$$

Now for $X \in \mathcal{X}$ and $\mathcal{G} \in \mathcal{G}$, we have

$$P[|E[X | \mathcal{G}]| \geq \alpha] \leq \alpha^{-1} E[|E[X | \mathcal{G}]|] \leq \alpha^{-1} E[E[|X| | \mathcal{G}]] = \alpha^{-1} E[|X|] \leq M/\alpha.$$

Hence, with $A := \{|E[X | \mathcal{G}]| \geq \alpha\}$, we have

$$E[|E[X | \mathcal{G}]|; A] \leq E[E[|X| | \mathcal{G}]\mathbf{1}_A] = E[E[|X|\mathbf{1}_A | \mathcal{G}]] = E[|X|\mathbf{1}_A] < \epsilon. \quad \blacksquare$$

Let $\langle \mathcal{F}_n \rangle$ be a filtration and τ a stopping time. What is the information up to time τ , \mathcal{F}_τ ? On the part of Ω where $\tau = n$, it is \mathcal{F}_n . Thus, we define

$$\mathcal{F}_\tau := \{A \in \mathcal{F}; \forall n A \cap [\tau = n] \in \mathcal{F}_n\}.$$

This is a σ -field and $\tau \in \mathcal{F}_\tau$.

Homework: Show that if $\langle X_n \rangle$ is adapted to $\langle \mathcal{F}_n \rangle$ and τ is a finite stopping time, then $X_\tau \in \mathcal{F}_\tau$.

Homework: Show that if τ_1 and τ_2 are finite stopping times, then $\mathcal{F}_{\tau_1 \wedge \tau_2} = \mathcal{F}_{\tau_1} \cap \mathcal{F}_{\tau_2}$.

Proposition. Let $\langle (X_n, \mathcal{F}_n) \rangle$ be a martingale (submartingale; supermartingale) and τ be a finite stopping time. Then for all n ,

$$E[X_n | \mathcal{F}_\tau] = X_{\tau \wedge n} \text{ a.s. } (\geq ; \leq)$$

Proof. It's sufficient to do the submartingale case. Note that $X_{\tau \wedge n} \in \mathcal{F}_\tau$. Since both sides of the desired inequality are in \mathcal{F}_τ , it suffices to show that for all $A \in \mathcal{F}_\tau$,

$$E[X_n; A] \geq E[X_{\tau \wedge n}; A].$$

This holds by the following calculation:

$$\begin{aligned} E[X_n \mathbf{1}_A] &= \sum_k E[X_n \underbrace{\mathbf{1}_A \mathbf{1}_{[\tau=k]}}_{\in \mathcal{F}_k}] \geq \sum_k E[X_{n \wedge k} \mathbf{1}_A \mathbf{1}_{[\tau=k]}] \\ &= \sum_k E[X_{n \wedge \tau} \mathbf{1}_A \mathbf{1}_{[\tau=k]}] = E[X_{n \wedge \tau} \mathbf{1}_A]. \quad \blacksquare \end{aligned}$$

Homework: Show that if $\langle (X_n, \mathcal{F}_n) \rangle$ is a submartingale and τ is a stopping time, then $\langle (X_{\tau \wedge n}, \mathcal{F}_n) \rangle$ is a submartingale.

Corollary. If $\langle X_n \rangle$ is a UI martingale and τ is a finite stopping time, then $\{X_{\tau \wedge n}\}$ is UI.

Proof. Apply the preceding two propositions. \blacksquare

Corollary. If $\langle X_n \rangle$ is a UI submartingale and τ is a finite stopping time, then $\{X_{\tau \wedge n}\}$ is UI.

Proof. Since $|X_{\tau \wedge n}| \leq |X_\tau| + |X_n|$, it suffices to show that $E[|X_\tau|] < \infty$. Now $\langle X_n^+ \rangle$ is a submartingale, so by the proposition, $E[X_{\tau \wedge n}^+] \leq E[X_n^+]$. Also, $E[X_0] \leq E[X_{\tau \wedge n}]$ by the corollary to Wald's equation. Hence,

$$E[|X_{\tau \wedge n}|] = 2E[X_{\tau \wedge n}^+] - E[X_{\tau \wedge n}] \leq 2E[X_n^+] - E[X_0] \leq 2E[|X_n|] - E[X_0].$$

Fatou's lemma gives the result. \blacksquare

Homework: 35.7.

We'll now usually just write our results for the submartingale case.

Proposition. If $\langle X_n \rangle$ is a submartingale and $\tau_1 \leq \tau_2 \leq j \in \mathbb{N}$ are bounded stopping times, then $E[X_{\tau_1}] \leq E[X_{\tau_2}]$.

Proof. Write $Z_n := X_n - X_{n-1}$. Then $X_n = X_0 + \sum_{k=1}^n Z_k$, so

$$E[X_{\tau_2} - X_{\tau_1}] = E\left[\sum_{k=\tau_1+1}^{\tau_2} Z_k\right] = E\left[\sum_{k=1}^j Z_k \mathbf{1}_{[\tau_1 < k \leq \tau_2]}\right] \geq 0$$

since $[\tau_1 < k \leq \tau_2] \in \mathcal{F}_{k-1}$. \blacksquare

The Optional Sampling Theorem. *If $\tau_1 \leq \tau_2$ are finite stopping times and $\langle (X_{\tau_2 \wedge n}, \mathcal{F}_n) \rangle$ is a UI submartingale, then*

$$X_{\tau_1} \leq E[X_{\tau_2} \mid \mathcal{F}_{\tau_1}] \text{ a.s.}$$

Note that by a preceding homework and corollary, $\langle (X_{\tau_2 \wedge n}, \mathcal{F}_n) \rangle$ is a UI submartingale if $\langle (X_n, \mathcal{F}_n) \rangle$ is a UI submartingale.

Proof. First, we show that $E[X_{\tau_1}] \leq E[X_{\tau_2}]$. Fix $j \in \mathbb{N}$. Then $\tau_1 \wedge j \leq \tau_2 \wedge j \leq j$ are bounded stopping times, whence the preceding proposition gives

$$E[X_{\tau_1 \wedge j}] = E[X_{\tau_2 \wedge (\tau_1 \wedge j)}] \leq E[X_{\tau_2 \wedge (\tau_2 \wedge j)}] = E[X_{\tau_2 \wedge j}].$$

Now $E[X_{\tau_2 \wedge j}] \rightarrow E[X_{\tau_2}]$ since $\{X_{\tau_2 \wedge j}\}$ is UI. Furthermore, by the preceding corollary, $\{X_{\tau_1 \wedge j}\} = \{X_{\tau_2 \wedge (\tau_1 \wedge j)}\}$ is UI, whence $E[X_{\tau_1 \wedge j}] \rightarrow E[X_{\tau_1}]$. Therefore $E[X_{\tau_1}] \leq E[X_{\tau_2}]$.

To show the finer conclusion desired, it suffices to show that $\forall A \in \mathcal{F}_{\tau_1}$ $E[X_{\tau_1}; A] \leq E[X_{\tau_2}; A]$. By the homework, $\mathcal{F}_{\tau_1} = \mathcal{F}_{\tau_1 \wedge \tau_2} \subseteq \mathcal{F}_{\tau_2}$. Given $A \in \mathcal{F}_{\tau_1}$, set $\tau := \tau_1 \mathbf{1}_A + \tau_2 \mathbf{1}_{A^c}$. This is a stopping time since for all n ,

$$[\tau \leq n] = \underbrace{([\tau_1 \leq n] \cap A)}_{\in \mathcal{F}_n} \cup \underbrace{([\tau_2 \leq n] \cap A^c)}_{\in \mathcal{F}_n}.$$

Also, $\tau \leq \tau_2$, so by the result proved in the first paragraph, we have

$$0 \leq E[X_{\tau_2} - X_{\tau}] = E[X_{\tau_2} - X_{\tau_1} \mathbf{1}_A - X_{\tau_2} \mathbf{1}_{A^c}] = E[X_{\tau_2} \mathbf{1}_A - X_{\tau_1} \mathbf{1}_A]. \quad \blacksquare$$

When does a submartingale converge a.s.? Note that the heuristic that it tends to increase suggests that it always has an extended-real limit. However, this is not so, even if it is a.s. bounded:

Example. Let $\langle x_n \rangle$ be any sequence in \mathbb{R} . Then there is a martingale $\langle X_n \rangle$ that behaves like $\langle x_n \rangle$, i.e., $X_n - x_n$ is eventually constant a.s. (stabilizes a.s.). To see this, let $z_n := x_n - x_{n-1}$ and $\langle Z_n \rangle$ be independent, mean-0 random variables with $\sum_n P[Z_n \neq z_n] < \infty$. Set $X_n := \sum_{k=1}^n Z_k$. Then $\langle X_n \rangle$ is a martingale and by the Borel–Cantelli lemma, $X_n - x_n$ stabilizes a.s.

Yet a mild hypothesis assures convergence:

Theorem 35.5 (Doob’s Martingale Convergence Theorem). *Let $\langle X_n \rangle$ be a submartingale (or supermartingale) with $K := \sup_n E[|X_n|] < \infty$. Then $X := \lim_{n \rightarrow \infty} X_n$ exists a.s. and $E[|X|] \leq K$.*

To prove this, we consider the number of upcrossings of intervals $[\alpha, \beta]$: Given X_1, \dots, X_n , set $\sigma_1 := \min\{j \in [1, n]; X_j \leq \alpha \text{ or } j = n\}$, $\tau_1 := \min\{j \in (\sigma_1, n]; X_j \geq \beta \text{ or } j = n\}$, $\sigma_2 := \min\{j \in (\tau_1, n]; X_j \leq \alpha \text{ or } j = n\}$, $\tau_2 := \min\{j \in (\sigma_2, n]; X_j \geq \beta \text{ or } j = n\}$, etc. These are clearly stopping times. The upcrossings are the intervals $[\sigma_i, \tau_i]$ for which $X_{\sigma_i} \leq \alpha$ and $X_{\tau_i} \geq \beta$. Note that each upcrossing is responsible for an increase of $\geq \beta - \alpha$.

Theorem 35.4 (Doob's Upcrossing Inequality). *If $\langle X_1, \dots, X_n \rangle$ is a submartingale and $\alpha < \beta$, then the number U of upcrossings of $[\alpha, \beta]$ satisfies*

$$E[U] \leq \frac{E[(X_n - \alpha)^+]}{\beta - \alpha}.$$

Proof. Since $\langle (X_j - \alpha)^+ \rangle$ is a submartingale, we may take $\langle X_j \rangle$ to be nonnegative and $\alpha = 0$. (Note that U doesn't change.) Because all $X_j \geq 0$, we have

$$X_n = X_{\tau_n} \geq X_{\tau_n} - X_{\sigma_1} = \sum_{k=1}^n (X_{\tau_k} - X_{\sigma_k}) + \sum_{k=2}^n (X_{\sigma_k} - X_{\tau_{k-1}}).$$

Since $n \geq \sigma_k \geq \tau_{k-1}$, the proposition (on optional sampling) gives $E[X_{\sigma_k} - X_{\tau_{k-1}}] \geq 0$. Also, $X_{\tau_k} - X_{\sigma_k} \geq \beta$ for each of the U upcrossings and is ≥ 0 otherwise. Hence $\sum_{k=1}^n (X_{\tau_k} - X_{\sigma_k}) \geq \beta U$, so $E[X_n] \geq E[\beta U]$, i.e., $E[U] \leq E[X_n]/\beta$. ■

Proof of Theorem 35.5. To show that $\lim X_n$ exists a.s. in $[-\infty, \infty]$, it suffices to show that the number of upcrossings of each $[\alpha, \beta]$ is finite a.s. for $\alpha, \beta \in \mathbb{Q}$ with $\alpha < \beta$. Let U_n be the number of upcrossings of $[\alpha, \beta]$ by time n . Then

$$E[U_n] \leq \frac{E[(X_n - \alpha)^+]}{\beta - \alpha} \leq \frac{E[|X_n|] + |\alpha|}{\beta - \alpha} \leq \frac{K + |\alpha|}{\beta - \alpha}.$$

Since $U_n \uparrow$, the MCT shows that $\lim U_n < \infty$ a.s., as desired.

Let $X := \lim X_n$. Then $|X| = \lim |X_n|$ and Fatou's lemma gives $E[|X|] \leq K$, so $|X| < \infty$ a.s. ■

Corollary. *If $\langle X_n \rangle$ is a nonnegative supermartingale, then $X := \lim_{n \rightarrow \infty} X_n$ exists a.s. and $E[X] \leq E[X_1]$.* ■

Homework: Let $\langle Y_n \rangle$ be IID with law $(\delta_1 + \delta_{-1})/2$.

- Find $B_n \in \sigma(Y_1, \dots, Y_n)$ such that $P(B_n) \rightarrow 0$ and $P(B_n \text{ i.o.}) = 1$.
- Define $X_0 := 0$, $X_{n+1} := X_n(1 + Y_{n+1}) + \mathbf{1}_{B_n} Y_{n+1}$ for $n \geq 0$. Show that $\langle X_n \rangle$ is a martingale.
- Show that $X_n \rightarrow 0$ in probability.
- Show that $P[X_n \text{ converges}] = 0$.

It is very important to know when $E[\lim X_n] = \lim E[X_n]$. Recall from the corollary to Theorem 16.14 that if $X_n \rightarrow X$ a.s., then $\{X_n\}$ is UI iff $X_n \rightarrow X$ in L^1 (in which case $E[X_n] \rightarrow E[X]$). Thus:

Proposition. *If $\langle X_n \rangle$ is a submartingale, then the following are equivalent:*

- $\{X_n\}$ is UI.
- $\lim X_n$ exists a.s. and in L^1 .
- $\lim E[X_n]$ exists in L^1 .

Proof. (ii) \Rightarrow (iii) is trivial. (ii) \Rightarrow (i) by the general corollary and (i) \Rightarrow (ii), (iii) \Rightarrow (ii) by Theorem 35.5. ■

However, there is a martingale $\langle X_n \rangle$ that is bounded in L^1 with $E[X_n] = 0 = E[\lim X_n]$, yet $E[|X_n|] \not\rightarrow 0$. For example, let $X_1 = X_2 = X_3 = \dots \sim (\delta_1 + \delta_{-1})/2$. It is also possible to have an example with $X_n \rightarrow 0$ a.s.; see the later extra credit problem.

We get an additional equivalence for martingales:

Theorem. *If $\langle (X_n, \mathcal{F}_n) \rangle$ is a martingale, then $X_n \rightarrow X$ in L^1 iff $X_n = E[X | \mathcal{F}_n]$ and X has a version in $\sigma(\bigcup_n \mathcal{F}_n)$.*

Write $\mathcal{F}_\infty := \sigma(\bigcup_n \mathcal{F}_n)$; we also write this as $\mathcal{F}_n \uparrow \mathcal{F}_\infty$ and as $\mathcal{F}_\infty = \bigvee_n \mathcal{F}_n$. We say that $\langle (X_n, \mathcal{F}_n) \rangle$ is **closed** by (X, \mathcal{F}_∞) .

Proof. \Rightarrow : Because $X_n \xrightarrow{L^1} X$, we have $X_n \mathbf{1}_A \xrightarrow{L^1} X \mathbf{1}_A$ for all A . Take $A \in \mathcal{F}_m$. Then

$$E[X; A] = \lim E[X_n; A] = E[X_m; A].$$

That is, X_m is a version of $E[X | \mathcal{F}_m]$. Since all $X_n \in \mathcal{F}_\infty$, also $X \in \mathcal{F}_\infty$.

\Leftarrow : We have seen that $\{X_n\}$ is UI, so let $X_\infty := \limsup X_n$. We must show that $X = X_\infty$. Since $\bigcup \mathcal{F}_n$ is a π -system, it suffices to check that $\forall A \in \bigcup \mathcal{F}_n$ $E[X; A] = E[X_\infty; A]$. Indeed, if $A \in \mathcal{F}_n$, then $E[X; A] = E[X_n; A] = E[X_\infty; A]$, as we have just seen in the other part of the proof. ■

Theorem 35.6. *If $\mathcal{F}_n \uparrow \mathcal{F}_\infty$ and $Z \in L^1(P)$, then $E[Z | \mathcal{F}_n] \rightarrow E[Z | \mathcal{F}_\infty]$ a.s.*

Proof. Apply the preceding to $X := E[Z | \mathcal{F}_\infty]$. ■

Corollary (Lévy's 0-1 Law). *If $\mathcal{F}_n \uparrow \mathcal{F}_\infty$ and $A \in \mathcal{F}_\infty$, then $P[A | \mathcal{F}_n] \rightarrow \mathbf{1}_A$ a.s.* ■

Corollary (Kolmogorov's 0-1 Law). *If $\{\mathcal{F}_n\}$ are independent, then the tail σ -field of $\{\mathcal{F}_n\}$ is trivial.*

Proof. Let $\mathcal{G}_n := \sigma(\mathcal{F}_1, \dots, \mathcal{F}_n)$. If A is in the tail σ -field, then $A \in \bigvee_n \mathcal{G}_n$, so $\mathbf{1}_A = \lim P(A | \mathcal{G}_n) = \lim P(A)$ since A is independent of \mathcal{G}_n . Thus, $P(A) \in \{0, 1\}$. ■

Homework: 35.4 (In fact, show that if $\exists \epsilon > 0$ $\sum_n P[Y_n > 1 + \epsilon] = \infty$, then $X = 0$ a.s.)

Homework: Let Z_n be IID, mean-0 random variables. Let $\theta \in L^1$ and $Y_n := Z_n - \theta$. Show that $E[\theta | Y_1, \dots, Y_n] \rightarrow \theta$ a.s.

Homework: (Pólya's urn) An urn contains 1 red ball and 1 green ball. Choose one ball at random and replace it with another of the same color. Repeat indefinitely. Show that the fraction of red balls in the urn has an a.s. limit. Extra credit: identify the limit distribution.

Extra Credit Homework: Give a martingale $\langle X_n \rangle$ such that $\sup E[|X_n|] < \infty$, for all k , $E[X_k] = 0 = \lim X_n$ a.s., but $E[|X_n|] \not\rightarrow 0$.

Recall Example 35.2: Let ν be a finite, signed measure on (Ω, P) and $\langle \mathcal{F}_n \rangle$ be a filtration. Suppose that $\nu \upharpoonright \mathcal{F}_n \ll P \upharpoonright \mathcal{F}_n$ for all n . Let $X_n := d(\nu \upharpoonright \mathcal{F}_n) / d(P \upharpoonright \mathcal{F}_n)$. We saw that $\langle (X_n, \mathcal{F}_n) \rangle$ is a martingale. Let $X := \limsup X_n$. Write $\mathcal{F}_\infty := \bigvee_{n \geq 1} \mathcal{F}_n$.

Theorem 35.7. (i) If $\nu \upharpoonright \mathcal{F}_\infty \ll P \upharpoonright \mathcal{F}_\infty$, then $X := d(\nu \upharpoonright \mathcal{F}_\infty) / d(P \upharpoonright \mathcal{F}_\infty)$.
(ii) If $\nu \upharpoonright \mathcal{F}_\infty \perp P \upharpoonright \mathcal{F}_\infty$, then $P[X = 0] = 1$.

This is not hard to deduce from Theorems 35.5 and 35.6—read p. 470.

Also recall Example 35.3, the specialization to P being Lebesgue measure on $(0, 1]$, $\mathcal{F}_n := \sigma(I_k^{(n)}; 0 \leq k < 2^n)$, where $I_k^{(n)} := (k2^{-n}, (k+1)2^{-n}]$. We always have that $\nu \upharpoonright \mathcal{F}_n \ll P \upharpoonright \mathcal{F}_n$ and

$$X_n = \sum_{k=0}^{2^n-1} \mathbf{1}_{I_k^{(n)}} \nu(I_k^{(n)}) 2^n.$$

We get that $X_n \rightarrow F'$ P -a.s., where $F(\omega) := \nu((0, \omega])$ is function of bounded variation: see Example 35.10 on p. 471.

Theorem 35.3 (Doob's Maximal Inequality). If $\langle X_1, \dots, X_n \rangle$ is a submartingale and $\alpha > 0$, then

$$P[\max_{i \leq n} X_i \geq \alpha] \leq \alpha^{-1} E[X_n; \max_{i \leq n} X_i \geq \alpha] \leq \alpha^{-1} E[|X_n|].$$

This extends Kolmogorov's inequality (Theorem 22.4), since if S_i are the partial sums of independent, mean-0 random variables, then $\langle S_i^2 \rangle$ is a submartingale.

Proof. Let $\tau_2 := n$ and $\tau_1 := \min\{i \leq n; X_i \geq \alpha \text{ or } i = n\}$. Let $M_k := \max_{i \leq k} X_i$. Then for all k ,

$$[M_n \geq \alpha] \cap [\tau_1 \leq k] = [M_k \geq \alpha] \in \mathcal{F}_k,$$

i.e., $[M_n \geq \alpha] \in \mathcal{F}_{\tau_1}$. Thus, the optional sampling theorem gives

$$\alpha P[M_n \geq \alpha] \leq E[X_{\tau_1}; M_n \geq \alpha] \leq E[X_{\tau_2}; M_n \geq \alpha] \leq E[|X_n|]. \quad \blacksquare$$

Homework: Show that if τ is a finite stopping time and $\langle (X_{\tau \wedge n}, \mathcal{F}_n) \rangle$ is a UI submartingale, then for all $\alpha > 0$,

$$P[\sup_n X_{\tau \wedge n} \geq \alpha] \leq \alpha^{-1} E[X_\tau; \sup_n X_{\tau \wedge n} \geq \alpha] \leq \alpha^{-1} E[|X_\tau|].$$

Doob's L^p Inequality. If $\langle X_1, \dots, X_n \rangle$ is a nonnegative submartingale and $p > 1$, then

$$\|\max_{i \leq n} X_i\|_p \leq \frac{p}{p-1} \|X_n\|_p.$$

Proof. Because of Theorem 35.3, it suffices to show that for any random variables $X, Y \geq 0$, if

$$\forall \alpha > 0 \quad P[X \geq \alpha] \leq \alpha^{-1} E[Y; X \geq \alpha],$$

then

$$\forall p > 1 \quad \|X\|_p \leq \frac{p}{p-1} \|Y\|_p.$$

Now

$$\begin{aligned}
\|X\|_p^p &= E[X^p] = \int_0^\infty P[X^p > s] ds = \int_0^\infty P[X > s^{1/p}] ds \\
&= \int_0^\infty pt^{p-1} P[X > t] dt \leq \int_0^\infty pt^{p-2} \int_\Omega Y \mathbf{1}_{[X \geq t]} dP dt \\
&= \int_\Omega pY \int_0^\infty t^{p-2} \mathbf{1}_{[X \geq t]} dt dP = \int_\Omega \frac{p}{p-1} Y X^{p-1} dP \\
&\leq \frac{p}{p-1} \|Y\|_p \|X^{p-1}\|_q,
\end{aligned}$$

where $q^{-1} = 1 - p^{-1} = (p-1)/p$. That is,

$$E[X^p] \leq \frac{p}{p-1} E[X^p]^{1/q} \|Y\|_p.$$

If $\|X\|_p < \infty$, then this gives the desired inequality. If not, then note that we may replace X by $X \wedge K$ in the hypothesis: $[X \wedge K \geq \alpha] = [X \geq \alpha]$ for $\alpha \leq K$ and $[X \wedge K \geq \alpha] = \emptyset$ for $\alpha > K$. This gives

$$\|X \wedge K\|_p \leq \frac{p}{p-1} \|Y\|_p,$$

and now we may let $K \rightarrow \infty$. ■

Doob's L^p Convergence Theorem. *If $\langle X_n \rangle$ is a martingale or nonnegative submartingale, $p > 1$, and $\sup_n E[|X_n|^p] < \infty$, then $\exists X$ such that $X_n \rightarrow X$ a.s. and in L^p .*

Problem 35.4 gave a counter-example for $p = 1$. Also, doubling bets on fair games gives a martingale counter-example.

Proof. Since $\|X_n\|_1 \leq \|X_n\|_p$, Theorem 35.5 gives $X_n \rightarrow X$ a.s. Since $|X_n|$ is a submartingale, Doob's L^p inequality and the MCT (or Fatou's lemma) give $\sup |X_n| \in L^p$. Since $|X_n - X|^p \leq 2^p \sup |X_n|^p \in L^1$, the LDCT gives $X_n \rightarrow X$ in L^p . ■

Homework: Let $p > 1$, A_n be independent with $P(A_n) = \left(\frac{n}{n+1}\right)^p$ for $n \geq 1$, and $X_n := n \mathbf{1}_{\bigcap_{k=1}^{n-1} A_k}$. Show that $\langle X_n \rangle$ is a supermartingale and $E[X_n^p] = 1$, but $X_n \rightarrow 0$ a.s. Thus, Doob's L^p convergence theorem does not extend to nonnegative supermartingales.

A sub/super/martingale indexed by the negative integers is called a **backwards** or **reversed** sub/super/martingale. This adjective refers only to the index set! A filtration can be defined with respect to any poset.

Example. Let $\langle Y_n \rangle$ be IID in $L^1(P)$, $S_n := \sum_{k=1}^n Y_k$, $X_{-n} := S_n/n$, and $\mathcal{F}_{-n} := \sigma(S_n, S_{n+1}, \dots)$. Then for $1 \leq k \leq n$, $E[Y_1 | \mathcal{F}_{-n}] = E[Y_k | \mathcal{F}_{-n}]$ by symmetry.

To prove this rigorously, we show that $E[Y_1; A] = E[Y_k; A]$ for all $A \in \mathcal{F}_{-n}$. Both sides here depend only on the joint distribution of Y_1 (or Y_k) and S_n, S_{n+1}, \dots : see the proposition below. So assume that P is a product measure, the law of $\langle Y_n \rangle$ on \mathbb{R}^∞ . Then there is a bijection of the probability space interchanging Y_1 and Y_k but preserving P and S_m for all $m \geq n$. This proves the equality.

Since $\sum_{k=1}^n E[Y_k | \mathcal{F}_{-n}] = E[S_n | \mathcal{F}_{-n}] = S_n$, it follows that $E[Y_1 | \mathcal{F}_{-n}] = S_n/n = X_{-n}$. Thus, $\langle (X_{-n}, \mathcal{F}_{-n}) \rangle$ is a UI reversed martingale.

Proposition. Let $X: (\Omega, \mathcal{F}, P) \rightarrow (\Omega', \mathcal{F}')$ and $Y: (\Omega, \mathcal{F}, P) \rightarrow (\Omega'', \mathcal{F}'')$ be random variables and $f: (\Omega', \mathcal{F}') \rightarrow \mathbb{R}$ be Borel with $f(X) \in L^1(P)$. Then $E[f(X) | Y]$ depends only on f and the law of (X, Y) . In other words, for every probability measure μ on $\mathcal{F}' \times \mathcal{F}''$ such that $\int |f(x)| d\mu(x, y) < \infty$, there is a Borel $g_\mu: (\Omega'', \mathcal{F}'') \rightarrow \mathbb{R}$ such that $g_{(X, Y)_*P}(Y)$ is a version of $E[f(X) | Y]$.

Proof. Write $\bar{f}(x, y) := f(x)$ and $\bar{\mathcal{F}}'' := \{\Omega' \times A; A \in \mathcal{F}''\}$. Set g_μ to be the Radon–Nikodym derivative of $(\bar{f}\mu) \upharpoonright \bar{\mathcal{F}}''$ with respect to $\mu \upharpoonright \bar{\mathcal{F}}''$. A check of the definition of conditional expectation shows that this works: if $\nu := (X, Y)_*P$, then for all Borel B ,

$$\int_{[Y \in B]} g_\nu(Y) dP = \int_{[y \in B]} g_\nu(y) d\nu(x, y) = \int_{[y \in B]} f(x) d\nu(x, y) = \int_{[Y \in B]} f(X) dP. \quad \blacksquare$$

Note: every reversed martingale is UI.

Write $\mathcal{F}_{-\infty} := \bigcap_n \mathcal{F}_{-n}$.

Theorem 35.8. Let $\langle (X_{-n}, \mathcal{F}_{-n}) \rangle$ be a reversed submartingale. Then its limit, $X_{-\infty} := \lim_{n \rightarrow \infty} X_{-n}$, exists a.s. in $[-\infty, \infty)$ and $X_{-\infty} \leq E[X_{-1} | \mathcal{F}_{-\infty}]$ a.s. If $\langle X_{-n} \rangle$ is a reversed martingale, then $X_{-\infty} \in (-\infty, \infty)$ a.s. and $X_{-\infty} = E[X_{-1} | \mathcal{F}_{-\infty}]$ a.s.

Proof. Given $[\alpha, \beta]$, let U_n be the number of upcrossings of $[\alpha, \beta]$ by $X_{-n}, X_{-n+1}, \dots, X_{-1}$. Then the upcrossing lemma gives

$$E[U_n] \leq \frac{E[|X_{-1}|] + |\alpha|}{\beta - \alpha}.$$

Hence $\sup U_n < \infty$ a.s., so $X_{-\infty} := \lim X_{-n}$ exists a.s. in $[-\infty, \infty)$.

Given $a < 0$, $\langle (X_{-n} \vee a, \mathcal{F}_{-n}) \rangle$ is a reversed submartingale. We claim it is UI. The rest then follows since $\forall A \in \mathcal{F}_{-\infty} \forall n \geq 1 \ E[X_{-n} \vee a; A] \leq E[X_{-1} \vee a; A]$; taking $n \rightarrow \infty$ gives $E[X_{-\infty} \vee a; A] \leq E[X_{-1} \vee a; A]$. Now take $a \rightarrow -\infty$.

To prove UI, let $c > 0$. Then

$$cP[X_{-n} \vee a \geq c] \leq E[X_{-n} \vee a; X_{-n} \vee a \geq c] \leq E[X_{-1} \vee a; X_{-n} \vee a \geq c] \leq E[X_{-1}^+].$$

Therefore $P[X_{-n} \vee a \geq c] \rightarrow 0$ uniformly in n as $c \rightarrow \infty$. Comparing the middle terms in these inequalities gives UI. \blacksquare

Theorem 35.9. If $Z \in L^1(P)$ and $\langle \mathcal{F}_{-n} \rangle$ is a filtration, then $E[Z | \mathcal{F}_{-n}] \rightarrow E[Z | \mathcal{F}_{-\infty}]$ a.s.

Proof. $\langle (E[Z | \mathcal{F}_{-n}], \mathcal{F}_{-n}) \rangle$ is a reversed martingale, so its limit is $Z_{-\infty} = E[E[Z | \mathcal{F}_{-1}] | \mathcal{F}_{-\infty}] = E[Z | \mathcal{F}_{-\infty}]$. \blacksquare

Example. With the previous example, we have $S_n/n \rightarrow X_{-\infty} = E[X_{-1} | \mathcal{F}_{-\infty}] = E[Y_1 | \mathcal{F}_{-\infty}]$ a.s. and in L^1 . On the other hand, by Kolmogorov's 0-1 law, $\lim S_n/n$ is a constant a.s. Hence $X_{-\infty} = E[X_{-\infty}] = E[E[Y_1 | \mathcal{F}_{-\infty}]] = E[Y_1]$ a.s. This gives a new derivation of the SLLN. While it doesn't give Etemadi's version that assumed only pairwise independence, it doesn't need full independence either—exchangeability is enough, though we don't then get that the limit is constant a.s.

Here, we call $\langle X_1, \dots, X_n \rangle$ *exchangeable* if for all permutations π of $\langle 1, 2, \dots, n \rangle$, $\langle X_{\pi(1)}, \dots, X_{\pi(n)} \rangle$ has the same distribution. We call $\langle X_i; i \geq 1 \rangle$ *exchangeable* if for all n , $\langle X_1, \dots, X_n \rangle$ is exchangeable.

Suppose that all X_i are indicators. It is not hard to give examples of exchangeable $\langle X_1, \dots, X_n \rangle$; for example, pick k uniformly out of n to be 1. In fact, every exchangeable sequence $\langle X_1, \dots, X_n \rangle$ of indicators is a mixture of these examples, because if $S_n := \sum_{k=1}^n X_k$, then the law of $\langle X_1, \dots, X_n \rangle$ conditional on S_n is the uniform measure on S_n out of n . The analogue for infinite exchangeable sequences is that every example is a mixture of Bernoulli trials:

Theorem 35.10 (de Finetti's Theorem). *If $\langle X_i; i \geq 1 \rangle$ is an exchangeable sequence of indicators, then there is a random variable θ with values in $[0, 1]$ such that given θ , $\langle X_i \rangle$ is Bernoulli(θ), i.e., for all m and $u_i \in \{0, 1\}$, $P[X_1 = u_1, \dots, X_m = u_m \mid \theta] = \theta^k (1 - \theta)^{m-k}$ a.s., where $k := \sum_{i=1}^m u_i$.*

Proof. This is just a limit of the finite case: Define $\mathcal{F}_{-n} := \sigma(S_n, S_{n+1}, \dots)$. Fix $m \geq 1$, and consider $n \geq m$. Because $\langle X_1, \dots, X_n \rangle$ is exchangeable, so is its law conditional on \mathcal{F}_{-n} by the preceding proposition. Given $u_i \in \{0, 1\}$, write $A := [X_1 = u_1, \dots, X_m = u_m]$. Then with $k := \sum_{i=1}^m u_i$, we have a.s.

$$P(A \mid \mathcal{F}_{-n}) = \frac{S_n}{n} \cdot \frac{S_n - 1}{n - 1} \cdots \frac{S_n - k + 1}{n - k + 1} \cdot \frac{n - S_n}{n - k} \cdot \frac{n - S_n - 1}{n - k - 1} \cdots \frac{n - S_n - m + k + 1}{n - m + 1}.$$

Taking $n \rightarrow \infty$ and using Theorem 35.9 gives $P(A \mid \mathcal{F}_{-\infty}) = \theta^k (1 - \theta)^{m-k}$ a.s. Because this is θ -measurable, the result follows by the tower property, Theorem 34.4. ■

Homework: (Pólya's urn) An urn contains r red balls and g green balls. Choose one ball at random and replace it with another of the same color. Repeat indefinitely. Let X_n be the indicator that the n th ball is red. Show that $\langle X_n \rangle$ is exchangeable. Extra credit: identify the distribution of the random variable θ that gives the mixture, in this case, of de Finetti's theorem.

Homework: With the notation of the proof of de Finetti's theorem, show that $\mathcal{F}_{-\infty} = \sigma(\theta)$.

Homework: Let $\langle X_i; i \geq 1 \rangle$ be simple random variables, each taking values in a finite set, U .

- Suppose that $\langle X_1, \dots, X_n \rangle$ is exchangeable. Define $Z_n(u) := |\{i \leq n; X_i = u\}|$. Describe the law of $\langle X_1, \dots, X_n \rangle$ conditional on $Z_n := \langle Z_n(u); u \in U \rangle$.
- Suppose that $\langle X_i; i \geq 1 \rangle$ is exchangeable. Show that $Z := \lim_{n \rightarrow \infty} Z_n/n$ exists a.s. *Hint:* Let $\mathcal{F}_{-n} := \sigma(Z_n, X_{n+1}, X_{n+2}, \dots)$.
- With the assumptions and notation of part (b), show that given Z , $\langle X_i \rangle$ is IID, namely, for all m and $u_i \in U$, $P[X_1 = u_1, \dots, X_m = u_m \mid W] = \prod_{i=1}^m Z(u_i)$.

You may omit pp. 474–478.