Martingales

Manjunath Krishnapur

Contents

CHAPTER 1

Conditional probability and expectation

1. Conditional expectation

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $\mathcal{G} \subseteq \mathcal{F}$ be a sub sigma algebra of $\mathcal{F}$. Let $X : \Omega \to \mathbb{R}$ be a real-valued random variable. The goal is to find the closest $\mathcal{G}$-measurable random variable to X. For example, if $\mathcal{G} = \sigma(Z)$, then we want the function g so that $g(Z)$ is the closest to X. This problem of predicting one (perhaps not easily observable) random variable in terms of other (observable) random variables is one of the fundamental problems of Statistics.

To say anything, we must first decide on the sense of closeness.

Square-integrable case. If $\mathbb{E}[X^2] < \infty$, then we can ask for $\mathcal{G}$-measurable Y that minimizes $\mathbb{E}[|X-Y|^2]$. Why does is exist, and is it unique? For this, we move to equivalence classes of random variables and regard $W = L^2(\Omega, \mathcal{G}, \mathbb{P}|_{\mathcal{G}})$ as a closed subspace of the Hilbert space $H = L^2(\Omega, \mathcal{F}, \mathbb{P})$. Hilbert space theory tell us that there is a projection map $P : H \to W$ such that for any u, the unique closest vector in W is Pu. An equivalent way of stating this is that $Pu \in W$ and $\langle u, v\rangle = \langle Pu, v\rangle$ for all $v \in W$.

Thus, if $[X] \in L^2$ is the equivalence class containing X, and Y is any member of the equivalence class $P[X]$ (any two choices agree up to a $\mathbb{P}$-null $\mathcal{G}$-measurable set), then $\mathbb{E}[|X - Y|^2] \leqslant \mathbb{E}[|X - Z|^2]$ for any $\mathcal{G}$-measurable square integrable random variable Z. Equivalently,

(1) Y is $\mathcal{G}$-measurable and square integrable,

(2) $\mathbb{E}[XZ] = \mathbb{E}[YZ]$ for any $\mathcal{G}$-measurable, square integrable Z.

For later purpose, we note that the projection operator that occurs above has a special property (which does not even make sense for a general orthogonal projection in a Hilbert space).

Exercise 1

If $X \geqslant 0$ a.s. and $\mathbb{E}[X^2] < \infty$, show that $P_W[X] \geqslant 0$ a.s. [**Hint:** If $[Y] = P_W[X]$, then $\mathbb{E}[(X - Y_+)^2] \leqslant \mathbb{E}[(X-Y)^2]$ with equality if and only if $Y \geqslant 0$ a.s.]

Integrable case. Now suppose we only assume that $\mathbb{E}[|X|] < \infty$. It is tempting to consider the closed subspace $W = L^1(\Omega, \mathcal{G}, \mathbb{P}|_{\mathcal{G}})$ of the Banach space $H = L^1(\Omega, \mathcal{F}, \mathbb{P})$. But there is no good projection theory in L^1, hence we cannot repeat what we did for the square integrable case.

Example 1

On $([-1,1],\mathcal{B},\lambda)$ (where λ is the uniform probability measure), let $\mathcal{G} = \sigma(x \mapsto |x|)$. Then $\mathcal{G}$-measurable random variables are precisely measurable even functions. If $X(t) = t$ and $Y = f(|t|)$, then $\mathbb{E}[|X-Y|] = \int_{-1}^{1} |t-f(|t|)| = \int_0^1 (|t-f(t)|+|t+f(t)|)dt$. The integrand is at least $2t$, and equality is achieved if $-t < f(t) < t$. There are infinitely many measurable f satisfying this for all $t \in [0,1]$, hence there is no unique closest $\mathcal{G}$-measurable random variable to X.
Find example where existence fails

The correct analogy with the square integrable case is to think of a random variable X as acting on other random variables by $Y \mapsto \mathbb{E}[XY]$. If $X \in L^2$, the correct space of Y is L^2, while if $X \in L^1$, then the correct space of Y is L^∞. This leads us to the following definition.

Definition 1: Conditional expectation

A random variable $Y : \Omega \to \mathbb{R}$ is said to be a conditional expectation of X given $\mathcal{G}$ if (a) Y is $\mathcal{G}$-measurable and integrable, and (b) $\mathbb{E}[YZ] = \mathbb{E}[XZ]$ for all bounded $\mathcal{G}$-measurable random variables Z. Any such Y is denoted $\mathbb{E}[X \,\big|\, \mathcal{G}]$.

Some remarks are in order.

(1) It suffices to check the second condition for indicator variables. That is, $\mathbb{E}[Y\mathbf{1}_A] = \mathbb{E}[X\mathbf{1}_A]$ for all $A \in \mathcal{G}$. If this holds, then $\mathbb{E}[YZ] = \mathbb{E}[XZ]$ for simple $\mathcal{G}$-measurable Z. For general bounded $\mathcal{G}$-measurable Z, find simple functions Z_n such that $|Z_n| \leqslant |Z|$ and $Z_n \xrightarrow{a.s.} Z$. DCT applies on both sides of $\mathbb{E}[XZ_n] = \mathbb{E}[YZ_n]$ to show that $\mathbb{E}[XZ] = \mathbb{E}[YZ]$.

(2) The same reasoning as in the previous point shows that if $\mathbb{E}[|X|^p] < \infty$ for some $p > 0$, and $\mathbb{E}[X\mathbf{1}_A] = \mathbb{E}[Y\mathbf{1}_A]$ for all $A \in \mathcal{G}$, then $\mathbb{E}[XZ] = \mathbb{E}[YZ]$ for the larger class of $Z \in L^q(\Omega,\mathcal{G},\mathbb{P}|_{\mathcal{G}})$, where $\frac{1}{p}+\frac{1}{q} = 1$.

(3) Taking $p = 2$ in the previous point, we see that for square integrable X, the conditional expectation exists and is the closest $L^2(\Omega,\mathcal{G},\mathbb{P}|_{\mathcal{G}})$ random variable to X.

The main question is whether a conditional expectation exists for integrable X, and if it is unique. Yes and Yes. Before giving a proof, let us see some examples.

Example 2

Let $B, C \in \mathcal{F}$. Let $\mathcal{G} = \{\emptyset, B, B^c, \Omega\}$ and let $X = \mathbf{1}_C$. Since $\mathcal{G}$-measurable random variables must be constant on B and on B^c, we must take $Y = \alpha\mathbf{1}_B + \beta\mathbf{1}_{B^c}$. Writing the condition for equality of integrals of Y and X over B and over B^c, we get $\alpha\mathbb{P}(B) = \mathbb{P}(C \cap B)$, $\beta\mathbb{P}(B^c) = \mathbb{P}(C \cap B^c)$. It is easy to see that the equality of integrals over $\emptyset$ and over Ω also hold. Hence, the unique choice for conditional expectation of X given $\mathcal{G}$ is

$$Y(\omega) = \begin{cases} \mathbb{P}(C \cap B)/\mathbb{P}(B) & \text{if } \omega \in B, \\ \mathbb{P}(C \cap B^c)/\mathbb{P}(B^c) & \text{if } \omega \in B^c. \end{cases}$$

This agrees with the notion that we learned in basic probability classes. If we get to know that B happened, we update our probability of C to $\mathbb{P}(C \cap B)/\mathbb{P}(B)$ and if we get to know that B^c happened, we update it to $\mathbb{P}(C \cap B^c)/\mathbb{P}(B^c)$.

Exercise 2

Let $B_1, \ldots, B_n$ be a measurable partition of Ω. Assume that $\mathbb{P}(B_k) > 0$ for each k. Show that the unique conditional expectation of $\mathbf{1}_C$ given $\mathcal{G}$ is

$$\sum_{k=1}^{n} \frac{\mathbb{P}(C \cap B_k)}{\mathbb{P}(B_k)} \mathbf{1}_{B_k}.$$

Example 3

Suppose Z is $\mathbb{R}^d$-valued and (X, Z) has density $f(x, z)$ with respect to Lebesgue measure on $\mathbb{R} \times \mathbb{R}^d$. Let $\mathcal{G} = \sigma(Z)$. Then, we claim that a version of $\mathbb{E}[X \mid \mathcal{G}]$ is given by

$$Y(\omega) = \begin{cases} \frac{\int_{\mathbb{R}} x f(x, Z(\omega)) dx}{\int_{\mathbb{R}} f(x, Z(\omega)) dx} & \text{if the denominator is positive,} \\ 0 & \text{otherwise.} \end{cases}$$

As Y is a function of Z, it is $\mathcal{G}$-measurable. Here, it is clear that the set of ω for which $\int f(x, Z(\omega)) dx$ is zero is a $\mathcal{G}$-measurable set. Hence, Y defined above is $\mathcal{G}$-measurable. We leave it as an exercise to check that Y is a version of $\mathbb{E}[X \mid \mathcal{G}]$.

Uniqueness of conditional expectation: Suppose Y_1, Y_2 are two versions of $\mathbb{E}[X \mid \mathcal{G}]$. Then $\int_A Y_1 d\mathbb{P} = \int_A Y_2 d\mathbb{P}$ for all $A \in \mathcal{G}$, since both are equal to $\int_A X d\mathbb{P}$. Let $A = \{\omega : Y_1(\omega) > Y_2(\omega)\}$. Then the equality $\int_A (Y_1 - Y_2) d\mathbb{P} = 0$ can hold if and only if $\mathbb{P}(A) = 0$ (since the integrand is positive on A).

Similarly $\mathbb{P}\{Y_2 - Y_1 > 0\} = 0$. This, $Y_1 = Y_2$ a.s. (which means that $\{Y_1 \neq Y_2\}$ is $\mathcal{G}$-measurable and has zero probability under $\mathbb{P}$).

Thus, conditional expectation, if it exists, is unique up to almost sure equality.

Existence of conditional expectation: We give two proofs.

First approach - Radon-Nikodym theorem: Let $X \geqslant 0$ and $\mathbb{E}[X] < \infty$. Then consider the measure $\mathbb{Q} : \mathcal{G} \to [0, 1]$ defined by $\mathbb{Q}(A) = \int_A X d\mathbb{P}$ (we assumed non-negativity so that $\mathbb{Q}(A) \geqslant 0$ for all $A \in \mathcal{G}$). Further, $\mathbb{P}$ is a probability measure when restricted to $\mathcal{G}$ (we continue to denote it by $\mathbb{P}$). It is clear that if $A \in \mathcal{G}$ and $\mathbb{P}(A) = 0$, then $\mathbb{Q}(A) = 0$. In other words, $\mathbb{Q}$ is absolutely continuous to $\mathbb{P}$ on $(\Omega, \mathcal{G})$. By the Radon-Nikodym theorem, there exists $Y \in L^1(\Omega, \mathcal{G}, \mathbb{P})$ such that $\mathbb{Q}(A) = \int_A Y d\mathbb{P}$ for all $A \in \mathcal{G}$. Thus, Y is $\mathcal{G}$-measurable and $\int_A Y d\mathbb{P} = \int_A X d\mathbb{P}$ (the right side is $\mathbb{Q}(A)$). Thus, Y is a version of $\mathbb{E}[X \mid \mathcal{G}]$.

For a general integrable random variable X, let $X = X_+ - X_-$ and let Y_+ and Y_- be versions of $\mathbb{E}[X_+ \mid \mathcal{G}]$ and $\mathbb{E}[X_- \mid \mathcal{G}]$, respectively. Then $Y = Y_+ - Y_-$ is a version of $\mathbb{E}[X \mid \mathcal{G}]$.

Remark 1

Where did we use the integrability of X in all this? When $X \geqslant 0$, we did not! In other words, for a non-negative random variable X (even if not integrable), there exists a Y taking values in $\mathbb{R}_+ \cup \{+\infty\}$ such that Y is $\mathcal{G}$-measurable and $\int_A Y d\mathbb{P} = \int_A X d\mathbb{P}$. However, it is worth noting that if X is integrable, so is Y.

In the more general case, if there is a set of positive measure on which both Y_+ and Y_- are both infinite, then $Y_+ - Y_-$ is ill-defined on that set. Therefore, it is best to assume that $\mathbb{E}[|X|] < \infty$ so that Y_+ and Y_- are finite a.s.

Second approach - Approximation by square integrable random variables: Let $X \geqslant 0$ be an integrable random variable. Let $X_n = X \wedge n$ so that X_n are square integrable (in fact bounded) and $X_n \uparrow X$. Let Y_n be versions of $\mathbb{E}[X_n \mid \mathcal{G}]$, defined by the projections $P_W[X_n]$ as discussed earlier.

Now, $X_{n+1} - X_n \geqslant 0$, hence by the exercise above $P_W[X_{n+1} - X_n] \geqslant 0$ a.s., hence by the linearity of projection, $P_W[X_n] \leqslant P_W[X_{n+1}]$ a.s. In other words, $Y_n(\omega) \leqslant Y_{n+1}(\omega)$ for all $\omega \in \Omega_n$ where $\Omega_n \in \mathcal{G}$ is such that $\mathbb{P}(\Omega_n) = 1$. Then, $\Omega' := \cap_n \Omega_n$ is in $\mathcal{G}$ and has probability 1, and for $\omega \in \Omega'$, the sequence $Y_n(\omega)$ is non-decreasing.

Define $Y(\omega) = \lim_n Y_n(\omega)$ if $\omega \in \Omega'$ and $Y(\omega) = 0$ for $\omega \notin \Omega'$. Then Y is $\mathcal{G}$-measurable. Further, for any $A \in \mathcal{G}$, by MCT we see that $\int_A Y_n d\mathbb{P} \uparrow \int_A Y d\mathbb{P}$ and $\int_A X_n d\mathbb{P} \uparrow X d\mathbb{P}$. If $A \in \mathcal{G}$, then

$\int_A Y_n d\mathbb{P} = \int_A X_n d\mathbb{P}$. Thus, $\int_A Y d\mathbb{P} = \int_A X d\mathbb{P}$. This proves that Y is a conditional expectation of X given $\mathcal{G}$.

2. Properties of conditional expectation

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. We write $\mathcal{G}, \mathcal{G}_i$ for sub sigma algebras of $\mathcal{F}$ and X, X_i for integrable $\mathcal{F}$-measurable random variables on Ω.

Properties specific to conditional expectations:

(1) If X is $\mathcal{G}$-measurable, then $\mathbb{E}[X \mid \mathcal{G}] = X$ a.s. In particular, this is true if $\mathcal{G} = \mathcal{F}$.

(2) If X is independent of $\mathcal{G}$, then $\mathbb{E}[X \mid \mathcal{G}] = \mathbb{E}[X]$. In particular, this is true if $\mathcal{G} = \{\emptyset, \Omega\}$.

(3) Tower property: If $\mathcal{G}_1 \subseteq \mathcal{G}_2$, then $\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}_2] \mid \mathcal{G}_1] = \mathbb{E}[X \mid \mathcal{G}_1]$ a.s. In particular (taking $\mathcal{G} = \{\emptyset, \Omega\}$), we get $\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}]] = \mathbb{E}[X]$.

(4) $\mathcal{G}$-measurable random variables are like constants for conditional expectation: For any bounded $\mathcal{G}$-measurable random variable Z, we have $\mathbb{E}[XZ \mid \mathcal{G}] = Z\mathbb{E}[X \mid \mathcal{G}]$ a.s.

The first statement is easy, since X itself satisfies the properties required of the conditional expectation. The second is easy too, since the constant random variable $\mathbb{E}[X]$ is $\mathcal{G}$-measurable and for any $A \in \mathcal{G}$ we have $\mathbb{E}[X\mathbf{1}_A] = \mathbb{E}[X]\mathbb{E}[\mathbf{1}_A]$.

Property (4): If $Z = \mathbf{1}_B$ for some $B \in \mathcal{G}$, it is the very definition of conditional expectation. From there, deduce the property when Z is a simple random variable, a non-negative random variable, and a general integrable random variable (but we also need XZ to be integrable, which is implied if Z is bounded). We leave the details as an exercise.

Property (3): Now consider the tower property which is of enormous importance to us. But the proof is straightforward. Let $Y_1 = \mathbb{E}[X \mid \mathcal{G}_1]$ and $Y_2 = \mathbb{E}[X \mid \mathcal{G}_2]$. If $A \in \mathcal{G}_1$, then by definition, $\int_A Y_1 d\mathbb{P} = \int_A X d\mathbb{P}$. Further, $\int_A Y_2 d\mathbb{P} = \int_A X d\mathbb{P}$ since $A \in \mathcal{G}_2$ too. This shows that $\int_A Y_1 d\mathbb{P} = \int_A Y_2 d\mathbb{P}$ for all $A \in \mathcal{G}_1$. Further, Y_1 is $\mathcal{G}_1$-measurable. Hence, it follows that $Y_1 = \mathbb{E}[Y_2 \mid \mathcal{G}_1]$. This is what is claimed there.

Properties akin to expectation:

(1) Linearity: For $\alpha, \beta \in \mathbb{R}$, we have $\mathbb{E}[\alpha X_1 + \beta X_2 \mid \mathcal{G}] = \alpha\mathbb{E}[X_1 \mid \mathcal{G}] + \beta\mathbb{E}[X_2 \mid \mathcal{G}]$ a.s.

(2) Positivity: If $X \geqslant 0$ a.s., then $\mathbb{E}[X \mid \mathcal{G}] \geqslant 0$ a.s. and $\mathbb{E}[X \mid \mathcal{G}]$ is zero a.s. if and only if $X = 0$ a.s. As a corollary, if $X_1 \leqslant X_2$, then $\mathbb{E}[X_1 \mid \mathcal{G}] \leqslant \mathbb{E}[X_2 \mid \mathcal{G}]$.

(3) Conditional MCT: If $0 \leqslant X_n \uparrow X$ a.s., then $\mathbb{E}[X_n \mid \mathcal{G}] \uparrow \mathbb{E}[X \mid \mathcal{G}]$ a.s. Here either assume that X is integrable or make sense of the conclusion using Remark 1.

(4) Conditional Fatou's: Let $0 \leqslant X_n$. Then, $\mathbb{E}[\liminf X_n \mid \mathcal{G}] \leqslant \liminf \mathbb{E}[X_n \mid \mathcal{G}]$ a.s.

(5) Conditional DCT: Let $X_n \overset{a.s.}{\to} X$ and assume that $|X_n| \leqslant Y$ for some Y with finite expectation, then $\mathbb{E}[X_n \mid \mathcal{G}] \overset{a.s.}{\to} \mathbb{E}[X \mid \mathcal{G}]$.

(6) Conditional Jensen's inequality: If $\phi : \mathbb{R} \to \mathbb{R}$ is convex and X and $\phi(X)$ are integrable, then $\mathbb{E}[\phi(X) \mid \mathcal{G}] \geqslant \phi(\mathbb{E}[X \mid \mathcal{G}])$. In particular, of $\mathbb{E}[|X|^p] < \infty$ for some $p \geqslant 1$, then $\mathbb{E}[|X|^p \mid \mathcal{G}] \geqslant (\mathbb{E}[|X| \mid \mathcal{G}])^p$ of which the special cases $\mathbb{E}[|X| \mid \mathcal{G}] \geqslant |\mathbb{E}[X \mid \mathcal{G}]|$ and $\mathbb{E}[X^2 \mid \mathcal{G}] \geqslant (\mathbb{E}[X \mid \mathcal{G}])^2$ are particularly useful.

(7) Conditional Cauchy-Schwarz: If $\mathbb{E}[X^2], \mathbb{E}[Y^2] < \infty$, then $(\mathbb{E}[XY \mid \mathcal{G}])^2 \leqslant \mathbb{E}[X^2 \mid \mathcal{G}]\mathbb{E}[Y^2 \mid \mathcal{G}]$.

If conditional expectation was expectation with respect to a measure, then no proofs are needed! As we shall see, in nice situations but not always, there is a *random probability measure* $\mu_{\mathcal{G},\omega} : \mathcal{F} \to \mathbb{R}$, called the conditional probability of $\mathbb{P}$ given $\mathcal{G}$, such that $\mathbb{E}[X \mid \mathcal{G}](\omega) = \int X(\omega')\mu_{G,\omega}(d\omega')$. Thus, the usual properties of expectation, applied ω by ω to $\mu_{\mathcal{G},\omega}$, give the above stated properties.

But the assumption that conditional probability exists is not necessary for the above properties to hold. We give direct proofs of the above statements.

Proofs of properties of conditional expectations:

(1) Let Y_i be versions of $\mathbb{E}[X_i \mid \mathcal{G}]$. Then for $A \in \mathcal{G}$,

$$\begin{aligned}\int_A (\alpha Y_1 + \alpha Y_2) d\mathbb{P} &= \alpha \int_A Y_1 d\mathbb{P} + \beta \int_A Y_2 d\mathbb{P} \\ &= \alpha \int_A X_1 d\mathbb{P} + \beta \int_A X_2 d\mathbb{P} = \int_A (\alpha X_1 + \beta X_2) d\mathbb{P}\end{aligned}$$

which shows that $\alpha Y_1 + \beta Y_2$ is a version of $\mathbb{E}[\alpha X_1 + \beta X_2 \mid \mathcal{G}]$.

(2) This is clear if you go back to the proof of the existence of conditional expectation. Here is a more direct proof. Let Y be a version of $\mathbb{E}[X \mid \mathcal{G}]$ and set $A = \{Y < 0\} \in \mathcal{G}$. Then $\int_A Y d\mathbb{P} = \int_A X d\mathbb{P} \geqslant 0$ (as $X \geqslant 0$ a.s.) but $Y < 0$ on A, hence $\mathbb{P}(A) = 0$.

(3) Choose versions Y_n of $\mathbb{E}[X_n \mid \mathcal{G}]$. By redefining them on a zero probability set we may assume that $Y_1 \leqslant Y_2 \leqslant \ldots$, hence $Y = \lim Y_n$ exists. For any $A \in \mathcal{G}$, by the usual MCT we have $\mathbb{E}[Y_n \mathbf{1}_A] \uparrow \mathbb{E}[Y \mathbf{1}_A]$ and $\mathbb{E}[X_n \mathbf{1}_A] \uparrow \mathbb{E}[X \mathbf{1}_A]$. But also $\mathbb{E}[Y_n \mathbf{1}_A] = \mathbb{E}[X_n \mathbf{1}_A]$ for each n, hence $\mathbb{E}[Y \mathbf{1}_A] = \mathbb{E}[X \mathbf{1}_A]$. This is what was claimed.

(4) Since $Z := \liminf X_n$ is the increasing limit of $Z_n := \inf_{k \geqslant n} X_k$, for any $A \in \mathcal{G}$ by the conditional MCT we have $\mathbb{E}[Z_n \mid \mathcal{G}] \uparrow \mathbb{E}[Z \mid \mathcal{G}]$. But $X_n \geqslant Z_n$, hence $\mathbb{E}[Z_n \mid \mathcal{G}] \leqslant \mathbb{E}[X_n \mid \mathcal{G}]$. Putting these together, we see that $\liminf \mathbb{E}[X_n \mid \mathcal{G}] \geqslant \mathbb{E}[Z \mid \mathcal{G}]$ which is what we wanted.

(5) Apply the conditional Fatou's lemma to $Y - X_n$ and $Y + X_n$.

(6) Fix a version of $\mathbb{E}[X \mid \mathcal{G}]$ and $\mathbb{E}[\phi(X) \mid \mathcal{G}]$. Write $\phi(t) = \sup_{i \in I}(a_i + b_i t)$, where I is countable (e.g., supporting lines at all rationals). For each $i \in I$, we have $\mathbb{E}[\phi(X) \mid \mathcal{G}] \geqslant \mathbb{E}[a_i + b_i X \mid \mathcal{G}] = a_i + b_i \mathbb{E}[X \mid \mathcal{G}]$. Take supremum over $i \in I$ to get $\phi(\mathbb{E}[X \mid \mathcal{G}])$ on the right.

(7) Observe that $\mathbb{E}[(X - tY)^2 \mid \mathcal{G}] \geqslant 0$ a.s. for any $t \in \mathbb{R}$. Hence $\mathbb{E}[X^2 \mid \mathcal{G}] + t^2\mathbb{E}[Y^2 \mid \mathcal{G}] - 2t\mathbb{E}[XY \mid \mathcal{G}] \geqslant 0$ a.s. The set of zero measure indicated by "a.s." depends on t, but we can choose a single set of zero measure such that the above inequality holds for all $t \in \mathbb{Q}$, a.s. (for a fixed version of $\mathbb{E}[X \mid \mathcal{G}]$ and $\mathbb{E}[Y \mid \mathcal{G}]$). By continuity in t, it holds for all $t \in \mathbb{R}$, a.s. Optimize over t to get the conditional Cauchy-Schwarz.

3. Conditional probability

Definition 2: Regular conditional probability

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\mathcal{G}$ be a sub sigma algebra of $\mathcal{F}$. By regular conditional probability of $\mathbb{P}$ given $\mathcal{G}$, we mean any function $Q : \Omega \times \mathcal{F} \to [0, 1]$ such that

(1) For $\mathbb{P}$-a.e. $\omega \in \Omega$, the map $A \to Q(\omega, A)$ is a probability measure on $\mathcal{F}$.

(2) For each $A \in \mathcal{F}$, then map $\omega \to Q(\omega, A)$ is a version of $\mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]$.

The second condition of course means that for any $A \in \mathcal{F}$, the random variable $Q(\cdot, A)$ is $\mathcal{G}$-measurable and $\int_B Q(\omega, A) d\mathbb{P}(\omega) = \mathbb{P}(A \cap B)$ for all $B \in \mathcal{G}$.

It is clear that if it exists, it must be unique (in the sense that if Q' is another conditional probability, then $Q'(\omega, \cdot) = Q(\omega, \cdot)$ for a.e. $\omega[\mathbb{P}]$. However, unlike conditional expectation, conditional probability does not necessarily exist.

Suppose we define $Q(\omega, B)$ to be a version of $\mathbb{E}[\mathbf{1}_B \mid \mathcal{G}]$ for each $B \in \mathcal{F}$. Can we not simply prove that Q is a conditional probability? The second property is satisfied by definition. But for $Q(\omega, \cdot)$ to be a probability measure, we require that for any $B_n \uparrow B$ it must hold that $Q(\omega, B_n) \uparrow Q(\omega, B)$. Although the conditional MCT assures us that this happens for a.e. ω, the exceptional set where it fails depends on B and B_ns. As there are uncountably many such sequences (unless $\mathcal{F}$ is finite) it may well happen that for each ω, there is some sequence for which it fails (an uncountable union of zero probability sets may have probability one). A concrete example where it does not exist is given at the end of the section. This is why, the existence of conditional probability is not trivial. But it does exist in all cases of interest.

Theorem 1

Let M be a complete and separable metric space and let $\mathcal{B}_M$ be its Borel sigma algebra. Then, for any Borel probability measure $\mathbb{P}$ on $(M, \mathcal{B}_M)$ and any sub sigma algebra $\mathcal{G} \subseteq \mathcal{B}_M$, a regular conditional probability Q exists. It is unique in the sense that if Q' is another regular conditional probability, then $Q(\omega, \cdot) = Q'(\omega, \cdot)$ for $\mathbb{P}$-a.e. $\omega \in M$.

In probability theory we generally do not ask for any structure on the probability space, but in this theorem we do. It is really a matter of language, since we always restrict our random variables to take values in complete and separable metric spaces. Thus, another way to state the above theorem is that in a general probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{G} \subseteq \mathcal{F}$, regular conditional probability w.r.t. $\mathcal{G}$ may not exist on $\mathcal{F}$, but it will exist on any sub sigma algebra $\mathcal{F}' \subseteq \mathcal{F}$ that is generated by a random variable taking values in a complete, separable metric space. We state this as a theorem.

Theorem 2

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\mathcal{G} \subseteq \mathcal{F}$ be a sub sigma algebra. Let $\mathcal{F}' = \sigma(X)$ be any sub sigma algebra of $\mathcal{F}$ generated by a random variable $X : \Omega \mapsto M$ where M is a complete and separable metric space (endowed with its Borel sigma algebra). Then a regular conditional probability for $\mathcal{G}$ exists on $\mathcal{F}'$. That is, there is a $Q : \Omega \times \mathcal{F}' \mapsto [0, 1]$ such that $Q(\omega, \cdot)$ is a probability measure on $(\Omega, \mathcal{F}', \mathbb{P})$ for each $\omega \in \Omega$ and $Q(\cdot, A)$ is a version of $\mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]$ for each $A \in \mathcal{F}'$.

In this situation of $\mathcal{F}' = \sigma(X)$, one can push forward $Q(\omega, \cdot)$ to M and get probability measures $\nu_\omega = Q(\omega, \cdot) \circ X^{-1}$. Then ν_ω is called the regular conditional distribution of X given $\mathcal{G}$. For example, if $X = (X_1, \ldots, X_d)$ is $\mathbb{R}^d$-valued, then ν_ω is the measure with the distribution function

$$F_\omega(t_1, \ldots, t_d) = Q(\omega, \{X_1 \leqslant t_1, \ldots, X_m \leqslant t_d\}).$$

We shall prove this for the special case when $\Omega = \mathbb{R}$. The same proof can be easily written for $\Omega = \mathbb{R}^d$, with only minor notational complication. The general fact can be deduced from the fact that any complete separable metric space M is isomorphic as a measure space to a Borel subset of the real line[1]. Of course, it should be noted that the metric plays little role, if the topology is preserved by changing the metric, we may do so. For example, $(0, 1)$ is not complete with the usual metric, but we can endow it with a complete metric.

---

[1] For a proof, see Chapter 13 of Dudley's book *Real analysis and probability* or this paper by B. V. Rao and S. M. Srivastava.

Proof of Theorem 1 when $M = \mathbb{R}$. We start with a Borel probability measure $\mathbb{P}$ on $(\mathbb{R}, \mathcal{B}_{\mathbb{R}})$ and $\mathcal{G} \subseteq \mathcal{B}_{\mathbb{R}}$. For each $t \in \mathbb{Q}$, let Y_t be a version of $\mathbb{E}[\mathbf{1}_{(-\infty,t]} \mid \mathcal{G}]$. For any rational $t < t'$, we know that $Y_t(\omega) \leqslant Y_{t'}(\omega)$ for all $\omega \notin N_{t,t'}$ where $N_{t,t'}$ is a Borel set with $\mathbb{P}(N_{t,t'}) = 0$. Further, by the conditional MCT, there exists a Borel set N_* with $\mathbb{P}(N_*) = 0$ such that for $\omega \notin N_*$, we have $\lim_{t\to\infty} Y_t(\omega) = 1$ and $\lim_{t\to-\infty} Y_t = 0$ where the limits are taken through rationals only.

Let $N = \bigcup_{t,t'} N_{t,t'} \cup N_*$ so that $\mathbb{P}(N) = 0$ by countable additivity. For $\omega \notin N$, the function $t \to Y_t(\omega)$ from $\mathbb{Q}$ to $[0, 1]$ is non-decreasing and has limits 1 and 0 at $+\infty$ and $-\infty$, respectively. Now define $F : \Omega \times \mathbb{R} \to [0, 1]$ by

$$F(\omega, t) = \begin{cases} \inf\{Y_s(\omega) : s > t, s \in \mathbb{Q}\} & \text{if } \omega \notin N, \\ 0 & \text{if } \omega \in N. \end{cases}$$

By exercise 3 below, for any $\omega \notin N$, we see that $F(\omega, \cdot)$ is the CDF of some probability measure μ_ω on $\mathbb{R}$, provided $\omega \notin N$. Define $Q : \Omega \times \mathcal{B}_{\mathbb{R}} \to [0, 1]$ by $Q(\omega, A) = \mu_\omega(A)$. We claim that Q is a conditional probability of $\mathbb{P}$ given $\mathcal{G}$.

The first condition, that $Q(\omega, \cdot)$ be a probability measure on $\mathcal{B}_{\mathbb{R}}$ is satisfied by construction. We only need to prove the first condition. To this end, define

$$\mathcal{H} = \{A \in \mathcal{B}_{\mathbb{R}} : Q(\cdot, A) \text{ is a version of } \mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]\}.$$

First we claim that $\mathcal{H}$ is a λ-system. Indeed, if $A_n \uparrow A$ and $Q(\cdot, A_n)$ is a version of $\mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]$, then by the conditional MCT, $Q(\cdot, A)$ which is the increasing limit of $Q(\cdot, A_n)$, is a version of $\mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]$. Similarly, if $A \subseteq B$ and $Q(\cdot, A)$, $A(\cdot, B)$ are versions of $\mathbb{E}[\mathbf{1}_A \mid \mathcal{G}]$ and $\mathbb{E}[\mathbf{1}_B \mid \mathcal{G}]$, then by linearity of conditional expectations, $Q(\cdot, B \setminus A) = Q(\cdot, B) - Q(\cdot, A)$ is a version of $\mathbb{E}[\mathbf{1}_{B\setminus A} \mid \mathcal{G}]$.

Next, we claim that $\mathcal{H}$ contains the π-system of all intervals of the form $(-\infty, t]$ for some $t \in \mathbb{R}$. For fixed t, by definition $Q(\omega, (-\infty, t])$ is the decreasing limit of $Y_s(\omega) = \mathbb{E}[\mathbf{1}_{(-\infty,s]} \mid \mathcal{G}](\omega)$ as $s \downarrow t$, whenever $\omega \notin N$. By the conditional MCT it follows that $Q(\cdot, (-\infty, t])$ is a version of $\mathbb{E}[\mathbf{1}_{(-\infty,t]} \mid \mathcal{G}]$.

An application of the π-λ theorem shows that $\mathcal{H} = \mathcal{B}_{\mathbb{R}}$. This completes the proof. ∎

The following exercise was used in the proof.

Exercise 3

Let $f : \mathbb{Q} \to [0, 1]$ be a non-decreasing function such that $f(t)$ converges to 1 or 0 according as $t \to +\infty$ or $t \to -\infty$, respectively. Then define $F : \mathbb{R} \to [0, 1]$ by $F(t) = \inf\{f(q) : t < q \in \mathbb{Q}\}$. Show that F is a CDF of a probability measure.

Proof of Theorem 1 for general M. Let $\phi : M \to \mathbb{R}$ be a Borel isomorphism. That is ϕ is bijective and ϕ, ϕ^{-1} are both Borel measurable. We are given a probability measure $\mathbb{P}$ on $(M, \mathcal{B}_M)$

and a sigma algebra $\mathcal{G} \subseteq \mathcal{B}_M$. Let $\mathbb{P}' = \mathbb{P} \circ \phi^{-1}$ be its pushforward probability measure on $(\mathbb{R}, \mathcal{B}_\mathbb{R})$. Let $\mathcal{G}' = \{\phi(A) : A \in \mathcal{G}\}$, clearly a sub sigma algebra of $\mathcal{B}_\mathbb{R}$.

From the already proved case, we get $Q' : \mathbb{R} \times \mathcal{B}_M \to [0,1]$, a conditional probability of $\mathbb{P}'$ given $\mathcal{G}'$. Define $Q : M \times \mathcal{B}_M \to [0,1]$ by $Q(\omega, A) = Q'(\phi(\omega), \phi(A))$. Check that Q' is a conditional probability of $\mathbb{P}$ given $\mathcal{G}$. ■

Now we give an example where regular conditional probability does not exist.

Example 4

Consider $([0,1], \mathcal{B}, \lambda)$. Let $E \subseteq [0,1]$ be a non-measurable set with $\lambda^*(E) = 1 = \lambda^*(E^c)$. Let $\mathcal{F} = \sigma\{\mathcal{B}, E\} = \{(A \cap E) \sqcup (B \cap E^c) : A, B \in \mathcal{B}\}$. On $\mathcal{F}$, define a measure by

$$\mu((A \cap E) \sqcup (B \cap E^c)) = \frac{1}{2}\lambda(A) + \frac{1}{2}\lambda(B).$$

This is well defined because $A \cap E = A' \cap E$ implies that $A \Delta A' \subseteq E^c$ and measurable subsets of E^c have zero measure (why?). Thus, $\lambda(A) = \lambda(A')$. Similarly, $\lambda(B \cap E^c)$ does not depend on the choice of B. That μ is a measure is then easy to see.

Suppose regular conditional probability Q of μ w.r.t. $\mathcal{B}$ were to exist. Then $Q(\cdot, A) = \mathbf{1}_A$ a.e., for any $A \in \mathcal{B}$, in particular for $A = [0,x]$ for any $x \in [0,1]$. Take intersection over $x \in \mathbb{Q}$ to see that $Q(\omega, [0,x]) = \mathbf{1}_{[0,x]}(\omega)$ for all $x \in \mathbb{Q} \cap [0,1]$, for a.e.$\omega$. As $x \mapsto Q(\omega, [0,x])$ is a distribution function, we see that $Q(\omega, \cdot) = \delta_\omega$, for a.e. ω. But then, $Q(\omega, E) = 1$ if $\omega \in E$ and $Q(\omega, E) = 0$ if $\omega \in E^c$. But this means that $Q(\cdot, E) = \mathbf{1}_E$ is not $\mathcal{B}$ measurable, contradicting a requirement of conditional probability.

4. Relationship between conditional probability and conditional expectation

Let M be a complete and separable metric space (or in terms introduced earlier, a Polish space). Let $\mathbb{P}$ be a probability measure on $\mathcal{B}_M$ and let $\mathcal{G} \subseteq \mathcal{B}_M$ be a sub sigma algebra. Let Q be a regular conditional probability for $\mathbb{P}$ given $\mathcal{G}$ which exists, as discussed in the previous section. Let $X : M \to \mathbb{R}$ be a Borel measurable, integrable random variable. We defined the conditional expectation $\mathbb{E}[X \mid \mathcal{G}]$ in the first section. We now claim that the conditional expectation is actually the expectation with respect to the conditional probability measure. In other words, we claim that

$$\mathbb{E}[X \mid \mathcal{G}](\omega) = \int_M X(\omega')dQ_\omega(\omega') \tag{1}$$

where $Q_\omega(\cdot)$ is a convenient notation probability measure $Q(\omega, \cdot)$ and $dQ_\omega(\omega')$ means that we use Lebesgue integral with respect to the probability measure Q_ω (thus ω' is a dummy variable which is integrated out).

To show this, it suffices to argue that the right hand side of (1) is $\mathcal{G}$-measurable, integrable and that its integral over $A \in \mathcal{G}$ is equal to $\int_A X d\mathbb{P}$.

Firstly, let $X = \mathbf{1}_B$ for some $B \in \mathcal{B}_M$. Then, the right hand side is equal to $Q_\omega(B) = Q(\omega, B)$. By definition, this is a version of $\mathbb{E}[\mathbf{1}_B \mid \mathcal{G}]$. By linearity, we see that (1) is valid whenever X is a simple random variable.

If X is a non-negative random variable, then we can find simple random variables $X_n \geqslant 0$ that increase to X. For each n

$$\mathbb{E}[X_n \mid \mathcal{G}](\omega) = \int_M X_n(\omega') dQ_\omega(\omega') \ \text{ a.e.}\omega[\mathbb{P}].$$

The left side increases to $\mathbb{E}[X \mid \mathcal{G}]$ for a.e.. ω by the conditional MCT. For fixed $\omega \notin N$, the right side is an ordinary Lebesgue integral with respect to a probability measure Q_ω and hence the usual MCT shows that it increases to $\int_M X(\omega')dQ_\omega(\omega')$. Thus, we get (1) for non-negative random variables.

For a general integrable random variable X, write it as $X = X_+ - X_-$ and use (1) individually for $X_\pm$ and deduce the same for X.

Remark 2

Here we explain the reasons why we introduced conditional probability. In most books on martingales, only conditional expectation is introduced and is all that is needed. However, when conditional probability exists, conditional expectation becomes an actual expectation with respect to a probability measure. This makes it simpler to not have to prove many properties for conditional expectation as we shall see in the following section. Also, it is aesthetically pleasing and psychologically satisfying to know that conditional probability exists in most circumstances of interest.

A more important point is that, for discussing Markov processes (as we shall do when we discuss Brownian motion), conditional probability is the more natural language in which to speak. This is explained next.

4.1. Specifying measures by conditional probabilities. Think of the following familiar objects in probability: Markov chains, Branching processes, Pólya's urn scheme. The last one will be defined later in the course, but the point here is that in all three cases and many others, the verbal description is of the form: "do something, and depending on what the outcome is, do this or that, ...". The very description contains the idea of conditioning. We explain with the example of Markov chains.

Markov chains: A discrete time Markov chain on a state space S with a sigma algebra $\mathcal{S}$ is specified by two ingredients: A probability measure ν on S and a stochastic kernel $\kappa : S \times \mathcal{S} \mapsto [0, 1]$ such that $\kappa(\cdot, A)$ is measurable for all $A \in \mathcal{S}$ and $\kappa(x, \cdot)$ is a probability measure on $(S, \mathcal{S})$.

Then, a Markov chain with initial distribution ν and transition kernel κ is a collection of random variables $(X_n)_{n\geqslant 0}$ (on some probability space) such that $X_0 \sim \nu$ and the conditional distribution of X_{n+1} given $X_0, \ldots, X_n$ is $\kappa(X_n, \cdot)$.

Does a Markov chain exist? It is easy to answer yes by defining probability measures μ_n on $(S^n, \mathcal{S}^{\otimes n})$ by

$$\mu_n(A_0 \times A_1 \times \ldots \times A_{n-1}) = \int_{A_0} \ldots \int_{A_{n-1}} \nu(dx_0)\kappa(x_0, dx_1) \ldots \kappa(x_{n-3}, dx_{n-2})\kappa(x_{n-2}, dx_{n-1})$$

for $A_i \in \mathcal{S}$. This does define a probability measure on S^n, and further, these measures are consistent (the projection of μ_{n+1} to the first n co-ordinates gives μ_n). By Kolmogorov's consistency theorem, there is a measure μ on $(S^{\mathbb{N}}, \mathcal{S}^{\mathbb{N}})$ whose projection to the first n co-ordinates is μ_n. On $(S^{\mathbb{N}}, \mathcal{S}^{\mathbb{N}}, \mu)$, the co-ordinate random variables form a Markov chain with the given initial distribution and transition kernel.

Remark 3

In fact, the Markov property is not needed. Suppose a measure ν on $(S, \mathcal{S})$, and stochastic kernels $\kappa_n : S^n \times \mathcal{S} \mapsto [0, 1]$ are given. Define

$$\begin{aligned}&\mu_n(A_0 \times A_1 \times \ldots \times A_{n-1})\\ &= \int_{A_0} \ldots \int_{A_{n-1}} \nu(dx_0)\kappa_1(x_0, dx_1)\kappa_2(x_0, x_1, dx_2) \ldots \kappa_{n-1}(x_0, \ldots x_{n-2}, dx_{n-1}).\end{aligned}$$

This is again a consistent family of distributions, and we can construct random variables $(X_k)_{k\geqslant 0}$ on a suitable probability space so that $(X_0, \ldots, X_{n-1}) \sim \mu_n$.

In that sequence, $X_0 \sim \nu$ and the conditional distribution of X_n given $(X_0, \ldots, X_{n-1}) = (x_0, \ldots, x_{n-1})$ is given by $\kappa_n(x_0, \ldots, x_{n-1}, dx)$. Thus, a sequence of random variables may be described by giving the distribution of X_0, and for each $n \geqslant 1$ specifying the distribution of X_n given the previous X_is.

5. Cautionary tales on conditional probability

Even when knows all the definitions in and out, it is easy to make mistakes with conditional probability. Extreme caution is advocated! Practising some explicit computations also helps. Two points are to be noted.

Always condition on a sigma-algebra: Always specify the experiment first and then the outcome of the experiment. From the nature of the experiment, we can work out the way probabilities and expectations are to be updated for every possible outcome of the experiment. Then we apply that to the outcome that actually occurs.

For example, suppose I tell you that the bus I caught today morning had a 4-digit registration number of which three of the digits were equal to 7, and ask you for the chance that the remaining digit is also a 7. You should refuse to answer that question, as it is not specified what experiment was conducted. Did I note down the first three digits and report them to you, or did I look for how many 7s there were and reported that to you? It is not enough to know what I observed, but also what else I could have observed.

Conditioning on zero probability events: If (X, Y) have a joint density $f(x, y)$, then $\mathbb{E}[X \mid Y] = \frac{\int xf(x,Y)dx}{\int f(x,Y)dx}$. If we set $Y = 0$ in this formula, we get $\mathbb{E}[X \mid Y = 0]$. However, since conditional expectation is only defined up to zero measure sets, we can also set $\mathbb{E}[X \mid Y = 0]$ to be any other value. Why this particular formula?

The point is the same as asking for the value of a measurable function at a point - changing the value at a point is of no consequence for most purposes. However, there may be some justification for choosing a particular value. For example, if 0 is a Lebesgue point of f, it makes sense to take $f(0)$ to be $\lim\limits_{\varepsilon\downarrow 0} \frac{1}{2\varepsilon}\int_{-\varepsilon}^{\varepsilon} f(x)dx$. This is true in particular if f is continuous at 0.

Similarly, if we have to specify a particular value for $\mathbb{E}[X \mid Y = 0]$, it has to be approached via some limits, for example we may define it as $\lim_{\varepsilon\downarrow 0} \mathbb{E}[X \mid |Y| < \varepsilon]$, if the limit exists (and $\mathbb{P}(|Y| < \varepsilon) > 0$ for any $\varepsilon > 0$). For instance, if the joint density $f(x, y)$ is continuous, this will limit will be equal to the formula we got earlier, i.e., $\frac{\int xf(x,0)dx}{\int f(x,0)dx}$.

Example 5

Here is an example that illustrates both the above points. Let (U, V) be uniform on $[0, 1]^2$. Consider the diagonal line segment $L = \{(u, v) \in [0, 1]^2 : u = v\}$. What is the expected value of U conditioned on the event that it lies on L? This question is ambiguous as the experiment is not specified and the event that (U, V) lies on L has zero probability. Here are two possible interpretations. See Figure 5.

(1) The experiment measured $Y = U - V$ and the outcome was 0. In this case we are conditioning on $\sigma\{Y\}$. If we take limits of $\mathbb{E}[U \bigm| |Y| < \varepsilon]$ as $\varepsilon \downarrow 0$, we get $\mathbb{E}[X \bigm| Y = 0] = \frac{1}{2}$.

(2) The experiment measured $Z = U/V$ and the outcome was 1. In this case we are conditioning on $\sigma\{Z\}$. If we take limits of $\mathbb{E}[U \bigm| |Z| < \varepsilon]$ as $\varepsilon \downarrow 0$, we get $\mathbb{E}[X \bigm| Z = 1] = \frac{2}{3}$ (do the calculation!).

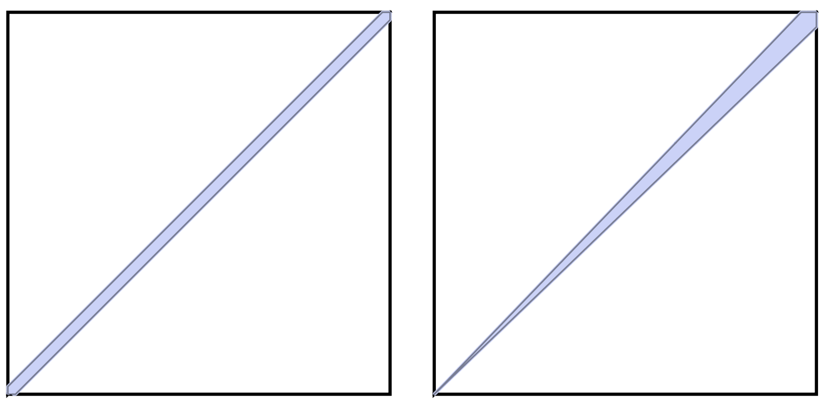

FIGURE 1. Two ways of conditioning a uniformly chosen point on the square to lie on the diagonal line. In the first case we condition on $|U - V| < \varepsilon$ and in the second case on $|\frac{U}{V} - 1| < \varepsilon$ (a value of $\varepsilon = 0.02$ was taken). Under that conditioning, the point is uniform in the shaded regions. In the first conditioning, U is almost uniform on $[0, 1]$ but not so in the second.

Conditional probability is the largest graveyard of mistakes in probability, hence it is better to keep these cautions in mind[2]. There are also other kinds of mistakes such as mistaking $\mathbb{P}(A \bigm| B)$

[2]There are many probability puzzles or "paradoxes", where the gist is some mistake in conditioning of the kind stated above (the famous *Tuesday brother problem* is an example of conditioning on an event without telling what the measurement is). A real-life example: No less a probabilist than Yuval Peres told us of a mistake he made once: In studying the random power series $f(z) = a_0 + a_1 z + a_2 z^2 + \dots$ where a_k are i.i.d. $N(0, 1)$, he got into a contradiction

for $\mathbb{P}(B \mid A)$, a standard example being the Bayes' paradox (given that you tested positive for a rare disease, what is the chance that you actually have the disease?) that we talked about in more basic courses. This same thing is called *representational fallacy* by Kahnemann and Tversky in their study of psychology of probability (a person who is known to be a doctor or a mathematician is given to be intelligent, systematic, introverted and absent-minded. What is the person more likely to be - doctor or mathematician?).

Here is a mind-bender for your entertainment[3]. If you like this, you can find many other such questions on Gil Kalai's blog under the category Test your intuition.

Elchanan Mossel's amazing dice paradox: A fair die is thrown repeatedly until a **6** turns up. Given that all the throws showed up even numbers, what is the expected number of throws (including the last throw)?

One intuitive answer is that it is like throwing a die with only three faces, $2, 4, 6$, till a 6 turns up, hence the number of throws is a Geometric random variable with mean 3. This is wrong!

6. Finer aspects of conditional probabilities (omit on first, second, third and fourth readings)

For those interested to go deeper into the subtleties of conditional probabilities, here are things I may expand on in the notes at some later time. You may safely skip all of this and increase the happiness in your life by a non-negative amount. The two aspects touched upon here are not of equal value. The first one is somewhat esoteric and I don't know if anyone cares about it anymore. The second one is fundamental to statistical physics, which forms a large part of probability theory.

6.1. Existence of regular conditional probabilities. Given $(\Omega, \mathcal{F}, \mathbb{P})$ and $\mathcal{G} \subseteq \mathcal{F}' \subseteq \mathcal{F}$, we want to know when a regular conditional probability $Q : \Omega \times \mathcal{F}' \mapsto [0, 1]$ for conditioning with respect to $\mathcal{G}$, exists. If $\mathcal{F}'$ can be taken equal to $\mathcal{F}$, all the better! The strongest statements I know of are from a paper of Jirina[4], of which I state the friendlier statement (for the more general one, see the paper). We need the notion of a *perfect measure*, introduced by Gnedenko and Kolmogorov.

A probability measure $\mathbb{P}$ on $(\Omega, \mathcal{F})$ is said to be perfect if for every $f : \Omega \mapsto \mathbb{R}$ that is Borel measurable, there is a Borel set $B \subseteq f(\Omega)$ such that $\mathbb{P} \circ f^{-1}(B^c) = 0$. Since forward images of

---

thinking that conditioning on $f(0) = 0$ is the same as conditioning the zero set of f to contain the origin! The two conditionings are different for reasons similar to the example given in the text.

[3] Thanks to P. Vasanth, Manan Bhatia and Gaurang Sriramanan for bringing this to my attention!

[4] Jirina, Conditional probabilities on σ-algebras with countable basis. **Czech. Math. J.** 4 (79), 372-380 (1954) [Selected Transitions in Mathematical Statistics and Probability, vol. 2 (Providence: American Mathematical Society, 1962), pp. 79-86]

measurable sets need not be measurable, it is not always true that $f(\Omega)$ is a Borel set, hence this definition which settles for something a little less.

Theorem: Assume that $\mathbb{P}$ is a perfect measure on $(\Omega, \mathcal{F})$, and that $\mathcal{F}$ is countably generated. Then for any $\mathcal{G} \subseteq \mathcal{F}$, a regular conditional probability exists on $\mathcal{F}$.

Corollary: Let $(X, \mathcal{B})$ be a metric space with its Borel sigma algebra. Assume that $\mathbb{P}$ is an inner regular probability Borel probability measure (i.e., $\mathbb{P}(A) = \sup\{\mathbb{P}(K) : K \subseteq A, \ K \text{ compact}\}$ for any $A \in \mathcal{B}$). Then, for any sub-sigma algebra $\mathcal{G} \subseteq \mathcal{B}$, a regular conditional probability exists $Q : X \times \mathcal{B} \mapsto [0, 1]$ exists.

One of the fundamental facts about complete, separable metric spaces is that every Borel probability measure is inner regular. Hence, our earlier theorem that regular conditional probabilities exist when working on Polish spaces is a consequence of the above theorem.

Perfect probabilties were introduced in the 1950s when the foundations of weak convergence laid down by Prokhorov were still fresh. Over decades, the emphasis in probability has shifted to studying interesting models coming from various applications, and the setting of complete separable metric spaces has proved adequate for all purposes. Modern books in probability often don't mention this concept (even Kallenberg does not!). A good reference (if you still want to wade into it) for all this and more is the highly educational book of K. R. Parthasarathy titled *Probability measures on metric spaces*.

6.2. Specifying a measure via conditional probabilities. We already saw that the joint distribution of a sequence of random variables $X_1, X_2, \ldots$ may be specified by giving the marginal distribution of X_1 and the conditional distribution of X_n given $\sigma\{X_1, \ldots, X_{n-1}\}$ for $n \geqslant 2$.

What if the specifications are more complicated? For example, suppose we want $\{X_i : i \in I\}$, where the conditional distribution of $\{X_i : i \in F\}$ given $\{X_i : i \notin F\}$ are given for each finite set F. Can we construct such a collection?

It is clear that some consistency conditions are needed.

Example 6

Let X_1, X_2, X_3 be integer-valued random variables such that $\mathbb{P}\{(X_1, X_2, X_3) = (i, j, k)\} > 0$ for all $(i, j, k) \in \mathbb{Z}^3$ (to avoid worries about division by zero). Then

$$\sum_j \mathbb{P}\{X_1 = i \mid X_2 = j, X_3 = k\}\mathbb{P}\{X_2 = j \mid X_3 = k\} = \sum_j \mathbb{P}\{X_1 = i, X_2 = j \mid X_3 = k\}$$

which is a consistency requirement among the conditional distributions. You may object that the second factor in the sum on the left is not quite in the form of conditional distribution of $\{X_i : i \in F\}$ given $\{X_i : i \notin F\}$. No problem, rewrite the above as

$$\sum_{j,\ell} \mathbb{P}\{X_1 = i \mid X_2 = j, X_3 = k\}\mathbb{P}\{X_1 = \ell, X_2 = j \mid X_3 = k\} = \sum_j \mathbb{P}\{X_1 = i, X_2 = j \mid X_3 = k\}$$

so that all terms are of that form.

Let us now formulate the question taking into account this kind of consistency requirement. The problem is already very interesting and non-trivial if the random variables take only two values and I is countable[5].

Specification: Let I be a countable set and let $\Omega = \{0, 1\}^I$ (a compact metric space) and $\mathcal{G} = \mathcal{B}(\Omega)$. Suppose that for each finite $F \subseteq I$ we are given a stochastic kernel $\lambda_F : \Omega \times \mathcal{F} \mapsto [0, 1]$ such that

(1) $\lambda_F(x, \cdot)$ is a Borel probability measure on $\mathcal{G}$.

(2) $\lambda_F(\cdot, A)$ is measurable w.r.t $\mathcal{G}_F := \sigma\{\omega_j : j \notin F\}$.

(3) $\lambda_F(\cdot, A) = \mathbf{1}_A$ if $A \in \mathcal{G}_F$.

(4) If $F_1 \subseteq F_2$, then $\lambda_{F_2} \circ \lambda_{F_1} = \lambda_{F_2}$ where

$$\lambda_{F_2} \circ \lambda_{F_1}(x, A) := \int_\Omega \lambda_{F_1}(y, A)\lambda_{F_2}(x, dy). \tag{2}$$

A collection $\{\lambda_F\}$ satisfying these conditions is called a *specification*.

Gibbs measure: Given a specification as above, does there exists a measure μ on $(\Omega, \mathcal{G})$ such that the regular conditional distribution given $\mathcal{G}_F$ is λ_F, for any finite $F \subseteq I$. Such a measure μ is called a *Gibbs measure*.

Equivalently, we may ask if there exist $\{0, 1\}$-valued random variables $(X_i)_{i\in I}$ (on some probability space) such that $\lambda_F(x, \cdot)$ is the conditional distribution of $(X_i)_{i\in F}$ given that $(X_i)_{i\in F^c} = (x_i)_{i\in F^c}$, for any finite $F \subseteq I$. Then μ is distribution of $(X_i)_{i\in I}$.

[5]The material below is taken from C. Preston, *Random Fields*, Springer, Berlin Heidelberg, 2006.

The result on existence of Gibbs measures: Unlike in the Kolmogorov consistency theorem, the obvious consistency conditions (2) are not sufficient to ensure the existence of Gibbs measures. We need more. The following fundamental forms the basis of the probabilistic study of Gibbs measures coming from statistical physics. The additional conditions imposed are not easy to interpret, but there are easy to check sufficient conditions that ensure they hold.

Theorem 3: Dobrushin-Lanford-Ruelle

Assume that the specification $\{\lambda_F\}$ satisfies the following conditions:

(1) There exists $x_0 \in \Omega$ such that given any finite $F \subseteq I$ and any $\varepsilon > 0$, there is a probability measure ν on $\{0,1\}^F$ such that for any $A \subseteq \{0,1\}^F$ satisfying $\nu(A) < \delta$ and any finite $F' \supseteq F$, we have $\lambda_{F'}(x_0, A) < \varepsilon$.

(2) For any finite dimensional cylinder set $A \in \mathcal{G}$ and any finite $F \subseteq I$ and any $\varepsilon > 0$, there is a finite $F' \subseteq I$ and a function $f : \{0,1\}^{F'} \mapsto \mathbb{R}$ such that $|\lambda_F(x, A) - f(x_{F'})| < \varepsilon$ for all $x \in \Omega$, where $x_{F'}$ is the projection of x to $\{0,1\}^{F'}$.

Then, a Gibbs measure exists for the given specification.

The question of uniqueness or non-uniqueness of Gibbs measure is one of the most fundamental questions in statistical physics, and underlies the mathematical study of *phase transitions*.

CHAPTER 2

Martingales in discrete time: theory

1. Martingales

1.1. The setting. Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $\mathcal{F}_\bullet = (\mathcal{F}_n)_{n\in\mathbb{N}}$ be a collection of sigma subalgebras of $\mathcal{F}$ indexed by natural numbers such that $\mathcal{F}_m \subseteq \mathcal{F}_n$ whenever $m < n$. Then we say that $\mathcal{F}_\bullet$ is a *filtration* and refer to the quadruple $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ as a *filtered probability space*. We refer to n as "time", and the sigma-algebra at a given time represents the complete knowledge at that instant.

A sequence of random variables $X = (X_n)_{n\in\mathbb{N}}$ defined on $(\Omega, \mathcal{F}, \mathbb{P})$ is said to be *adapted* to the filtration $\mathcal{F}_\bullet$ if $X_n \in \mathcal{F}_n$ for each n.

Definition 3: Martingales, Submartingales, Supermartingales

Let $X = (X_n)_{n\in\mathbb{N}}$ be an adapted process on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$. Assume that each X_n is integrable. We say that X is a

(1) *super-martingale* if $\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] \leqslant X_{n-1}$ a.s. for each $n \geqslant 1$,

(2) *sub-martingale* if $-X$ is a super-martingale,

(3) *martingale* is it is both a super-martingale and a sub-martingale.

When we want to explicitly mention the filtration, we write $\mathcal{F}_\bullet$-martingale or $\mathcal{F}_\bullet$-super-martingale etc.

Unlike say Markov chains, the definition of martingales does not appear to put too strong a restriction on the distributions of X_n, it is only on a few conditional expectations. Nevertheless, very power theorems can be proved at this level of generality, and there are any number of examples to justify making a definition whose meaning is not obvious on the surface.

1.2. Examples. In this section we give classes of examples.

Example 7: Random walk

Let ξ_n be independent random variables with finite mean and let $\mathcal{F}_n = \sigma\{\xi_1, \dots, \xi_n\}$ (so $\mathcal{F}_0 = \{\emptyset, \Omega\}$). Define $X_0 = 0$ and $X_n = \xi_1 + \dots + \xi_n$ for $n \geqslant 1$. Then, X is $\mathcal{F}_\bullet$-adapted, X_n have finite mean, and

$$\begin{aligned} \mathbb{E}[X_n \mid \mathcal{F}_{n-1}] &= \mathbb{E}[X_{n-1} + \xi_n \mid \mathcal{F}_{n-1}] \\ &= \mathbb{E}[X_{n-1} \mid \mathcal{F}_{n-1}] + \mathbb{E}[\xi_n \mid \mathcal{F}_{n-1}] \\ &= X_{n-1} + \mathbb{E}[\xi_n] \end{aligned}$$

since $X_{n-1} \in \mathcal{F}_{n-1}$ and ξ_n is independent of $\mathcal{F}_{n-1}$. Thus, if $\mathbb{E}[\xi_n]$ is positive for all n, then X is a sub-martingale; if $\mathbb{E}[\xi_n]$ is negative for all n, then X is a super-martingale; if $\mathbb{E}[\xi_n] = 0$ for all n, then X is a martingale.

Example 8: Product martingale

Let ξ_n be independent, non-negative random variables and let $X_n = \xi_1\xi_2 \dots \xi_n$ and $X_0 = 1$. Then, with $\mathcal{F}_n = \sigma\{\xi_1, \dots, \xi_n\}$, we see that X is $\mathcal{F}_\bullet$-adapted and $\mathbb{E}[X_n]$ exists (equals the product of $\mathbb{E}[\xi_k]$, $k \leqslant n$). Lastly,

$$\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] = \mathbb{E}[X_{n-1}\xi_n \mid \mathcal{F}_{n-1}] = X_{n-1}\mu_n$$

where $\mu_n = \mathbb{E}[\xi_n]$. Hence, if $\mu_n \geqslant 1$ for all n, then X is a sub-martingale, if $\mu_n = 1$ for all n, then X is a martingale, and if $\mu_n \leqslant 1$ for all n, then X is a super-martingale.
In particular, replacing ξ_n by ξ_n/μ_n, we see that $Y_n := \frac{X_n}{\mu_1 \dots \mu_n}$ is a martingale.

Example 9: Log-likelihood function

Let $S = \{1, 2, \dots, m\}$ be a finite set with a probability mass function $p(i)$, $1 \leqslant i \leqslant m$. Suppose $X_1, X_2, \dots$ are i.i.d. samples from this distribution. The likelihood-function of the first n samples is defined as

$$L_n = \prod_{k=1}^{n} p(X_k).$$

Its logarithm, $\ell_n := \log L_n = \sum_{k=1}^{n} \log p(X_k)$, is called the log-likelihood function. This is a sum of i.i.d. random variables $\log p(X_k)$, and they have finite mean $H := \mathbb{E}[\log p(X_k)] = \sum_{i=1}^{m} p(i) \log p(i)$ (if $p(i) = 0$ for some i, interpret $p(i) \log p(i)$ as zero). Hence $\ell_n - nH$ is a martingale (with respect to the filtration given by $\mathcal{F}_n = \sigma\{X_1, \dots, X_n\}$), by the same logic as in the first example.

Example 10: Doob martingale

Here is a very general way in which any (integrable) random variable can be put at the end of a martingale sequence. Let X be an integrable random variable on $(\Omega, \mathcal{F}, \mathbb{P})$ and let $\mathcal{F}_\bullet$ be any filtration. Let $X_n = \mathbb{E}[X \mid \mathcal{F}_n]$. Then, (X_n) is $\mathcal{F}_\bullet$-adapted, integrable and

$$\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] = \mathbb{E}\left[\mathbb{E}\left[X \mid \mathcal{F}_n\right] \mid \mathcal{F}_{n-1}\right] = \mathbb{E}[X \mid \mathcal{F}_{n-1}] = X_{n-1}$$

by the tower property of conditional expectation. Thus, (X_n) is a martingale. Such martingales got by conditioning one random variable w.r.t. an increasing family of sigma-algebras is called a *Doob martingale*.

Often $X = f(\xi_1, \ldots, \xi_m)$ is a function of independent random variables ξ_k, and we study X be sturying the evolution of $\mathbb{E}[X \mid \xi_1, \ldots, \xi_k]$, revealing the information of ξ_ks, one by one. This gives X as the end-point of a Doob martingale. The usefulness of this construction will be clear in a few lectures.

Example 11: Increasing process

Let A_n, $n \geqslant 0$, be a sequence of random variables such that $A_0 \leqslant A_1 \leqslant A_2 \leqslant \ldots$ a.s.Assume that A_n are integrable. Then, if $\mathcal{F}_\bullet$ is any filtration to which A is adapted, then

$$\mathbb{E}[A_n \mid \mathcal{F}_{n-1}] - A_{n-1} = \mathbb{E}[A_n - A_{n-1} \mid \mathcal{F}_{n-1}] \geqslant 0$$

by positivity of conditional expectation. Thus, A is a sub-martingale. Similarly, a decreasing sequence of random variables is a super-martingale[a].

[a] An interesting fact that we shall see later is that any sub-martingale is a sum of a martingale and an increasing process. This seems reasonable since a sub-martingale increases on average while a martingale stays constant on average.

Example 12: Harmonic functions

Let $R = (R_n)_{n\geqslant 0}$ be a simple random walk on a graph G with a countable vertex set V where each vertex has finite degree. This means that R is a Markov chain with transition probabilities $p_{i,j} = \frac{1}{\deg(i)}$ if $j \sim i$, and $p_{i,j} = 0$ otherwise. Let $\phi : V \mapsto \mathbb{R}$ be a harmonic function, i.e., $\phi(i) = \frac{1}{\deg(i)} \sum_{j:j\sim i} \phi(j)$, for all $i \in V$. Then, $X_n = \phi(V_n)$ is a martingale. Indeed,

$$\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] = \frac{1}{\deg(V_{n-1})} \sum_{j:j\sim V_{n-1}} \phi(j) = \phi(V_{n-1}) = X_{n-1}.$$

We say that ϕ is subharmonic if $\phi(i) \leqslant \frac{1}{\deg(i)} \sum_{j:j\sim i} \phi(j)$, for all $i \in V$ and that ϕ is superharmonic if the inequality goes the other way. Correspondingly, X will be a submartingale or a supermartingale.

Example 13: Branching process

Consider a Galton-Watson process with offspring variable L with $\mathbb{P}\{L = k\} = p_k$ for $k \in \mathbb{N}$. Recall the informal description: At generation 0 there is one individual, who gives birth to a random number of offsprings according to the distribution of L. These offsprings belong to the first generation, and each of them give birth to offsprings according to the same distribution, independently of each other and their ancestors. And so on.

Let Z_n is the number of individuals in the nth generation. A precise construction of Z_ns can by setting $Z_0 = 1$ and

$$Z_n := \begin{cases} L_{n,1} + \ldots + L_{n,Z_{n-1}} & \text{if } Z_{n-1} \geqslant 1, \\ 0 & \text{if } Z_{n-1} = 0. \end{cases}$$

where $L_{n,k}$, $n, k \geqslant 1$, are i.i.d. copies of L.

The natural filtration to consider here is $\mathcal{F}_n = \sigma\{L_{m,k} : m \leqslant n, k \geqslant 1\}$. Clearly $Z = (Z_n)_{n\geqslant 0}$ is adapted to $\mathcal{F}_\bullet$. Assume that $\mathbb{E}[L] = m < \infty$. Then, (see the exercise below to justify the steps)

$$\begin{aligned} \mathbb{E}[Z_n \mid \mathcal{F}_{n-1}] &= \mathbb{E}[\mathbf{1}_{Z_{n-1}\geqslant 1}(L_{n,1} + \ldots + L_{n,Z_{n-1}}) \mid \mathcal{F}_{n-1}] \\ &= \mathbf{1}_{Z_{n-1}\geqslant 1} Z_{n-1} m \\ &= Z_{n-1} m. \end{aligned}$$

Thus, $\frac{Z_n}{m^n}$ is a martingale.

Exercise 4

If N is a $\mathbb{N}$-valued random variable independent of ξ_m, $m \geqslant 1$, and ξ_m are i.i.d. with mean μ, then $\mathbb{E}[\sum_{k=1}^{N} \xi_k \mid N] = \mu N$.

Example 14: Pólya's urn scheme

An urn has $b_0 > 0$ black balls and $w_0 > 0$ white balls to start with. A ball is drawn uniformly at random and returned to the urn with an additional new ball of the same colour. Draw a ball again and repeat. The process continues forever. A basic question about this process is what happens to the contents of the urn? Does one colour start dominating, or do the proportions of black and white equalize?

In precise notation, the above description may be captured as follows. Let U_n, $n \geqslant 1$, be i.i.d. Uniform$[0,1]$ random variables. Let $b_0 > 0$, $w_0 > 0$, be given. Then, define $B_0 = b_0$ and $W_0 = w_0$. For $n \geqslant 1$, define (inductively)

$$\xi_n = \mathbf{1}\left(U_n \leqslant \frac{B_{n-1}}{B_{n-1} + W_{n-1}}\right), \quad B_n = B_{n-1} + \xi_n, \quad W_n = W_{n-1} + (1 - \xi_n).$$

Here, ξ_n is the indicator that the nth draw is a black, B_n and W_n stand for the number of black and white balls in the urn before the $(n+1)$st draw. It is easy to see that $B_n + W_n = b_0 + w_0 + n$ (since one ball is added after each draw).

Let $\mathcal{F}_n = \sigma\{U_1, \ldots, U_n\}$ so that ξ_n, B_n, W_n are all $\mathcal{F}_n$ measurable. Let $X_n = \frac{B_n}{B_n + W_n} = \frac{B_n}{b_0 + w_0 + n}$ be the proportion of balls after the nth draw (X_n is $\mathcal{F}_n$-measurable too). Observe that

$$\mathbb{E}[B_n \mid \mathcal{F}_{n-1}] = B_{n-1} + \mathbb{E}[\mathbf{1}_{U_n \leqslant X_{n-1}} \mid \mathcal{F}_{n-1}] = B_{n-1} + X_{n-1} = \frac{b_0 + w_0 + n}{b_0 + w_0 + n - 1} B_{n-1}.$$

Thus,

$$\begin{aligned}
\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] &= \frac{1}{b_0 + w_0 + n} \mathbb{E}[B_n \mid \mathcal{F}_{n-1}] \\
&= \frac{1}{b_0 + w_0 + n - 1} B_{n-1} \\
&= X_{n-1}
\end{aligned}$$

showing that (X_n) is a martingale.

1.3. New martingales out of old. Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space.

- Suppose $X = (X_n)_{n \geqslant 0}$ is a $\mathcal{F}_\bullet$-martingale and $\phi : \mathbb{R} \to \mathbb{R}$ is a convex function. If $Y_n = \phi(X_n)$ has finite expectation for each n, then $Y = (Y_n)_{n \geqslant 0}$ is a sub-martingale. If X was a sub-martingale to start with, and if ϕ is increasing and convex, then Y is a sub-martingale.

Indeed, $\mathbb{E}[\phi(X_n) \mid \mathcal{F}_{n-1}] \geqslant \phi(\mathbb{E}[X_n \mid \mathcal{F}_{n-1}])$ by conditional Jensen's inequality. If X is a martingale, then the right hand side is equal to $\phi(X_{n-1})$ and we get the sub-martingale property for $(\phi(X_n))_{n\geqslant 0}$.

If X was only a sub-martingale, then $\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] \geqslant X_{n-1}$ and hence the increasing property of ϕ is required to conclude that $\phi(\mathbb{E}[X_n \mid \mathcal{F}_{n-1}]) \geqslant \phi(X_{n-1})$.

- If $t_0 < t_1 < t_2 < \ldots$ is a subsequence of natural numbers, and X is a martingale (or sub-martingale or super-martingale), then $X_{t_0}, X_{t_1}, X_{t_2}, \ldots$ is also a martingale (respectively sub-martingale or super-martingale). Of course, we must take the new filtration $\mathcal{F}_{t_0}, \mathcal{F}_{t_1}, \ldots$. This follows from the tower property of conditional expectation: If $n > m$,

$$\mathbb{E}[X_n \mid \mathcal{F}_m] = \mathbb{E}[\mathbb{E}[\ldots\mathbb{E}[X_n \mid \mathcal{F}_{n-1}] \mid \mathcal{F}_{n-2}] \ldots \mid \mathcal{F}_m].$$

But it is a very interesting question that we shall ask later as to whether the same is true if t_i are random times.

If we had a continuous time-martingale $X = (X_t)_{t\geqslant 0}$, then again $X(t_i)$ would be a discrete time martingale for any $0 < t_1 < t_2 < \ldots$. Results about continuous time martingales can in fact be deduced from results about discrete parameter martingales using this observation and taking closely spaced points t_i. If we get to continuous-time martingales at the end of the course, we shall explain this fully.

- Let X be a martingale and let $H = (H_n)_{n\geqslant 1}$ be a predictable sequence. This just means that $H_n \in \mathcal{F}_{n-1}$ for all $n \geqslant 1$. Then, define $(H.X)_n = \sum_{k=1}^{n} H_k(X_k - X_{k-1})$. Assume that $(H.X)_n$ is integrable for each n (true for instance if H_n is a bounded random variable for each n). Then, $(H.X)$ is a martingale. If X was a sub-martingale to start with, then $(H.X)$ is a sub-martingale provided H_n are non-negative, in addition to being predictable.

Proof. $\mathbb{E}[(H.X)_n - (H.X)_{n-1} \mid \mathcal{F}_{n-1}] = \mathbb{E}[H_n(X_n - X_{n-1}) \mid \mathcal{F}_{n-1}] = H_n\mathbb{E}[X_n - X_{n-1} \mid \mathcal{F}_{n-1}]$. If X is a martingale, the last term is zero. If X is a sub-martingale, then $\mathbb{E}[X_n - X_{n-1} \mid \mathcal{F}_{n-1}] \geqslant 0$ and because $H_n \geqslant 0$, the sub-martingale property of $(H.X)$ follows. ■

1.4. Continuous time martingales? For continuous time processes, we must change the setting. But most of what we said in this section goes through, with appropriate modifications.

- A filtration indexed by a totally ordered set I such as $\mathbb{R}_+ = [0, \infty)$ or $\mathbb{Z}$ or $\{0, 1, \ldots, n\}$ etc. is just a family of sub sigma algebras $\mathcal{F}_t$, $t \in I$ such that $\mathcal{F}_t \subseteq \mathcal{F}_s$ if $t \leqslant s$.
- $X = (X_t)_{t\in I}$ is said to be adapted to $\mathcal{F}_\bullet = (\mathcal{F}_t)_{t\in I}$ if $X_t \in \mathcal{F}_t$ for all $t \in I$.
- Defining martingales may look tricky as there may be no "previous instant" $n - 1$. However, as we saw above, for a discrete time martingale, $\mathbb{E}[X_n \mid \mathcal{F}_m] = X_m$ for any $n > m$.

This can be taken as the definition in general. That is $X = (X_t)_{t\in I}$ is defined to be a supermartingale if $\mathbb{E}[X_t \mid \mathcal{F}_s] \leqslant X_s$ for any $s < t$.

- If $J \subseteq I$ and $(X_t)_{t\in I}$ is a submartingale w.r.t. $(\mathcal{F}_t)_{t\in I}$, then so is $(X_t)_{t\in J}$ w.r.t $(\mathcal{F}_t)_{t\in J}$.
- By the previous remark, if $(X_t)_{t\in\mathbb{R}_+}$ is a martingale, then so is any sequence $X_{q_1}, X_{q_2}, \ldots$ provided $q_1 < q_2 < \ldots$. Theorems that we prove for discrete time martingales apply to all such subsequences, and putting them together, one can get analogous theorems for the original process indexed by $\mathbb{R}_+$. This will be explained at the end. For now, this is just to say that studying discrete time martingales is sufficient.

Here are two examples of martingales in continuous time.

Exercise 5: Brownian motion

Let $W = (W_t)_{t\geqslant 0}$ be a standard Brownian motion. Recall that this means that for any $t_1 < \ldots < t_n$, the variables $W_{t_1}, W_{t_2} - W_{t_1} \ldots, W_{t_n} - W_{t_{n-1}}$ are independent Gaussians with zero means and variances $t_1, t_2 - t_1, \ldots, t_n - t_{n-1}$, respectively (the sample path continuity of $t \mapsto W_t$ is not needed for the following exercise).

Show that W_t and $W_t^2 - t$ are martingales.

Exercise 6: Poisson process

Let $N = (N_t)_{t\geqslant 0}$ be a homogenous Poisson process with intensity 1. Recall that this means that for any $t_1 < \ldots < t_n$, the variables $N_{t_1}, N_{t_2} - N_{t_1}, \ldots, N_{t_n} - N_{t_{n-1}}$ are independent Poisson random variables with parameters $t_1, t_2 - t_1, \ldots, t_n - t_{n-1}$ respectively.

Show that $N_t - t$ and $(N_t - t)^2 - t$ are martingales.

2. A short preview of things to come

Here are three distinct themes in the theory of martingales that we shall explore in some detail. The usefulness can only be appreciated when one sees the variety of problems that can be solved using these ideas.

(1) Several theorems that we have seen for sums of independent random variables go through for martingales, with a few modifications. Two examples are Hoeffding's inequality and Kolmogorov's maximal inequality. While the proofs are not all that different from the one for independent sums, the applicability is greatly expanded by this generalisation. With greater effort and imposing suitable conditions one can (we do not touch upon these) also prove central limit theorems, laws of iterated logarithm, etc.

(2) Optional stopping/sampling theorems. By definition of martingales, $\mathbb{E}[M_T] = \mathbb{E}[M_0]$ for any t, but the extension to certain kinds of random times T, known as stopping times, brings with it many amazing consequences. The closest things one might have seen are perhaps Wald's identities in sequential analysis (where the idea of stopping at a random time is the key point).

(3) Convergence theorems for martingales. Martingales stay constant on average. Turns out, each sample path of a martingale must either converge or oscillate wildly. Convergence theorems restrict the possibility of oscillating wildly to conclude convergence. The single most useful statement is that any martingale sequence that is uniformly integrable must converge almost surely and in L^1. This has innumerable consequences.

3. Hoeffding's inequality

Theorem 4: Hoeffding's inequality

Let $X = (X_0, \ldots, X_n)$ be a martingale. Assume that $|X_k - X_{k-1}| \leqslant d_k$ a.s. for $1 \leqslant k \leqslant n$. Then $\mathbb{P}\{X_n - X_0 \geqslant t\} \leqslant e^{-\frac{t^2}{2D^2}}$ for any $t > 0$, where $D^2 = d_1^2 + \ldots + d_n^2$.

Earlier we proved this for the martingale $X_k = \xi_1 + \ldots + \xi_k$, where ξ_i are independent random variables with $|\xi_k| \leqslant d_k$. A key step in the proof is that for a zero mean random variable Y with $|Y| \leqslant d$,

$$\mathbb{E}[e^{\theta Y}] \leqslant e^{\frac{1}{2}\theta^2 d^2}. \tag{3}$$

Recall that this is proved by writing Y as the convex combination $\frac{Y+d}{2d}(-d) + \frac{d-Y}{2d}(d)$ and using convexity of exponential to get $\mathbb{E}[e^{\theta Y}] \leqslant \frac{e^{\theta d}+e^{-\theta d}}{2}$. It is an elementary fact that the latter is

$$\sum_{k\geqslant 0} \frac{\theta^{2k} d^{2k}}{(2k)!} \leqslant \sum_{k\geqslant 0} \frac{(\theta^2 d^2/2)^k}{k!} = e^{\frac{1}{2}\theta^2 d^2}.$$

Proof of Hoeffding's inequality. Fix $t > 0$ and $\theta > 0$ and use Markov's inequality:

$$\mathbb{P}\{X_n - X_0 \geqslant t\} = e^{-\theta t}\mathbb{E}[e^{\theta(X_n - X_0)}].$$

Conditioning on $\mathcal{F}_{n-1}$, the right side is just $\mathbb{E}[e^{\theta(X_{n-1}-X_0)}\mathbb{E}[e^{\theta(X_n - X_{n-1})} \mid \mathcal{F}_{n-1}]]$. Conditional on $\mathcal{F}_{n-1}$, the random variable $X_n - X_{n-1}$ has zero mean (martingale property) and is bounded by d_n anyway. By (3), the inner conditional expectation is at most $e^{\theta^2 d_n^2/2}$. Thus,

$$\mathbb{E}[e^{\theta(X_n - X_0)}] \leqslant e^{\theta d_n^2/2}\mathbb{E}[e^{\theta(X_{n-1}-X_0)}].$$

Continuing, we get $\mathbb{E}[e^{\theta(X_n - X_0)}] \leqslant e^{\frac{1}{2}\theta^2 D^2}$. Thus, $\mathbb{P}\{X_n - X_0 \geqslant t\} \leqslant e^{-\theta t + \frac{1}{2}\theta^2 D^2}$. The optimal choice is $\theta = \frac{t}{D}$, which gives the claimed bound. ∎

The only difference in the proof is that for independent random variables we factored $\mathbb{E}[e^{\theta S_n}]$ as a product of $\mathbb{E}[e^{\theta X_k}]$, but here we do it by conditioning on the previous step. While there is not much novelty in the proof of this extension, it greatly enhances the applicability.

3.1. Concentration for functions of independent random variables. Let $\xi_1, \ldots, \xi_n$ be i.i.d. random variables and let $f : \mathbb{R}^n \to \mathbb{R}$ be a fixed function. Assume that $|f(x) - f(y)| \leqslant d$ if $x, y \in \mathbb{R}^n$ differ in at most one co-ordinate (i.e., $x_i = y_i$ for all $i \neq j$, for some j).

Theorem 5: McDiarmid's inequality

In the above setting, let $Y = f(\xi_1, \ldots, \xi_n)$. Then

$$\mathbb{P}\{|Y - \mathbb{E}[Y]| \geqslant t\} \leqslant 2e^{-\frac{t^2}{2nd^2}}.$$

Proof. Create the Doob martingale $X_k = \mathbb{E}[Y \mid \mathcal{F}_k]$, where $\mathcal{F}_k = \sigma\{\xi_1, \ldots, \xi_k\}$. Then $X_n = Y$ while $X_0 = \mathbb{E}[Y]$. If we argue that $|X_k - X_{k-1}| \leqslant d$ a.s., then by Hoeffding's inequality (apply to X and to $-X$ and add up the inequalities), we get the claimed result, since $D^2 = nd^2$.

To see the bound on $X_{k+1} - X_k$, consider ξ_k', an independent copy of ξ_k. Then

$$X_k = \mathbb{E}[f(\xi_1, \ldots, \xi_{k-1}, \xi_k, \xi_{k+1}, \ldots, \xi_n) \mid \xi_1, \ldots, \xi_{k-1}, \xi_k],$$
$$X_{k-1} = \mathbb{E}[f(\xi_1, \ldots, \xi_{k-1}, \xi_k', \xi_{k+1}, \ldots, \xi_n) \mid \xi_1, \ldots, \xi_{k-1}, \xi_k].$$

Note that the conditioning is on the same set in both cases, but f acts on vectors that differ only in the kth co-ordinate. Therefore,

$$|X_k - X_{k-1}| \leqslant \mathbb{E}[|f(\xi_1, \ldots, \xi_{k-1}, \xi_k, \xi_{k+1}, \ldots, \xi_n) - f(\xi_1, \ldots, \xi_{k-1}, \xi_k', \xi_{k+1}, \ldots, \xi_n)| \mid \mathcal{F}_k].$$

As the random variable is bounded by d, so is its conditional expectation. ■

Innumerable problems of probabilistic combinatorial optimization are of the form given here. For example,

(1) Let ξ be a uniformly sampled binary string. Let Y_n be the number of times the segment $s = 1011$ occurs in ξ. That is $Y_n = \sum_{i=1}^{n-3} \mathbf{1}_{(\xi_i, \ldots, \xi_{i+3}) = s}$. While we can compute the mean ($\mathbb{E}[Y_n] \sim n/4$) and variance of Y_n ($\mathrm{Var}(Y_n) \asymp n$), the dependence causes difficulties in studying the variable. But Hoeffding's inequality tells us that

$$\mathbb{P}\{|Y_n - \mathbb{E}[Y_n]| \geqslant t\sqrt{n}\} \leqslant e^{-ct^2}.$$

That is, Y_n has sub-Gaussian tails around its mean, on the scale of $\sqrt{n}$, the same behaviour (up to the constant c in the exponent) as if Y_n were a sum of i.i.d. random variables.

(2) Pick two independent binary strings U, V of length n uniformly and independently at random. Let

$$Y_n = \max\{k : \exists i_1 < \ldots < i_k,\ j_1 < \ldots < j_k, \text{ such that } U_{i_r} = V_{j_r} \text{ for all } r\},$$

be the longest length of a common subsequence. By some elementary arguments one can show that $\mathbb{E}[Y_n] \sim cn$ for some $0 < c < 1$, but the value of c is unknown. Remarkably, without any knowledge of the mean and variance, we can still say that $\mathbb{P}\{|Y_n - \mathbb{E}[Y_n]| \geqslant t\sqrt{n}\} \leqslant 2e^{-t^2/2}$ (use Theorem 5 with $\xi_i = (U_i, V_i)$).

Of course, Hoeffding's gives a one-way inequality. The window length one gets in specific problems may not be optimal. The power is in the generality.

4. Stopping times

Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space. Let $T : \Omega \to \mathbb{N} \cup \{+\infty\}$ be a random variable. If $\{T \leqslant n\} \in \mathcal{F}_n$ for each $n \in \mathbb{N}$, we say that T is a *stopping time*.

Equivalently we may ask for $\{T = n\} \in \mathcal{F}_n$ for each n. The equivalence with the definition above follows from the fact that $\{T = n\} = \{T \leqslant n\} \setminus \{T \leqslant n-1\}$ and $\{T \leqslant n\} = \cup_{k=0}^{n}\{T = k\}$. The way we defined it, it also makes sense for continuous time. For example, if $(\mathcal{F}_t)_{t \geqslant 0}$ is a filtration and $T : \Omega \to [0, +\infty]$ is a random variable, then we say that T is a stopping time if $\{T \leqslant t\} \in \mathcal{F}_t$ for all $t \geqslant 0$.

Example 15

Let X_k be random variables on a common probability space and let $\mathcal{F}^X$ be the natural filtration generated by them. If $A \in \mathcal{B}(\mathbb{R})$ and $\tau_A = \min\{n \geqslant 0 : X_n \in A\}$, then τ_A is a stopping time. Indeed, $\{\tau_A = n\} = \{X_0 \notin A, \ldots, X_{n-1} \notin A, X_n \in A\}$ which clearly belongs to $\mathcal{F}_n$.
On the other hand, $\tau'_A := \max\{n : X_n \notin A\}$ is not a stopping time as it appears to require future knowledge. One way to make this precise is to consider $\omega_1, \omega_2 \in \Omega$ such that $\tau'_A(\omega_1) = 0 < \tau'_A(\omega_2)$ but $X_0(\omega_1) = X_0(\omega_2)$. I we can find such ω_1, ω_2, then any event in $\mathcal{F}_0$ contains both of them or neither. But $\{\tau'_A \leqslant 0\}$ contains ω_1 but not ω_2, hence it cannot be in $\mathcal{F}_0$. In a general probability space we cannot guarantee the existence of ω_1, ω_2 (for example Ω may contain only one point or X_k may be constant random variables!), but in sufficiently rich spaces it is possible. See the exercise below.

Exercise 7

Let $\Omega = \mathbb{R}^{\mathbb{N}}$ with $\mathcal{F} = \mathcal{B}(\mathbb{R}^{\mathbb{N}})$ and let $\mathcal{F}_n = \sigma\{\Pi_0, \Pi_1, \ldots, \Pi_n\}$ be generated by the projections $\Pi_k : \Omega \to \mathbb{R}$ defined by $\Pi_k(\omega) = \omega_k$ for $\omega \in \Omega$. Give an honest proof that τ'_A defined as above is not a stopping time (let A be a proper subset of $\mathbb{R}$).

Suppose T, S are two stopping times on a filtered probability space. Then $T \wedge S$, $T \vee S$, $T + S$ are all stopping times. However cT and $T - S$ need not be stopping times (even if they take values in $\mathbb{N}$). This is clear, since $\{T \wedge S \leqslant n\} = \{T \leqslant n\} \cup \{S \leqslant n\}$ etc. More generally, if $\{T_m\}$ is a countable family of stopping times, then $\max_m T_m$ and $\min_m T_m$ are also stopping times.

Small digression into continuous time: We shall use filtrations and stopping times in the Brownian motion class too. There the index set is continuous and complications can arise. For example, let $\Omega = C[0, \infty)$, $\mathcal{F}$ its Borel sigma-algebra, $\mathcal{F}_t = \sigma\{\Pi_s : s \leqslant t\}$. Now define $\tau, \tau' : C[0, \infty) \to [0, \infty)$ by $\tau(\omega) = \inf\{t \geqslant 0 : \omega(t) \geqslant 1\}$ and $\tau'(\omega) = \inf\{t \geqslant 0 : \omega(t) > 1\}$ where the infimum is interpreted to be $+\infty$ is the set is empty. In this case, τ is an $\mathcal{F}_\bullet$-stopping time but τ' is not (why?). In discrete time there is no analogue of this situation. When we discuss this in Brownian motion, we shall enlarge the sigma-algebra $\mathcal{F}_t$ slightly so that even τ' becomes a stopping time. This is one of the reasons why we do not always work with the natural filtration of a sequence of random variables.

4.1. The sigma algebra at a stopping time. If T is a stopping time for a filtration $\mathcal{F}_\bullet$, then we want to define a sigma-algebra $\mathcal{F}_T$ that contains all information up to and including the random time T.

To motivate the idea, assume that $\mathcal{F}_n = \sigma\{X_0, \ldots, X_n\}$ for some sequence $(X_n)_{n \geqslant 0}$. One might be tempted to define $\mathcal{F}_T$ as $\sigma\{X_0, \ldots, X_T\}$ but a moment's thought shows that this does not make sense as written since T itself depends on the sample point. One way to fix this is to "freeze the process" at time T to get the *stopped process* $Y_n := X_{T \wedge n}$, $n \geqslant 0$ and define

$$\mathcal{F}_T := \sigma\{Y_n : n \geqslant 0\}. \tag{4}$$

This is well-defined, and it is clear that it captures all knowledge up to the stopping time T.

Another way to think of this is to partition the sample space as $\Omega = \sqcup_{n \geqslant 0}\{T = n\}$ and on the portion $\{T = n\}$ we consider the sigma-algebra generated by $\{X_0, \ldots, X_n\}$. Putting all these together we get a sigma-algebra that we call $\mathcal{F}_T$. To summarize, we say that $A \in \mathcal{F}_T$ if and only if $A \cap \{T = n\} \in \mathcal{F}_n$ for each $n \geqslant 0$. Observe that this condition is equivalent to asking for $A \cap \{T \leqslant n\} \in \mathcal{F}_n$ for each $n \geqslant 0$ (check!). Thus, we arrive at the definition

$$\mathcal{F}_T := \{A \in \mathcal{F} : A \cap \{T \leqslant n\} \in \mathcal{F}_n\} \quad \text{for each } n \geqslant 0. \tag{5}$$

The two definitions (4) and (5) are equivalent when the filtration is the one generated by X. We leave this as an exercise. For general filtrations, we take (5) as the definition. Some basic observations.

(1) It does not make a difference if we wrote $\{T=n\}$ instead of $\{T\leqslant n\}$ in (5). But in continuous time setting, it makes sense to define $\mathcal{F}_T=\{A\in\mathcal{F}: A\cap\{T\leqslant t\}\in\mathcal{F}_t \text{ for all } t\}$.

(2) $\mathcal{F}_T$ is a sigma-algebra. Indeed,

$$A^c\cap\{T\leqslant n\}=\{T\leqslant n\}\setminus(A\cap\{T\leqslant n\}),$$
$$\Big(\bigcup_{k\geqslant 1}A_k\Big)\cap\{T\leqslant n\}=\bigcup_{k\geqslant 1}(A_k\cap\{T\leqslant n\}).$$

From these it follows that $\mathcal{F}_T$ is closed under complements and countable unions. As $\Omega\cap\{T\leqslant n\}=\{T\leqslant n\}\in\mathcal{F}_n$, we see that $\Omega\in\mathcal{F}_T$. Thus $\mathcal{F}_T$ is a sigma-algebra.

(3) The idea behind the definition of $\mathcal{F}_T$ is that it somehow encapsulates all the information we have up to the random time T. The following lemma is a sanity check that this intuition is captured in the definition (i.e., if the lemma were not true, we would have to change our definition!). Here and later, note that X_T is the random variable $\omega\mapsto X_{T(\omega)}(\omega)$. But this makes sense only if $T(\omega)<\infty$, hence we assume finiteness below. Alternately, we can fix some random variable X_∞ that is $\mathcal{F}$-measurable and use that to define X_T when $T=\infty$.

Lemma 1

Let $X=(X_n)_{n\geqslant 0}$ be adapted to the filtration $\mathcal{F}_\bullet$ and let T be a finite $\mathcal{F}_\bullet$-stopping time. Then X_T is $\mathcal{F}_T$-measurable.

Proof. $\{X_T\leqslant u\}\cap\{T\leqslant n\}=\{X_n\leqslant u\}\cap\{T\leqslant n\}$ which is in $\mathcal{F}_n$, since $\{X_n\leqslant u\}$ and $\{T\leqslant n\}$ both are. Therefore, $\{X_T\leqslant u\}\in\mathcal{F}_T$ for any $u\in\mathbb{R}$, meaning that X_T is $\mathcal{F}_T$-measurable. ■

(4) Another fact is that T is $\mathcal{F}_T$-measurable (again, it would be a strange definition if this was not true - after all, by time T we know that value of T). To show this we just need to show that $\{T\leqslant m\}\in\mathcal{F}_m$ for any $m\geqslant 0$. But that is true because for every $n\geqslant 0$ we have

$$\{T\leqslant m\}\cap\{T\leqslant n\}=\{T\leqslant m\wedge n\}\in\mathcal{F}_{m\wedge n}\subseteq\mathcal{F}_n.$$

(5) If T, S are stopping times and $T\leqslant S$ (caution! here we mean $T(\omega)\leqslant S(\omega)$ for every $\omega\in\Omega$), then $\mathcal{F}_T\subseteq\mathcal{F}_S$. To see this, suppose $A\in\mathcal{F}_T$. Then $A\cap\{T\leqslant n\}\in\mathcal{F}_n$ for each n.

If $A \in \mathcal{F}_T$, then $A \cap \{T \leqslant n\} \in \mathcal{F}_n$ and hence $(A \cap \{T \leqslant n\}) \cap \{S \leqslant n\} \in \mathcal{F}_n$, as S is a stopping time. But if $T \leqslant S$, then $A \cap \{S \leqslant n\} = A \cap \{S \leqslant n\} \cap \{T \leqslant n\}$ and hence $A \cap \{S \leqslant n\} \in \mathcal{F}_n$. This shows that $A \in \mathcal{F}_S$.

All these should make it clear that that the definition of $\mathcal{F}_T$ is sound and does indeed capture the notion of information up to time T.

For the sake of completeness: In the last property stated above, suppose we only assumed that $T \leqslant S$ a.s. Can we still conclude that $\mathcal{F}_T \subseteq \mathcal{F}_S$? Let $C = \{T > S\}$ so that $C \in \mathcal{F}$ and $\mathbb{P}(C) = 0$. If we try to repeat the proof as before, we end up with

$$A \cap \{S \leqslant n\} = [(A \cap \{T \leqslant n\}) \cap \{S \leqslant n\}] \ \cup \ (A \cap \{S \leqslant n\} \cap C).$$

The first set belongs to $\mathcal{F}_n$ but there is no assurance that $A \cap C$ does, since we only know that $C \in \mathcal{F}$.

One way to get around this problem (and many similar ones) is to complete the sigma-algebras as follows. Let $\mathcal{N}$ be the collection of all null sets in $(\Omega, \mathcal{F}, \mathbb{P})$. That is,

$$\mathcal{N} = \{A \subseteq \Omega : \exists\, B \in \mathcal{F} \text{ such that } B \supseteq A \text{ and } \mathbb{P}(B) = 0\}.$$

Then define $\bar{\mathcal{F}}_n = \sigma\{\mathcal{F}_n \cup \mathcal{N}\}$. This gives a new filtration $\bar{\mathcal{F}}_\bullet = (\bar{\mathcal{F}}_n)_{n \geqslant 0}$ which we call the completion of the original filtration (strictly speaking, this completion depended on $\mathcal{F}$ and not merely on $\mathcal{F}_\bullet$. But we can usually assume without loss of generality that $\mathcal{F} = \sigma\{\cup_{n \geqslant 0} \mathcal{F}_n\}$ by decreasing $\mathcal{F}$ if necessary. In that case, it is legitimate to call $\bar{\mathcal{F}}_\bullet$ the completion of $\mathcal{F}_\bullet$ under $\mathbb{P}$).

It is to be noted that after enlargement, $\mathcal{F}_\bullet$-adapted processes remain adapted to $\bar{\mathcal{F}}_\bullet$, stopping times for $\mathcal{F}_\bullet$ remain stopping times for $\bar{\mathcal{F}}_\bullet$, etc. Since the enlargement is only by $\mathbb{P}$-null sets, it can be see that $\mathcal{F}_\bullet$-super-martingales remain $\bar{\mathcal{F}}_\bullet$-super-martingales, etc. Hence, there is no loss in working in the completed sigma algebras.

Henceforth we shall simply assume that our filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ is such that all $\mathbb{P}$-null sets in $(\Omega, \mathcal{F}, \mathbb{P})$ are contained in $\mathcal{F}_0$ (and hence in $\mathcal{F}_n$ for all n). Let us say that $\mathcal{F}_\bullet$ is complete to mean this.

Exercise 8

Let T, S be stopping times with respect to a complete filtration $\mathcal{F}_\bullet$. If $T \leqslant S$ a.s (w.r.t. $\mathbb{P}$), show that $\mathcal{F}_T \subseteq \mathcal{F}_S$.

Exercise 9

Let $T_0 \leqslant T_1 \leqslant T_2 \leqslant \ldots$ (a.s.) be stopping times for a complete filtration $\mathcal{F}_\bullet$. Is the filtration $(\mathcal{F}_{T_k})_{k\geqslant 0}$ also complete?

5. Optional stopping or sampling

Let $X = (X_n)_{n\geqslant 0}$ be a super-martingale on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$. We know that (1) $\mathbb{E}[X_n] \leqslant \mathbb{E}[X_0]$ for all $n \geqslant 0$ and (2) $(X_{n_k})_{k\geqslant 0}$ is a super-martingale for any subsequence $n_0 < n_1 < n_2 < \ldots$.

Optional stopping theorems are statements that assert that $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$ for a stopping time T. *Optional sampling theorems* are statements that assert that $(X_{T_k})_{k\geqslant 0}$ is a super-martingale for an increasing sequence of stopping times $T_0 \leqslant T_1 \leqslant T_2 \leqslant \ldots$. Usually one is not careful to make the distinction and OST could refer to either kind of result.

Neither of these statements is true without extra conditions on the stopping times. But they are true when the stopping times are bounded, as we shall prove in this section. In fact, it is best to remember only that case, and derive more general results whenever needed by writing a stopping time as a limit of bounded stopping times. For example, $T \wedge n$ are bounded stopping times and $T \wedge n \stackrel{a.s.}{\to} T$ as $n \to \infty$.

Now we state the precise results for bounded stopping times.

Theorem 6: Optional stopping theorem

Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space and let T be a stopping time for $\mathcal{F}_\bullet$. If $X = (X_n)_{n\geqslant 0}$ is a $\mathcal{F}_\bullet$-super-martingale, then $(X_{T\wedge n})_{n\geqslant 0}$ is a $\mathcal{F}_\bullet$-super-martingale. In particular, $\mathbb{E}[X_{T\wedge n}] \leqslant \mathbb{E}[X_0]$ for all $n \geqslant 0$.

Proof. Let $H_n = \mathbf{1}_{n\leqslant T}$. Then $H_n \in \mathcal{F}_{n-1}$ because $\{T \geqslant n\} = \{T \leqslant n-1\}^c$ belongs to $\mathcal{F}_{n-1}$. By the observation earlier, $(H.X)_n$, $n \geqslant 0$, is a super-martingale. But $(H.X)_n = X_{T\wedge n} - X_0$ and this proves that $(X_{T\wedge n})_{n\geqslant 0}$ is an $\mathcal{F}_\bullet$-super-martingale. Then of course $\mathbb{E}[X_{T\wedge n}] \leqslant \mathbb{E}[X_0]$. ■

Optional stopping theorem is a remarkably useful tool. The way it is applied is to strengthen the above statement to say that $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$ (equality if it is a martingale) for a stopping time T. This would follow if we could show that $\mathbb{E}[X_{T\wedge n}] \to \mathbb{E}[X_T]$ as $n \to \infty$. This seems reasonable as $X_{T\wedge n} \stackrel{a.s.}{\to} X_T$ for any finite stopping time T. However, it is not always true as the following example shows.

Example 16

Let (X_n) be a simple symmetric random walk on integers started at the origin and let T be the first time when the random walk visits the state 1 (it is well-known that $T < \infty$ a.s.). Then X is a martingale, $X_T = 1$ a.s. but $X_0 = 0$ a.s., hence the expectations do not match.

If (X_n) is a positive supermartingale, then for any finite stopping time T, almost sure convergence of $X_{T\wedge n}$ to X_T and Fatou's lemma imply that $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$. We give more conveniently applicable conditions in the theorem below (the conditions hold for X if and only if they hold for its negative, which is convenient to apply to martingales to get the conclusion $\mathbb{E}[X_T] = \mathbb{E}[X_0]$).

Theorem 7: Optional stopping theorem - an extension

Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space and let T be a finite stopping time for $\mathcal{F}_\bullet$. If $X = (X_n)_{n\geqslant 0}$ is a $\mathcal{F}_\bullet$-sub-martingale. Then $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$ if any one of the following conditions are met.

(1) $\{X_{T\wedge n}\}_{n\geqslant 1}$ is uniformly integrable.

(2) $\{X_{T\wedge n}\}_{n\geqslant 1}$ is either (a) uniformly bounded or (b) dominated by an integrable random variable or (c) bounded in L^2.

(3) T is uniformly bounded.

(4) The differences $\{X_{n+1} - X_n\}$ are uniformly bounded by a constant, and $\mathbb{E}[T] < \infty$.

Proof. Since T is finite, $X_{T\wedge n} \stackrel{a.s.}{\to} X_T$. Hence, $X_T \stackrel{L^1}{\to} X_T$ if and only if $\{X_{T\wedge n}\}$ is uniformly integrable. Convergence in L^1 implies convergence of expectations. This proves the first statement.

Each of the three conditions in the second statement is sufficient for uniform integrability, hence the second follows from the first.

If $T \leqslant N$ a.s. then $|X_{T\wedge n}| \leqslant |X_0| + \ldots + |X_N|$ which is an integrable random variable. Therefore, the sequence $\{X_{T\wedge n}\}_{n\geqslant 1}$ is dominated and hence uniformly integrable.

If $|X_{n+1} - X_n| \leqslant b$ a.s. for all n, then $|X_n| \leqslant |X_0| + nb$. Hence $|X_{T\wedge n}| \leqslant |X_0| + bT$. As $|X_0| + bT$ is integrable, the domination condition holds and thus $\{X_{T\wedge n}\}$ is uniformly integrable. ■

Although the conditions given here may be worth remembering, it is much better practise to always write $\mathbb{E}[X_{T\wedge n}] \leqslant \mathbb{E}[X_0]$ and then think of ways in which to let $n \to \infty$ and get $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$. While uniform integrability is necessary and sufficient, it is hard to check, but there may be other situation-specific ways to interchange limit and expectation.

Needless to say, we just stated the result for super-martingales. From this, the reverse inequality holds for sub-martingales (by applying the above to $-X$) and hence equality holds for martingales.

In Theorem 6 we think of stopping a process at a stopping time. There is a variant where we sample the process at an increasing sequence of stopping times and the question is whether the observed process retains the martingale/super-martingale property. This can be thought of as a non-trivial extension of the trivial statement that if $(X_n)_n$ is a super-martingale w.r.t. $(\mathcal{F}_n)_n$, then for any $n_0 \leqslant n_1 \leqslant n_2 \leqslant \ldots$, the process $(X_{n_k})_k$ is a super-martingale w.r.t. $(\mathcal{F}_{n_k})_k$.

Theorem 8: Optional sampling theorem

Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space and let $X = (X_n)_{n \geqslant 0}$ be a $\mathcal{F}_\bullet$-super-martingale. Let T_n, $n \geqslant 0$ be bounded stopping times for $\mathcal{F}_\bullet$ such that $T_0 \leqslant T_1 \leqslant T_2 \leqslant \ldots$ Then, $(X_{T_k})_{k \geqslant 0}$ is a super-martingale with respect to the filtration $(\mathcal{F}_{T_k})_{k \geqslant 0}$.
If we only assume that $T_0 \leqslant T_1 \leqslant T_3 \leqslant \ldots$ a.s., then the conclusion remains valid if we assume that the given filtration is complete.

The condition of boundedness of the stopping times can be replaced by the condition that $\{X_{T_k \wedge n}\}_{n \geqslant 0}$ is uniformly integrable for any k. The reasons are exactly the same as those that went into the proof of Theorem 7.

Proof. Since X is adapted to $\mathcal{F}_\bullet$, it follows that X_{T_k} is $\mathcal{F}_{T_k}$-measurable. Further, if $|T_k| \leqslant N_k$ w.p.1. for a fixed number N_k, then $|X_{T_k}| \leqslant |X_0| + \ldots + |X_{N_k}|$ which shows the integrability of X_{T_k}. The theorem will be proved if we show that if $S \leqslant T \leqslant N$ where S, T are stopping times and N is a fixed number, then

$$\mathbb{E}[X_T \mid \mathcal{F}_S] \leqslant X_S \quad \text{a.s.} \tag{6}$$

Since X_S and $\mathbb{E}[X_T \mid \mathcal{F}_S]$ are both $\mathcal{F}_S$-measurable, (6) follows if we show that $\mathbb{E}[(X_T - X_S)\mathbf{1}_A] \leqslant 0$ for every $A \in \mathcal{F}_S$.

Now fix any $A \in \mathcal{F}_S$ and define $H_k = \mathbf{1}_{S+1 \leqslant k \leqslant T}\mathbf{1}_A$. This is the indicator of the event $A \cap \{S \leqslant k-1\} \cap \{T \geqslant k\}$. Since $A \in \mathcal{F}_S$ we see that $A \cap \{S \leqslant k-1\} \in \mathcal{F}_{k-1}$ while $\{T \geqslant k\} = \{T \leqslant k-1\}^c \in \mathcal{F}_{k-1}$. Thus, H is predictable. In words, this is the betting scheme where we bet 1 rupee on each game from time $S+1$ to time T, but only if A happens (which we know by time S). By the gambling lemma, we conclude that $\{(H.X)_n\}_{n \geqslant 1}$ is a super-martingale. But $(H.X)_n = (X_{T \wedge n} - X_{S \wedge n})\mathbf{1}_A$. Put $n = N$ and get $\mathbb{E}[(X_T - X_S)\mathbf{1}_A] \leqslant 0$ since $(H.X)_0 = 0$. Thus (6) is proved. ■

An alternate proof of Theorem 8 is outlined below.

Second proof of Theorem 8. As in the first proof, it suffices to prove (6).

First assume that $S \leqslant T \leqslant S+1$ a.s. Let $A \in \mathcal{F}_S$. On the event $\{S = T\}$ we have $X_T - X_S = 0$. Therefore,

$$\int_A (X_T - X_S)dP = \int_{A \cap \{T=S+1\}} (X_{S+1} - X_S)dP$$

$$(7) \qquad = \sum_{k=0}^{N-1} \int_{A \cap \{S=k\} \cap \{T=k+1\}} (X_{k+1} - X_k)dP.$$

For fixed k, we see that $A \cap \{S = k\} \in \mathcal{F}_k$ since $A \in \mathcal{F}_S$ and $\{T = k+1\} = \{T \leqslant k\}^c \cap \{S = k\} \in \mathcal{F}_k$ because $T \leqslant S+1$. Therefore, $A \cap \{S = k\} \cap \{T = k+1\} \in \mathcal{F}_k$ and the super-martingale property of X implies that $\int_B (X_{k+1} - X_k)dP \leqslant 0$ for any $B \in \mathcal{F}_k$. Thus, each term in (7) is non-positive. Hence $\int_A X_S dP \geqslant \int_A X_T dP$ for every $A \in \mathcal{F}_T$. This just means that $\mathbb{E}[X_S \mid \mathcal{F}_T] \leqslant X_T$. This completes the proof assuming $S \leqslant T \leqslant S+1$.

In general, since $S \leqslant T \leqslant N$, let $S_0 = S$, $S_1 = T \wedge (S+1)$, $S_2 = T \wedge (S+2), \ldots, S_N = T \wedge (S+N)$ so that each S_k is a stopping time, $S_N = T$, and for each k we have $S_k \leqslant S_{k+1} \leqslant S_k + 1$ a.s. Deduce from the previous case that $\mathbb{E}[X_T \mid \mathcal{F}_S] \leqslant X_S$ a.s. ■

We end this section by giving an example to show that optional sampling theorems can fail if the stopping times are not bounded.

Example 17

Let ξ_i be i.i.d. $\text{Ber}_{\pm}(1/2)$ random variables and let $X_n = \xi_1 + \ldots + \xi_n$ (by definition $X_0 = 0$). Then X is a martingale. Let $T_1 = \min\{n \geqslant 1 : X_n = 1\}$.

A theorem of Pólya asserts that $T_1 < \infty$ w.p.1. But $X_{T_1} = 1$ a.s. while $X_0 = 0$. Hence $\mathbb{E}[X_{T_1}] \neq \mathbb{E}[X_0]$, violating the optional stopping property (for bounded stopping times we would have had $\mathbb{E}[X_T] = \mathbb{E}[X_0]$). In gambling terminology, if you play till you make a profit of 1 rupee and stop, then your expected profit is 1 (an not zero as optional stopping theorem asserts).

If $T_j = \min\{n \geqslant 0 : X_n = j\}$ for $j = 1, 2, 3, \ldots$, then it again follows from Pólya's theorem that $T_j < \infty$ a.s. and hence $X_{T_j} = j$ a.s. Clearly $T_0 < T_1 < T_2 < \ldots$ but $X_{T_0}, X_{T_1}, X_{T_2}, \ldots$ is not a super-martingale (in fact, being increasing it is a sub-martingale!).

This example shows that applying optional sampling theorems blindly without checking conditions can cause trouble. But the boundedness assumption is by no means essential. Indeed, if the above example is tweaked a little, optional sampling is restored.

Example 18

In the previous example, let $-A < 0 < B$ be integers and let $T = \min\{n \geqslant 0 : X_n = -A \text{ or } X_n = B\}$. Then T is an unbounded stopping time. In gambling terminology, the gambler has capital A and the game is stopped when he/she makes a profit of B rupees or the gambler goes bankrupt. If we set $B = 1$ we are in a situation similar to before, but with the somewhat more realistic assumption that the gambler has finite capital.
By the optional sampling theorem $\mathbb{E}[X_{T\wedge n}] = 0$. By a simple argument (or Pólya's theorem) one can prove that $T < \infty$ w.p.1. Therefore, $X_{T\wedge n} \stackrel{a.s.}{\to} X_T$ as $n \to \infty$. Further, $|X_{T\wedge n}| \leqslant B+A$ from which by DCT it follows that $\mathbb{E}[X_{T\wedge n}] \to \mathbb{E}[X_T]$. Therefore, $\mathbb{E}[X_T] = 0$. In other words optional stopping property is restored.

6. Wald's identities

If $X_0, X_1, \ldots$ are i.i.d. random variables with finite mean, and T is an $\mathbb{N}$-valued random variable independent of X_is, then $\mathbb{E}[S_T] = \mathbb{E}[X_1]\mathbb{E}[T]$. To see this, write $S_T = \sum_n X_n \mathbf{1}_{T\geqslant n}$ and take expectations to get

$$\mathbb{E}[S_T] = \sum_{n=0}^{\infty} \mathbb{E}[X_n \mathbf{1}_{T\geqslant n}] = \mathbb{E}[X_1] \sum_{n=0}^{\infty} \mathbf{1}_{T\geqslant n} = \mathbb{E}[X_1]\mathbb{E}[T].$$

The application of Fubini's theorem to interchange of Expectation and sum in the first equality is justified by first working out the same kind of expression with $|X_n|$ in place of X_n.

Motivated by applications in statistics, Wald worked out conditions under which T could be allowed to depend on the sequence (X_n). These are known as Wald's identities. We give two of them.

Theorem 9: Wald's first identity

Let X_n be i.i.d. with finite expectation. Let T be a stopping time for the natural filtration of (X_n) with $\mathbb{E}[T] < \infty$. Then $\mathbb{E}[S_T] = \mathbb{E}[X_1]\mathbb{E}[T]$.

Theorem 10: Wald's second identity

Let X_n be i.i.d. with zero mean and finite variance σ^2. Let T be a stopping time for the natural filtration of (X_n) with $\mathbb{E}[T] < \infty$. Then $\mathbb{E}[S_T^2] = \sigma^2\mathbb{E}[T]$.

Proof of Wald's first identity. As $S_n - n\mathbb{E}[X_1]$ is a martingale, hence $\mathbb{E}[S_{T\wedge n} - (T\wedge n)\mathbb{E}[X_1]] = 0$ by optional stopping theorem. Write this as $\mathbb{E}[S_{T\wedge n}] = \mathbb{E}[T\wedge n]\mathbb{E}[X_1]$. By MCT, the right side converges to $\mathbb{E}[T]\mathbb{E}[X_1]$. On the left side, $S_{T\wedge n} \stackrel{a.s.}{\to} S_T$, and $S_{T\wedge n}$ is dominated by $Y = \sum_{k=0}^{\infty} |X_k| \mathbf{1}_{T\geqslant k}$.

As $\{T \geqslant k\} = (\{T \leqslant k-1\})^c \in \mathcal{F}_{k-1}$ and X_k is independent of $\mathcal{F}_{k-1}$,

$$\mathbb{E}[|X_k|\mathbf{1}_{T\geqslant k}] = \mathbb{E}[\mathbb{E}[|X_k|\mathbf{1}_{T\geqslant k} \mid \mathcal{F}_{k-1}]] = \mathbb{E}[\mathbf{1}_{T\geqslant k}\mathbb{E}[X_k]] = \mathbb{E}[|X_k|]\mathbb{P}\{T \geqslant k\}.$$

Summing over k shows that $\mathbb{E}[Y] = \mathbb{E}[|X_1|]\mathbb{E}[T]$. ■

The application of optional stopping here could be avoided easily. The same argument that showed $\mathbb{E}[Y] = \mathbb{E}[|X_1|]\mathbb{E}[T]$ also shows that $\mathbb{E}[S_T] = \mathbb{E}[X_1]\mathbb{E}[T]$. Applying Fubini's now requires us to have already shown that $\mathbb{E}[Y] < \infty$.

Proof of Wald's second identity. We try to mimic the previous proof, replacing S_n by the martingale $S_n^2 - n\sigma^2$. Thus, $\mathbb{E}[S^2_{T\wedge n}] = \sigma^2\mathbb{E}[T \wedge n]$. By MCT, the right side converges to $\sigma^2\mathbb{E}[T]$. Now $S_{T\wedge n} \overset{a.s.}{\to} S_T$. If we show that $S_{T\wedge n}$ converges in L^2, the limit must be S_T again, and hence $\mathbb{E}[S^2_{T\wedge n}] \to \mathbb{E}[S^2_T]$, completing the proof. To show the required convergence in L^2, observe that for $m < n$

$$\begin{aligned}\mathbb{E}[|S_{T\wedge n} - S_{T\wedge m}|^2] &= \mathbb{E}\left[\left(\sum_{k=m+1}^{n} X_k\mathbf{1}_{T\geqslant k}\right)^2\right] \\ &= \sum_{k=m+1}^{n} \mathbb{E}[X_k^2\mathbf{1}_{T\geqslant k}] + 2\sum_{m+1\leqslant k<\ell\leqslant n} \mathbb{E}[X_kX_\ell\mathbf{1}_{T\geqslant \ell}].\end{aligned}$$

For $k < \ell$,

$$\mathbb{E}[X_kX_\ell\mathbf{1}_{T\geqslant\ell}] = \mathbb{E}[X_k\mathbf{1}_{T\geqslant\ell}\mathbb{E}[X_\ell \mid \mathcal{F}_{\ell-1}]] = 0$$

while

$$\mathbb{E}[X_k^2\mathbf{1}_{T\geqslant k}] = \mathbb{E}[\mathbf{1}_{T\geqslant k}\mathbb{E}[X_k^2 \mid \mathcal{F}_{k-1}]] = \mathbb{E}[X_1^2]\mathbb{P}\{T \geqslant k\}.$$

Thus, $\mathbb{E}[|S_{T\wedge n} - S_{T\wedge m}|^2] = \mathbb{E}[X_1]^2 \sum_{k=m+1}^{n} \mathbb{P}\{T \geqslant k\}$, which goes to zero as $m, n \to \infty$ (as $\mathbb{E}[T] < \infty$). This shows that $S_{T\wedge n}$ is Cauchy in L^2, and hence convergent. This completes the proof. ■

7. Applications of the optional stopping theorem

7.1. Gambler's ruin problem. Let $S_n = \xi_1 + \ldots + \xi_n$ be simple symmetric random walks, where ξ_i are i.i.d. $\mathrm{Ber}_\pm(1/2)$. Fix $-a < 0 < b$. What is the probability that S hits b before $-a$? With $T = T_{-a} \wedge T_b$ where $T_x = \min\{n \geqslant 0 : X_n = x\}$ we know that $T < \infty$ a.s.[1] and hence $\mathbb{E}[X_{T\wedge n}] = 0$ for all n. Since $|X_{T\wedge n}| \leqslant a + b$, we can let $n \to \infty$ and use DCT to conclude that $\mathbb{E}[X_T] = 0$. Hence,

[1] If you don't know this, here is a simple argument - Divide the coin tosses into disjoint blocks of length $\ell = a + b$, and observe that with probability $2^{-\ell}$, all tosses in a block are heads. Hence, there is some block which has all heads. If the random walk is not to the left of $-a$ at the beginning of this block, then it will be to the right of b at the end of the block.

if $\alpha = \mathbb{P}\{X_T = b\}$ then $1 - \alpha = \mathbb{P}\{X_T = -a\}$ and

$$0 = \mathbb{E}[X_T] = \alpha b - (1-\alpha)a$$

which gives $\alpha = \frac{a}{a+b}$.

Exercise 10

Let ξ_i be i.i.d. with $\xi_1 = +1$ w.p. p and $\xi_1 = -1$ w.p. $q = 1-p$. Let $X_n = \xi_1 + \ldots + \xi_n$. Find the probability that X hits B before $-A$ (for $A, B > 0$, of course).

One can get more information about the time T as follows. Recall that $\{S_n^2 - n\}$ is a martingale, hence $\mathbb{E}[S_{T\wedge n}^2 - (T \wedge n)] = 0$ for all n. To interchange expectation with limit as $n \to \infty$, we rewrite this as $\mathbb{E}[S_{T\wedge n}^2] = \mathbb{E}[T \wedge n]$. The left side converges to $\mathbb{E}[S_T^2]$ by DCT (as $|S_{T\wedge n}| \leqslant a + b$) and the right side converges to $\mathbb{E}[T]$ (by MCT). Hence

$$E[T] = \mathbb{E}[S_T^2] = (-a)^2 \frac{b}{a+b} + b^2 \frac{a}{a+b} = ab.$$

In particular, when $a = b$, we get b^2, which makes sense in view of the fundamental fact that a random walk moves distance $\sqrt{t}$ in time t.

7.2. Waiting times for patterns in coin tossing. Let $\xi_1, \xi_2, \ldots$ be i.i.d. Ber$(1/2)$ variables (fair coin tosses). Let $\tau_{1011} = \min\{n \geqslant 1 : (\xi_{n-3}, \ldots, \xi_n) = (1,0,1,1)\}$ and similarly define τ_ε for any patter $\varepsilon \in \{0,1\}^k$ for some k. Clearly these are stopping times for the filtration $\mathcal{F}_n = \sigma\{\xi_1, \ldots, \xi_n\}$. We would like to understand the distribution or the mean of these stopping times.

Clearly τ_1 and τ_0 are Geometric random variables with mean 2. Things are less simple for other patterns. Since this is written out in many places and was explained in class and is given as exercise to write out a proper proof, will skip the explanation here. The final answer depends on the overlaps in the pattern. If $\varepsilon = (\varepsilon_1, \ldots, \varepsilon_k)$, then

$$\mathbb{E}[\tau_\varepsilon] = \sum_{j=1}^{k} 2^j \mathbf{1}_{(\varepsilon_1,\ldots,\varepsilon_j)=(\varepsilon_{k-j+1},\ldots,\varepsilon_k)}.$$

In particular, $\mathbb{E}[\tau_{11}] = 6$ while $\mathbb{E}[\tau_{10}] = 4$. You may remember that the usual proof of this involves setting up a gambling game where the kth gambler enters with 1 rupee in hand, just before the kth toss, and bets successively on the kth toss being ε_1 (at the same time the $(k-i)$th gambler, if still in the game, is betting on it being ε_{i+1}). If instead, the kth gambler come with k rupees in hand, one can find the second moment of τ_ε and so on. If the kth gambler comes with $e^{\theta k}$ rupees, where θ is sufficiently small, then the moment generating function of τ_ε can also be found.

Aside from the proof of this claim that uses optional stopping theorem, what is the reason for different waiting times for different patterns of the same length? This can be understood qualitatively in terms of the waiting time paradox.

The waiting time "paradox": If buses come regularly with inter-arrival times of one hour, but a person who has no watch goes at random to the bus stop, her expected waiting time is 30 minutes. However, if inter-arrival times are random with equal chance of being 90 minutes or 30 minutes (so one hour on average still), then the expected waiting time jumps to 37.5 minutes! The reason is that the person is 3 times more likely to have entered in a 90 minute interval than in a 30 minute interval.

What does this have to do with the waiting times in patterns. The "buses" $\underline{10}$ and $\underline{11}$ are equally frequent (chance $1/4$ at any $(n-1, n)$ slot), but 10 is more regularly spaced than $\underline{11}$. In fact $\underline{11}$ buses can crowd together as in the string $\underline{11111}$ which has 4 occurrences of $\underline{11}$. But to get four $\underline{10}$ buses we need at least 8 tosses. Thus, the waiting time for the less regular bus is more!

8. Random walks on graphs

Let $G = (V, E)$ be a graph with a countable vertex set V. We shall always assume that each vertex has finite degree and that the graph is connected. Simple random walk on G (usually written SRW) is a markov chain $X = (X_n)_{n \geqslant 0}$ with transition probabilities $p_{u,v} = \frac{1}{\deg(u)}$ for $v \sim u$, and $p_{u,v} = 0$ if v is not adjacent to u. Usually we fix X_0 to be some vertex w (in which case we write $\mathbb{P}_w$, $\mathbb{E}_w$ to indicate that).

Recall that a function $f : V \mapsto \mathbb{R}$ is said to be harmonic at a vertex u if $\frac{1}{\deg(u)} \sum_{v:v\sim u} f(v) = f(u)$ (this is called the *mean value property*). We saw that if f if harmonic on the whole of V, then $(f(X_n))_{n\geqslant 0}$ is a martingale. In such a situation, optional sampling theorem tells us that $(f(X_{\tau\wedge n}))_{n\geqslant 0}$ is also a martingale for any stopping time τ. Here is an extension of this statement.

Theorem 11

Let X be SRW on G. Let B be a proper subset of V and let τ denote the hitting time of B by the random walk. Suppose $f : V \mapsto \mathbb{R}$ is harmonic (or sub-harmonic) at all vertices of $V \setminus B$. Then, $(f(X_{\tau\wedge n}))_{n\geqslant 0}$ is a martingale (respectively, sub-martingale) with respect to the filtration $(\mathcal{F}_{\tau\wedge n})_{n\geqslant 0}$.

Note that this is not obvious and does not follow from the earlier statement. If we define $M_n = f(X_n)$, then M may not a martingale, since f need not harmonic on B. Therefore, $f(X_{\tau\wedge n})$ is not got by stopping a martingale (in which case OST would have implied the theorem), it is just that this stopped process is a martingale!

Proof. Let f be harmonic and set $M_n = f(X_{\tau\wedge n})$ and let $\mathcal{G}_n = \mathcal{F}_{\tau\wedge n}$. We want to show that $\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = M_n$. Clearly M_n is $\mathcal{G}_n$ measurable (since $X_{\tau\wedge n}$ is). Let $A \in \mathcal{G}_n$ and let $A' = A \cap \{\tau \leqslant n\}$ and $A'' = A \cap \{\tau > n\}$.

On A' we have $M_{n+1} = M_n = f(X_\tau)$ and hence $\mathbb{E}[M_{n+1}\mathbf{1}_{A'}] = \mathbb{E}[M_n\mathbf{1}_{A'}]$.

On A'', we have that X_{n+1} is a uniformly chosen random neighbour of X_n (independent of all the conditioning) and hence,

$$\mathbb{E}[M_{n+1}\mathbf{1}_{A''}] = \mathbf{1}_{A''}\frac{1}{\deg(X_n)}\sum_{v:v\sim X_n} f(u) = \mathbf{1}_{A''}f(X_n)$$

where the last equality holds because $X_n \notin B$ and f is harmonic there. But $f(X_n) = M_n$ on A'', (since $\tau > n$), and hence we see that $\mathbb{E}[M_{n+1}\mathbf{1}_{A''}] = \mathbb{E}[M_n\mathbf{1}_{A''}]$.

Adding the two we get $\mathbb{E}[M_{n+1}\mathbf{1}_A] = \mathbb{E}[M_n\mathbf{1}_A]$ for all $A \in \mathcal{G}_n$, hence $\mathbb{E}[M_{n+1} \mid \mathcal{G}_n] = M_n$. ■

Remark 4: Reversible Markov chains

Can the discussions of this section be carried over to general Markov chains? Not quite, but it can be to *reversible* Markov chains. Let X be a Markov chain on a countable state space S with transition matrix P. We shall assume that the chain is irreducible. Recall that the chain is said to be reversible if there is a $\pi = (\pi_i)_{i\in S}$ on S (called the stationary measure) such that $\pi(i)p_{i,j} = \pi(j)p_{j,i}$ for all $i, j \in S$.

If the chain is reversible, we can make a graph G with vertex set S and edges $i \sim j$ whenever $p_{i,j} > 0$ (reversibility forces $p_{j,i} > 0$, hence the graph is undirected). For any $i \sim j$, define the conductance of the corresponding edge as $C_{i,j} = \pi(i)p_{i,j}$. By reversibility, $C_{j,i} = C_{i,j}$, hence the conductance is associated to the edge, not the direction. Then the given Markov chain is a random walk on this graph, except that the transitions are not uniform. They are given by

$$p_{i,j} = \begin{cases} \frac{C_{i,j}}{C_{i,\cdot}} & \text{if } j \sim i, \\ 0 & \text{otherwise} \end{cases}$$

where $C_{i,\cdot} = \sum_{k:k\sim i} C_{i,k}$. Conversely, for any graph $G = (V, E)$ with specified conductances on edges, if we define transition probabilities as above, we get a reversible Markov chain.

All the discussions in the section can be taken over to general reversible chains, with appropriate modifications. If a chain is not reversible, for example suppose there are two states i, j such that $p_{i,j} > 0$ but $p_{j,i} = 0$, are quite different.

8.1. Discrete Dirichlet problem and gambler's ruin. Let $G = (V, E)$ be a connected graph with vertex set V and edge set E and every vertex having finite degrees. Let X denote the simple random walk on G. We consider two problems.

Gambler's ruin problem: Let A, C be disjoint proper subsets of V. Find $\mathbb{P}_x\{\tau_A < \tau_C\}$ for any $x \in V$. Here τ_A is the hitting time of the set A by the SRW X.

Discrete Dirichlet problem: Let B be a proper subset of V. Fix a function $\phi : B \mapsto \mathbb{R}$. Find a function $f : V \mapsto \mathbb{R}$ such that (a) $f(x) = \phi(x)$ for all $x \in B$, (b) f is harmonic on $V \setminus B$. This is a system of linear equations, one for each $v \in V \setminus B$, and in the variables $f(x)$, $x \in V \setminus B$.

These two problems are intimately related. To convey the main ideas without distractions, we restrict ourselves to finite graphs now.

(1) Observe that the solution to the Dirichlet problem, if it exists, is unique. Indeed, if f, g are two solutions, then $h = f - g$ is harmonic on $V \setminus B$ and $h = 0$ on B. Now let x_0 be a point where h attains its maximum (here finiteness of the graph is used). If $x_0 \notin B$, then $h(x_0)$ is the average of the values of h at the neighbours of x_0, hence each of those values must be equal to $h(x_0)$. Iterating this, we get a point $x \in B$ such that $h(x) = h(x_0)$ (connectedness of the graph is used here). Therefore, the maximum of h is zero. Similarly the minimum is zero and we get $f = g$.

(2) Let $f(x) = P_x\{\tau_A < \tau_B\}$ in the gambler's ruin problem. We claim that f is harmonic at every $x \notin B := A \cup C$. Indeed, for any $x \notin B$, condition on the first step of the Markov chain to see that

$$f(x) = \mathbb{E}_x[\mathbb{P}\{\tau_A < \tau_B \mid X_1\}] = \mathbb{E}_x[\mathbb{P}_{X_1}\{\tau_A < \tau_B\}] = \frac{1}{\deg(x)} \sum_{y:y\sim x} f(y).$$

Further, f is 1 on A and 0 on C. Hence f is just the solution to the discrete Dirichlet problem with $B = A \cup C$ and $\phi = \mathbf{1}_A$. rst

(3) Conversely, suppose a set B is given and for every $x \in B$ we solve the gambler's ruin problem with $A = \{x\}$ and $C = B \setminus \{x\}$. Let $\mu_x(y) = \mathbb{P}_y\{\tau_x = \tau_B\}$ denote the solution. Then, given any $\phi : B \mapsto \mathbb{R}$, it is easy to see that $f(\cdot) = \sum_{x \in B} \phi(x)\mu_x(\cdot)$ is a solution to the discrete Dirichlet problem (linear combinations of harmonic functions is harmonic).

(4) The solution in the previous point may be rewritten as (with $M_n = f(X_{\tau \wedge n})$)

$$f(y) = \sum_{x \in B} \phi(x)\mathbb{P}_y\{\tau_B = \tau_x\} = \sum_{x \in B} \phi(x)\mathbb{P}_y\{M_\tau = x\} = \mathbb{E}_y[M_\tau].$$

(5) Here is another way to see that the solution f to the Dirichlet problem must be given like this. From Theorem 11 we know that M_n is a martingale. Hence $\mathbb{E}[f(X_{\tau\wedge n})] = \mathbb{E}[f(X_0)]$, in particular, if $X_0 = v$ then $\mathbb{E}_v[f(X_{\tau\wedge n})] = f(v)$. Let $n \to \infty$ and DCT ($f(X_{\tau\wedge n})$ is of course uniformly bounded) to conclude that $f(v) = \mathbb{E}[f(X_\tau)] = \mathbb{E}[M_\tau]$.

To summarize, we have shown the existence and uniqueness of the solution to the discrete Dirichlet problem, and related it to the solution to the gambler's ruin problem. This can be summarized as follows.

Theorem 12

Let $G = (V, E)$ be a finite connected graph and let B be a proper subset of vertices. Given $\phi : B \mapsto \mathbb{R}$, the unique solution to the discrete Dirichlet problem with boundary data ϕ is given by $f(x) = \mathbb{E}_x[\phi(X_\tau)]$ where X is the simple random walk on B and τ is its first hitting time of the set B.

Electrical networks: With the above discussion, we have related the gambler's ruin problem to the Dirichlet problem, without being able to solve either of them! Indeed, in general it is hopeless to expect an explicit solution. However, it is worth noting that the discrete Dirichlet problem arises in a different area that looks unrelated, namely that of electrical networks (a more sophisticated name is discrete potential theory). This will not bring any miracles, but the intuition from electrical networks can be of use in studying random walks and vice versa. Now we describe the electrical network formulation.

Imagine that G is an electric network where each edge is replaced by a unit resistor. The vertices in A are connected to batteries and the voltages at these points are maintained at $\phi(x), x \in A$. Then, electric current flows through the network and at each vertex a voltage is established. According to Kirchoff's law, the voltages at the vertices are precisely the solution to the discrete Dirichlet problem. Here is an example where we use knowledge of reduction in electrical networks to find the voltage at one vertex. This reduction is very special, and for general graphs there is not much one can do.

Example 19

Let G be the tree shown in the picture below. It is a tree with one vertex of degree 3 which we call the root O, and three leaves A, B, C. Let the distance (the number of edges between the root and the leaf) to the three leaves be a, b, c respectively. What is the probability that a simple random walk starting from O hits A before $\{B, C\}$?

As we have seen, the answer is given by $f(O)$, where $f : V \mapsto \mathbb{R}$ is a function that is harmonic except at the leaves and $f(A) = 1$, $f(B) = f(C) = 0$. As discussed above, this is the same problem in electrical networks (with each edge replaced by a unit resistor) of finding the voltage function when batteries are connected so that A is maintained at voltage 1 and B, C at voltage 0. In high school, we have seen ruled for resistors in series and parallel, so this is the same problem as a graph with four vertices O', A', B', C', where $O'A', OB', OC'$ have resistances a, b, c,, respectively. Then the effective resistance between A' and $\{B', C'\}$ is $a + \frac{1}{\frac{1}{b}+\frac{1}{c}}$, hence the effective current is the reciprocal of this. Therefore, the voltage at O' is $\frac{a^{-1}}{a^{-1}+b^{-1}+c^{-1}}$.

Exercise: Show this by solving for the harmonic function on the tree (without using this network reduction business!).

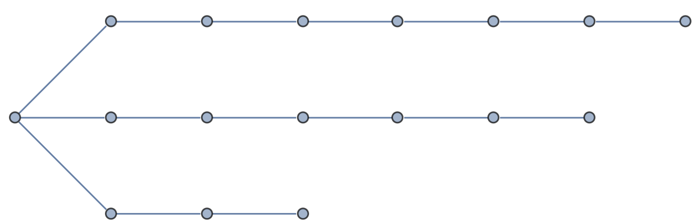

Variational principles: Using the terminology of electrical networks will not really help solve any problem. What do help are variational principles. Here is an easy exercise.

Exercise 11

Given a finite graph G and a proper subset of vertices B and a function $\phi : B \mapsto \mathbb{R}$, consider the functional $\mathcal{L}[f] := \sum_{u\sim v}(f(u) - f(v))^2$ on $\mathcal{H} = \{f : V \mapsto \mathbb{R} : f(x) = \phi(x) \text{ for all } x \in B\}$. Then the unique minimizer of L on $\mathcal{H}$ is the solution to the discrete Dirichlet problem with boundary data ϕ.

To illustrate the point, we now go to infinite graphs $G = (V, E)$ (again V is countable, each vertex has finite degree and G is connected). Recall that simple random walk X on G is recurrent if $\mathbb{P}_v\{\tau_v^+ < \infty\} = 1$ for some v (in which case it follows that $\mathbb{P}_w\{\tau_u < \infty\} = 1$ for all $w \neq u \in V$) where

$\tau_v^+ = \min\{n \geqslant 1 : X_n = v\}$ (observe the condition $n \geqslant 1$, as opposed to $n \geqslant 0$ in the definition of τ_v). If not recurrent, it is called transient.

Again, fixing v and consider $f(x) = \mathbb{P}_x\{\tau_v < \infty\}$. If the graph is recurrent, then $f = 1$ everywhere, whereas if it is transient, we may prove that $f(x) \to 0$ as $x \to \infty$ (i.e., given $\varepsilon > 0$, there is a finite $F \subseteq V$ such that $|f(x)| < \varepsilon$ for all $x \notin F$). This way, one may expect to prove a theorem (this statement is not quite true as stated) that the graph is transient if and only if there is a harmonic function $f : V \mapsto \mathbb{R}$ such that $f(v) = 1$, $f(x) \to 0$ as $x \to \infty$ and f is harmonic on $V \setminus \{v\}$. But this is still hard to use, because finding harmonic functions may be delicate. This is where the variational principle is useful. We state the following theorem without proof[2]. For a fixed $v \in V$, a cut-set is any collection of edges such that every infinite simple path starting from v must use one of the edges in Π.

Theorem 13

Let G be an infinite connected network and let v be a fixed vertex. The following are equivalent.

(1) SRW on G is transient.

(2) There exists $W : E \mapsto \mathbb{R}_+$ such that $\sum_{e \in \Pi} W(e) \geqslant 1$ for every cut-set Π and $\sum_{e \in E} W(e)^2 < \infty$.

To illustrate the usefulness of this theorem, let us prove Pólya's theorem for random walks on $\mathbb{Z}^d$. Let us fix the vertex 0 and consider the existence of a W as required in the theorem.

$d = 1$: Any pair of edges $\{[n, n+1], [-m-1, -m]\}$ where $n, m > 0$, is a cut-set. From that it is easy to see that $W([n, n+1]) \geqslant 1$ for infintiely many n (in fact for all positive n or for all negative n or both). But then $\sum W(e)^2 = \infty$, showing that the random walk must be recurrent.

$d = 2$: For any n, let $B(n) = \{-n, \ldots, n\}^2$. Let Π_n be the collection of edges that are in $B(n+1)$ but not in $B(n)$. There are $4(n+1)$ edges in $\Pi(n)$, and if the sum $\sum_{e \in \Pi_n} W(e) \geqslant 1$, then $\sum_{e \in \Pi_n} W(e)^2 \geqslant \frac{1}{4(n+1)}$ by Cauchy-Schwarz. As Π_ns are pairwise disjoint, this shows that $\sum_e W(e)^2 = \infty$.

$d \geqslant 3$. Define $W(e) = \frac{1}{|e|}$ where $|e|$ is the Euclidean distance from the origin to the mid-point of e. There are about n^{d-1} edges having $|e| \in [n, n+1]$, so the total sum of squares is like $\sum_n \frac{1}{n^{d-1}}$ which is finite. But is the condition $\sum_{e \in \Pi} W(e) \geqslant 1$ satisfied? For cut-sets of the form $B(n+1) \setminus B(n)$ where $B(n) = \{-n, \ldots, n\}^d$, this is clear. We leave the general case as an exercise.

The power of this theorem is in its robustness (as opposed to criteria such as $\sum_n p_{u,u}^{(n)} < \infty$ that we see in Markov chain class). If finitely many edges are added to the graph, it does not make a

[2]Chapter 2 of the book *Probability on trees and networks* by Lyons and Peres is an excellent resource for this subject. Another important resource is the paper *The extremal length of a network* by R. J. Duffin.

difference to the existence of W (also for finitely many deletions, provided it does not disconnect v from infinity) and hence to the question of recurrence or transience!

9. Maximal inequality

Kolmogorov's proof of his famous inequality was perhaps the first proof using martingales, although the term did not exist then!

Lemma 2: Kolmogorov's maximal inequality

Let ξ_k be independent random variables with zero means and finite variances. Let $S_n = \xi_1 + \ldots + \xi_n$. Then,

$$\mathbb{P}\left\{\max_{k\leqslant n}|S_k| \geqslant t\right\} \leqslant \frac{1}{t^2}\text{Var}(S_n).$$

Proof. We know that $(S_k)_{k\geqslant 0}$ is a martingale and $(S_k^2)_{k\geqslant 0}$ is a sub-martingale. Let $T = \min\{k : |S_k| \geqslant t\}\wedge n$ (i.e., T is equal to n or to the first time S exits $(-t, t)$, whichever is earlier). Then T is a bounded stopping time and $T \leqslant n$. By OST, $\{S_T^2, S_n^2\}$ is a sub-martingale and thus $\mathbb{E}[S_T^2] \leqslant \mathbb{E}[S_n^2]$. By Chebyshev's inequality,

$$\mathbb{P}\left\{\max_{k\leqslant n}|S_k| \geqslant t\right\} = \mathbb{P}\{S_T^2 \geqslant t^2\} \leqslant \frac{1}{t^2}\mathbb{E}[S_T^2] \leqslant \frac{1}{t^2}\mathbb{E}[S_n^2].$$

Thus the inequality follows. ■

This is an amazing inequality that controls the supremum of the entire path $S_0, S_1, \ldots, S_n$ in terms of the end-point alone! It takes a little thought to realize that the inequality $\mathbb{E}[S_T^2] \leqslant \mathbb{E}[S_n^2]$ is not a paradox. One way to understand it is to realize that if the path goes beyond $(-t, t)$, then there is a significant probability for the end point to be also large. This intuition is more clear in certain precursors to Kolmogorov's maximal inequality. In the following exercise you will prove one such, for symmetric, but not necessarily integrable, random variables.

Exercise 12

Let ξ_k be independent symmetric random variables and let $S_k = \xi_1 + \ldots + \xi_k$. Then for $t > 0$, we have

$$\mathbb{P}\left\{\max_{k\leqslant n} S_k \geqslant t\right\} \leqslant 2\mathbb{P}\{S_n \geqslant t\}.$$

Hint: Let T be the first time k when $S_k \geqslant t$. Given everything up to time $T = k$, consider the two possible future paths formed by $(\xi_{k+1}, \ldots, \xi_n)$ and $(-\xi_{k+1}, \ldots, -\xi_n)$. If $S_T \geqslant t$, then clearly for at least one of these two continuations, we must have $S_n \geqslant t$. Can you make this reasoning precise and deduce the inequality?

For a general super-martingale or sub-martingale, we can write similar inequalities that control the running maximum of the martingale in terms of the end-point.

Lemma 3: Doob's inequalities

Let X be a super-martingale. Then for any $t > 0$ and any $n \geqslant 1$,

(1) $\mathbb{P}\Big\{\max_{k \leqslant n} X_k \geqslant t\Big\} \leqslant \frac{1}{t}\{\mathbb{E}[X_0] + \mathbb{E}[(X_n)_-]\}$,

(2) $\mathbb{P}\Big\{\min_{k \leqslant n} X_k \leqslant -t\Big\} \leqslant \frac{1}{t}\mathbb{E}[(X_n)_-]$.

Proof. Let $T = \min\{k : X_k \geqslant t\} \wedge n$. By OST $\{X_0, X_T\}$ is a super-martingale and hence $\mathbb{E}[X_T] \leqslant \mathbb{E}[X_0]$. But

$$\begin{aligned}\mathbb{E}[X_T] &= \mathbb{E}[X_T \mathbf{1}_{X_T \geqslant t}] + \mathbb{E}[X_T \mathbf{1}_{X_T < t}] \\ &= \mathbb{E}[X_T \mathbf{1}_{X_T \geqslant t}] + \mathbb{E}[X_n \mathbf{1}_{X_T < t}] \\ &\geqslant \mathbb{E}[X_T \mathbf{1}_{X_T \geqslant t}] - \mathbb{E}[(X_n)_-]\end{aligned}$$

since $\mathbb{E}[X_n \mathbf{1}_A] \geqslant -\mathbb{E}[(X_n)_-]$ for any A. Thus, $\mathbb{E}[X_T \mathbf{1}_{X_T \geqslant t}] \leqslant \mathbb{E}[X_0] + \mathbb{E}[(X_n)_-]$. Now use Chebyshev's inequality to write $\mathbb{P}\{X_T \geqslant t\} \leqslant \frac{1}{t}\mathbb{E}[X_T \mathbf{1}_{X_T \geqslant t}]$ to get the first inequality.

For the second inequality, define $T = \min\{k : X_k \leqslant -t\} \wedge n$. By OST $\{(X_T), (X_n)\}$ is a super-martingale and hence $\mathbb{E}[X_T] \geqslant \mathbb{E}[X_n]$. But

$$\begin{aligned}\mathbb{E}[X_T] &= \mathbb{E}[X_T \mathbf{1}_{X_T \leqslant -t}] + \mathbb{E}[X_n \mathbf{1}_{X_T > -t}] \\ &\leqslant -t\mathbb{P}\{X_T \leqslant -t\} + \mathbb{E}[(X_n)_+].\end{aligned}$$

Hence $\mathbb{P}\{X_T \leqslant -t\} \leqslant \frac{1}{t}\{\mathbb{E}[(X_n)_+] - \mathbb{E}[X_n]\} = \frac{1}{t}\mathbb{E}[(X_n)_-]$. ■

For convenience, let us write down the corresponding inequalities for sub-martingales (which of course follow by applying Lemma 3 to $-X$): If $X_0, \ldots, X_n$ is a sub-martingale, then for any $t > 0$ we have

$$\mathbb{P}\{\max_{k \leqslant n} X_k \geqslant t\} \leqslant \frac{1}{t}\mathbb{E}[(X_n)_+], \tag{8}$$

$$\mathbb{P}\{\min_{k \leqslant n} X_k \leqslant -t\} \leqslant \frac{1}{t}\{-\mathbb{E}[X_0] + \mathbb{E}[(X_n)_+]\}. \tag{9}$$

If ξ_i are independent with zero mean and finite variances and $S_n = \xi_1 + \ldots + \xi_n$ is the corresponding random walk, then the above inequality when applied to the sub-martingale S_k^2 reduces to Kolmogorov's maximal inequality.

Maximal inequalities are useful in proving the Cauchy property of partial sums of a random series with independent summands. Here is an exercise.

Exercise 13

Let ξ_n be independent random variables with zero means. Assume that $\sum_n \mathrm{Var}(\xi_n) < \infty$. Show that $\sum_k \xi_k$ converges almost surely. [Extra: If interested, extend this to independent ξ_ks taking values in a separable Hilbert space H such that $\mathbb{E}[\langle \xi_k, u\rangle] = 0$ for all $u \in H$ and such that $\sum_n \mathbb{E}[\|\xi_n\|^2] < \infty$.]

10. Doob's up-crossing inequality

For a real sequence $x_0, x_1, \dots, x_n$ and any $a < b$, define the number of up-crossings of the sequence over the interval $[a, b]$ as the maximum number k for which there exist indices $0 \leqslant i_1 < j_1 < i_2 < j_2 < \dots < i_k < j_k \leqslant n$ such that $x_{i_r} \leqslant a$ and $x_{j_r} \geqslant b$ for all $r = 1, 2, \dots, k$. Intuitively it is the number of times the sequence crosses the interval in the upward direction. Similarly we can define the number of down-crossings of the sequence (same as the number of up-crossings of the sequence $(-x_k)_{0\leqslant k\leqslant n}$ over the interval $[-b, -a]$).

Lemma 4: Doob's up-crossing inequality

Let $X_0, \dots, X_n$ be a sub-martingale. Let $U_n[a, b]$ denote the number of up-crossings of the sequence $X_0, \dots, X_n$ over the interval $[a, b]$. Then,

$$\mathbb{E}[U_n[a, b] \mid \mathcal{F}_0] \leqslant \frac{\mathbb{E}[(X_n - a)_+ \mid \mathcal{F}_0] - (X_0 - a)_+}{b - a}.$$

What is the importance of this inequality? It will be in showing the convergence of martingales or super-martingales under some mild conditions. In continuous time, it will yield regularity of paths of martingales (existence of right and left limits).

The basic point is that a real sequence $(x_n)_n$ converges if and only if the number of up-crossings of the sequence over any interval is finite. Indeed, if $\liminf x_n < a < b < \limsup x_n$, then the sequence has infinitely many up-crossings and down-crossings over $[a, b]$. Conversely, if $\lim x_n$ exists, then the sequence is Cauchy and hence over any interval $[a, b]$ with $a < b$, there can be only finitely many up-crossings. In these statements the limit could be $\pm\infty$.

Proof. First assume that $X_n \geqslant 0$ for all n and that $a = 0$. Let $T_0 = 0$ and define the stopping times

$$T_1 := \min\{k \geqslant T_0 : X_k = 0\}, \quad T_3 := \min\{k \geqslant T_2 : X_k = 0\}, \quad \dots$$

$$T_2 := \min\{k \geqslant T_1 : X_k \geqslant b\}, \quad T_4 := \min\{k \geqslant T_3 : X_k \geqslant b\}, \quad \dots$$

where the minimum of an empty set is defined to be n. T_i are strictly increasing up to a point when T_k becomes equal to n and then the later ones are also equal to n. In what follows we only need T_k for $k \leqslant n$ (thus all the sums are finite sums).

$$
\begin{aligned}
X_n - X_0 &= \sum_{k\geqslant 0} X(T_{2k+1}) - X(T_{2k}) + \sum_{k\geqslant 1} X(T_{2k}) - X(T_{2k-1}) \\
&\geqslant \sum_{k\geqslant 0} (X(T_{2k+1}) - X(T_{2k})) + bU_n[0,b].
\end{aligned}
$$

The last inequality is because for each k for which $X(T_{2k}) \geqslant b$, we get one up-crossing and the corresponding increment $X(T_{2k}) - X(T_{2k-1}) \geqslant b$.

Now, by the optional sampling theorem (since $T_{2k+1} \geqslant T_{2k}$ are both bounded stopping times), we see that

$$
\mathbb{E}[X(T_{2k+1}) - X(T_{2k}) \mid \mathcal{F}_0] = \mathbb{E}\left[\mathbb{E}[X(T_{2k+1}) - X(T_{2k}) \mid \mathcal{F}_{T_{2k}}] \mid \mathcal{F}_0\right] \geqslant 0.
$$

Therefore, $\mathbb{E}[X_n - X_0 \mid \mathcal{F}_0] \geqslant b\mathbb{E}[U_n[0,b] \mid \mathcal{F}_0]$. This gives the up-crossing inequality when $a = 0$ and $X_n \geqslant 0$.

In the general situation, just apply the derived inequality to the sub-martingale $(X_k - a)_+$ (which crosses $[0, b-a]$ whenever X crosses $[a,b]$) to get

$$
\mathbb{E}[(X_n - a)_+ \mid \mathcal{F}_0] - (X_0 - a)_+ \geqslant (b-a)\mathbb{E}[U_n[a,b] \mid \mathcal{F}_0]
$$

which is what we claimed. ■

The break up of $X_n - X_0$ over up-crossing and down-crossings was okay, but how did the expectations of increments over down-crossings become non-negative? There is a distinct sense of something suspicious about this! The point is that $X(T_3) - X(T_2)$, for example, is not always non-negative. If X never goes below a after T_2, then it can be positive too. Indeed, the sub-martingale property somehow ensures that this positive part off sets the $-(b-a)$ increment that would occur if $X(T_3)$ did go below a.

We invoked OST in the proof. Optional sampling was in turn proved using the gambling lemma. It is an instructive exercise to write out the proof of the up-crossing inequality directly using the gambling lemma (start betting when below a, stop betting when reach above b, etc.).

11. Convergence theorem for L^2-bounded martingales

One of the fundamental results about martingales is that if it is uniformly integrable, then it converges almost surely and in L^1. The almost sure convergence holds under the weaker assumption that the martingale is L^1-bounded, i.e., $\sup_n \mathbb{E}|X_n| < \infty$. However, it is the L^1 convergence, in particular the convergence of expectations, that is most useful.

In this section, we give a proof of the convergence under the stronger assumption of L^2-boundedness. The conclusion is also strengthened to convergence in L^2. In many applications, this is sufficient, but we also prove the more general theorem in a later section.

Theorem 14: Square integrable martingales

Let X be a martingale on $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ such that $\sup_n \mathbb{E}[X_n^2] < \infty$. Then there is an L^2 random variable X_∞ such that $X_n \to X_\infty$ a.s. and in L^2.

First, a basic observation about a square integrable martingale X. Assume $\mathbb{E}[X_n^2] < \infty$ for each n (no need for a uniform bound). By the projection interpretation of conditional expectations, $X_{n+1} - X_n$ is an orthogonal to $L^2(\Omega, \mathcal{F}_n, \mathbb{P})$. In particular, $\{X_{k+1} - X_k\}_{k \geqslant 0}$ is an orthogonal set in $L^2(\Omega, \mathcal{F}, \mathbb{P})$ and hence for any $m > n$, we have

$$\mathbb{E}[(X_m - X_n)^2] = \sum_{k=n}^{m-1} \mathbb{E}[(X_{k+1} - X_k)^2]. \tag{10}$$

Proof of Theorem 14. Apply (10) with $n = 0$ and let $m \to \infty$ to see that

$$\sum_{k=0}^{\infty} \mathbb{E}[(X_{k+1} - X_k)^2] \leqslant \sup_m \mathbb{E}[(X_m - X_0)^2].$$

Under the L^2-boundedness assumption, the series on the left converges. Hence, $\mathbb{E}[(X_m - X_n)^2] \to 0$ as $m, n \to \infty$ by using (10) again, since the right side is the tail of a convergent series. Thus, $\{X_n\}$ is a Cauchy sequence in L^2 and hence there is some $X_\infty \in L^2$ such that $X_n \to X_\infty$ in L^2.

We now show almost sure convergence. Applying Doob's maximal inequality to the sub-martingale $\{|X_k - X_n|\}_{k \geqslant n}$, we get for any $m > n$,

$$\mathbb{P}\{\max_{n \leqslant k \leqslant m} |X_k - X_n| \geqslant \varepsilon\} \leqslant \frac{\mathbb{E}[|X_m - X_n|]}{\varepsilon} \leqslant \frac{1}{\varepsilon}\sqrt{\mathbb{E}[(X_m - X_n)^2]}.$$

As the latter goes to zero as $m, n \to \infty$, we see that

$$\mathbb{P}\{|X_k - X_j| \geqslant 2\varepsilon \text{ for some } k > j > n\} \to 0 \text{ as } n \to \infty.$$

Let $\varepsilon = \frac{1}{\ell}$ and choose N_ℓ so that for $n \geqslant N_\ell$, the probability of the event on the left is less than $\frac{1}{\ell^2}$. By Borel-Cantelli lemma, almost surely, only finitely many of these events occur. Therefore, the sequence $\{X_n\}$ is a Cauchy sequence, almost surely. Thus $X_n \overset{a.s.}{\to} X'_\infty$ for some X'_∞. However, the L^2 limit is X_∞, therefore $X_\infty = X'_\infty$ a.s. ■

Can we deduce the martingale convergence theorem for L^1-bounded martingales, by approximating them with L^2-bounded martingales? This is a tempting approach, but the naive way of doing it will give the result only under additional restrictions.

Theorem 15: Convergence for uniformly integrable martingales with uniformly bounded differences

Let $X = (X_n)_{n \geqslant 0}$ be a uniformly integrable martingale. Assume that $|X_{n+1} - X_n| \leqslant b$ a.s., for all n for some $b < \infty$. Then, X_n converges almost surely and in L^1 to an integrable random variable X_∞.

PROOF. Fix a positive integer M and let $\tau_M = \min\{k : |X_k| \geqslant M\}$. Then $\{X(\tau_M \wedge n)\}_{n\geqslant 0}$ is a martingale. Further, $|X(\tau_M \wedge n)| \leqslant M + b$, since the jumps are bounded by b, and the martingale is within $[-M, M]$ at time $\tau_M - 1$. Thus, $\{X(\tau_M \wedge n)\}$ is an L^2-bounded martingale and hence by Theorem 14, there is some $Z_M \in L^2$ such that $X(\tau_M \wedge n) \to Z_M$ a.s. and in L^2, as $n \to \infty$.

Further, applying Doob's maximal inequality, if $C = \sup_n \mathbb{E}[|X_n|]$, then

$$\mathbb{P}\{\tau_M < \infty\} = \lim_{n\to\infty} \mathbb{P}\{\tau_M \leqslant n\} \leqslant \frac{1}{M}\mathbb{E}[|X_n|] \leqslant \frac{C}{M}.$$

As $\tau_M \leqslant \tau_{M+1}$, it follows that $A = \cup_M\{\tau_M = \infty\}$ has probability 1. Further, on the event $\{\tau_M = \infty\}$, it is clear that $Z_{M'} = Z_M$ for all $M' > M$ (in fact, $X(\tau_M \wedge n) = X(\tau_{M'} \wedge n) = X(n)$ for all n). Therefore, we may consistently define a random variable Z by setting it equal to Z_M on the event $\{\tau_M = \infty\}$. It is then clear that $X_n \stackrel{a.s.}{\to} Z$ on the event A. Since $\mathbb{P}(A) = 1$, we have proved that $X_n \stackrel{a.s.}{\to} Z$.

The integrability of Z follows by Fatou's lemma and the remaining parts of the martingale convergence theorem (that uniform integrability implies L^1 convergence etc.) are general facts that follow once we have almost sure convergence. ■

12. Convergence theorem for super-martingales

In this and the next section, we present the general results on martingale convergence. Unlike the square integrable case where we used only the maximal inequality, here the proofs use the upcrossing inequality.

Theorem 16: Super-martingale convergence theorem

Let X be a super-martingale on $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$. Assume that $\sup_n \mathbb{E}[(X_n)_-] < \infty$.

(1) Then, $X_n \stackrel{a.s.}{\to} X_\infty$ for some integrable (hence finite) random variable X_∞.

(2) In addition, $X_n \to X_\infty$ in L^1 if and only if $\{X_n\}$ is uniformly integrable. If this happens, we also have $\mathbb{E}[X_\infty \mid \mathcal{F}_n] \leqslant X_n$ for each n.

In other words, when a super-martingale does not explode to $-\infty$ (in the mild sense of $\mathbb{E}[(X_n)_-]$ being bounded), it must converge almost surely!

PROOF. Fix $a < b$. Let $D_n[a, b]$ be the number of down-crossings of $X_0, \ldots, X_n$ over $[a, b]$. By applying the up-crossing inequality to the sub-martingale $-X$ and the interval $[-b, -a]$, and taking expectations, we get

$$\begin{aligned}\mathbb{E}[D_n[a, b]] &\leqslant \frac{\mathbb{E}[(X_n - b)_-] - \mathbb{E}[(X_0 - b)_-]}{b - a} \\ &\leqslant \frac{1}{b-a}(\mathbb{E}[(X_n)_-] + |b|) \leqslant \frac{1}{b-a}(M + |b|)\end{aligned}$$

where $M = \sup_n \mathbb{E}[(X_n)_-]$. Let $D[a,b]$ be the number of down-crossings of the whole sequence (X_n) over the interval $[a,b]$. Then $D_n[a,b] \uparrow D[a,b]$ and hence by MCT we see that $\mathbb{E}[D[a,b]] < \infty$. In particular, $D[a,b] < \infty$ w.p.1.

Consequently, $\mathbb{P}\{D[a,b] < \infty \text{ for all } a < b, \quad a, b \in \mathbb{Q}\} = 1$. Thus, X_n converges w.p.1., and we define X_∞ as the limit (for ω in the zero probability set where the limit does not exist, define X_∞ as 0). Thus $X_n \stackrel{a.s.}{\to} X_\infty$.

We observe that $\mathbb{E}[|X_n|] = \mathbb{E}[X_n] + 2\mathbb{E}[(X_n)_-] \leqslant \mathbb{E}[X_0] + 2M$. By Fatou's lemma, $\mathbb{E}[|X_\infty|] \leqslant \liminf \mathbb{E}[|X_n|] \leqslant 2M + \mathbb{E}[X_0]$. Thus X_∞ is integrable.

This proves the first part. The second part is very general - whenever $X_n \stackrel{a.s.}{\to} X$, we have L^1 convergence if and only if $\{X_n\}$ is uniformly integrable. Lastly, $\mathbb{E}[X_{n+m} \mid \mathcal{F}_n] \leqslant X_n$ for any $n, m \geqslant 1$. Let $m \to \infty$ and use L^1 convergence of X_{n+m} to X_∞ to get $\mathbb{E}[X_\infty \mid \mathcal{F}_n] \leqslant X_n$.

This completes the proof. ∎

A direct corollary that is often used is

Corollary 1

A non-negative super-martingale converges almost surely to a finite random variable.

13. Convergence theorem for martingales

Now we deduce the consequences for martingales.

Theorem 17: Martingale convergence theorem

Let $X = (X_n)_{n \geqslant 0}$ be a martingale with respect to $\mathcal{F}_\bullet$. Assume that X is L^1-bounded.

(1) Then, $X_n \stackrel{a.s.}{\to} X_\infty$ for some integrable (in particular, finite) random variable X_∞.

(2) In addition, $X_n \stackrel{L^1}{\to} X_\infty$ if and only if X is uniformly integrable. In this case, $\mathbb{E}[X_\infty \mid \mathcal{F}_n] = X_n$ for all n.

(3) If X is L^p bounded for some $p > 1$, then $X_\infty \in L^p$ and $X_n \stackrel{L^p}{\to} X_\infty$.

Observe that for a martingale the condition of L^1-boundedness, $\sup_n \mathbb{E}[|X_n|] < \infty$, is equivalent to the weaker looking condition $\sup_n \mathbb{E}[(X_n)_-] < \infty$, since $\mathbb{E}[|X_n|] - 2\mathbb{E}[(X_n)_-] = \mathbb{E}[X_n] = \mathbb{E}[X_0]$ is a constant.

Proof. The first two parts of the proof are immediate since a martingale is also a super-martingale. To conclude $\mathbb{E}[X_\infty \mid \mathcal{F}_n] = X_n$, we apply the corresponding inequality in the super-martingale convergence theorem to both X and to $-X$.

For the third part, if X is L^p bounded, then it is uniformly integrable and hence $X_n \to X_\infty$ a.s. and in L^1. To get L^p convergence, consider the non-negative sub-martingale $\{|X_n|\}$ and let $X^* = \sup_n |X_n|$. From Lemma 5 we conclude that $X^* \in L^p$. Of course, X^* dominates $|X_n|$ and $|X_\infty|$. Hence,

$$|X_n - X_\infty|^p \leqslant 2^{p-1}(|X_n|^p + |X_\infty|^p) \leqslant 2^p (X^*)^p$$

by the inequality $|a+b|^p \leqslant 2^{p-1}(|a|^p + |b|^p)$ by the convexity of $x \mapsto |x|^p$. Thus, $|X_n - X_\infty|^p \overset{a.s.}{\to} 0$ and the sequence is dominated by $2^p(X^*)^p$ which is integrable. Dominated convergence theorem shows that $\mathbb{E}[|X_n - X_\infty|^p] \to 0$. ■

We used the following lemma in the last part of the above proof[3]. This lemma is similar in spirit and in its use to maximal inequalities in analysis, such as the famous one of Hardy and Littlewood.

Lemma 5: A maximal inequality

Let $(Y_n)_{n\geqslant 0}$ be an L^p-bounded non-negative sub-martingale. Then $Y^* := \sup_n Y_n$ is in L^p and in fact $\mathbb{E}[(Y^*)^p] \leqslant C_p \sup_n \mathbb{E}[Y_n^p]$ where $C_p = \left(\frac{p}{p-1}\right)^p$.

Proof. Let $Y_n^* = \max_{k\leqslant n} Y_k$. Fix $\lambda > 0$ and let $T = \min\{k \geqslant 0 : Y_k \geqslant \lambda\}$. By the optional sampling theorem, for any fixed n, the sequence of two random variables $\{Y_{T\wedge n}, Y_n\}$ is a sub-martingale. Hence, $\int_A Y_n dP \geqslant \int_A Y_{T\wedge n} dP$ for any $A \in \mathcal{F}_{T\wedge n}$. Let $A = \{Y_{T\wedge n} \geqslant \lambda\}$ so that $\mathbb{E}[Y_n \mathbf{1}_A] \geqslant \mathbb{E}[Y_{T\wedge n}\mathbf{1}_{Y_{T\wedge n}\geqslant\lambda}] \geqslant \lambda\mathbb{P}\{Y_n^* \geqslant \lambda\}$. On the other hand, $\mathbb{E}[Y_n\mathbf{1}_A] \leqslant \mathbb{E}[Y_n \mathbf{1}_{Y^*\geqslant\lambda}]$ since $Y_n^* \leqslant Y^*$. Thus, $\lambda\mathbb{P}\{Y_n^* > \lambda\} \leqslant \mathbb{E}[Y_n\mathbf{1}_{Y^*\geqslant\lambda}]$.

Let $n \to \infty$. Since $Y_n^* \uparrow Y^*$, we get

$$\lambda\mathbb{P}\{Y^* > \lambda\} \leqslant \limsup_{n\to\infty} \lambda\mathbb{P}\{Y_n^* \geqslant \lambda\} \leqslant \limsup_{n\to\infty} \mathbb{E}[Y_n \mathbf{1}_{Y^*\geqslant\lambda}] = \mathbb{E}[Y_\infty \mathbf{1}_{Y^*\geqslant\lambda}].$$

where Y_∞ is the a.s. and L^1 limit of Y_n (exists, because $\{Y_n\}$ is L^p bounded and hence uniformly integrable). To go from the tail bound to the bound on pth moment, we use the identity $\mathbb{E}[(Y^*)^p] = \int_0^\infty p\lambda^{p-1}\mathbb{P}\{Y^* \geqslant \lambda\}d\lambda$ valid for any non-negative random variable in place of Y^*. Using the tail bound, we get

$$\begin{aligned}\mathbb{E}[(Y^*)^p] &\leqslant \int_0^\infty p\lambda^{p-2}\mathbb{E}[Y_\infty\mathbf{1}_{Y^*\geqslant\lambda}]d\lambda \leqslant \mathbb{E}\left[\int_0^\infty p\lambda^{p-2} Y_\infty \mathbf{1}_{Y^*\geqslant\lambda} d\lambda\right] \quad \text{(by Fubini's)}\\ &= \frac{p}{p-1}\mathbb{E}[Y_\infty \cdot (Y^*)^{p-1}].\end{aligned}$$

[3]Arghydeep Chatterjee suggested the following argument that does not require this lemma to prove L^p-convergence in Theorem 17. We already know that $X_n \to X_\infty$ a.s. and in L^1 and that $X_n = \mathbb{E}[X_\infty \mid \mathcal{F}_n]$. From $|X_n|^p \overset{a.s.}{\to} |X_\infty|^p$ and Fatou's lemma and the assumption of L^p-boundedness, we see that $\mathbb{E}[|X_\infty|^p] < \infty$. By the conditional Jensen's inequality, $|X_n|^p \leqslant \mathbb{E}[|X_\infty|^p \mid \mathcal{F}_n]$, hence $\{|X_n|^p\}$ is a uniformly integrable sequence. As $|X_n - X_\infty|^p \leqslant 2^p(|X|_n^p + |X_\infty|^p)$, the sequence $\{|X_n - X_\infty|^p\}$ is also uniformly integrable. But $|X_n - X_\infty|^p \overset{a.s.}{\to} 0$, hence we also get $\mathbb{E}[|X_n - X_\infty|^p] \to 0$.

Let q be such that $\frac{1}{q}+\frac{1}{p}=1$. By Hölder's inequality, $E[Y_\infty \cdot (Y^*)^{p-1}] \leqslant \mathbb{E}[Y_\infty^p]^{\frac{1}{p}}\mathbb{E}[(Y^*)^{q(p-1)}]^{\frac{1}{q}}$. Since $q(p-1)=p$, this gives us $\mathbb{E}[(Y^*)^p]^{1-\frac{1}{q}} \leqslant \frac{p}{p-1}\mathbb{E}[Y_\infty^p]^{\frac{1}{p}}$. Hence, $\mathbb{E}[(Y^*)^p] \leqslant C_p\mathbb{E}[Y_\infty^p]$ with $C_p = (p/(1-p))^p$. By virtue of Fatou's lemma, $\mathbb{E}[Y_\infty^p] \leqslant \liminf \mathbb{E}[Y_n^p] \leqslant \sup_n \mathbb{E}[Y_n^p]$. Thus, $\mathbb{E}[(Y^*)^p] \leqslant C_p \sup_n \mathbb{E}[Y_n^p]$. ■

Alternately, from the inequality $\lambda\mathbb{P}\{Y_n^* > \lambda\} \leqslant \mathbb{E}[Y_n \mathbf{1}_{Y^* \geqslant \lambda}]$ we could have (by similar steps, but without letting $n \to \infty$) arrived at a bound of the form $\mathbb{E}[(Y_n^*)^p] \leqslant C_p\mathbb{E}[Y_n^p]$. The right hand side is bounded by $C_p \sup_n \mathbb{E}[Y_n^p]$ while the left hand side increases to $\mathbb{E}[(Y^*)^p]$ by monotone convergence theorem. This is another way to complete the proof.

One way to think of the martingale convergence theorem is that we have extended the martingale from the index set $\mathbb{N}$ to $\mathbb{N} \cup \{+\infty\}$ retaining the martingale property. Indeed, the given martingale sequence is the Doob martingale given by the limit variable X_∞ with respect to the given filtration.

While almost sure convergence is remarkable, it is not strong enough to yield useful conclusions. Convergence in L^1 or L^p for some $p \geqslant 1$ are much more useful. In this context, it is important to note that L^1-bounded martingales do not necessarily converge in L^1.

Example 20: Critical branching process

Consider a Galton-Watson tree (branching process) with mean off-spring distribution equal to 1 (any non-degenerate distribution will do, eg., Poisson(1)). Then if Z_n denotes the number of individuals in the nth generation (we start with $Z_0 = 1$), then Z_n is a non-negative martingale, and $\mathbb{E}[Z_n] = 1$, hence it is L^1-bounded. But $Z_\infty = 0$ (either recall this fact from previous classes, or prove it from the martingale convergence theorem!). Thus $\mathbb{E}[Z_n] \not\to \mathbb{E}[Z_\infty]$, showing that L^1-convergence fails.

14. Reverse martingales

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Let $\mathcal{F}_i$, $i \in I$ be sub-sigma algebras of $\mathcal{F}$ indexed by a partially ordered set $(I, \leqslant)$ such that $\mathcal{F}_i \subseteq \mathcal{F}_j$ whenever $i \leqslant j$. Then, we may define a martingale or a sub-martingale etc., with respect to this "filtration" $(\mathcal{F}_i)_{i \in I}$. For example, a martingale is a collection of integrable random variables X_i indexed by $i \in I$ such that X_i is $\mathcal{F}_i$-measurable and $\mathbb{E}[X_j \mid \mathcal{F}_i] = X_i$ whenever $i \leqslant j$.

If the index set is $-\mathbb{N} = \{0, -1, -2, \ldots\}$ with the usual order, we say that X is a reverse martingale or a reverse sub-martingale etc.

What is different about reverse martingales as compared to martingales is that our questions will be about the behaviour as $n \to -\infty$, towards the direction of decreasing information. It turns out that the results are even cleaner than for martingales!

Theorem 18: Reverse martingale convergence theorem

Let $X = (X_n)_{n \leqslant 0}$ be a reverse martingale. Then $\{X_n\}$ is uniformly integrable. Further, there exists a random variable $X_{-\infty}$ such that $X_n \to X_{-\infty}$ almost surely and in L^1.

Proof. Since $X_n = \mathbb{E}[X_0 \mid \mathcal{F}_n]$ for all n, the uniform integrability follows from Exercise **??**.

Let $U_n[a, b]$ be the number of down-crossings of $X_n, X_{n+1}, \ldots, X_0$ over $[a, b]$. The up-crossing inequality (applied to $X_n, \ldots, X_0$ over $[a, b]$) gives $\mathbb{E}[U_n[a, b]] \leqslant \frac{1}{b-a}\mathbb{E}[(X_0-a)_+]$. Thus, the expected number of up-crossings $U_\infty[a, b]$ by the full sequence $(X_n)_{n \leqslant 0}$ has finite expectation, and hence is finite w.p.1.

As before, w.p.1., the number of down-crossings over any interval with rational end-points is finite. Hence, $\lim_{n \to -\infty} X_n$ exists almost surely. Call this $X_{-\infty}$. Uniform integrability shows that convergence also takes place in L^1. ■

What about reverse super-martingales or reverse sub-martingales? Although we shall probably have no occasion to use this, here is the theorem which can be proved on the same lines.

Theorem 19

Let $(X_n)_{n \leqslant 0}$ be a reverse super-martingale. Assume that $\sup_n \mathbb{E}[X_n] < \infty$. Then $\{X_n\}$ is uniformly integrable and X_n converges almost surely and in L^1 to some random variable $X_{-\infty}$.

Proof. Exercise. ■

This covers almost all the general theory that we want to develop. The rest of the course will consist in milking these theorems to get many interesting consequences.

CHAPTER 3

Martingales: applications

1. Lévy's forward and backward laws

Let X be an integrable random variable on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

Question 1: If $\mathcal{F}_n$, $n \geqslant 0$, is an increasing sequence of sigma-algebras, then what happens to the sequence $\mathbb{E}[X \mid \mathcal{F}_n]$ as $n \to \infty$?

Question 2: If $\mathcal{G}_n$, $n \geqslant 0$ is a decreasing sequence of sigma-algebras, then what happens to $\mathbb{E}[X \mid \mathcal{G}_n]$ as $n \to \infty$.

Note that the question here is different from conditional MCT. The random variable is fixed and the sigma-algebras are changing. A natural guess is that the limit might be $\mathbb{E}[X \mid \mathcal{F}_\infty]$ and $\mathbb{E}[X \mid \mathcal{G}_\infty]$ respectively, where $\mathcal{F}_\infty = \sigma\{\bigcup_n \mathcal{F}_n\}$ and $\mathcal{G}_\infty = \bigcap_n \mathcal{G}_n$. We shall prove that these guesses are correct.

Forward case: The sequence $X_n = \mathbb{E}[X \mid \mathcal{F}_n]$ is a martingale because of the tower property $\mathbb{E}[\mathbb{E}[X \mid \mathcal{F}_n] \mid \mathcal{F}_m] = \mathbb{E}[X \mid \mathcal{F}_m]$ for $m < n$. Recall that such martingales are called Doob martingales.

Being conditional expectations of a given X, the martingale is uniformly integrable and hence X_n converges a.s.and in L^1 to some X_∞. We claim that $X_\infty = \mathbb{E}[X \mid \mathcal{F}_\infty]$ a.s..

Indeed, X_n is $\mathcal{F}_\infty$-measurable for each n and hence the limit X_∞ is $\mathcal{F}_\infty$-measurable (since the convergence is almost sure, there is a null set issue which can be dealt with by completing the sigma-algebras. Alternately, define $X_\infty(\omega) = \lim X_n(\omega)$ when the limit exists, and $X_\infty(\omega) = 0$ when $\lim X_n(\omega)$ does not exist. Then X_∞ is $\mathcal{F}_\infty$-measurable).

Define the measure μ and ν on $\mathcal{F}_\infty$ by $\mu(A) = \int_A X dP$ and $\nu(A) = \int_A X_\infty dP$ for $A \in \mathcal{F}_\infty$. What we want to show is that $\mu(A) = \nu(A)$ for all $A \in \mathcal{F}_\infty$. If $A \in \mathcal{F}_m$, then for any $n > m$, we have

$$\int_A X dP = \int_A X_m P = \int_A X_n dP \overset{n\to\infty}{\longrightarrow} \int_A X_\infty dP.$$

The first inequality holds because $X_m = \mathbb{E}[X \mid \mathcal{F}_m]$ and the second equality holds because $X_m = \mathbb{E}[X_n \mid \mathcal{F}_m]$ for $n > m$. The last convergence holds because $X_n \to X$ in L^1. Comparing the first and last quantities in the above display, we see that $\mu(A) = \nu(A)$ for all $A \in \bigcup_m \mathcal{F}_m$.

Thus, $\bigcup_n \mathcal{F}_n$ is a π-system on which μ and ν agree. By the $\pi - \lambda$ theorem, they agree of $\mathcal{F}_\infty = \sigma\{\bigcup_n \mathcal{F}_n\}$. This completes the proof that $\mathbb{E}[X \mid \mathcal{F}_n] \xrightarrow{\text{a.s.}, L^1} \mathbb{E}[X \mid \mathcal{F}_\infty]$.

Backward case: Write $X_{-n} = \mathbb{E}[X \mid \mathcal{G}_n]$ for $n \in \mathbb{N}$. Then X is a reverse martingale w.r.t the filtration $\mathcal{G}_{-n}$, $n \in \mathbb{N}$. By the reverse martingale convergence theorem, we get that X_n converges almost surely and in L^1 to some X_∞.

We claim that $X_\infty = \mathbb{E}[X \mid \mathcal{G}_\infty]$. Since X_∞ is $\mathcal{G}_n$ measurable for every n (being the limit of X_k, $k \geqslant n$), it follows that X_∞ is $\mathcal{G}_\infty$-measurable. Let $A \in \mathcal{G}_\infty$. Then $A \in \mathcal{G}_n$ for any n and hence $\int_A X dP = \int_A X_n dP$ which converges to $\int_A X_\infty dP$. Thus, $\int_A X dP = \int_A X_\infty dP$ for all $A \in \mathcal{F}_\infty$.

2. Kolmogorov's zero-one law

As a corollary of the forward law, we may prove Kolmogorov's zero-one law.

Theorem 20: Kolmogorov's zero-one law

Let ξ_n, $n \geqslant 1$ be independent random variables and let $\mathcal{T} = \bigcap_n \sigma\{\xi_n, \xi_{n+1}, \ldots\}$ be the tail sigma-algebra of this sequence. Then $\mathbb{P}(A)$ is 0 or 1 for every $A \in \mathcal{T}$.

Proof. Let $\mathcal{F}_n = \sigma\{\xi_1, \ldots, \xi_n\}$. Then $\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}_n] \to \mathbb{E}[\mathbf{1}_A \mid \mathcal{F}_\infty]$ in L^1 and almost surely. But $\mathcal{F}_\infty = \sigma\{\xi_1, \xi_2, \ldots\}$. Thus if $A \in \mathcal{T} \subseteq \mathcal{F}_\infty$ then $\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}_\infty] = \mathbf{1}_A$ a.s. On the other hand, $A \in \sigma\{\xi_{n+1}, \xi_{n+2}, \ldots\}$ from which it follws that A is independent of $\mathcal{F}_n$ and hence $\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}_n] = \mathbb{E}[\mathbf{1}_A] = \mathbb{P}(A)$. The conclusion is that $\mathbf{1}_A = \mathbb{P}(A)$ a.s., which is possible if and only if $\mathbb{P}(A)$ equals 0 or 1. ■

3. Strong law of large numbers

The strong law of large number under first moment condition is an easy consequence of the reverse martingale theorem.

Theorem 21

Let ξ_n, $n \geqslant 1$ be i.i.d. real-valued random variables with zero mean and let $S_n = \xi_1 + \ldots + \xi_n$. Then $\frac{1}{n} S_n \stackrel{\text{a.s.}}{\to} 0$.

Proof. Let $\mathcal{G}_n = \sigma\{S_n, S_{n+1}, \ldots\} = \sigma\{S_n, \xi_{n+1}, \xi_{n+2}, \ldots\}$, a decreasing sequence of sigma-algebras. Hence $M_{-n} := \mathbb{E}[\xi_1 \mid \mathcal{G}_n]$ is a reverse martingale and hence converges almost surely and in L^1 to some $M_{-\infty}$.

But $\mathbb{E}[\xi_1 \mid \mathcal{G}_n] = \frac{1}{n} S_n$ (why?). Thus, $\frac{1}{n} S_n \to M_{-\infty}$ almost surely and in L^1. But the limit of $\frac{1}{n} S_n$ is clearly a tail random variable of ξ_ns and hence must be constant. Thus, $M_{-\infty} = \mathbb{E}[M_{-\infty}] = \lim \frac{1}{n} \mathbb{E}[S_n] = 0$. In conclusion, $\frac{1}{n} S_n \stackrel{\text{a.s.}}{\to} 0$. ■

4. Critical branching process

Let Z_n, $n \geqslant 0$ be the generation sizes of a Galton-Watson tree with offspring distribution $p = (p_k)_{k\geqslant 0}$. If $m = \sum_k kp_k$ is the mean, then Z_n/m^n is a martingale (we saw this earlier).

If $m < 1$, then $\mathbb{P}\{Z_n \geqslant 1\} \leqslant \mathbb{E}[Z_n] = m^n \to 0$ and hence, the branching process becomes extinct w.p.1. For $m = 1$ this argument fails. We show using martingales that extinction happens even in this cases.

Theorem 22

If $m = 1$ and $p_1 \neq 1$, then the branching process becomes extinct almost surely.

Proof. If $m = 1$, then Z_n is a non-negative martingale and hence converges almost surely to a finite random variable Z_∞. But Z_n is integer-valued. Thus,

$$Z_\infty = j \Leftrightarrow Z_n = j \text{ for all } n \geqslant n_0 \text{ for some } n_0.$$

But if $j \neq 0$ and $p_1 < 1$, then it is easy to see that $\mathbb{P}\{Z_n = j \text{ for all } n \geqslant n_0\} = 0$ (since conditional on $\mathcal{F}_{n-1}$, there is a positive probability of p_0^j that $Z_n = 0$). Thus, $Z_n = 0$ eventually. ∎

In the supercritical case we know that there is a positive probability of survival. If you do not know this, prove it using the second moment method as follows.

Exercise 14

By conditioning on $\mathcal{F}_{n-1}$ (or by conditioning on $\mathcal{F}_1$), show that (1) $\mathbb{E}[Z_n] = m^n$, (2) $\mathbb{E}[Z_n^2] \asymp (1+\sigma^2)m^{2n}$. Deduce that $\mathbb{P}\{Z_n > 0\}$ stays bounded away from zero. Conclude positive probability of survival.

We also have the martingale Z_n/m^n. By the martingale convergence theorem $W := \lim Z_n/m^n$ exists, a.s. On the event of extinction, clearly $W = 0$. On the event of survival, is it necessarily the case that $W > 0$ a.s.? If yes, this means that whenever the branching process survives, it does so by growing exponentially, since $Z_n \sim W\, m^n$. The answer is given by the famous Kesten-Stigum theorem.

Theorem 23: Kesten-Stigum theorem

Assume that $\mathbb{E}[L] > 1$ and that $p_1 \neq 1$. Then, $W > 0$ almost surely on the event of survival if and only if $\mathbb{E}[L \log_+ L] < \infty$.

We now prove a weaker form of this, that if $\mathbb{E}[L^2] < \infty$, then $W > 0$ on the event of survival (this was in fact proved by Kolmogorov earlier).

Kesten-Stigum under finite variance condition. Assume $\sigma^2 = \mathbb{E}[L^2] < \infty$. Then by Exercise 14, $\frac{Z_n}{m^n}$ is an L^2 bounded martingale. Therefore it converges to W almost surely and in L^2.

In particular, $\mathbb{P}(W = 0) < 1$. However, by conditioning on the first generation, we see that $q = \mathbb{P}\{W = 0\}$ satisfies the equation $q = \mathbb{E}[q^L]$ (if the first generation has L children, in each of the trees under these individuals, the corresponding $W_i = 0$ and these W_i are independent). But the usual proof of the extinction theorem shows that there are only two solutions to the equation $q = \mathbb{E}[q^L]$, namely 1 and the extinction probability of the tree. Since we have see that $q < 1$, it must be equal to the extinction probability. That is $W > 0$ a.s. on the event of survival. ■

5. Pólya's urn scheme

Initially the urn contain b black and w white balls. Let B_n be the number of black balls after n steps. Then $W_n = b + w + n - B_n$. We have seen that $X_n := B_n/(B_n + W_n)$ is a martingale. Since $0 \leqslant X_n \leqslant 1$, uniform integrability is obvious and $X_n \to X_\infty$ almost surely and in L^1. Since X_n are bounded, the convergence is also in L^p for every $p < \infty$.

Theorem 24

X_∞ bas Beta(b, w) distribution.

Proof. Let V_k be the colour of the kth ball drawn. It takes values 1 (for black) and 0 (for white). It is an easy exercise to check that

$$\mathbb{P}\{V_1 = \varepsilon_1, \ldots, V_m = \varepsilon_m\} = \frac{b(b+1)\ldots(b+r-1)w(w+1)\ldots(w+s-1)}{(b+w)(b+w+1)\ldots(b+w+n-1)}$$

if $r = \varepsilon_1 + \ldots + \varepsilon_m$ and $s = n - r$. The key point is that the probability does not depend on the order of ε_is. In other words, any permutation of $(V_1, \ldots, V_n)$ has the same distribution as $(V_1, \ldots, V_n)$, a property called *exchangeability*.

From this, we see that for any $0 \leqslant r \leqslant n$, we have

$$\mathbb{P}\left\{X_n = \frac{b+r}{b+w+n}\right\} = \binom{n}{r}\frac{b(b+1)\ldots(b+r-1)w(w+1)\ldots(w+(n-r)-1)}{(b+w)(b+w+1)\ldots(b+w+n-1)}.$$

In the simplest case of $b = w = 1$, the right hand side is $\frac{1}{n+1}$. That is, X_n takes the values $\frac{r+1}{n+2}$, $0 \leqslant r \leqslant n$, with equal probabilities. Clearly then $X_n \xrightarrow{d}$ Unif$[0, 1]$. Hence, $X_\infty \sim$ Unif$[0, 1]$.

In general, we write the above probability as

$$\begin{aligned}\mathbb{P}\left\{X_n = \frac{b+r}{b+w+n}\right\} &= \frac{n!}{r!(n-r)!}\frac{\Gamma(b+r)\Gamma(w+n-r)\Gamma(b+w)}{\Gamma(b)\Gamma(w)\Gamma(b+w+n)} \\ &\sim \frac{\Gamma(b+w)}{\Gamma(b)\Gamma(w)}\frac{n^{n+\frac{1}{2}}(b+r)^{b+r-\frac{1}{2}}(w+n-r)^{w+n-r-\frac{1}{2}}}{r^{r+\frac{1}{2}}(n-r)^{n-r+\frac{1}{2}}(b+w+n)^{b+w+n-\frac{1}{2}}}\end{aligned}$$

using Stirlings' approximation (valid if both r and $n-r$ are large). Now $b+r \sim r$ and $w+n-r \sim n-r$ and $b+w+n \sim n$ and hence we arrive at

$$\mathbb{P}\left\{X_n = \frac{b+r}{b+w+n}\right\} \sim \frac{\Gamma(b+w)}{\Gamma(b)\Gamma(w)}\left(\frac{r}{n}\right)^{b-1}\left(1-\frac{r}{n}\right)^{w-1}\frac{1}{n}$$

from which it is easy to see that $X_n \stackrel{d}{\to} \text{Beta}(b,w)$. Hence $X_\infty \sim \text{Beta}(b,w)$. ■

Here is a possibly clever way to avoid computations in the last step.

Exercise 15

For each initial value of b, w, let $\mu_{b,w}$ be the distribution of X_∞ when the urn starts with b black and w white balls. Each $\mu_{b,w}$ is a probability measure on $[0,1]$.

(1) Show that $\mu_{b,w} = \frac{b}{b+w}\mu_{b+1,w} + \frac{w}{b+w}\mu_{b,w+1}$.

(2) Check that $\text{Beta}(b,w)$ distributions satisfy the above recursions.

(3) Assuming $(b,w) \mapsto \mu_{b,w}$ is continuous, deduce that $\mu_{b,w} = \text{Beta}(b,w)$ is the only solution to the recursion.

One can introduce many variants of Pólya's urn scheme. For example, whenever a ball is picked, we may add r balls of the same color and q balls of the opposite color. That changes the behaviour of the urn greatly and in a typical case, the proportions of black balls converges to a constant.

Here is a muti-color version which shares all the features of Pólya's urn above.

Multi-color Pólya's urn scheme: We have ℓ colors denoted $1, 2, \ldots, \ell$. Initially an urn contains $b_k > 0$ balls of color k (b_k need not be integers). At each step of the process, a ball is drawn uniformly at random from the urn, its color noted, and returned to the urn with another ball of the same color. Let $B_k(n)$ be the number of balls of kth color after n draws. Let ξ_n be the color of the ball drawn in the nth draw.

Exercise 16

(1) Show that $\frac{1}{n+b_1+\ldots+b_\ell}(B_1(n), \ldots, B_\ell(n))$ converges almost surely (and in L^p for any p) to some random vector $(Q_1, \ldots, Q_\ell)$.

(2) Show that $\xi_1, \xi_2, \ldots$ is an exchangeable sequence.

(3) For $b_1 = \ldots = b_\ell = 1$, show that $(Q_1, \ldots, Q_\ell)$ has Dirichlet$(1, 1, \ldots, 1)$ distribution. In general, it has Dirichlet$(b_1, \ldots, b_\ell)$ distribution.

This means that $Q_1 + \ldots + Q_\ell = 1$ and $(Q_1, \ldots, Q_{\ell-1})$ has density

$$\frac{\Gamma(b_1 + \ldots + b_\ell)}{\Gamma(b_1) \ldots \Gamma(b_\ell)} x_1^{b_1-1} \ldots x_{\ell-1}^{b_{\ell-1}-1}(1 - x_1 - \ldots - x_{\ell-1})^{b_\ell - 1}$$

on $\Delta = \{(x_1, \ldots, x_{\ell-1}) : x_i > 0 \text{ for all } i \text{ and } x_1 + \ldots + x_{\ell-1} < 1\}$.

Blackwell-Macqueen urn scheme: Here is a generalization of Pólya's urn scheme to infinitely many colours. Start with the unit line segment $[0,1]$, each point of which is thought of as a distinct colour. Pick a uniform random variable V_1, after which we add a line segment of length 1 that has colour V_1 (you may imagine that the new segment is attached to the old one at the point V_1). Now we have the original line segment and a new line segment, and we draw a point uniformly at random from the union of the two line segments. If it falls in the original segment at location V_2, a new line segment of colour V_2 is added and if it falls in the segment of colour V_1, then a new line segment of length 1 having colour V_1 is added. The process continues.

If one considers the situation after the first step, the colour V_1 is like the black in a Pólya's urn scheme with $b = 1 = w$. Hence the proportion of V_1 converges almost surely to $P_1 \sim \text{unif}[0,1]$. When the kth colour appears, it appears with a line segment of length 1 and the original line segment has length 1. If we ignore all the points that fall in the other coloured segments that have appeared before, then again we have a Pólya urn with $b = w = 1$. This leads to the following conclusion: The proportions of the colours that appear, in the order of appearance, converges almost surely to $(P_1, P_2, \ldots)$ where $P_1 = U_1$, $P_2 = (1-U_1)U_2$, $P_3 = (1-U_1)(1-U_2)U_3, \ldots$ where U_i are i.i.d. uniform random variables on $[0,1]$.

The random vector P has a distribution on the infinite simplex $\Delta = \{(p_1, p_2, \ldots) : p_i \geqslant 0, \ \sum_i p_i = 1\}$ that is known as a GEM distribution (for Griffiths-Engel-McCloskey) and random vector $P^{\downarrow}$ got from P by ranking the co-ordinates in decreasing order is said to have Poisson-Dirichlet distribution (on the ordered simplex $\Delta^{\downarrow} = \{(p_1, p_2, \ldots) : p_1 \geqslant p_2 \geqslant \ldots \geqslant 0 \text{ and } \sum_i p_i = 1\}$. If we allow the initial stick to have length $\theta > 0$ (the segments added still have length 1), then the resulting distribution on Δ and $\Delta^{\downarrow}$ are called $\text{GEM}(0,\theta)$ and $\text{PD}(0,\theta)$ distributions.

6. Liouville's theorem

Recall that a harmonic function on $\mathbb{Z}^2$ is a function $f : \mathbb{Z}^2 \to \mathbb{R}$ such that $f(x) = \frac{1}{4}\sum_{y:y\sim x} f(y)$ for all $x \in \mathbb{Z}^2$.

Theorem 25: Liouville's theorem

If f is a non-constant harmonic function on $\mathbb{Z}^2$, then $\sup f = +\infty$ and $\inf f = -\infty$.

Proof. If not, by negating and/or adding a constant we may assume that $f \geqslant 0$. Let X_n be simple random walk on $\mathbb{Z}^2$. Then $f(X_n)$ is a martingale. But a non-negative super-martingale converges almost surely. Hence $f(X_n)$ converges almost surely.

But Pólya's theorem says that X_n visits every vertex of $\mathbb{Z}^2$ infinitely often w.p.1. This contradicts the convergence of $f(X_n)$ unless f is a constant. ■

Observe that the proof shows that a non-constant super-harmonic function on $\mathbb{Z}^2$ cannot be bounded below. The proof uses recurrence of the random walk. But in fact the same theorem holds on $\mathbb{Z}^d$, $d \geqslant 3$, although the simple random walk is transient there.

For completeness, here is a quick proof of Pólya's theorem in two dimensions.

Exercise 17

Let S_n be simple symmetric random walk on $\mathbb{Z}^2$ started at $(0,0)$.

(1) Show that $\mathbb{P}\{S_{2n} = (0,0)\} = \frac{1}{4^{2n}} \sum_{k=0}^{n} \frac{(2n)!}{k!^2 (n-k)!^2}$ and that this expression reduces to $\left(\frac{1}{2^{2n}}\binom{2n}{n}\right)^2$.

(2) Use Stirling's formula to show that $\sum_n \mathbb{P}\{S_{2n} = (0,0)\} = \infty$.

(3) Conclude that $\mathbb{P}\{S_n = (0,0) \text{ i.o.}\} = 1$.

The question of existence of bounded or positive harmonic functions on a graph (or in the continuous setting) is important. Here are two things that we may cover if we get time.

- There are no bounded harmonic functions on $\mathbb{Z}^d$ (Blackwell).
- Let μ be a probability measure on $\mathbb{R}$ and let f be a harmonic function for the random walk with step distribution μ. This just means that f is continuous and $\int_{\mathbb{R}} f(x+a)d\mu(x) = f(a)$. Is f necessarily constant? We shall discuss this later (under the heading "Choquet-Deny theorem").

To prove Blackwell's theorem, we first prove a lemma for general Markov chains. Let $P = (p_{i,j})_{i,j \in S}$ be a Markov transition matrix on a countable state space S. A function $f : S \to \mathbb{R}$ is said to be P-harmonic if $f(i) = \sum_{j \in S} p_{i,j} f(j)$ for all $i \in S$. If $X = (X_n)_{n \geqslant 0}$ is a Markov chain having transition P, then the P-harmonicity of f can be equivalently stated as $\mathbb{E}[f(X_{n+1}) \mid X_0, \ldots, X_n] = f(X_n)$. In other words, $(f(X_n))_n$ is a martingale.

Lemma 6

In the above setting, assume that for any $i, j \in S$, there is a coupling (X_n, Y_n) such that individually X and Y are Markov chains with transition matrix P and initial states i, j respectively, and such that $\tau = \min\{n : X_n = Y_n\} < \infty$ a.s. Then, any bounded P-harmonic function on S is constant.

Proof. Fix $i, j \in S$ and a coupling (X_n, Y_n) as assumed. Let f be a P-harmonic function such that $|f| \leqslant M$. As already observed, $f(X_n)$ and $f(Y_n)$ are both martingales and hence by the optional stopping theorem $f(i) = \mathbb{E}[f(X_{\tau \wedge n})]$ and $f(j) = \mathbb{E}[f(Y_{\tau \wedge n})]$. As $|f| \leqslant M$ and $\tau \wedge n \to \tau$ a.s., DCT implies that $\mathbb{E}[f(X_{\tau\wedge n})] \to \mathbb{E}[f(X_\tau)]$ and $\mathbb{E}[f(Y_{\tau\wedge n})] \to \mathbb{E}[f(Y_\tau)]$ as $n \to \infty$. Hence $f(i) = \mathbb{E}[f(X_\tau)] = f(j)$ which implies that f is a constant. ∎

Remark 5

It was not really necessary to use the optional stopping theorem in the proof above. Let $Z_n = Y_n \mathbf{1}_{n\leqslant\tau} + X_n \mathbf{1}_{n>\tau}$. Then (X_n, Z_n) is also a coupling with the same meeting time τ, but they stick together after that. By the martingale property, $f(i) = \mathbb{E}[f(X_n)]$ and $f(j) = \mathbb{E}[f(Z_n)]$. Therefore,

$$|f(i) - f(j)| = |\mathbb{E}[f(X_n) - f(Z_n)]| \leqslant \mathbb{E}[|f(X_n) - f(Z_n)|] \quad \leqslant 2M\mathbb{P}\{\tau > n\}$$

which converges to zero as $n \to \infty$. This shows that $f(i) = f(j)$ for any i, j.

Harmonic function on a graph is the same as the P-harmonic function where P is the transition for simple random walk on the graph. Observe that a P-harmonic function is also Q-harmonic if $Q = (P + I)/2$. The transition matrix Q is a *lazy version* of P, wherein at each step it stays put with probability $1/2$ and when it moves, it moves according to P.

Theorem 26: Blackwell: No bounded harmonic functions on $\mathbb{Z}^d$

If $f : \mathbb{Z}^d \to \mathbb{R}$ is bounded and harmonic, then f is constant.

Proof. Let P be the transition matrix for the lazy version of simple symmetric random walk on $\mathbb{Z}^d$. Then f is P-harmonic. By Lemma 6, it suffices to show that for any $x, y \in \mathbb{Z}^d$, we can couple the chains starting at x, y so that they meet eventually. This is achieved by as follows.

Let $X_0 = x$ and $Y_0 = y$. Here is how the steps are coupled at time n. Pick $U \sim \text{Unif}\{1, \dots, d\}$ and $\xi, \eta \sim \text{Unif}\{0, 1\}$, all independent and independent of the chains up to time n.

(1) If $U = i$ and $X_n(i) = Y_n(i)$, then let $(X_{n+1}, Y_{n+1}) - (X_n, Y_n)$ be equal to 0 if $\xi = 0$ and equal to (e_i, e_i) if $\xi = 1, \eta = 0$ and equal to $(-e_i, -e_i)$ if $\xi = 1, \eta = 1$.

(2) If $U = i$ and $X_n(i) < Y_n(i)$, then let $(X_{n+1}, Y_{n+1}) - (X_n, Y_n)$ be equal to $(e_i, 0)$ if $\xi = 0$ and equal to $(0, -e_i)$ if $\xi = 1$.

(3) If $U = i$ and $X_n(i) > Y_n(i)$, then let $(X_{n+1}, Y_{n+1}) - (X_n, Y_n)$ be equal to $(-e_i, 0)$ if $\xi = 0$ and equal to $(0, e_i)$ if $\xi = 1$.

We leave it as an exercise to check that X and Y are lazy simple symmetric random walks. Further, for any co-ordinate i, $X_n(i) - Y_n(i)$ is a lazy simple symmetric random walk in 1-dimension (with lazyness probability $1 - \frac{1}{2d}$) that gets absorbed at 0. Hence it will eventually get absorbed, in other words $X_n(i) = Y_n(i)$ for large enough n. This completes the proof of coupling. ∎

6.1. Regular trees. There are non-constant, bounded harmonic functions on certain graphs. For example, let $\mathcal{T}_b$ denote the infinite regular tree where each vertex has b neighbours. Assume that $b \geqslant 3$. Fix a vertex o as the root vertex and label its neighbours as $u_1, \ldots, u_b$. For $1 \leqslant k \leqslant b$, write V_k for the set of vertices from which the path to the root passes through u_k. For a vertex u, let $|u|$ denote its distance from the root.

$$f(u) = \begin{cases} \frac{1}{(b-1)^{|u|}b} & \text{if } u \in V_2 \cup \ldots \cup V_b, \\ \frac{1}{b} & \text{if } u = o, \\ 1 - \frac{1}{(b-1)^{|u|-1}b} & \text{if } u \in V_1. \end{cases} \tag{11}$$

Exercise 18

Show that f is harmonic. In particular, there exist bounded harmonic functions on the b-regular tree for $b \geqslant 3$.

The checking is routine, but more interesting is to explain where this function came from. For this, consider any graph $G = (V, E)$ and fix two disjoint subsets of vertices $A, B \subseteq V$. Define $g : V \to \mathbb{R}$ by

$$g(x) = \mathbb{P}_x\{\text{RW hits } A \text{ before } B\}.$$

Then by conditioning on the first step, it is easy to see that $g(x) = \frac{1}{\deg(x)} \sum_{y \sim x} g(y)$, provided $x \notin A$. However, on A, the mean value property fails. To get a function that is harmonic on the whole graph, we should try to push the sets A and B out of the graph. For the b-regular tree, we take $A = V_1 \cap \{|u| = N\}$ and $B = (V_2 \cup \ldots \cup V_b) \cap \{|u| = N\}$. Then defining g as above, we get a harmonic function on the complement of $A \cup B$, in particular on $\{u : |u| < N\}$. By reducing the problem to that of a biased random walk on integers, it is easy to obtain an explicit form for $g = g_N$. Then let $N \to \infty$ to get the function f defined above. It is harmonic everywhere.

An alternate way to explain this is to note that the SRW on the b-regular tree (for $b \geqslant 3$) is transient. Because of the tree structure, this implies that eventually the RW stays in one of the V_ks forever. The function f is precisely

$$f(x) = \mathbb{P}_x\{\text{RW eventually stays in } V_1\}. \tag{12}$$

Exercise 19

Consider a biased RW on $\mathbb{Z}$ that goes one step to the right with probability p and one step to the left with probability q where $p < q$. Show that $\mathbb{P}_x\{\text{RW hits } 0\} = (q/p)^x$ for all $x \geqslant 0$. Hence deduce that the function defined in (12) agrees with the one defined in (11).

7. Hewitt-Savage zero-one law

There are many zero-one laws in probability, asserting that a whole class of events are trivial. For a sequence of random variables, here are three important classes of such events.

Below, ξ_n, $n \geqslant 1$, are random variables on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and taking values in $(X, \mathcal{F})$. Then $\xi = (\xi_n)_{n\geqslant 1}$ is a random variable taking values in $(X^{\mathbb{N}}, \mathcal{F}^{\otimes \mathbb{N}})$. These definitions can be extended to two sided-sequences $(\xi_n)_{n\in\mathbb{Z}}$ easily.

(1) The *tail sigma-algebra* is defined as $\mathcal{T} = \cap_n \mathcal{T}_n$ where $\mathcal{T}_n = \sigma\{\xi_n, \xi_{n+1}, \ldots\}$.

(2) The *exchangeable sigma-algebra* $\mathcal{S}$ is the sigma-algebra of those events that are invariant under finite permutations.

More precisely, let G be the sub-group (under composition) of all bijections $\pi : \mathbb{N} \to \mathbb{N}$ such that $\pi(n) = n$ for all but finitely many n. It is clear how G acts on $X^{\mathbb{N}}$:

$$\pi((\omega_n)) = (\omega_{\pi(n)}).$$

Then

$$\mathcal{S} := \{\xi^{-1}(A) : A \in \mathcal{F}^{\otimes \mathbb{N}} \text{ and } \pi(A) = A \text{ for all } \pi \in G\}.$$

If G_n is the sub-group of $\pi \in G$ such that $\pi(k) = k$ for every $k > n$ and $\mathcal{S}_n := \{\xi^{-1}(A) : A \in \mathcal{F}^{\otimes \mathbb{N}}$ and $\pi(A) = A$ for all $\pi \in G_n\}$, then $\mathcal{S}_n$ are sigma-algebras that decrease to $\mathcal{S}$.

(3) The *translation-invariant sigma-algebra* $\mathcal{I}$ is the sigma-algebra of all events invariant under translations.

More precisely, let $\theta_n : X^{\mathbb{N}} \to X^{\mathbb{N}}$ be the translation map $[\theta_n(\omega)]_k = \omega_{n+k}$. Then, $\mathcal{I} = \{A \in \mathcal{F}^{\otimes \mathbb{N}} : \theta_n(A) = A \text{ for all } n \in \mathbb{N}\}$ (these are events invariant under the action of the semi-group $\mathbb{N}$).

Kolmogorov's zero-one law asserts that under and product measure $\mu_1 \otimes \mu_2 \otimes \ldots$, the tail sigma-algebra is trivial. Ergodicity is the statement that $\mathcal{I}$ is trivial and it is true for i.i.d. product measures $\mu^{\otimes \mathbb{N}}$. The exchangeable sigma-algebra is also trivial under i.i.d. product measure, which is the result we prove in this section. First an example.

Example 21

The event $A = \{\omega \in \mathbb{R}^{\mathbb{N}} : \lim \omega_n = 0\}$ is an invariant event. In fact, every tail event is an invariant event. But the converse is not true. For example,

$$A = \{\omega \in \mathbb{R}^{\mathbb{N}} : \lim_{n\to\infty} (\omega_1 + \ldots + \omega_n) \text{ exists and is at most } 0\}$$

is an invariant event but not a tail event. This is because $\omega = (-1, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots)$ belongs to A and so does every finite permutation of ω as the sum does not change. But changing the first co-ordinate to 0 gives $\omega' = (0, \frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots)$, which is not in A.

Theorem 27: Hewitt-Savage 0-1 law

Let μ be a probability measure on $(X, \mathcal{F})$. Then the invariant sigma-algebra $\mathcal{S}$ is trivial under the product measure $\mu^{\otimes \mathbb{N}}$.

In terms of random variables, we may state this as follows: Let ξ_n be i.i.d. random variables taking values in X. Let $f : X^{\mathbb{N}} \mapsto \mathbb{R}$ be a measurable function such that $f \circ \pi = f$ for all $\pi \in G$. Then, $f(\xi_1, \xi_2, \ldots)$ is almost surely a constant.

We give a proof using reverse martingale theorem. There are also more direct proofs.

Proof. For any integrable Y (that is measurable w.r.t $\mathcal{F}^{\otimes \mathbb{N}}$), the sequence $\mathbb{E}[Y \mid \mathcal{S}_n]$ is a reverse martingale and hence $\mathbb{E}[Y \mid \mathcal{S}_n] \xrightarrow{\text{a.s.}, L^1} \mathbb{E}[Y \mid \mathcal{S}]$.

Now fix $k \geq 1$ and let $\phi : X^k \to \mathbb{R}$ be a bounded measurable function. Take $Y = \phi(X_1, \ldots, X_k)$. We claim that

$$\mathbb{E}[\phi(X_1, \ldots, X_k) \mid \mathcal{S}_n] = \frac{1}{n(n-1)\ldots(n-k+1)} \sum_{\substack{1 \leq i_1, \ldots, i_k \leq n \\ \text{distinct}}} \phi(X_{i_1}, \ldots, X_{i_k}).$$

To see this, observe that by symmetry (since $\mathcal{S}_n$ does no distinguish between $X_1, \ldots, X_n$), we have $\mathbb{E}[\phi(X_{i_1}, \ldots, X_{i_k}) \mid \mathcal{S}_n]$ is the same for all distinct $i_1, \ldots, i_k \leq n$. When you add all these up, we get

$$\mathbb{E}\left[\sum_{\substack{1 \leq i_1, \ldots, i_k \leq n \\ \text{distinct}}} \phi(X_{i_1}, \ldots, X_{i_k}) \,\middle|\, \mathcal{S}_n \right] = \sum_{\substack{1 \leq i_1, \ldots, i_k \leq n \\ \text{distinct}}} \phi(X_{i_1}, \ldots, X_{i_k})$$

since the latter is clearly $\mathcal{S}_n$-measurable. There are $n(n-1)\ldots(n-k+1)$ terms on the left, each of which is equal to $\mathbb{E}[\phi(X_1, \ldots, X_k) \mid \mathcal{S}_n]$. This proves the claim.

Together with the reverse martingale theorem, we have shown that

$$\frac{1}{n(n-1)\ldots(n-k+1)} \sum_{\substack{1 \leq i_1, \ldots, i_k \leq n \\ \text{distinct}}} \phi(X_{i_1}, \ldots, X_{i_k}) \xrightarrow{\text{a.s.}, L^1} \mathbb{E}[\phi(X_1, \ldots, X_k) \mid \mathcal{S}].$$

The number of summands on the left in which X_1 participates is $k(n-1)(n-2)\dots(n-k+1)$. If $|\phi| \leqslant M_\phi$, then the total contribution of all terms containing X_1 is at most

$$M_\phi \frac{k(n-1)(n-2)\dots(n-k+1)}{n(n-1)(n-2)\dots(n-k+1)} \to 0$$

as $n \to \infty$. Thus, the limit is a function of $X_2, X_3, \dots$. By a similar reasoning, the limit is a tail-random variable for the sequence $X_1, X_2, \dots$. By Kolmogorov's zero-one law it must be a constant (then the constant must be its expectation). Hence,

$$\mathbb{E}[\phi(X_1, \dots, X_k) \,\big|\, \mathcal{S}] = \mathbb{E}[\phi(X_1, \dots, X_k)].$$

As this is true for every bounded measurable ϕ, we see that $\mathcal{S}$ is independent of $\sigma\{X_1, \dots, X_k\}$. As this is true for every k, $\mathcal{S}$ is independent of $\sigma\{X_1, X_2, \dots\}$. But $\mathcal{S} \subseteq \sigma\{X_1, X_2, \dots\}$ and therefore $\mathcal{S}$ is independent of itself. This implies that for any $A \in \mathcal{S}$ we must have $\mathbb{P}(A) = \mathbb{P}(A \cap A) = \mathbb{P}(A)^2$ which implies that $\mathbb{P}(A)$ equals 0 or 1. ■

8. Exchangeable random variables

Let ξ_n, $n \geqslant 1$, be any sequence of random variables. Recall that this means that

$$(\xi_{\pi(1)}, \xi_{\pi(2)}, \dots) \stackrel{d}{=} (\xi_1, \xi_2, \dots)$$

for any bijection (permutation) $\pi : \mathbb{N} \mapsto \mathbb{N}$ that fixes all but finitely many elements. Since distribution of infinitely many random variables is nothing but the collection of all finite dimensional distributions, this is equivalent to saying that

$$(\xi_{i_1}, \dots, \xi_{i_n}) \stackrel{d}{=} (\xi_1, \dots, \xi_n)$$

for any $n \geqslant 1$ and any distinct $i_1, \dots, i_n$.

We have seen an example of an exchangeable sequence in Pólya's urn scheme, namely the successive colours drawn.

Example 22

If ξ_n are i.i.d., then they are exchangeable. More generally, consider finitely many probability measures $\mu_1, \dots, \mu_k$ on some $(\Omega, \mathcal{F})$ and let $p_1, \dots, p_k$ be positive numbers that add up to 1. Pick $L \in \{1, \dots, k\}$ with probabilities $p_1, \dots, p_k$, and conditional on L, pick an i.i.d. sequence $\xi_1, \xi_2, \dots$ from μ_L. Then (unconditionally) ξ_is are exchangeable but not independent.

The above example essentially covers everything, according to a fundamental theorem of de Finetti! Before stating it, let us recall the exchangeable sigma-algebra $\mathcal{S}$ of the collection of all sets in the product sigma-algebra $\mathcal{F}^{\otimes \mathbb{N}}$ that are invariant under finite permutations of co-ordinates. Let us also define $\mathcal{S}_n$ as the collection of all events invariant under the permutations of the first n co-ordinates. The $\mathcal{S} = \cap_n \mathcal{S}_n$.

Theorem 28: de Finetti

Let $\xi_1, \xi_2, \ldots$ be an exchangeable sequence of random variables taking values in $(X, \mathcal{F})$. Then, they are i.i.d. conditional on $\mathcal{S}$. By this we mean that

$$\mathbb{E}\left[\phi_1(\xi_1)\ldots\phi_k(\xi_k) \;\middle|\; \mathcal{S}\right] = \prod_{j=1}^{k} \mathbb{E}[\phi_j(\xi_1) \mid \mathcal{S}]$$

for any $k \geqslant 1$ and any bounded measurable $\phi_j : X \mapsto \mathbb{R}$.

If the situation is nice enough that a regular conditional probability given $\mathcal{S}$ exists, then the statement is equivalent to saying that the conditional distribution is (almost surely) a product of identical probability distributions on $\mathcal{F}$.

Before proving this, let us prove a lemma very similar to the one we used in the proof of the Hewitt-Savage zero one law.

Lemma 7

Let $\xi_1, \xi_2, \ldots$ be an exchangeable sequence taking values in $(X, \mathcal{F})$. Fix $k \geqslant 1$ and any bounded measurable $\psi : X^k \mapsto \mathbb{R}$. Then, as $n \to \infty$

$$\frac{1}{n^k} \sum_{1 \leqslant i_1, \ldots, i_k \leqslant n} \psi(\xi_{i_1}, \ldots, \xi_{i_k}) \overset{\text{a.s.}}{\to} \mathbb{E}\left[\psi(\xi_1, \ldots, \xi_k) \;\middle|\; \mathcal{S}\right].$$

Proof. We claim that

$$\mathbb{E}\left[\psi(\xi_1, \ldots, \xi_k) \;\middle|\; \mathcal{S}_n\right] = \frac{1}{n(n-1)\ldots(n-k+1)} \sideset{}{'}\sum_{i_1, \ldots, i_k \leqslant n} \psi(\xi_{i_1}, \ldots, \xi_{i_k}) \tag{13}$$

where $\sum'_{i_1, \ldots, i_k \leqslant n}$ denotes summation over distinct $i_1, \ldots, i_k \leqslant n$. The reason is that the right hand side is clearly in $\mathcal{S}_n$ (since it is a symmetric function of $\xi_1, \ldots, \xi_n$). Further, if $Z = g(\xi_1, \ldots, \xi_n)$ where g is a symmetric measurable bounded function from X^n to $\mathbb{R}$, then for any permutation π of $[n]$,

$$\begin{aligned}\mathbb{E}[Z\psi(\xi_1, \ldots, \xi_k)] &= \mathbb{E}[g(\xi_{\pi(1)}, \ldots, \xi_{\pi(n)})\psi(\xi_{\pi(1)}, \ldots, \xi_{\pi(k)})] \\ &= \mathbb{E}[g(\xi_1, \ldots, \xi_n)\psi(\xi_{\pi(1)}, \ldots, \xi_{\pi(k)})]\end{aligned}$$

where the first line used the exchangeability of ξ_is and the second used the symmetry of g. By such symmetric functions generate the sigma-algebra $\mathcal{S}_n$, hence this shows that $\mathbb{E}[\psi(\xi_{\pi(1)}, \ldots, \xi_{\pi(k)}) \mid \mathcal{S}_n]$ is the same for all permutations π of $[n]$. Therefore the expectation of the right hand side of (13) is also the same.

Now, by Lévy's backward law (or reverse martingale theorem) we know that $\mathbb{E}[\phi(\xi_1, \ldots, \xi_k) \mid \mathcal{S}_n]$ converges to $\mathbb{E}[\phi(\xi_1, \ldots, \xi_k) \mid \mathcal{S}]$. On the right hand side, we may replace $n(n-1)\ldots(n-k+1)$ by n^k

(the ratio goes to 1 as $n \to \infty$) and extend the sum to all $i_1, \ldots, i_k$ since the number of terms with at least two equal indices is of order n^{k-1} and its contribution is at most $\|\psi\|_{\sup}$ (thus the contribution gets washed away when divided by n^k). ■

Now we prove de Finetti's theorem.

Proof of de Finetti's theorem. By the lemma applied to $\psi(x_1, \ldots, x_k) = \phi_1(x_1) \ldots \phi_k(x_k)$,

$$\frac{1}{n^k} \sum_{i_1,\ldots,i_k \leqslant n} \phi_1(\xi_{i_1}) \ldots \phi_k(\xi_{i_k}) \overset{a.s.}{\to} \mathbb{E}\left[\phi_1(X_1) \ldots \phi_k(X_k) \,\middle|\, \mathcal{S}\right].$$

On the other hand, the left hand side factors into a product of $\frac{1}{n}\sum_{i=1}^{n} \phi_\ell(x_i)$ over $\ell = 1, 2, \ldots, k$, and again by the Lemma the ℓth factor converges almost surely to $\mathbb{E}[\phi_\ell(\xi_1) \mid \mathcal{S}]$. This proves the theorem. ■

There are many alternate ways to state the thorem of de Finetti. One is to say that every exchangeable measure is a convex combination of i.i.d. product measures. Another way is this:

If $(\xi_n)_n$ is an exchangeable sequence of random variables taking values in a Polish space X, then there exists a Borel measurable function $f : [0,1] \times [0,1] \mapsto X$ such that

$$(\xi_1, \xi_2, \xi_3, \ldots) \overset{d}{=} (f(V, V_1), f(V, V_2), f(V, V_3), \ldots)$$

where $V, V_1, V_2, \ldots$ are i.i.d.uniform$[0,1]$ random variables. Here V represents the common information contained in $\mathcal{S}$, and conditional on that, the variables are i.i.d.

8.1. About the exchangeable sigma algebra. Suppose X_i are i.i.d. By the Hewitt-Savage zero-one law, the exchangeable sigma algebra $\mathcal{S}$ is trivial. What is it in the case of a general exchangeable sequence $(X_n)_n$? To get an idea, first consider the case where X_ns take values in a finite set A. Then, by the lemma above, $\frac{1}{n}\sum_{k=1}^{n} \mathbf{1}_{X_k = a}$ converges almost surely to some $\theta(a)$ for each $a \in A$. Then θ is a random probability vector on A. Further, for any fixed n, it is clear that $\mathcal{S}_n$ is precisely the sigma algebra generated by $\frac{1}{n}\sum_{k=1}^{n} \mathbf{1}_{X_k = a}$, $a \in A$. This suggests that the exchangeable sigma-algebra $\mathcal{S}$ must be just the sigma-algebra generated by θ (i.e., by $\theta(a)$, $a \in A$). To fill up with a precise statement

This also gives a way to think of de Finetti's theorem (in fact this was implicit in the proof). Think of an exchangeable sequence of random variables taking values in a finite set A. Then when we condition on $\mathcal{S}_n$, we know the number of times each $a \in A$ appears among $X_1, \ldots, X_n$. In other words, we know the multi-set $\{X_1, \ldots, X_n\}$. By exchangeablity, the conditional distribution of $(X_1, \ldots, X_n)$ is uniform distribution on all sequences in A^n that are consistent with these frequencies. Put another way, from the multi-set $\{X_1, \ldots, X_n\}$, sample n times without replacement, and place the elements in the order that they are sampled. If we fix a k and consider $X_1, \ldots, X_k$, then for large n sampling without replacement and sampling with replacement are essentially the same, which is the statement that $X_1, \ldots, X_k$, given $\mathcal{S}_n$, are approximately i.i.d.

9. Absolute continuity and singularity of product measures

Consider a sequence of independent random variables X_n (they may take values in different spaces). We are told that either (1) $X_n \sim \mu_n$ for each n or (2) $X_n \sim \nu_n$ for each n.Here μ_n and ν_n are given probability distributions. From one realization of the sequence $(X_1, X_2, \ldots)$, can we tell whether the first situation happened or the second?

In measure theory terms, the question may be formulated as follows.

Question: Let μ_n, ν_n be probability measures on $(\Omega_n, \mathcal{G}_n)$. Let $\Omega = \times_n \Omega_n$, $\mathcal{F} = \otimes_n \mathcal{G}_n$ and $\mu = \otimes_n \mu_n$ and $\nu = \otimes \nu_n$. Then, μ, ν are probability measures on $(\Omega, \mathcal{F})$. Assume that $\nu_n \ll \mu_n$ for each n. Can we say whether (1) $\nu \ll \mu$, (2) $\nu \perp \mu$ or (3) neither of the previous two options?

Let us consider a concrete example where direct calculations settle the above question. It also serves to show that both $\nu \perp \mu$ and $\nu \ll \mu$ are possibilities.

Example 23

Let $\mu_n = \text{unif}[0,1]$ and $\nu_n = \text{unif}[0, 1+\delta_n]$. Then, $\nu[0,1]^{\mathbb{N}} = \prod_n \frac{1}{1+\delta_n}$. Thus, if $\prod_n (1+\delta_n) = \infty$, then $\mu[0,1]^{\mathbb{N}} = 1$ while $\nu[0,1]^{\mathbb{N}} = 0$. Hence, $\mu \perp \nu$.

On the other hand, if $\prod_n (1+\delta_n) < \infty$, then we claim that $\nu \ll \mu$. To see this, pick U_n, V_n be i.i.d. unif$[0,1]$. Define $X_n = (1+\delta_n) U_n \sim \nu_n$. Further, set

$$Y_n = \begin{cases} X_n & \text{if } X_n \leqslant 1, \\ V_n & \text{if } X_n > 1. \end{cases}$$

Check that V_n are i.i.d with uniform distribution on $[0,1]$. In short, $(X_1, X_2, \ldots) \sim \nu$ and $(Y_1, Y_2, \ldots) \sim \mu$. Now,

$$\mathbb{P}\{X_n = Y_n \text{ for all } n\} = \mathbb{P}\{X_n \leqslant (1+\delta_n)^{-1} \text{ for all } n\} = \prod_{n=1}^{\infty} \frac{1}{1+\delta_n}$$

which is positive by assumption. Thus, there is a way to construct $X \sim \mu$ and $Y \sim \nu$ such that $X = Y$ with positive probability. Then we cannot possibly have $\mu \perp \nu$ (in itself this is not enough to say that $\nu \ll \mu$).

We used the special properties of uniform distribution to settle the above example. In general it is not that easy, but Kakutani provided a complete answer.

Theorem 29: Kakutani's theorem

Let μ_n, ν_n be probability measures on $(\Omega_n, \mathcal{F}_n)$ and assume that $\mu_n \ll \nu_n$ with Radon-Nikodym theorem f_n Let $\mu = \otimes_n \mu_n$ and $\nu = \otimes_n \nu_n$, probability measures on $\Omega = \times_n \Omega_n$ with the product sigma algebra. Let $a_n = \int_{\Omega_n} \sqrt{f_n} d\nu_n$. Then,$f(x) := \prod_{k=1}^{\infty} f_k(x_k)$ converges ν-almost surely

(1) If $\prod_{k=1}^{\infty} a_k > 0$, then $\mu \ll \nu$ and and $d\mu(x) = f(x)\, d\nu(x)$.

(2) If $\prod_{k=1}^{\infty} a_k = 0$, then $\mu \perp \nu$.

To see the relevance of martingales in this problem, let $\Pi_k : \Omega \to \Omega_k$ denote the projection maps and let $\mathcal{G}_n = \sigma\{\Pi_1, \ldots, \Pi_n\}$. Define $\xi_k = f_k \circ \Pi_k$, i.e., $\xi_k(\omega) = f_k(\omega_k)$. Then ξ_k are independent random variables under ν (as ν is a product measure), with

$$\mathbb{E}_\nu[\xi_n] = \int_{\Omega_n} f_n d\nu_n = \mu_n(\Omega_n) = 1.$$

Therefore, $X_n := \xi_1 \ldots \xi_n$ is a ν-martingale. ince $f_1(\omega_1) \ldots f_n(\omega_n)$ is the Radon-Nikodym derivative of $\mu_1 \otimes \ldots \otimes \mu_n$ w.r.t. $\nu_1 \otimes \ldots \otimes \nu_n$, it follows that

$$\int_A X_n d\nu = \mu(A) \text{ for } A \in \mathcal{G}_n. \tag{14}$$

This shows the relevance of the martingale X_n to the question. Similarly, $Y_n = \frac{\sqrt{\xi_1} \ldots \sqrt{\xi_n}}{a_1 \ldots a_n}$ is a ν-martingale as $\sqrt{\xi_k}/a_k$ are independent random variables with unit mean (observe that $a_k \leqslant 1$ by Cauchy-Schwarz inequality). A general result on *product martingales* (products of independent positive random variables with unit mean) is given in the next section, but here we go directly to the proof of Kakutani's theorem .

Proof of Kakutani's theorem. If $\prod_n a_n > 0$, then $\{Y_n\}$ is L^2 bounded as

$$\mathbb{E}[Y_n^2] = \frac{1}{a_1^2 \ldots a_n^2} \leqslant \frac{1}{\prod_k a_k^2}.$$

Therefore, $Y_n \to Y_\infty$ in $L^2(\nu)$ and hence $Y_n^2 \to Y_\infty^2$ in $L^1(\nu)$. As $X_n = Y_n^2 a_1^2 \ldots a_n^2$, this shows that $X_\infty = Y_\infty \prod_k a_k^2$ a.s.[ν] and that $X_n \to X_\infty$ in $L^1(\nu)$. Therefore, for any $A \in \mathcal{B}_{\mathbb{R}^\mathbb{N}}$, we have $\int X_n \mathbf{1}_A d\nu \to \int X_\infty \mathbf{1}_A d\nu$. But from (14), we see that $\int X_n \mathbf{1}_A d\nu = \mu(A)$ if $A \in \mathcal{G}_m$ and $n \geqslant m$. Thus if $d\theta = X_\infty d\nu$, then $\theta(A) = \mu(A)$ for all $A \in \cup_k \mathcal{G}_k$ (all finite dimensional cylinder sets). This is a π-system that generates $\mathcal{B}_{\mathbb{R}^\mathbb{N}}$, hence $\theta = \mu$. In other words, $\mu \ll \nu$ and X_∞ is the Radon-Nikodym derivative of μ w.r.t. ν. This proves the first part.

Suppose $\prod_n a_n = 0$. Let $A_n = \{\omega \in \Omega_1 \times \ldots \times \Omega_n : f_1(\omega_1)\ldots f_n(\omega_n) > 1\}$ and let $B_n = A_n \times \Omega_{n+1} \times \Omega_{n+2} \times \ldots$. Then

$$\begin{aligned}\nu(B_n) &= \int_{A_n} d\nu_1(\omega_1)\ldots d\nu_n(\omega_n) \\ &\leqslant \int_{A_n} \sqrt{f_1(\omega_1)}\ldots\sqrt{f_n(\omega_n)}d\nu_1(\omega_1)\ldots d\nu_n(\omega_n) \text{ (the integrand is } \geqslant 1 \text{ on } A_n) \\ &\leqslant a_1 \ldots a_n\end{aligned}$$

by expanding the integral to $\Omega_1 \times \ldots \times \Omega_n$. On the other hand, $B_n^c = A_n^c \times \Omega_{n+1} \times \ldots$ and hence

$$\begin{aligned}\mu(B_n^c) &= \int_{A_n^c} f_1(\omega_1)\ldots f_n(\omega_n)d\nu_1(\omega_1)\ldots d\nu_n(\omega_n) \\ &\leqslant \int_{A_n^c} \sqrt{f_1(\omega_1)}\ldots\sqrt{f_n(\omega_n)}d\nu_1(\omega_1)\ldots d\nu_n(\omega_n) \text{ (as } x \leqslant \sqrt{x} \text{ for } 0 \leqslant x \leqslant 1) \\ &\leqslant a_1 \ldots a_n\end{aligned}$$

again by expanding the integral to $\Omega_1 \times \ldots \times \Omega_n$.

If $\prod_k a_k = 0$, we have found sets B_n such that $\mu(B_n^c) \to 0$ and $\nu(B_n) \to 0$. Choose a subsequence $\{n_k\}$ such that $\mu(B_{n_k}^c)$ and $\nu(B_{n_k})$ are both summable. By Borel-Cantelli lemma, (here f.o. stands for "finitely often")

$$\mu(B_{n_k}^c \text{ f.o.}) = 1 \text{ and } \nu(B_{n_k} \text{ f.o.}) = 1.$$

The two sets are clearly disjoint, hence $\mu \perp \nu$. ■

Now we come to a more general question, going beyond the product structure. The answer is less explicit in general, but very useful.

9.1. Question. Let μ, ν be probability measures on $(\Omega, \mathcal{F})$. Suppose $\mathcal{F}_1 \subseteq \mathcal{F}_2 \subseteq \ldots$ are sub sigma-algebras of $\mathcal{F}$ such that $\sigma\{\cup_n \mathcal{F}_n\} = \mathcal{F}$. Let $\mu_n = \mu\big|_{\mathcal{F}_n}$ and $\nu_n = \nu\big|_{\mathcal{F}_n}$ be the restrictions of μ and ν to $\mathcal{F}_n$. Assume that $\nu_n \ll \mu_n$ for each n. Is $\nu \ll \mu$. If not are there conditions?

This subsumes the question of product measures by taking $\Omega = \times_n \Omega_n$ and $\mathcal{F}_n = \sigma\{\Pi_1, \ldots, \Pi_n\}$, the sigma algebra generated by the first n projections. The answer to this question is as follows.

Theorem 30

Let X_n be the Radon-Nikodym derivative of μ_n w.r.t. ν_n. Let $X = \limsup X_n$. Then the Lebesgue decomposition of μ w.r.t. ν is given by $\mu = X d\nu + \mathbf{1}_{X=\infty}\mu$.

Proof. Let $\theta = \frac{1}{2}\mu + \frac{1}{2}\nu$. Then μ, ν are absolutely continuous to θ, so we denote their Radon-Nikodym derivatives w.r.t. θ as V, W respectively. Then for any set $A \in \mathcal{F}$, we have

$$\mu(A) = \mu(A \cap \{W > 0\}) + \mu(A \cap \{W = 0\})$$
$$= \int_A V\mathbf{1}_{W>0} d\theta + \mu(A \cap \{W = 0\}) = \int_A \frac{V}{W}\mathbf{1}_{W>0} d\nu + \mu(A \cap \{W = 0\})$$

because $d\theta = \frac{1}{W}d\nu$ on the set $W > 0$. Thus, the Lebesgue decomposition of μ w.r.t. ν is given by

$$d\mu = h d\nu + d\gamma \text{ where } h = \frac{V}{W}\mathbf{1}_{W>0} \text{ and } \gamma(\cdot) = \mu(\cdot \cap \{W = 0\}). \tag{15}$$

Observe that $\gamma \perp \nu$ because $\nu\{W = 0\} = 0$.

Further, $d\mu_n = V_n d\theta_n$ where $V_n = \mathbb{E}_\theta[V \mid \mathcal{F}_n]$. As a Doob martingale[1], $V_n \to V$ a.s.$[\theta]$ and in $L^1(\theta)$. Similarly, $W_n \to W$ a.s.$[\theta]$ and in $L^1(\theta)$.

As $d\mu_n = X_n d\nu_n$, we see that $X_n W_n = V_n$ a.s.$[\theta]$. Therefore, $X = \frac{V}{W}$ a.s.$[\theta]$ on the event $\{W > 0\}$, and $X = \infty$ a.s.$[\theta]$ on the event $\{W = 0\}$. Comparing to (15) completes the proof. ■

Does the theorem on product measures follow from here? Under the assumption that $\prod_n a_n > 0$ (in the notation of Theorem 29, the proof we gave precisely showed that $X < \infty$ a.s. w.r.t. ν and μ, and hence $d\mu = X d\nu$ (here $X = f_1 \otimes f_2 \otimes \ldots$). Under the assumption $\prod_n a_n = 0$, our earlier proof did not show that $X = \infty$ a.s.$[\mu]$. That follows from the result on product martingales stated and proved in the next section.

10. Product martingales

If ξ_n are *independent* positive random variables with unit mean, then $X_n := \xi_1 \ldots \xi_n$ is a martingale. Such martingales are called *product martingales* and occur in many situations. As non-negative martingales, $X_n \overset{a.s.}{\to} X_\infty$ for some finite random variable X_∞ (in fact $\mathbb{E}[X_\infty] \leqslant 1$ by Fatou's lemma. The question is of uniform integrability of $\{X_n\}$ which is equivalent to "$\mathbb{E}[X_\infty] = 1$" or that X is a Doob martingale.

Lemma 8

In the above setting, let $a_n = \mathbb{E}[\sqrt{\xi_n}]$. Then there are two possibilities.

(1) $\prod_n a_n > 0$. In this case, $\{X_n\}$ is uniformly integrable, $\mathbb{E}[X_\infty] = 1$. If $\xi_n > 0$ a.s. for all n, then $X_\infty > 0$ a.s.

(2) $\prod_n a_n = 0$. In this case, $\{X_n\}$ is not uniformly integrable and $X_\infty = 0$ a.s.

---

[1]Reason: Let V_∞ be the limit guaranteed by the martingale convergence theorem (in $L^1(\theta)$ and a.s.$[\theta]$). Then $V_n = \mathbb{E}[V_\infty \mid \mathcal{F}_n]$. But then, $\mathbb{E}_\theta[V_\infty \mid \mathcal{F}_n] = \mathbb{E}_\theta[V \mid \mathcal{F}_n]$ for all n, which implies that $\int_A V_\infty d\theta = \int_A V d\theta$ for all $A \in \bigcup_n \mathcal{F}_n$. By the $\pi - \lambda$ theorem, it follows that $V_\infty = V$ a.s.$[\theta]$.

Observe that $a_k \leqslant \sqrt{\mathbb{E}[\xi_k]} = 1$ for all k. Hence the partial products $\prod_{j=1}^n a_j$ are decreasing in n and have a limit in $[0,1]$, which is what we mean by $\prod_n a_n$.

PROOF OF LEMMA 8. Let $Y_n = \prod_{j=1}^n \frac{\xi_j}{\sqrt{a_j}}$. Then X_n and Y_n are both martingales (w.r.t. $\mathcal{F}_n = \sigma\{\xi_1, \ldots, \xi_n\}$) and are related as $X_n = Y_n^2 a_1^2 \ldots a_n^2$. As they are non-negative and have mean 1, we also know that $X_n \overset{a.s.}{\to} X_\infty$ and $Y_n \overset{a.s.}{\to} Y_\infty$ where X_∞ and Y_∞ are integrable (hence finite almost surely).

(1) Suppose $\prod_n a_n > 0$. Then $\mathbb{E}[Y_n^2] = \frac{1}{a_1^2 \ldots a_n^2}$ is uniformly bounded. As an L^2-bounded martingale, $Y_n \to Y_\infty$ in L^2. In particular, Y_n^2 converges to Y_∞^2 almost surely and in L^1, which implies that $\{Y_n^2\}$ must be uniformly integrable. But $X_n \leqslant Y_n^2$ (as $a_j \leqslant 1$ for all j), which means that $\{X_n\}$ is uniformly integrable. In particular, we also have $\mathbb{E}[X_\infty] = \lim \mathbb{E}[X_n] = 1$. In particular, $\mathbb{P}\{X_\infty > 0\} > 0$. But if ξ_ns are strictly positive, then the event $\{X_\infty > 0\}$ is a tail event of $(\xi_n)_n$, hence by Kolmogorov's zero one law it must have probability 1.

(2) Suppose $\prod_n a_n = 0$. Observe that $X_\infty = Y_\infty^2 \prod_j a_j^2$ and Y_∞ is a finite random variable. Hence $X_\infty = 0$ a.s.. ■

11. The Haar basis and almost sure convergence

Consider $L^2[0,1]$ with respect to the Lebesgue measure. Abstract Hilbert space theory says that L^2 is a Hilbert space, it has an orthonormal basis, and that for any orthonormal basis $\{\phi_n\}$ and any $f \in L^2$, we have

$$f \overset{L^2}{=} \sum_n \langle f, \phi_n \rangle \phi_n$$

which means that the L^2-norm of the difference between the left side and the nth partial sum on the right side converges to zero as $n \to \infty$.

But since L^2 consists of functions, it is possible to ask for convergence in other senses. In general, there is no almost-sure convergence in the above series.

Theorem 31

Let $H_{n,k}$, $n \geqslant 1, 0 \leqslant k \leqslant 2^n - 1$ be the Haar basis for L^2. Then, for any $f \in L^2$, the convergence holds almost surely.

PROOF. On the probability space $([0,1], \mathcal{B}, \lambda)$, define the random variables

$$X_n(t) = \sum_{m \leqslant n} \sum_{k \leqslant 2^m - 1} \langle f, H_{m,k} \rangle H_{m,k}(t).$$

We claim that X_n is a martingale. Indeed, it is easy to see that if $\mathcal{F}_n := \sigma\{H_{n,0}, \ldots, H_{n,2^n-1}\}$ (which is the same as the sigma algebra generated by the intervals $[k/2^n, (k+1)/2^n)$, $0 \leqslant k \leqslant 2^n - 1$), then $X_n = \mathbb{E}[f \mid \mathcal{F}_n]$. Thus, $\{X_n\}$ is the Doob-martingale of f with respect to the filtration $\mathcal{F}_\cdot$.

Further, $\mathbb{E}[X_n^2] = \sum_{m \leqslant n} \sum_{k \leqslant 2^m - 1} |\langle f, H_{m,k} \rangle|^2 \leqslant \|f\|_2^2$. Hence $\{X_n\}$ is an L^2-bounded martingale. It converges almost surely and in L^2. But in L^2 it converges to f. Hence $X_n \stackrel{a.s.}{\to} f$. ■

12. Karlin-McGregor formula

Consider n independent simple random walks on $\mathbb{Z}$, each going up with probability p and down with probability $q = 1 - p$ at each step. If they start at locations $a_1, \ldots, a_n$ and time 0, the probability that they are at locations $b_1, \ldots, b_n$ at time t (in some order) is

$$\sum_{\sigma \in \mathcal{S}_n} \prod_{i=1}^{n} P_{a_i, b_{\sigma(i)}}(t) \tag{16}$$

where $P_{a,b}(t)$ is the probability that a simple random walk started at location a at time 0 is at location b at time t. Explicitly,

$$P_{a,b}(t) = \binom{t}{\frac{t+b-a}{2}} p^{\frac{t+b-a}{2}} q^{\frac{t-b+a}{2}} \tag{17}$$

where, for x a positive integer, $\binom{x}{y}$ is interpreted as zero unless y is an integer and $0 \leqslant y \leqslant x$. If instead we specify which random walk should end where, then we get only one term corresponding to the specified permutation.

Here is a different question. In the same setting, what is the probability that none of the random walks hit each other in the meantime? This is a much harder question, but there is an amazing explicit answer that turns out to be deep and important (we cannot explain the latter here). It holds in somewhat greater generality.

Definition 4

Let $\mu_0, \mu_1, \ldots$ be probability distributions on $\mathbb{Z}$. A random walk on $\mathbb{Z}$ with step distributions $(\mu_k)_{k \geqslant 0}$ is a sequence of random variables $\{S_0, S_1, \ldots\}$ such that $S_{k+1} - S_k \sim \mu_k$ for $k \geqslant 0$. If $S_0 = a$ w.p.1., we say that the random walk starts at a.

We say that the *skip-free* condition is satisfied if for independent random walks $S^{(1)}, \ldots, S^{(n)}$ started at $a_1 < \ldots < a_n$, if whenever $S^{(i)}(t) > S^{(j)}$ for some t, then there must be an $s \leqslant t$ such that $S^{(i)}(s) = S^{(j)}(s)$. Note that the skip-free condition depends only on the step-distributions and the initial states. Two cases where it is satisfied are:

(1) $\mu_k(-1) = 0$ for all k.

(2) $\mu_k(0) = 0$ for all k and $a_1, \ldots, a_n$ are all even or all odd.

We introduce the notation for the transition probabilities

$$P_{a,b}(s,t) = \mathbb{P}\{S_t = b \,\big|\, S_s = a\}$$

which can be written in terms of the μ_ks. Note that $P_{a,b}(s,t)$ depends on a, b only through $b-a$, and in the special case when μ_k does not depend on k (time-homogeneity), it depends on s, t only through $t-s$. When $\mu_k(+1) = p$ and $\mu_k(-1) = q$, this reduces to (17).

Theorem 32: Karlin–McGregor formula

Consider independent random walks $S^{(1)}, \ldots, S^{(n)}$ with common step distributions $(\mu_k)_{k\geqslant 0}$ and started at $a_1 < \ldots < a_n$. Assume that the skip-free condition is satisfied. Then, the probability that they end up at locations $b_1 < \ldots < b_n$ at time t without any two of them intersecting up to that time, is equal to

$$\sum_{\sigma\in\mathcal{S}_n} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} P_{a_i,b_{\sigma(i)}}(0,t) = \det\left[P_{a_i,b_j}(0,t)\right]_{1\leqslant i,j\leqslant n}.$$

Proof. Denote that random walks as $S^{(1)}, \ldots, S^{(n)}$, where $S_0^{(j)} = j$. For any $\sigma \in \mathcal{S}_n$, consider

$$M^{\sigma}(s) = \prod_{i=1}^{n} \mathbb{P}\left\{S^{(j)}(t) = b_{\sigma(j)} \text{ for each } j \,\big|\, \mathcal{F}_s\right\} \;=\; \mathbb{E}\left[\prod_{j=1}^{n} \mathbf{1}_{S^{(j)}} = b_{\sigma(j)} \,\big|\, \mathcal{F}_s\right]$$

where $\mathcal{F}_s = \sigma\{S^{(j)}(r) : 0 \leqslant r \leqslant s,\ 1 \leqslant j \leqslant n\}$ is the natural filtration. Clearly M^{σ} is an $\mathcal{F}_{\bullet}$ martingale, and hence so is

$$M(s) := \sum_{\sigma\in\mathcal{S}_s} \operatorname{sgn}(\sigma) M^{\sigma}(s) \;=\; \mathbb{E}\left[\det\left(P_{S^{(j)}(s),b_k}(s,t)\right)_{1\leqslant j,k\leqslant n}\right].$$

Let $\tau = \inf\{s : S^{(j)}(s) = S^{(k)}(s) \text{ for some } j \neq k\}$ be the first time two of the random walks meet. The optional stopping theorem gives $\mathbb{E}[M(\tau \wedge t)] = \mathbb{E}[M(0)]$.

(1) $M(0)$ is precisely the quantity on the right side of the statement of the theorem.

(2) $M(\tau\wedge t) = 0$ if $\tau \leqslant t$ (if $S^{(j)}(s) = S^{(i)}(s)$, then the jth and ith rows of $\left(P_{S^{(j)}(s),b_k}(s,t)\right)_{1\leqslant j,k\leqslant n}$ are equal). But if $\tau > t$, then $M(\tau \wedge t) = M(t) = \sum_{\sigma} \operatorname{sgn}(\sigma) \prod_{i=1}^{n} \mathbf{1}_{S^{(j)}(t)=b_{\sigma(j)}}$, since $P_{a,b}(t,t) = \delta_{a,b}$. Thus, $\mathbb{E}[M(\tau \wedge t)] = \mathbb{P}\{\tau > t,\ \{S^{(1)}(t), \ldots, S^{(p)}(t)\} = \{b_1, \ldots, b_m\}\}$.

Thus we arrive at the identity

$$\sum_{\sigma} \operatorname{sgn}(\sigma)\mathbb{P}\{\{S^{(j)}(t) = b_{\sigma(j)} \text{ for each } j,\ \text{ no intersection up to time } t\} = \det\left[P_{a_i,b_j}(0,t)\right]_{1\leqslant i,j\leqslant n}.$$

If the transitions are such that two paths cannot cross each other without touching, then only the term when σ is the identity permutation survives on the left and we arrive at

$$\mathbb{P}\{\{S^{(j)}(t) = b_j \text{ for each } j,\ \text{ no intersection up to time } t\} = \det\left[P_{a_i,b_j}(0,t)\right]_{1\leqslant i,j\leqslant n}.$$

That was the claim. ■

The Karlin-McGregor formula can be used to solve many counting problems such as the following.

A problem of counting lattice paths: On $\mathbb{Z}^2$, an oriented lattice path is one of the form $\ldots, \mathbf{u}_k, \mathbf{u}_{k+1} \ldots$ such that $\mathbf{u}_{k+1} - \mathbf{u}_k$ is $(1,0)$ or $(0,1)$.

Proposition 1

Let $(a_i, b_i), (c_i, d_i) \in \mathbb{Z}^2$ for $\leqslant i \leqslant n$, and assume that $a_i + b_i = 0$ and $c_i + d_i = L$ for all i and that $a_1 < \ldots < a_n$ and $c_1 < \ldots < c_n$. The number of packets of *non-intersecting* oriented lattice paths that lead from (a_i, b_i) to (c_i, d_i), $i \leqslant n$, is

$$\det\left[\binom{L}{j + d_j - i - b_i}\right]_{i,j \leqslant n}.$$

This follows directly from Karlin-McGregor, just rotate the lattice by $45°$ so that the starting points are on one vertical line and the ending points are on a parallel vertical line. After this rotation, lattice paths become simple random walk paths.

A generalization of the ballot problem: Suppose there are p candidates $C_1, \ldots, C_p$ in an election who get $N_1, \ldots, N_p$ votes respectively. As the votes are counted one by one, what is the chance that throughout the counting process, the true order of the candidates is maintained?

We may take $N_1 \geqslant N_2 \geqslant \ldots \geqslant N_p$ without loss of generality. The question is to find the chance that throughout the counting, C_1 leads, then C_2, then C_3 etc. When $p = 2$, this is the famous ballot problem, for which the answer (usually got by the reflection principle) is $\frac{N_1 - N_2 + 1}{N_1 + 1}$. The answer to the general question is

$$\det\left[\frac{N_j!}{(N_j + i - j)!}\right]_{1 \leqslant i,j \leqslant p} \tag{18}$$

and can be derived from the Karlin-McGregor formula, though the derivation is a little less obvious (try first!).

Proof. Consider random walk with $\text{Geo}(r)$ steps, where $0 < r < 1$. Here $\text{Geo}(r)$ takes the value k with probability $(1-r)^k r$, for $k \geqslant 0$. Consider p independent random walks $X^{(k)}$ started at $-k$, for $1 \leqslant k \leqslant p$. We claim that

$$\mathbb{P}\left\{X^{(k)}(T) = N_k - k \text{ for each } k \text{ and they do not intersect ever} \,\middle|\, X^{(k)}(T) = N_k - k \text{ for each } k\right\}$$

$$= \lim_{Tr=1, r\downarrow 0} \frac{\det\left[\left(P_{-j, N_k - k}(T)\right)_{j,k \leqslant p}\right]}{\prod_{k=1}^{p} P_{-k, N_k - k}(T)}.$$

This is because, as $r \to 0$ and $rT \to 1$, with probability converging to 1, all the jumps are of size at most 1 and no two random walks jump at the same time. That alloww us to apply the Karlin-McGregor formula (for fixed $r > 0$, skip-free condition is violated). By writing out the negative binomial coefficient or general Poisson convergence of rare events, it follows that

$$\lim_{Tr=1, r\downarrow 0} P_{a,b}(T) = \mathbb{P}\{\text{Pois}(1) = b - a\} = \frac{e^{-1}}{(b-a)!}.$$

This leads to the simplification

$$\lim_{Tr=1,r\downarrow 0} \frac{\det\left[\left(P_{-j,N_k-k}(T)\right)_{j,k\leqslant p}\right]}{\prod_{k=1}^{p} P_{-k,N_k-k}(T)} = \frac{\det\left[\left(\frac{e^{-1}}{(N_k+j-k)!}\right)_{j,k\leqslant p}\right]}{\prod_{k=1}^{p} \frac{e^{-1}}{N_k!}} = \det\left[\frac{N_j!}{(N_j+i-j)!}\right]_{1\leqslant i,j\leqslant p},$$

the expression that appears in (18).

But what does this have to do with the ballot problem? When all jumps happen at different times, conditional on then event $X^{(k)}(T) = N_k - k$, the walk $X^{(k)}$ jumps a total of N_k times. Further, the the $N_1 + \ldots + N_p$ jumps occur in a uniform random order, just as they should to correspond to counting votes at random. Lastly, the ballot problem asked for weak inequalities, hence we converted to strict inequality by starting $X^{(k)}$ at $-k$. ■

Remark 6

Karlin-McGregor formula can be stated for skip-free random walks in continuous time (and even continuous space, if sample paths are continuous, for example for Brownian motions), with essentially the same proof, except that we need optional stopping theorem for continuous time martingales. If we did that, the proof above can be stated more naturally by considering independent Poisson processes (which is essentially the scaling limit of the random walks with Geometric steps, as $r \to 0$ and $rT \to 1$).

13. Kahane's multiplicative cascade

A dyadic interval in $[0, 1]$ is one of the form $I_{n,k} = [k2^{-n}, (k+1)2^{-n}]$ for some $n \geqslant 0$, $0 \leqslant k \leqslant 2^n - 1$. The interesting property of these intervals is that for any two of them, either one contains the other or the two have disjoint interiors. It is best to view this via the regular binary tree $\mathcal{T}$ that has root $\emptyset$ and where every vertex has two children.

We may index the vertices with the dyadic intervals $I_{n,k}$. The root vertex is indexed by $I_{0,0} = [0, 1]$ and the vertex indexed by vertex $I_{n,k}$ has two children, namely $I_{n+1,2k}$ (left-child) and $I_{n+1,2k+1}$ (right child), the two dyadic intervals of the next generation that are contained in it.

Now let $W, W_{n,k}$, $n \geqslant 0$, $0 \leqslant k \leqslant 2^n - 1$, be i.i.d. strictly positive random variables with a distribution μ. We construct a sequence of random measures as follows: $dM_0(x) = dx$ and for

$n \geqslant 0$ we set $dM_{n+1}(x) = g_{n+1}(x)dM_n(x)$, where $g_{n+1}(x) = g_n(x)W_{n+1,k}$ if $x \in I_{n+1,k}$. In other words, $dM_n(x) = f_n(x)dx$ where ($g_0(x) = 1$ and $k_0 = 0$ by definition)

$$f_n(x) = \prod_{j=0}^{n} g_j(x) = \prod_{m=0}^{n} W_{m,k_m} \quad \text{if } x \in I_{n,k_n} \subseteq I_{n-1,k_{n-1}} \subseteq \ldots \subseteq I_{1,k_0} \subseteq I_{0,0}.$$

Theorem 33: Kahane

The sequence of random measures M_n converge almost surely to a random measure M on $[0,1]$. Further, if $\mathbb{E}[W \log_2 W] < 1$, then $M \neq 0$ a.s. and M is almost surely a singular measure with no atoms.

Of course, the convergence here is in the Lévy metric on the space of probability measures on the line.

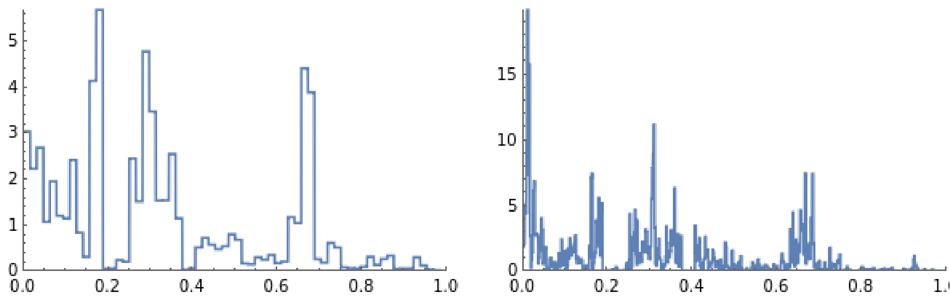

FIGURE 1. Figures of the densities f_7 and f_{10} when $W = e^{tZ-\frac{1}{2}t^2}$ where $Z \sim N(0,1)$ and $t = \frac{1}{2}$. It can be believed that the limit measure M is singular (note the markings on the Y-axis)

PROOF OF CONVERGENCE OF MEASURES. Fix $x \in [0,1]$. As it is a product of positive unit mean independent random variables, $f_n(x)$ is a positive martingale and hence converges a.s. Taking intersection over $x \in \mathbb{Q} \cap [0,1]$, we see that $f_n(x) \overset{a.s.}{\to} f(x)$ for all $x \in \mathbb{Q} \cap [0,1]$, almost surely, for some $f : \mathbb{Q} \cap [0,1] \to \mathbb{R}_+$.

For the same reason, $M_n[0,x]$ is a martingale, as

$$\begin{aligned}
\mathbb{E}[M_{n+1}[0,x] \mid \mathcal{F}_n] &= \mathbb{E}\left[\int_{[0,x]} f_{n+1}(t)dt \,\Big|\, \mathcal{F}_n\right] \\
&= \int_{[0,x]} \mathbb{E}[f_{n+1}(t) \mid \mathcal{F}_n]dt \\
&= \int_{[0,x]} f_n(t)dt \\
&= M_n[0,x].
\end{aligned}$$

Therefore, $M_n[0,x]$is also a martingale for each x. Again taking intersection over $x \in \mathbb{Q} \cap [0,1]$, we see that $M_n[0,x] \to G(x)$ for all $x \in Q \cap [0,1]$ a.s., for some function $G : \mathbb{Q} \cap [0,1] \to \mathbb{R}_+$. Then $\bar{G}(x) = \inf\{G(u) : u > x\}$ defines a CDF of a measure on $[0,1]$, which we call M. It is now clear that $M_n \to M$, a.s. ■

For the remaining properties, to simplify the proofs we shall make the assumption that $2\mathbb{E}[W^2] < 1$. This is a stronger assumption, as $x \log_2 x \leqslant 2x^2$ for all $x > 0$. The proof under the assumption that $\mathbb{E}[W \log_2 W] < 1$ is along similar lines, but more involved.

STRICT POSITIVITY OF THE LIMITING MEASURE UNDER THE STRONGER ASSUMPTION. By conditioning on $W_{0,0}$, we see that

$$M_n[0,1] \stackrel{d}{=} W(M'_n[0,1] + M''_n[0,1])$$

where W, M'_n, M''_n are independent. Hence, $b_n = \mathbb{E}[M_n[0,1]^2]$ satisfies the recursion

$$\begin{aligned} b_{n+1} &= \mathbb{E}[W^2]\mathbb{E}[M'_n[0,1]^2 + M''_n[0,1]^2 + 2M'_n[0,1]M''_n[0,1]] \\ &= \mathbb{E}[W^2](2\mathbb{E}[M_n[0,1]^2 + 2\mathbb{E}[M'_n[0,1]]^2) \\ &= 2\mathbb{E}[W^2](1 + b_n) \end{aligned}$$

because $\mathbb{E}[M_n[0,1]] = 1$ by the martingale property. Writing $\beta = 2\mathbb{E}[W^2]$ and repeating, we see that

$$b_n = \beta + \beta^2 + \ldots + \beta^{n-1} + \beta^n b_0$$

which converges to $1/(1-\beta)$ if $\beta < 1$. Thus, the martingale $M_n[0,1]$ is L^2-bounded and hence converges in L^1. Thus $\mathbb{E}[M[0,1]] = 1$, showing that $\mathbb{P}\{M[0,1] > 0\} > 0$. As the event $M[0,1] > 0$ is clearly a tail event of the $W_{n,k}$s, it follows that $M[0,1] > 0$ a.s. ■

CONTINUITY PROPERTIES OF THE LIMITING MEASURE. ■

14. Waiting for patterns

We have seen how to use optional stopping theorem to compute expected waiting times for patterns in a fair coin toss sequence. Here we deal with a more general situation[2].

Let $X_1, X_2, \ldots$ be an i.i.d. sequence of random variables having a distribution μ on a finite alphabet $\mathcal{A}$. We fix sequences $A = (a_1, \ldots, a_m) \in \mathcal{A}^m$ and $B = (b_1, \ldots, b_n) \in \mathcal{A}^n$. The following notation (just for this section) is useful:

$$A \odot B := \sum_{k=1}^{m \wedge n} \frac{\mathbf{1}_{(a_{m-k+1},\ldots,a_m)=(b_1,\ldots,b_k)}}{\mu(b_1)\ldots\mu(b_k)}.$$

[2]This is taken from the paper of Shuo-Yen Robert Li from (Annals of Probability, Vol 8, No. 6, 1171–1176) in which he discovered the martingale approach to computing such expectations and probabilities.

Theorem 34

Consider the sequence $(a_1, \dots, a_m, X_{m+1}, X_{m+2}, X_{m+3}, \dots)$ and let τ be the smallest t such that the segment of length n ending at X_{m+t} is equal to B. Then

$$\mathbb{E}[\tau] = B \odot B - A \odot B.$$

Proof. Consider the sequence $(X_1, X_2, \dots)$ and let $\hat{\tau} = \min\{t > m : (X_{t-n+1}, \dots, X_t) = B\}$. The theorem is equivalent to the claim that

$$\mathbb{E}[\hat{\tau} - m \mid (X_1, \dots, X_m) = A] = B \odot B - A \odot B.$$

To prove this, set up a gambling game, where, for $t \geqslant m+1$, the gambler G_t deposits 1 rupee just before X_t is revealed. The betting proceeds as follows:

G_t bets on $X_t = b_1$. If she loses she leaves. If she wins he gets $1/\mu(b_1)$ and bets it on $X_{t+1} = b_2$. If she loses she leaves, else she gets $1/\mu(b_1)\mu(b_2)$ and so on. The game is stopped at $\hat{\tau}$.

Let M_t denote the total profit of the team of gamblers after time $t \geqslant m$. The gamblers spend a total of $t - m$ rupees. The total earning is

$$\sum_{k=1}^{n} \frac{\mathbf{1}_{(X_{t-k+1},\dots,X_t)=(b_{n-k+1},\dots,b_n)}}{\mu(b_1)\dots\mu(b_k)} = (X_{t-m+1}, \dots, X_t) \odot B.$$

At time $t = \hat{\tau}$, then the right side is $B \odot B$ if $\tau \geqslant m+n$ and equal to $B \odot B - A \odot B$ if $\tau < m+n$.

By the optional sampling theorem (why does it apply?)

$$\begin{aligned}\mathbb{E}[\hat{\tau} - m \mid (X_1, \dots, X_m) = A] &= \mathbb{E}\left[\sum_{k=1}^{\hat{\tau}} \frac{\mathbf{1}_{(X_{t-k+1},\dots,X_t)=(b_{n-k+1},\dots,b_n)}}{\mu(b_1)\dots\mu(b_k)} \,\Big|\, (X_1, \dots, X_m) = A\right] \\ &= B \odot B - A \odot A.\end{aligned}$$

■

Theorem 35

Let $A = (a_1, \dots, a_m) \in \mathcal{A}^m$ and $B = (b_1, \dots, b_n) \in \mathcal{A}^n$. Assume that neither sequence is an extension of the other. Let $X_1, X_2, \dots$ be i.i.d. with distribution μ on Σ. Then

$$\frac{\mathbb{P}\{\tau_A < \tau_B\}}{\mathbb{P}\{\tau_B < \tau_A\}} = \frac{A \odot A - A \odot B}{B \odot B - B \odot A}$$

CHAPTER 4

Continuous time martingales

Let $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$ be a filtered probability space. Here the filtration $\mathcal{F}_\bullet = (\mathcal{F}_t)_{t\geqslant 0}$ is a collection of sigma-algebra indexed by "continuous time" $t \geqslant 0$. All this means is that $\mathcal{F}_t$ are sub sigma-algebras of $\mathcal{F}$ and that $\mathcal{F}_s \subseteq \mathcal{F}_t$ if $s < t$. Often we write $\mathcal{F}_\infty := \sigma(\cup_t \mathcal{F}_t)$.

Let $X = (X_t)_{t\geqslant 0}$ be a stochastic process on this probability space that is adapted to $\mathcal{F}_\bullet$ (i.e., X_t is $\mathcal{F}_t$-measurable for each $t \geqslant 0$) and such that each $X_t \in L^1(\mathbb{P})$. We say that X is a $\mathcal{F}_\bullet$-martingale if $\mathbb{E}[X_t \mid \mathcal{F}_s] = X_s$ for any $s < t$. Submartingales and supermartingales are similarly defined.

For much of the theory, we shall have to assume more about the filtration and/or the martingale.

(1) *Standard/Usual conditions* on the filtration: The filtration must be right-continuous ($\mathcal{F}_t^+ := \cap_{s>t}\mathcal{F}_s$ is equal to $\mathcal{F}_t$) and complete (all null sets for $\mathcal{F}_\infty := \sigma(\cup_t \mathcal{F}_t)$ are contained in $\mathcal{F}_0$).

(2) *Regularity of sample paths*: The sample paths $t \mapsto X_t(\omega)$ of the process must be continuous or at least RCLL (right continuous at all $t \geqslant 0$ and left limits exist at all $t > 0$) for a.e. ω.

It is safe to make these blanket assumptions and proceed. However, we shall indicate in the next section that if one makes either of these assumptions, the other can be had for free! Before that, here are some basic examples of martingales.

Example 24: Brownian motion

Let $W = (W_t)_{t\geqslant 0}$ be a $\mathcal{F}_\bullet$-Brownian motion. This means that $W_t - W_s \sim N(0, t-s)$ is independent of $\mathcal{F}_s$ for any $s < t$. Therefore,

$$\mathbb{E}[W_t \mid \mathcal{F}_s] = W_s + \mathbb{E}[W_t - W_s \mid \mathcal{F}_s] \;=\; W_s.$$

Thus, W is a martingale. On the other hand, W_t^2 is a sub-martingale, as

$$\begin{aligned}\mathbb{E}[W_t^2 \mid \mathcal{F}_s] &= \mathbb{E}[W_s^2 + (W_t - W_s)^2 + 2W_s(W_t - W_s) \mid \mathcal{F}_s] \\ &= W_s^2 + (t-s) \;\geqslant\; W_s^2.\end{aligned}$$

From this we also see that $W_t^2 - t$ is a martingale.

Exercise 20

Show that for any $\theta \in \mathbb{R}$, the process $M_\theta(t) = e^{\theta W(t) - \frac{1}{2}\theta^2 t}$ is a martingale. Argue that the θ-derivatives $M_\theta^{(k)}|_{\theta=0}$ are martingales and find them explicitly for $k \leqslant 4$.

Example 25: Poisson process

Let $N(\cdot)$ be a Poisson process on with intensity $\lambda > 0$ on the positive real line[a].

Then with $\mathcal{F}_\bullet$ denoting the natural filtration of $N(\cdot)$,

$$\mathbb{E}[N(t) \mid \mathcal{F}_s] = N(s) + \mathbb{E}[N(t) - N(s) \mid \mathcal{F}_s] = N(s) + \lambda(t-s)$$

from which we see that $N(t) - \lambda t$ is a martingale. Similarly, $\mathbb{E}[(N(t) - \lambda t)^2 \mid \mathcal{F}_s]$ is equal to

$$\begin{aligned}
&(N(s)-\lambda s)^2 + \mathbb{E}[(N(t) - N(s) - \lambda(t-s))^2 \mid \mathcal{F}_s] + 2(N(s)-\lambda s)\mathbb{E}[N(t) - N(s) - \lambda(t-s) \mid \mathcal{F}_s] \\
&= (N(s)-\lambda s)^2 + \lambda(t-s)
\end{aligned}$$

as the third summand vanishes. From this we see that $(N(t) - \lambda t)^2 - \lambda t$ is a martingale.

[a] By definition, this means that $(N(t))_{t \geqslant 0}$ is a stochastic process such that (1) $N(t_j) - N(t_{j-1})$ are independent for any $t_0 < t_1 < \ldots < t_n$ for any $n \geqslant 1$, (2) $N(t) - N(s) \sim \text{Pois}(\lambda(t-s))$ for $s < t$, (3) the sample paths are RCLL. One way to show the existence is to take i.i.d. $\text{Exp}(\lambda)$ random variables $(\xi_k)_k$, define $S_k = \xi_1 + \ldots + \xi_k$ (with $S_0 = 0$) and let $N(t) = k$ for $t \in [S_k, S_{k+1})$ be the count of the number of S_is "that arrived before time t". We leave it for you to check that $N(\cdot)$ satisfies the conditions of a Poisson process.

Exercise 21

Show that for any $\theta \in \mathbb{R}$, the process $M_\theta(t) = e^{\theta N(t) - \lambda t(e^\theta - 1)}$ is a martingale. Argue that the θ-derivatives $M_\theta^{(k)}|_{\theta=0}$ are martingales and find them explicitly for $k \leqslant 4$.

If $Y \in L^1(\Omega, \mathcal{F}, \mathbb{P})$ and we define $X_t = \mathbb{E}[Y \mid \mathcal{F}_t]$, then $(X_t)_{t \geqslant 0}$ is a martingale. Such martingales are known as Doob martingales. The nicest thing that could happen to a martingale is that it turns out to be a Doob martingale. A natural necessary condition for this is that $\{X_t\}$ is uniformly integrable. For example, in the Brownian motion example, (W_t) is not a Doob martingale because W_t are not uniformly integrable. In fact, as $W_t \sim N(0,t)$, we see that $\mathbb{E}[|W_t|\mathbf{1}_{|W_t|>M}] \to \infty$ as $t \to \infty$ for any $M < \infty$. However, if we restrict the time interval, then $(W_t)_{0 \leqslant t \leqslant T}$ is uniformly integrable. It is also a Doob martingale by taking $Y = W_T$.

1. Augmentation of filtrations and Regularization of martingales

1.1. Augmented filtration. For any filtration $\mathcal{F}_\bullet$, there is a minimal enlargement that satisfies the usual conditions, known. To define it, let $\mathcal{N}$ be the collection of all null sets in $\mathcal{F}_\infty := \sigma(\cup_t \mathcal{F}_t)$

and define

$$\overline{\mathcal{F}}_t^+ := \bigcap_{s>t} \sigma(\mathcal{F}_s \cup \mathcal{N}) = \sigma\left(\bigcap_{s>t} \mathcal{F}_s \cup \mathcal{N}\right).$$

The equality is left as an exercise. The danger of enlarging the filtration is that the martingale property can be lost. However, if the sample paths are RCLL , then the martingale property remains true w.r.t. the augmented filtration.

Lemma 9

Let X be a $\mathcal{F}_\bullet$-martingale. If X has RCLL paths, then X remains a martingale w.r.t. $\overline{\mathcal{F}_\bullet}^+$.

Proof. Let $s < t$ and $A \in \overline{\mathcal{F}}_s^+$. Then we can find $B \in \mathcal{F}_s^+$ and a null set N such that $A = B \sqcup N$. Let $s_1 > s_2 > \ldots \to s$. Since $B \in \mathcal{F}_{s_n}$ and X is a martingale w.r.t. $\mathcal{F}_\bullet$, we have $\mathbb{E}[X_t \mathbf{1}_B] = \mathbb{E}[X_{s_n} \mathbf{1}_B]$. But $(X_{s_n})_n$ is a reverse martingale (that is uniformly integrable because $X_{s_n} = \mathbb{E}[X_{s_1} \mid \mathcal{F}_{s_n}]$) and hence converges to X_s a.s. and in L^1. Thus, $\mathbb{E}[X_t \mathbf{1}_B] = \mathbb{E}[X_s \mathbf{1}_B]$. Now $\mathbb{E}[X_t \mathbf{1}_N] = 0 = \mathbb{E}[X_s \mathbf{1}_N]$ as N is null. Adding, we get $\mathbb{E}[X_t \mathbf{1}_A] = \mathbb{E}[X_s \mathbf{1}_A]$. Thus, $\mathbb{E}[X_t \mid \overline{\mathcal{F}}_s^+] = X_s$ a.s. ∎

1.2. Regularized process. The more surprising thing (since nothing in the definition of martingale leads us to expect so much!) is that almost any martingale can be modified to have RCLL paths. By modification of $X = (X_t)_{t \geqslant 0}$ we mean any $Y = (Y_t)_{t \geqslant 0}$ on the same probability space such that $X_t = Y_t$ a.s., for each $t \geqslant 0$. In other words, X and Y are the same, when viewed as random variables taking values in $\mathbb{R}^{[0,\infty)}$ (endowed with the cylinder sigma-algebra).

Lemma 10

Suppose X is a $\mathcal{F}_\bullet$-martingale. Assume that $\mathcal{F}_\bullet$ satisfies the usual conditions. Then X has a modification Y that is also a $\mathcal{F}_\bullet$-martingale and has RCLL sample paths.

This proof as well as many later ones will use the results for discrete time martingales. To do that, let $D_n = 2^{-n}\mathbb{Z}_+$ and $D = \cup_{n \geqslant 0} D_n$ denote the set of dyadic rationals. Fix $T \in \mathbb{N}$ and consider $(X_t)_{t \in D_n \cap [0,T]}$ which is the finite martingale sequence $(X(0), X(1/2^n), \ldots, X(2^n T/2^n))$. Doob's maximal inequality implies that

$$\mathbb{P}\left\{\max_{t \in D_n \cap [0,T]} |X_t| \geqslant u\right\} \leqslant \frac{2\mathbb{E}|X_T|}{u}. \tag{19}$$

while Doob's upcrossing inequality implies that

$$\mathbb{E}[U_{D_n,T}[a,b]] \leqslant \frac{\mathbb{E}|X_0 - a| + \mathbb{E}|X_T - a|}{b-a} \tag{20}$$

where $U_{D_n,T}[a,b]$ is the number of upcrossings of $[a,b]$ by $(X_t)_{t\in D_n\cap[0,T]}$. The key point to note is that the bounds in (19) and (20) do not depend on n. Thus, by letting $n\uparrow\infty$, we get

$$\mathbb{P}\left\{\sup_{t\in D\cap[0,T]}|X_t|\geqslant u\right\}\leqslant\frac{2\mathbb{E}|X_T|}{u}. \tag{21}$$

and

$$\mathbb{E}[U_{D,T}[a,b]]\leqslant\frac{\mathbb{E}|X_0-a|+\mathbb{E}|X_T-a|}{b-a} \tag{22}$$

where $U_{D,T}[a,b]$ is the number of upcrossings of $[a,b]$ by $(X_t)_{t\in D\cap[0,T]}$. To be precise, $U_{D,T}[a,b]$ is the supremum of all k for which there exist dyadic rational times $0<t_1<s_1<\ldots<t_k<s_k\leqslant T$ such that $X_{t_j}\leqslant a$ and $X_{s_j}\geqslant b$ for $j\leqslant k$.

Proof of Lemma 10. Consider $(X_t)_{t\in D}$. From (21) we see that the process if bounded on $D\cap[0,T]$, a.s. From (22) we see that up to time T, the number of upcrossings across $[a,b]$ is finite, a.s. We may take intersection over $T\in\mathbb{Z}_+$ and $a,b\in\mathbb{Q}$ to say that almost surely, the paths of $(X_t)_{t\in D}$ are locally bounded and make only finitely many upcrossings in any bounded interval of time. Let Ω_0 denote the good set of full probability on which this happens.

If $\omega\in\Omega_0$, then the sequence $X_{t_n}(\omega)$ has a finite limit for any bounded increasing or decreasing sequence (t_n) (as it makes only finitely many upcrossings across any interval). Obviously the limit cannot depend on the sequence (as seen by intermixing sequences). Thus we may define

$$Y_t(\omega)=\begin{cases}\lim_{s\downarrow t}X_s(\omega) & \text{if } \omega\in\Omega_0, \text{ where the limit is taken over } s\in D,\ s>t,\\ 0 & \text{if } \omega\notin\Omega_0.\end{cases}$$

Note that even for $t\in D$ and $\omega\in\Omega_0$, we may have $Y_t(\omega)\neq X_t(\omega)$. The proof will be completed by checking that Y has all the desired properties.

(1) Y is a modification of X. Fix $t\in[0,\infty)$. Take a strictly decreasing sequence $t_n\in D$ that converges to t. Then $X_{t_n}=\mathbb{E}[X_{t_1}\mid\mathcal{F}_{t_n}]$ is a reverse martingale and hence converges almost surely to $X_t\mathbb{E}[X_{t_1}\mid\mathcal{F}_t]$ as $\mathcal{F}_t=\cap_n\mathcal{F}_{t_n}$. But Y_t is also an almost sure limit of X_{t_n}. Hence $X_t=Y_t$ a.s.

(2) Y has RCLL paths. We only need to consider $\omega\in\Omega_0$. If $t_1>t_2>\ldots\to t$, then we may pick $s_k\in(t_k,t_k+\frac{1}{k})$ such that $s_k\in D$ and $|X_{s_k}(\omega)-Y_{t_k}(\omega)|\leqslant\frac{1}{k}$. Then $s_k\downarrow t$ and hence $X_{s_k}(\omega)\to Y_t(\omega)$. But then $Y_{t_k}(\omega)$ must also converge to $Y_t(\omega)$. As for left limits, observe that Y directly inherits from X the property that the paths are locally bounded and make finitely many upcrossings across any interval in a bounded interval of time (for $\omega\in\Omega_0$). Thus, if $t_n\uparrow t$, then $Y_{t_n}(\omega)$ must have a finite limit.

(3) Y is a $\mathcal{F}_\bullet$-martingale. From the definition, we see that Y_t is $\mathcal{F}_t^+$-measurable, and by assumption that $\mathcal{F}_\bullet$ is right continuous it follows that Y is adapted to $\mathcal{F}_\bullet$. To prove that it is a martingale, we must check that for any $s < t$ and any $A \in \mathcal{F}_s$, we have $\int_A Y_t d\mathbb{P} = \int_A Y_s d\mathbb{P}$. But as $Y_t = X_t$ and $Y_s = X_s$ a.s. and X is a martingale, this follows strightaway. ■

Remark 7

The space of RCLL trajectories is also known as $D[0,\infty)$, and it is the most general space in which we have a satisfactory theory of stochastic processes. It is more general than $C[0,\infty)$, and like the latter it can be endowed with the *Skorokhod metric* that makes it complete and separable. We do not go into these issues[a], as the course will be concerned with continuous martingales.

[a] Billingsley's book *Convergence of probability measures* is an excellent reference for it, as is K. R. Parthasarathy's *Probability measures on metric spaces*.

2. Basic tools of continuous time martingales

Here we collect some of the basic tools. We have see the use of analogous results for discrete time martingales. They will in fact be used to prove the theorems for continuous time martingales.

Throughout the section, our processes are on a filtered probability space $(\Omega, \mathcal{F}, \mathcal{F}_\bullet, \mathbb{P})$. The filtration is assumed to satisfy the usual conditions. The martingale $X = (X_t)_{t \geqslant 0}$ is adapted to $\mathcal{F}_\bullet$ and is assumed to have RCLL sample paths.

2.1. Maximal inequalities. For any $T < \infty$, we have

$$\mathbb{P}\left\{\sup_{t \leqslant T} |X_t| \geqslant u\right\} \leqslant \frac{2\mathbb{E}|X_T|}{u} \quad \text{for any } u > 0. \tag{23}$$

Further, for any $p > 1$, we have

$$\mathbb{E}\left[\left(\sup_{t \leqslant T} |X_t|\right)^p\right] \leqslant C_p \sup_{t \leqslant T} \mathbb{E}[|X_t|^p]. \tag{24}$$

The inequality (24) follows immediately from (21), as the supremum over $[0,T]$ and over $[0,T] \cap D$ are equal for RCLL paths. In exactly the same way, we can get (24) from the L^p-maximal inequality.

2.2. Upcrossing inequality. For $T < \infty$, let $U_T[a,b]$ denote the number of upcrossings of $(X_t)_{t \leqslant T}$ over the interval $[a,b]$. Then,

$$\mathbb{E}[U_T[a,b]] \leqslant \frac{\mathbb{E}|X_0 - a| + \mathbb{E}|X_T - a|}{b - a} \tag{25}$$

This follows from (22), as the number of upcrossings by $(X_t)_{t \in [0,T]}$ is equal to the number of upcrossings by $(X_t)_{t \in D \cap [0,T]}$ for RCLL paths.

2.3. Stopping times. Recall that a $\mathcal{F}_\bullet$-stopping time τ is a $\mathbb{R}_+ \cup \{+\infty\}$ valued random variable such that $\{\tau \leqslant t\} \in \mathcal{F}_t$ for all $t \geqslant 0$. The sigma-algebra at τ is

$$\mathcal{F}_\tau := \{A \in \mathcal{F} : A \cap \{\tau \leqslant t\} \in \mathcal{F}_t \text{ for all } t \geqslant 0\}.$$

To see that this is a sigma-algebra, observe that

$$A^c \cap \{\tau \leqslant t\} = \{\tau \leqslant t\} \setminus (A \cap \{\tau \leqslant t\}),$$

$$(\cup_n A_n) \cap \{\tau \leqslant t\} = \cup_n (A_n \cap \{\tau \leqslant t\}),$$

which show that $\mathcal{F}_\tau$ is closed under complements and countable unions. The following exercise makes the definition of $\mathcal{F}_\tau$ as the "information up to time τ" more palatable.

Exercise 22

Assume that $\mathcal{F}_\bullet$ is the natural filtration of some stochastic process $X = (X_t)_{t \geqslant 0}$. If τ is an $\mathcal{F}_\bullet$-stopping time, then $\mathcal{F}_\tau = \sigma\{X_{\tau \wedge t} : t \geqslant 0\}$.

Here are a few useful observations about stopping times.

(1) If $\tau \leqslant \tau'$, then $\mathcal{F}_\tau \subseteq \mathcal{F}_{\tau'}$. To see this, suppose $A \in \mathcal{F}_\tau$. Then

$$A \cap \{\tau' \leqslant t\} = A \cap \{\tau \leqslant t\} \cap \{\tau' \leqslant t\}$$

which is in $\mathcal{F}_t$ because $A \cap \{\tau \leqslant t\} \in \mathcal{F}_t$ (as $A \in \mathcal{F}_\tau$) and $\{\tau' \leqslant t\}$ (as τ' is a stopping time).

(2) τ is $\mathcal{F}_\tau$-measurable. This means that $\{\tau \leqslant s\} \in \mathcal{F}_\tau$ for any s. But that is true, since

$$\{\tau \leqslant s\} \cap \{\tau \leqslant t\} = \{\tau \leqslant s \wedge t\} \in \mathcal{F}_{s \wedge t} \subseteq \mathcal{F}_t.$$

(3) There exist stopping times τ_n that take only discrete values and decrease to τ. To see this, set $\tau_n = \frac{\lfloor 2^n \tau \rfloor + 1}{2^n}$.

(4) If $\tau_n \downarrow \tau$, then $\cap_n \mathcal{F}_{\tau_n} = \mathcal{F}_\tau$.

2.4. Optional stopping/sampling theorems. Assume that the filtration satisfies the usual conditions.

Theorem 36: Optional stopping theorem

Let X be a $\mathcal{F}_\bullet$-martingale having RCLL sample paths.

(1) If τ is a $\mathcal{F}_\bullet$-stopping time, then $X^\tau = (X_{\tau \wedge t})_{t \geqslant 0}$ is a $\mathcal{F}_\bullet$-martingale. In particular, $\mathbb{E}[X_{\tau \wedge t}] = \mathbb{E}[X_0]$.

(2) If $\tau \leqslant \tau'$ are two bounded stopping times, then $\mathbb{E}[X_{\tau'} \mid \mathcal{F}_\tau] = \mathbb{E}[X_\tau]$. In particular, $\mathbb{E}[X_{\tau'}] = \mathbb{E}[X_\tau]$.

Although the two statements are phrased differently, they are about bounded stopping times (as $\tau \wedge t$ is a bounded stopping time for each t). This is too restrictive, and the boundedness condition can be removed under uniform integrability assumptions. For convenience, we state it under a stronger condition that is easier to check and suffices for almost all our purposes.

Corollary 2: Optional stopping theorem

In the setting of the theorem, assume in addition that X is L^2 bounded. Then

(1) $\mathbb{E}[X_\tau] = \mathbb{E}[X_0]$ for any finite stopping time τ.

(2) $\mathbb{E}[X_{\tau'}] = \mathbb{E}[X_\tau]$ for any two finite stopping times $\tau \leqslant \tau'$.

Proof. By Doob's inequality, $X^* = \sup_t |X_t|$ satisfies $\mathbb{E}[(X^*)^2] \leqslant C \sup_t \mathbb{E}[X_t^2]$. As, $X_{\tau\wedge t}$, $t \geqslant 0$, are dominated by X^*, and converge almost surely to X_τ (as τ is finite a.s.), by DCT we conclude that $X_{\tau\wedge t} \to X_\tau$ in L^2. In particular, $\mathbb{E}[X_\tau] = \lim_{t\to\infty} \mathbb{E}[X_{\tau\wedge t}] = \mathbb{E}[X_0]$.

The second is similar. The theorem applies to the bounded stopping times $\tau \wedge t \leqslant \tau' \wedge t$ and gives $\int_A X_{\tau'\wedge t} d\mathbb{P} = \int X_{\tau\wedge t} d\mathbb{P}$ for any $A \in \mathcal{F}_{\tau\wedge t}$. Fix $A \in \mathcal{F}_{\tau\wedge s}$ and let $s < t \uparrow \infty$ to get $\int_A X_{\tau'} d\mathbb{P} = \int X_\tau d\mathbb{P}$ (by DCT, using domination by X^*). As $\cup_s \mathcal{F}_{\tau\wedge s}$ is an algebra that generates $\mathcal{F}_\tau$, it follows that the same identity holds for all $A \in \mathcal{F}_\tau$. ■

Theorem 36 will be proved by discretization. To go from the discrete to the continuous, we shall again need uniform integrability, but only of $(X_t)_{t\leqslant T}$ for any fixed $T < \infty$. This is because $X_t = \mathbb{E}[X_T \mid \mathcal{F}_t]$ for $t \leqslant T$. In particular, if $X_T^* = \sup_{t\leqslant T} |X_t|$, then by Doob's inequality ,

Proof of Theorem 36. Let $\tau_n = \frac{\lfloor 2^n \tau \rfloor + 1}{2^n}$ be the discrete valued stopping times that decrease to τ. Fix $s < t$ and let $s_n = \frac{\lfloor 2^n t \rfloor + 1}{2^n}$ and $t_n = \frac{\lfloor 2^n s \rfloor + 1}{2^n}$. If $A \in \mathcal{F}_s$, then $A \in \mathcal{F}_{s_n}$ and by the discrete time OST applied to $(X_{k/2^n})_{k\geqslant 0}$ we see that $\int_A X_{t_n\wedge\tau_n} d\mathbb{P} = \int_A X_{s_n\wedge\tau_n} d\mathbb{P}$. Now, $X_{t_n\wedge\tau_n} \overset{a.s.}{\to} X_{\tau\wedge t}$ and $X_{s_n\wedge\tau_n} \overset{a.s.}{\to} X_{s\wedge\tau}$ as $n \to \infty$. If we prove that these sequences are uniformly integrable, we get $\int_A X_{t\wedge\tau} d\mathbb{P} = \int_A X_{s\wedge\tau} d\mathbb{P}$, which shows that $(X_{t\wedge\tau})_{t\geqslant 0}$ is a martingale.

To get uniform integrability, fix $T > t$ and recall that $X[0,T]$ is uniformly integrable (as $X_s = \mathbb{E}[X_T \mid \mathcal{F}_s]$ for all $s \leqslant T$). Now $X_{\tau_n\wedge t_n}$ is a ■

3. Square integrable continuous martingales

These are the main objects of interest for us. By definition, a square integrable, continuous martingale is a stochastic process $X = (X_t)_{t\geqslant 0}$ that is a martingale on a filtered probability space (with a filtration satisfying the usual conditions) and having continuous sample paths[1]. We shall think of X as a $C[0,\infty)$ valued random variable.

[1] In some books such as Bass's *Stochastic processes*, uniform integrability is also assumed. We shall keep that separate.

If X is a square integrable martingale, then for any $s \leqslant t \leqslant u \leqslant v$, we have $X_t - X_s \perp X_v - X_u$ in $L^2(\mathbb{P})$. The proof is the same as in the discrete case.

Theorem 37: Doob-Meyer decomposition

Let X be a continuous sub-martingale w.r.t. $\mathcal{F}_\bullet$. Then there is a unique way t
$M+A$ where M is a continuous martingale w.r.t $\mathcal{F}_\bullet$ and A is an adaptd increasing
process with $A_0 = 0$.

3.1. Quadratic variation.

We shall assume this theorem, see one of the reference books for a proof. The key application is as follows. Let M be a continuous martingale. Then M^2 is a continuous sub-martingale, and hence has a Doob-Meyer decomposition. The increasing process in the decomposition is denoted $\langle M \rangle$ and called the *quadratic variation* process of M. Its defining property is that $\langle M \rangle_0 = 0$ and $M_t^2 - \langle M \rangle_t$ is a martingale.

This only means that X_t are random variables on a common probbaility space or equivalently that X is a $\mathbb{R}^{\mathbb{R}_+} = \{f : \mathbb{R}_+ \to \mathbb{R}\}$ valued random variable. As the cylinder sigma-algebra on $\mathbb{R}^{\mathbb{R}_+}$ contains only events that are describable in terms of countably many co-ordinates, this space is too large. We prefer that our random variables take values in a smaller space such as $C[0,\infty)$, which is a complete, separable metric space. But $C[0,\infty)$ is inadequate to cover all examples of interest, for instance the Poisson process $N(\cdot)$ above has jumps. The most general class of functions that is sufficient for all our purposes is the *Skorokhod space*

$$D[0,\infty) = \{f : \mathbb{R}_+ \to \mathbb{R} : f(t+) = f(t) \text{ and } f(t-) \text{ exists and is finite, for all } t \geqslant 0\}.$$

Here $f(t\pm)$ are the limits of $f(t \pm h)$ as $h \downarrow 0$. This class of functions is also called RCLL (right continuous with left limits) or *càdlàg* which is an abbreviation of a French phrase with the same meaning. It is useful to also have at hand the space $D[0,1]$, where the functions have the domain $[0,1]$.

To talk of random variables taking values in $D[0,\infty)$ or $D[0,1]$, we endow these spaces with a topology and consequently the corresponding Borel sigma algebras. Even if we stick to $D[0,1]$, sup-norm is not a good way to measure distances. For example, we want $\mathbf{1}_{[1/n,1]}$ to converge to $\mathbf{1}_{[0,1]}$ as $n \to \infty$, but the sup-norm of the difference is 1 for all n. This is reminiscent of the same issues that lead to Lévy metric on probability measures on $\mathbb{R}$. What we need is to allow a leeway in the domain space, not just in the range. For this purpose, let us introduce the legal set of reparameterizations

$$\mathcal{R} = \{\phi : [0,1] \to [0,1] : \phi : [0,1] \to [0,1] \text{ is strictly increasing, continuous and bijective}\}.$$

Let $I \in \mathcal{R}$ denote the identity map from $[0,1] \to [0,1]$. For $f, g \in D[0,1]$ define

$$d(f,g) = \inf_{\phi \in \mathcal{R}} \|\phi - I\|_{\sup} + \|f \circ \phi - g\|_{\sup}.$$

Observe that $D[0,1]$ functions are bounded (if $f(t_n) \to \pm\infty$, we can find a decreasing or increasing subsequence t_{n_k}. If t is the limit of that subsequence, either the left limit or right limit at t fails to exist) and therefore this $d(f,g)$ is well-defined and finite. To illustrate the role of the two terms in the definition of d, consider

(1) f, g continuous: We can take $\phi = I$ to see that $d(f,g) \leqslant \|f - g\|_{\sup}$.

(2) $f = \mathbf{1}_{[t,1]}$ and $g = \mathbf{1}_{[s,1]}$: We can take a ϕ that maps t to s to see that $d(f,g) \leqslant |s - t|$.

It is not difficult to see that d is a metric (Exercise!). The fact that $d(f,g) = 0$ implies $f = g$ requires them to be elements of $D[0,1]$ (e.g., if $f = \mathbf{1}_{[0,.5]}$ and $g = \mathbf{1}_{[0,0.5)}$, then $d(f,g) = 0$, but $f \notin D[0,1]$). The corresponding topology is called *Skorokhod topology*. It can be shown that $D[0,1]$ is a complete, separable metric space[2]. It was realized in 1950s that complete separable metric spaces are the natural settings for spaces where random variables take values. Thus, $D[0,1]$ (and $D[0,\infty)$) are the natural spaces to model random trajectories in one dimension.

[2]But not under the above given metric! It can be modified to become complete. This is similar to how $(0,1]$ is not a complete metric space under the usual metric, but it is complete under the metric $|\frac{1}{x}-\frac{1}{y}|$