



Stochastic processes I

Sommer 2009, Uni Bonn

Anton Bovier

*Institut für Angewandte Mathematik
Rheinische Friedrich-Wilhelms-Universität Bonn
Endenicher Allee 60
53115 Bonn*

Version: March 25, 2011

Contents

1	A review of measure theory	<i>page</i> 1
1.1	Probability spaces	1
1.2	Construction of measures	6
1.3	Random variables	14
1.4	Integrals	16
1.5	\mathcal{L}^p and L^p spaces	20
1.6	Fubini's theorem	23
1.7	Densities, Radon-Nikodým derivatives	23
2	Conditional expectations and conditional probabilities	31
2.1	Conditional expectations	31
2.2	Elementary properties of conditional expectations	35
2.3	The case of random variables with absolutely continuous distributions	37
2.4	The special case of L^2 -random variables	38
2.5	Conditional probabilities and conditional probability measures	39
3	Stochastic processes	41
3.1	Definition of stochastic processes	41
3.2	Construction of stochastic processes; Kolmogorov's theorem	44
3.3	Examples of stochastic processes	47
3.3.1	Independent random variables	47
3.3.2	Gaussian processes	48
3.3.3	Markov processes	52
3.3.4	Gibbs measures	55
4	Martingales	57
4.1	Definitions	57

4.2	Upcrossings and convergence	60
4.3	Inequalities	67
4.4	Doob decomposition	69
4.5	A discrete time Itô formula.	72
4.6	Central limit theorem for martingales	73
4.7	Stopping times, optional stopping	78
5	Markov processes	83
5.1	Markov processes with stationary transition probabilities	83
5.2	The strong Markov property	85
5.3	Markov processes and martingales	85
5.4	Harmonic functions and martingales	89
5.5	Dirichlet problems	90
5.6	Doob's h -transform	92
5.7	Markov chains with countable state space	95
6	Random walks and Brownian motion	99
6.1	Random walks	99
6.2	Construction of Brownian motion	101
6.3	Donsker's invariance principle	105
6.4	Martingale and Markov properties	108
6.5	Sample path properties	111
6.6	The law of the iterated logarithm	114
	<i>Bibliography</i>	119
	<i>Index</i>	121

A review of measure theory



In this first chapter I review the main concepts of measure theory that we will need. I will not give proofs in most cases. Those familiar with my W-Theorie 1 lecture will find that most of this material was covered there, except that here we will take a somewhat more abstract point of view, replacing the space of real numbers by arbitrary metric spaces. One will see, however, that this implies very few changes. For more details, there is a wealth of references on measure theory. See e.g. [2, 14, 10, 7, 5, 1].

1.1 Probability spaces

A *space*, Ω , is an arbitrary non-empty set. Elements of a space Ω will be denoted by ω . If $A \subset \Omega$ is a subset of Ω , we denote by $\mathbb{1}_A$ the indicator function of the set A , i.e.

$$\mathbb{1}_A(\omega) = \begin{cases} 1 & , \text{ if } \omega \in A, \\ 0 & , \text{ if } \omega \in A^c \equiv \Omega \setminus A. \end{cases} \quad (1.1)$$

Definition 1.1.1 Let Ω be a space. A family $\mathcal{A} \equiv \{A_\lambda\}_{\lambda \in I}$, $A_\lambda \subset \Omega$, with I an arbitrary set, is called a *class* of Ω . A non-empty class of Ω is called an *algebra*, if:

- (i) $\Omega \in \mathcal{A}$.
- (ii) For all $A \in \mathcal{A}$, $A^c \in \mathcal{A}$.
- (iii) For all $A, B \in \mathcal{A}$, $A \cup B \in \mathcal{A}$.

If \mathcal{A} is an algebra, and moreover

- (iv) $\bigcup_{n=1}^{\infty} A_n \in \mathcal{A}$, whenever for all $n \in \mathbb{N}$, $A_n \in \mathcal{A}$,

then \mathcal{A} is called a σ -algebra.

Definition 1.1.2 A space, Ω , together with a σ -algebra, \mathcal{F} , of subsets of Ω is called a *measurable space*, (Ω, \mathcal{F}) .

Definition 1.1.3 Let (Ω, \mathcal{F}) be a measurable space. A map $\mu : \mathcal{F} \rightarrow [0, \infty]$ from \mathcal{F} the non-negative real numbers (and infinity) is called a (positive) *measure*, if

- (i) $\mu(\emptyset) = 0$.
- (ii) For any countable family $\{A_n\}_{n \in \mathbb{N}}$ of mutually disjoint elements of \mathcal{F} ,

$$\mu \left(\bigcup_{n \in \mathbb{N}} A_n \right) = \sum_{n \in \mathbb{N}} \mu(A_n). \quad (1.2)$$

A measure, μ , is called finite, if $\mu(\Omega) < \infty$. A measure is called σ -finite, if there exists a countable class, Ω_n , of subsets of Ω , such that $\Omega = \bigcup_{n \in \mathbb{N}} \Omega_n$, such that, for all $n \in \mathbb{N}$, $\mu(\Omega_n) < \infty$.

A triple, $(\Omega, \mathcal{F}, \mu)$, is called a *measure space*.

Definition 1.1.4 Let (Ω, \mathcal{F}) be a measurable space. A positive measure, \mathbb{P} , on (Ω, \mathcal{F}) that satisfies $\mathbb{P}[\Omega] = 1$ is called a *probability measure*. A triple $(\Omega, \mathcal{F}, \mathbb{P})$, where Ω is a set, \mathcal{F} a σ -algebra of subsets of Ω , and \mathbb{P} a probability measure on (Ω, \mathcal{F}) , is called a *probability space*.

Probability spaces provide the scenery where probability theory takes place. The set of sceneries is huge, since we have so far not made any restriction on the allowable spaces Ω . In most instances, we will, however, want to stay on reasonable grounds. Fortunately, there is a quite canonical setting where everything we ever want to do can be constructed. This is the realm where Ω is a topological space and $\mathcal{F} = \mathcal{B}(\Omega)$ is the Borel- σ -algebra of Ω .

We recall the definition of a topological space.

Definition 1.1.5 A space, E , is called a *topological space*, if for every point $p \in E$ there exists a collection, \mathcal{U}_p , of subsets of E , called *neighborhoods*, with the following properties:

- (i) For every point, p , $\mathcal{U}_p \neq \emptyset$.
- (ii) Every neighborhood of p contains p .
- (iii) If $U_1, U_2 \in \mathcal{U}_p$, then there exists $U_3 \in \mathcal{U}_p$ such that $U_3 \subset U_1 \cap U_2$.
- (iv) If $U \in \mathcal{U}_p$ and $q \in U$, then there exists $V \in \mathcal{U}_q$ such that $V \subset U$.

Recall that in a topological space, one can define the notions such *open sets* and *closed sets*; open sets have the property that any of its points has a neighborhood that is contained in the set, and closed sets are the complements of open sets. Note that the empty set is also considered as an open set by default. Since the entire space E is also open, the empty set is, however, also a closed set. Two topological spaces are considered *equivalent*, or have the same *topology*, if they contain the same open sets. In particular, given two sets of collections of neighborhoods, \mathcal{U}_p , and \mathcal{V}_p , on a space E , then they generate the same topology, if for any $p \in E$, and any $U \in \mathcal{U}_p$, there exists $V \in \mathcal{V}_p$ such that $V \subset U$ and for any $V \in \mathcal{V}_p$ there exists $U \in \mathcal{U}_p$ such that $U \subset V$.

- Definition 1.1.6**(i) A topological space, Ω , is called *Hausdorff*, if any two distinct points in Ω have disjoint neighborhoods.
- (ii) A topological space, E , is called *separable*, if there exists a countable subset, $E_0 \subset E$ whose closure¹ is E .

Definition 1.1.7 Let E be a topological space. The *Borel- σ -algebra*, $\mathcal{B}(E)$ of E is the smallest σ -Algebra that contains all open sets of E .

The point behind the notion of the Borel- σ -algebra is that it is big enough to satisfy our needs, but small enough to ensure that it is possible to construct measures on it. Larger σ -algebras, such as the power set on any uncountable topological space, do not usually allow to define measures with nice properties on them.

One says that the Borel- σ -algebra is *generated* by the open sets of E . This notion will be used quite frequently. We say in general that a class, \mathcal{A} , of a space Ω generates a σ -Algebra, $\sigma(\mathcal{A})$, defined as the smallest σ -algebra that contains \mathcal{A} ,

$$\sigma(\mathcal{A}) \equiv \bigcap_{\substack{\mathcal{F} \supset \mathcal{A} \\ \mathcal{F} \text{ is } \sigma\text{-algebra}}} \mathcal{F}.$$

Even more structure appears if we work on so-called *metric spaces*.

Definition 1.1.8 Let E be a set. A map, $\rho : M \times M \rightarrow [0, +\infty]$, is called a *metric*, if

- (i) $\rho(x, y) = 0$, if and only if $x = y$;
- (ii) $\rho(x, y) = \rho(y, x)$;
- (iii) $\rho(x, z) \leq \rho(x, y) + \rho(y, z)$, for all $x, y, z \in E$.

¹ The closure of a subset, A , of a topological space is the intersection of all closed subsets containing A .

The set $B_r(x) \equiv \{y \in E : \rho(x, y) < r\}$ is called the (open) ball of radius r .

The set of neighborhoods obtained from the open balls associated to a metric, ρ , is called the *metric topology*. A topological space endowed with a metric and its metric topology is called a *metric space*.

A sequence $x_n \in E$, $n \in \mathbb{N}$ is called a *Cauchy sequence*, if for any $\epsilon > 0$, there exists $n_0 \in \mathbb{N}$, such that for all $n, m \geq n_0$, $\rho(x_n, x_m) < \epsilon$.

A metric space, E , is called complete, if any Cauchy sequence in E converges.

A further useful specialisation is the restriction to so called *Polish spaces*, that is *separable, completely metrisable spaces*. A completely metrisable space is a space that is homeomorphic to a complete metric space. That is, a Polish space is essentially a complete, separable metric space up to the fact that the metric may not have been fixed. Recall that \mathbb{R}^d is a Polish space, and so is $\mathbb{R}^{\mathbb{N}}$ when equipped with the product topology.

Note that in many cases, different families of sets generate the same σ -algebra. For instance, if E is not only a topological space, but a metric space with the topology given by the metric topology, then the set of open balls generates the Borel- σ -algebra. But also the set of closed balls will generate. If E is the real line, the half-lines also generate the Borel- σ -algebra.

An advantage in Ω being a Polish space lies in the fact that one can choose as a generator of the Borel- σ -algebra a *countable* collection of sets. E.g., in the case of the real line, the Borel- σ -algebra is already generated by the half-lines $(-\infty, q]$, with $q \in \mathbb{Q}$ (just observe that if x is any real number, there exists a sequence $q_n \downarrow x$, and the set $\bigcap_{n \in \mathbb{N}} (-\infty, q_n] = (-\infty, x]$ is also contained in the σ -algebra generated by these half-lines).

A related, but more general class of spaces are sometimes useful. These are called *Lousin spaces*. These are spaces that are homeomorphic to a Borel subset of a compact metric space.

Two notions of special types of classes are very useful in this context.

Definition 1.1.9 Let Ω be a space. A class of Ω , \mathcal{T} , is called a Π -system, if \mathcal{T} is closed under finite intersections; a class, \mathcal{G} , is called a λ -system, if

- (i) $\Omega \in \mathcal{G}$,
- (ii) If $A, B \in \mathcal{G}$, and $A \supset B$, then $A \setminus B \in \mathcal{G}$,
- (iii) If $A_n \in \mathcal{G}$ and $A_n \subset A_{n+1}$, then $\lim_{n \uparrow \infty} A_n \in \mathcal{G}$.

The following useful observation is called Dynkin's theorem.

Theorem 1.1.1 *If \mathcal{T} is a Π -system and \mathcal{G} is a λ -system, then $\mathcal{G} \supset \mathcal{T}$ implies that \mathcal{G} contains the smallest σ -algebra containing \mathcal{T} .*

The most useful application of Dynkin's theorem is the observation that, if two probability measures are equal on a Π -system that generates the σ -algebra, then they are equal on the σ -algebra. (since the set on which the two measures coincide forms a λ -system containing \mathcal{T}).

Examples. The general setup allows to treat many important examples on the same footing.

Countable spaces. If Ω is a countable space, the natural topology is the *discrete topology*. Here the set of neighborhoods of a point p is just the set $\{p\}$ itself. Clearly this is a topology, and all sets are open and closed with respect to it. The Borel- σ -algebra consists of the power set of Ω .

Euclidean space. \mathbb{R}^d equipped with the Euclidean metric is a metric space. Choosing as sets of neighborhoods the set of all open balls, $B_r(p) \equiv \{x \in \mathbb{R}^d : \|x - p\| < r\}$ turns this into a topological space. The corresponding Borel- σ -algebra is the smallest σ -algebra containing all these balls.

Note that, since on \mathbb{R}^d the Euclidean norm and the sup-norm are equivalent, the Borel- σ -algebra is also generated by open (or closed) rectangles.

Infinite product spaces. If E is a topological space, then the infinite Cartesian product space, E^∞ , can also be turned into a topological space through the *product topology*. Here the set of neighborhoods of a point $p \equiv (p_1, p_2, p_3, \dots)$ is given by the collection of sets

$$U_{p_1} \times U_{p_2} \times U_{p_k} \times E \times E \times \dots, \quad (1.3)$$

where $k \in \mathbb{N}$, and $U_{p_i} \in \mathcal{U}_{p_i}$. If $\mathcal{B}(E)$ is the Borel- σ -algebra of E , then the Borel- σ -algebra of E^∞ is the product σ -algebra, $\mathcal{B}(E^\infty) = \mathcal{B}(E)^{\otimes \infty}$, i.e. the σ -Algebra that is generated by the family of sets $A_1 \times \dots \times A_k \times E \times \dots$, $k \in \mathbb{N}$, $A_i \in \mathcal{B}(E)$ (where of course it also suffices to choose the sets $E \times \dots \times E \times A_k \times E \times \dots$, $k \in \mathbb{N}$, and A_k running through a generator of $\mathcal{B}(E)$).

If E is a metric space, then one can also turn E^∞ into a metric space, such that the associated metric topology is equivalent to the product topology. This is done, e.g. by setting

$$\|p - q\| \equiv \sum_{n=1}^{\infty} 2^{-n} \|p_n - q_n\|. \quad (1.4)$$

Note that this implies that, if E is a Polish space, then the infinite product space E^∞ equipped with the product topology is also Polish.

Infinite product spaces will be the scenario to discuss stochastic processes with discrete time, the main topic of this course.

Function spaces. Important examples of metric spaces are normed function spaces, such as the space of bounded, real-valued functions on \mathbb{R} (or subsets $I \subset \mathbb{R}$), equipped with the supremum norm

$$\|f - g\|_\infty \equiv \sup_{t \in I} |f(t) - g(t)|.$$

We will begin to deal with such examples in the later parts of this course, when we introduce Gaussian random processes with continuous time.

Spaces of measures. Another space we are often encountering in probability theory is that of *measures* on a Borel- σ -algebra. There are various ways to introduce topologies on spaces of measures, but a very common one is the so-called *weak topology*. Let E be the topological space in question, and $C_0(E, \mathbb{R})$ the space of real-valued, bounded, and continuous functions on E . We denote by $\mathcal{M}_+(E, \mathcal{B}(E))$ the set of all positive measures on $(E, \mathcal{B}(E))$. One can then define neighborhoods of a measure μ of the form

$$B_{\epsilon, k, f_1, \dots, f_k}(\mu) \equiv \left\{ \nu \in \mathcal{M}_+(E, \mathcal{B}(E)) : \max_{i=1}^k |\mu(f_i) - \nu(f_i)| < \epsilon \right\}, \quad (1.5)$$

where $\epsilon > 0$, $k \in \mathbb{N}$, and $f_i \in C_0(E, \mathbb{R})$.

If E is a Polish space, then the weak topology can also be derived from a suitably defined metric.

1.2 Construction of measures

The problem of the construction of measures in the general context of topological spaces is not entirely trivial. This is due to the richness of a Borel- σ -algebra and the hidden subtlety associated with the requirement of σ -additivity. The general strategy is to construct a “measure” first on a simpler set, an algebra or a semi-algebra, and then to use a powerful theorem ensuring the unique extendibility to the σ -algebra.

To do this we first define the notion of a σ -additive set function.

Definition 1.2.1 Let \mathcal{A} be a class of subset of some set Ω . A function $\nu : \mathcal{A} \rightarrow [0, \infty]$ is called a positive, σ - additive (or countably additive) set function, if

- (i) $\nu(\emptyset) = 0$;
- (ii) For any sequence, $A_k, k \in \mathbb{N}$, of mutually disjoint elements of \mathcal{A} such that $\bigcup_{k \in \mathbb{N}} A_k \in \mathcal{A}$,

$$\nu \left(\bigcup_{k \in \mathbb{N}} A_k \right) = \sum_{k \in \mathbb{N}} \nu(A_k) \quad (1.6)$$

The aim of this section is to prove the following version of Carathéodory's theorem.

Theorem 1.2.2 Let Ω be a set and let \mathcal{S} be an algebra on Ω . Let μ_0 be a countably additive map $\mathcal{S} \rightarrow [0, \infty]$. Then there exists a measure, μ , on $(\Omega, \sigma(\mathcal{S}))$, such that $\mu = \mu_0$ on \mathcal{S} .

If μ_0 is σ -finite, then μ is unique.

Proof. We begin by defining the notion of an *outer measure*.

Definition 1.2.2 Let Ω be a set. A map $\mu^* : \mathcal{P}(\Omega) \rightarrow [0, \infty]$ is called an outer measure if,

- (i) $\mu^*(\emptyset) = 0$;
- (ii) If $A \subset B$, then $\mu^*(A) \leq \mu^*(B)$ (increasing);
- (iii) for any sequence $A_n \in \mathcal{P}(\Omega)$, $n \in \mathbb{N}$,

$$\mu^* \left(\bigcup_{n \in \mathbb{N}} A_n \right) \leq \sum_{n \in \mathbb{N}} \mu^*(A_n) \quad (1.7)$$

(σ -sub-additivity).

Note that an outer measure is far less constraint than a measure; this is why it can be defined on any set, not just on σ -algebras.

Example. If $(\Omega, \mathcal{F}, \mu)$ is a measure space, we can define an extension of μ that will be an outer measure on $\mathcal{P}(\Omega)$ as follows: For any $D \subset \Omega$, let

$$\mu^*(D) \equiv \inf\{\mu(F) : F \in \mathcal{F}; F \supset D\}. \quad (1.8)$$

This is of course not how we want to proceed when constructing a measure. Rather, we will construct an outer measure from a σ -additive function on an algebra (that is also a Π -system), and then use this to construct a measure.

Next we define the notion of μ^* -measurability of sets.

Definition 1.2.3 A subset $B \subset \Omega$ is called μ^* -measurable, if, for all subsets $A \subset \Omega$,

$$\mu^*(A) = \mu^*(A \cap B) + \mu^*(A \cap B^c). \quad (1.9)$$

The set of μ^* -measurable sets is called $\mathcal{M}(\mu^*)$.

Theorem 1.2.3 (i) $\mathcal{M}(\mu^*)$ is a σ -algebra that contains all subsets $B \subset \Omega$ such that $\mu^*(B) = 0$.

(ii) The restriction of μ^* to $\mathcal{M}(\mu^*)$ is a measure.

Proof. Note first that in general, by sub-additivity,

$$\mu^*(A) \leq \mu^*(A \cap B) + \mu^*(A \cap B^c). \quad (1.10)$$

If $\mu^*(B) = 0$, we have also that

$$\mu^*(A) \geq \mu^*(A \cap B^c) = \mu^*(A \cap B) + \mu^*(A \cap B^c). \quad (1.11)$$

Thus, $\mathcal{M}(\mu^*)$ contains all set B with $\mu^*(B) = 0$. This implies in particular that $\emptyset \in \mathcal{M}(\mu^*)$. Also, by the symmetry of the definition, $\mathcal{M}(\mu^*)$ contains all its sets together with their complements. Thus the only non-trivial thing to show (i) is the stability under countable unions. Let B_1, B_2 be in $\mathcal{M}(\mu^*)$. Then

$$\begin{aligned} \mu^*(A \cap (B_1 \cup B_2)) &= \mu^*(A \cap (B_1 \cup B_2) \cap B_1) + \mu^*(A \cap (B_1 \cup B_2) \cap B_1^c) \\ &= \mu^*(A \cap B_1) + \mu^*(A \cap B_2 \cap B_1^c), \end{aligned} \quad (1.12)$$

where we used that $B_1 \in \mathcal{M}(\mu^*)$ for the first equality. Then

$$\begin{aligned} &\mu^*(A \cap (B_1 \cup B_2)) + \mu^*(A \cap (B_1 \cup B_2)^c) \\ &= \mu^*(A \cap B_1) + \mu^*(A \cap B_2 \cap B_1^c) + \mu^*(A \cap B_1^c \cap B_2^c) \\ &= \mu^*(A \cap B_1) + \mu^*(A \cap B_1^c) = \mu^*(A). \end{aligned} \quad (1.13)$$

Thus $B_1 \cup B_2 \in \mathcal{M}(\mu^*)$. This implies that $\mathcal{M}(\mu^*)$ is closed under finite union. Since it is also closed under passage to the complement, it is closed under finite intersection. Thus it is enough to show that countable unions of pairwise disjoint sets, $B_k \in \mathcal{M}(\mu^*)$, $k \in \mathbb{N}$, are in $\mathcal{M}(\mu^*)$. To show this, we show that, for all $m \in \mathbb{N}$,

$$\mu(A) = \sum_{n=1}^m \mu^*(A \cap B_n) + \mu^* \left(A \cap \bigcap_{n=1}^m B_n^c \right). \quad (1.14)$$

This holds for $m = 1$ by definition, and if it holds for m , then

$$\begin{aligned}\mu^* \left(A \cap \bigcap_{n=1}^m B_n^c \right) &= \mu^* \left(A \cap \bigcap_{n=1}^m B_n^c \cap B_{m+1} \right) + \mu^* \left(A \cap \bigcap_{n=1}^{m+1} B_n^c \right) \\ &= \mu^* (A \cap B_{m+1}) + \mu^* \left(A \cap \bigcap_{n=1}^{m+1} B_n \right),\end{aligned}$$

so, inserting this into (1.14), it holds for $m + 1$. Hence, by induction, it is true for all $m \in \mathbb{N}$.

From (1.14) we deduce further that

$$\mu(A) \geq \sum_{n=1}^m \mu^*(A \cap B_n) + \mu^* \left(A \cap \bigcap_{n=1}^{\infty} B_n^c \right). \quad (1.15)$$

Now we let m tend to infinity, and use sub-additivity:

$$\begin{aligned}\mu(A) &\geq \sum_{n=1}^{\infty} \mu^*(A \cap B_n) + \mu^* \left(A \cap \bigcap_{n=1}^{\infty} B_n^c \right) \\ &\geq \mu^* \left(A \cap \left(\bigcup_{n=1}^{\infty} B_n \right) \right) + \mu^* \left(A \cap \bigcap_{n=1}^{\infty} B_n^c \right).\end{aligned} \quad (1.16)$$

Since the converse inequality follows by sub-additivity, equality holds in (1.16) and thus the union $\bigcup_{n=1}^{\infty} B_n \in \mathcal{M}(\mu^*)$.

It remains to prove that μ^* restricted to $\mathcal{M}(\mu^*)$ is a measure. We know already that $\mu^*(\emptyset) = 0$. Let now B_n be disjoint as above. Let us choose in the first line of (1.16) $A = \bigcup_{n=1}^{\infty} B_n$. This gives

$$\mu^* \left(\bigcup_{n=1}^{\infty} B_n \right) \geq \sum_{n=1}^{\infty} \mu^*(B_n). \quad (1.17)$$

Since the converse inequality holds by sub-additivity, equality holds and the result is proven. \square

The preceding theorem provides a clear strategy for proving Carathéodory's theorem. All we need is to prescribe a σ -additive function, μ , on the algebra. Then construct an outer measure μ^* . This can be done in the following way: If \mathcal{S} is an algebra, set

$$\mu^*(D) = \inf\{\mu(A) : A \in \mathcal{S}; A \supset D\} \quad (1.18)$$

One needs to show that this is sub-additive and defines an outer measure. Once this is done, it remains to show that $\mathcal{M}(\mu^*)$ contains $\sigma(\mathcal{S})$. This is done by showing that it contains \mathcal{S} , since $\mathcal{M}(\mu^*)$ is a σ -algebra.

Let us now conclude our proof by carrying out these steps.

Lemma 1.2.4 *Let \mathcal{S} be an algebra, μ a σ -additive function on \mathcal{S} , and μ^* defined by (1.18). Then μ^* is an outer measure.*

Proof. First, note that the first two conditions for μ^* to be an outer measure are trivially satisfied. To prove sub-additivity, let A_n , $n \in \mathbb{N}$ be a family of subsets of Ω . For each n , we can choose sets $F_n \in \mathcal{S}$, such that $A_n \subset F_n$, and $\mu(F_n) \leq \mu^*(A_n) + \epsilon 2^{-n}$, for any $\epsilon > 0$ (since $\mu^*(A_n) = \inf\{\mu(F) : F \in \mathcal{S}; F \supset A_n\}$). Then, since $\bigcup_n F_n \supset \bigcup_n A_n$, and μ is σ -additive,

$$\mu^*\left(\bigcup_{n \in \mathbb{N}} A_n\right) \leq \mu\left(\bigcup_{n \in \mathbb{N}} F_n\right) \leq \sum_{n \in \mathbb{N}} \mu(F_n) \leq \sum_{n \in \mathbb{N}} \mu^*(A_n) + 2\epsilon, \quad (1.19)$$

which proves the claim since $\epsilon > 0$ is arbitrary. \square

Lemma 1.2.5 *Let μ^* be the outer measure defined by (1.18). Let $\mathcal{M}(\mu^*)$ be the σ -algebra of μ^* -measurable sets. Then $\sigma(\mathcal{S}) \subset \mathcal{M}(\mu^*)$.*

Proof. We must show that $\mathcal{M}(\mu^*)$ contains a family that generates $\sigma(\mathcal{S})$. In fact, we will show that it contains all the elements of the algebra \mathcal{S} . To see this, let $A \subset \Omega$ be arbitrary. Then (if $\mu^*(A) < \infty$), for any $\epsilon > 0$, there is a set $F \in \mathcal{S}$, such that $A \subset F$ and $\mu^*(A) \geq \mu(F) - \epsilon$. But then, for $B \in \mathcal{S}$,

$$\mu^*(A \cap B) \leq \mu(F \cap B) \quad (1.20)$$

and also

$$\mu^*(A \cap B^c) \leq \mu(F \cap B^c) \quad (1.21)$$

But the two sets on the right-hand sides are disjoint and in \mathcal{S} . Thus

$$\mu^*(A \cap B) + \mu^*(A \cap B^c) \leq \mu(F \cap B) + \mu(F \cap B^c) = \mu(F) \leq \mu^*(A) + \epsilon. \quad (1.22)$$

This proves

$$\mu^*(A) \geq \mu^*(A \cap B) + \mu^*(A \cap B^c) \quad (1.23)$$

and since the opposite inequality follows by sub-additivity, $B \in \mathcal{M}(\mu^*)$. \square

Thus we have in fact constructed an outer measure that is a measure on $\sigma(\mathcal{S})$ and that extends μ on \mathcal{S} . The uniqueness of the extension in the finite case follows from Dynkin's theorem. Assume that there are two extensions, μ and ν that coincide on \mathcal{S} . One verifies easily that the class of sets where $\mu(B) = \nu(B)$ is a λ -system which contains the Π -system

\mathcal{S} ; by Dynkin's theorem this λ -system must be $\sigma(\mathcal{S})$. Finally, if μ is σ -finite, one uses the following standard argument (that allows to carry many results from finite to σ -finite measures): By σ -finiteness, there exists a sequence of increasing sets, $\Omega_n \uparrow \Omega$, with $\mu(\Omega_n) < \infty$. Then the measure $\mu_n \equiv \mu \mathbb{I}_{\Omega_n} \uparrow \mu$. So if there are two extensions of a given σ -additive set function, then their restrictions to all Ω_n are finite measures and must coincide. But then so must their limits. This concludes the proof of Carathéodory's theorem. \square

Remark 1.2.1 Carathéodory's theorem should appear rather striking at first by its generality. It makes no assumptions on the nature of the space Ω whatsoever. Does this mean that the construction of a measure is in general trivial? The answer is of course no, but Carathéodory's theorem separates clearly the topological aspects from the algebraic aspects of measure theory. Namely, it shows that in a concrete situation, to construct a measure one needs to construct a σ -additive set function on an algebra that contains a Π -system that will generate the desired σ -algebra. The proof of Carathéodory's theorem shows that the extension to a measure is essentially a matter of algebra and completely general. We will see later how topological aspects enter into the construction of additive set functions, and why aspects like separability and metric topologies become relevant.

Remark 1.2.2 The σ -algebra $\mathcal{M}(\mu^*)$ is in general not equal to the σ -algebra generated by \mathcal{S} . In particular, we have seen that $\mathcal{M}(\mu^*)$ contains all sets of μ^* -measure zero, all of which need not be in $\sigma(\mathcal{S})$. This observation suggests to consider in general extensions of a given σ -algebra with respect to a measure that ensures that all sets of measure zero are measurable. Let $(\Omega, \mathcal{F}, \mu)$ be a measure space. Define the outer measure, μ^* , as in (1.8), and define the inner measure, μ_* , as

$$\mu_*(D) \equiv \sup\{\mu(F) : F \in \mathcal{F}; F \subset D\}. \quad (1.24)$$

Then

$$\mathcal{M}(\mu) \equiv \{A \subset \Omega : \mu_*(A) = \mu^*(A)\}. \quad (1.25)$$

One can easily check that $\mathcal{M}(\mu)$ is a σ -algebra that contains \mathcal{F} and all sets of outer measure zero.

Terminology. A measure, μ , defined on a Borel- σ -algebra is sometimes called a *Borel measure*. The measure space $(\Omega, \mathcal{M}(\mu), \mu)$ is called the *completion* of $(\Omega, \mathcal{F}, \mu)$.

It is a nice feature of null-sets that not only can they be added, but they can also be gotten rid off. This is the content of the next lemma.

Lemma 1.2.6 *Let $(\Omega, \mathcal{F}, \mu)$ be a probability space and assume that $G \subset \Omega$ is such that $\mu^*(G) = 1$. Then for any $A \in \mathcal{F}$, $\mu^*(G \cap A) = \mu(A)$ and if $\mathcal{G} \equiv \mathcal{F} \cap G$ (that is the set of all subsets of G of the form $G \cap A$, $A \in \mathcal{F}$), then (G, \mathcal{G}, μ^*) is a probability space.*

Proof. Exercise. □

Lebesgue measure. The prime example of the construction of a measure using Carathéodory's theorem is the Lebesgue measure on \mathbb{R} . Consider the algebra, \mathcal{S} , of all sets that can be written as finite unions of semi-open, disjoint intervals of the form $(a, b]$, and $(a, +\infty)$, $a \in \mathbb{R} \cup \{-\infty\}$, $b \in \mathbb{R}$. Clearly, the function λ , defined by

$$\lambda\left(\bigcup_i (a_i, b_i]\right) = \sum_i (b_i - a_i) \quad (1.26)$$

provides a countably additive set function (this needs a proof!!). Then we know that this can be extended to $\sigma(\mathcal{S}) = \mathcal{B}(\mathbb{R})$; more precisely, one actually constructs a measure on the σ -algebra $\mathcal{M}(\lambda^*)$, and strictly speaking it is this measure on the complete measure space $(\mathbb{R}, \mathcal{M}(\lambda), \lambda)$ that is called the Lebesgue measure.

Of course the same construction can be carried out on any finite non-empty interval, $I \subset \mathbb{R}$; the corresponding measures are finite and thus unique. It is easy to see that λ as a measure on \mathbb{R} is σ -finite and hence also unique.

The construction carries over, with obvious modifications, to \mathbb{R}^d : just replace half-open intervals by half-open rectangles. The key is that we have a natural notion of volume for the elementary objects, and that this provides a σ -additive function on the corresponding algebra.

On topological spaces, one can ask for a number of continuity related properties of measures that occasionally will come very handy.

Definition 1.2.4 Let Ω be a Hausdorff space and $\mathcal{B}(\Omega)$ the corresponding Borel- σ -algebra. A measure, μ , on $(\Omega, \mathcal{B}(\Omega))$, is called Then:

- (0) *Borel measure*, if for any compact set, $C \in \mathcal{F}$, $\mu(C) < \infty$;
- (i) *inner regular* or *tight*, if, for all $B \in \mathcal{F}$, $\mu(B) = \sup_{C \subset B} \mu(C)$, where the supremum is over all compact sets contained in B ;
- (ii) *outer regular*, if for all $B \in \mathcal{F}$, $\mu(B) = \inf_{O \supset B} \mu(O)$, where the infimum is over all open sets containing B .

- (iii) *locally finite*, if any point in Ω contains a neighborhood, U , such that $\mu(U) < \infty$.
- (iv) *Radon measure*, if it is inner regular and locally finite.

A very important result is that on a compact metrisable spaces, all probability measures are inner regular.

Theorem 1.2.7 *Let Ω be a compact metrisable space, and let \mathbb{P} be a probability measure on $(\Omega, \mathcal{B}(\Omega))$. Then \mathbb{P} is inner regular.*

Proof. Let \mathcal{A} be the class of elements, B , of $\mathcal{B}(\Omega)$, such that, for all $\epsilon > 0$, there exists a compact set, $K \subset B$, and an open set, $G \supset B$, such that $\mathbb{P}(B \setminus K) < \epsilon$, and $\mathbb{P}(G \setminus B) < \epsilon$. We want to show that \mathcal{A} is an algebra. First, if $B \in \mathcal{A}$, then its complement, B^c , will also be in \mathcal{A} (for G^c is closed, $B^c \supset G^c$, and $B^c \setminus G^c = G \setminus B$, and vice versa). Next, if $B_1, B_2 \in \mathcal{A}$. Then there are $K_i \subset B_i$ and $G_i \supset B_i$, such that $\mathbb{P}(B_i \setminus K_i) < \epsilon/2$ and $\mathbb{P}(G_i \setminus B_i) < \epsilon/2$. Then $K = K_1 \cup K_2$ and $G = G_1 \cup G_2$ are the desired sets for B . Thus \mathcal{A} is an algebra.

Now let B_n be an increasing sequence of elements of \mathcal{A} such that $\bigcup_{n \in \mathbb{N}} B_n = B$. Choosing set K_n and G_n as before, but with $\epsilon/2$ replaced by $\epsilon 2^{-n-1}$. Then there will be $N < \infty$ such that $\mathbb{P}\left(B \setminus \bigcup_{n=1}^N K_n\right) < \epsilon$. Namely, $\mathbb{P}\left(B \setminus \bigcup_{n \in \mathbb{N}} K_n\right) < \epsilon/2$, while $\mathbb{P}\left(\bigcup_{n \in \mathbb{N}} K_n \setminus \bigcup_{n=1}^N K_n\right) < \epsilon/2$. Therefore, there exists a compact set $K \equiv \bigcup_{n=1}^N K_n$ such that $\mathbb{P}(B \setminus K) < \epsilon$. The same construction works for the corresponding open sets, and so $B \in \mathcal{A}$. Thus \mathcal{A} is a σ -algebra.

Now Ω is metrisable, so there exists a metric, ρ , such that the topology of Ω is equivalent to the metric topology. Then, if K is a closed and thus compact subset of Ω , then K is the intersection of a sequence of open sets $G_n \equiv \{\omega \in \Omega : \rho(\omega, K) < 1/n\}$,

$$K = \bigcap_{n \in \mathbb{N}} G_n \tag{1.27}$$

Since $G_n \downarrow K$, and \mathbb{P} is finite, it follows that $\mathbb{P}[G_n] \downarrow \mathbb{P}[K]$ (recall that $G_n \setminus K \downarrow \emptyset$, and hence $\mathbb{P}[G_n \setminus K] \downarrow 0$, due to the σ -additivity. This means that $K \in \mathcal{A}$. Since \mathcal{A} is a σ -algebra that contains all closed sets, $\mathcal{B}(\Omega)$ contains \mathcal{A} . But since $\mathcal{B}(\Omega)$ is the smallest σ -algebra that contains all closed sets, $\mathcal{B}(\Omega) = \mathcal{A}$.

Now for any $B \in \mathcal{B}(\Omega)$, and $K \subset B$ compact, $\mathbb{P}(B) = \mathbb{P}(K) + \mathbb{P}(B \setminus K)$. But since, for any $B \in \mathcal{A}$, and for any $\epsilon > 0$, by definition, there exists

K such that $\mathbb{P}(B \setminus K) < \epsilon$. Thus $\sup\{\mathbb{P}(K) : K \subset B\} = \mathbb{P}(B)$, so \mathbb{P} is inner regular. \square

Remark 1.2.3 Note that the proof shows that \mathbb{P} is also outer regular. Measures that are both inner and outer regular are sometimes called regular.

1.3 Random variables

Definition 1.3.1 Let (Ω, \mathcal{F}) and (E, \mathcal{G}) be two measurable spaces. A map $f : \Omega \rightarrow E$ is called measurable from (Ω, \mathcal{F}) to (E, \mathcal{G}) , if, for all $A \in \mathcal{G}$, $f^{-1}(A) \equiv \{\omega \in \Omega : f(\omega) \in A\} \in \mathcal{F}$.

The notion of measurability implies that a measurable map is capable of transporting a measure from one space to another. Namely, if $(\Omega, \mathcal{F}, \mathbb{P})$ is a probability space, and f is a measurable map from (Ω, \mathcal{F}) to (E, \mathcal{G}) , then

$$\mathbb{P}_f \equiv \mathbb{P} \circ f^{-1}$$

defines a probability measure on (E, \mathcal{G}) , called the *induced measure*. Namely, for any $B \in \mathcal{G}$, by definition

$$\mathbb{P}_f(B) = \mathbb{P}[f^{-1}(B)]$$

is well defined, since $f^{-1}(B) \in \mathcal{F}$.

The standard notion of a *random variable* refers to a measurable function from some measurable space to the space $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$. We will generally extend this notion and call any measurable map from a measurable space (Ω, \mathcal{F}) to a measurable space $(E, \mathcal{B}(E))$, where E is a topological, respectively metric space, a *E -valued random variable* or a *E -valued Borel function*. Our privileged picture is then that we have an unspecified, so called *abstract probability space* $(\Omega, \mathcal{F}, \mathbb{P})$ on which all kinds of random variables, be it, reals, infinite sequences, functions, or measures, are defined, possibly simultaneously.

An important notion is then that of the σ -algebra generated by random variables.

Definition 1.3.2 Let (Ω, \mathcal{F}) be a measurable space, and let $(E, \mathcal{B}(E))$ be a topological space equipped with its Borel- σ -algebra. Let f be an E -valued random variable. We say that $\sigma(f)$ is the smallest σ -algebra such that f is measurable from $(\Omega, \sigma(f))$ to $(E, \mathcal{B}(E))$.

Note that $\sigma(f)$ depends on the set of values f takes. E.g., if f is real valued, but takes only finitely many values, the σ -algebra generated by f has just finitely many elements. If f is the constant function, then $\sigma(f) = \{\Omega, \emptyset\}$, the trivial σ -algebra. This notion is particularly useful, if several random variables are defined on the same probability space.

Dynkin's lemma has a sometimes useful analogue for so-called monotone classes of functions.

Theorem 1.3.8 *Let \mathcal{H} be a class of bounded functions on Ω to \mathbb{R} . Assume that*

- (i) \mathcal{H} is a vector space over \mathbb{R} ;
- (ii) $1 \in \mathcal{H}$;
- (iii) if $f_n \geq 0$ are in \mathcal{H} , and $f_n \uparrow f$, where f is bounded, then $f \in \mathcal{H}$.

Then, if \mathcal{H} contains the indicator functions of every element of a Π -system \mathcal{S} , then \mathcal{H} contains any bounded $\sigma(\mathcal{S})$ -measurable function.

Proof. Let \mathcal{D} be the class of subsets D of Ω such that $\mathbb{1}_D \in \mathcal{H}$. Then \mathcal{D} is a λ -system. Since by hypothesis it contains \mathcal{S} , by Dynkin's theorem, \mathcal{D} contains the σ -algebra generated by \mathcal{S} . Now let f be a $\sigma(\mathcal{S})$ -measurable function s.t. $0 \leq f \leq K < \infty$. Set

$$D(n, i) \equiv \{\omega \in \Omega : i2^{-n} \leq f(\omega) < (i+1)2^{-n}\}, \quad (1.28)$$

and set

$$f_n(\omega) \equiv \sum_{i=0}^{K2^n} i2^{-n} \mathbb{1}_{D(n, i)}(\omega). \quad (1.29)$$

Every $D(n, i)$ is $\sigma(\mathcal{S})$ -measurable, and so $\mathbb{1}_{D(n, i)} \in \mathcal{H}$, and so by (i), $f_n \in \mathcal{H}$. Since $f_n \uparrow f$, $f \in \mathcal{H}$.

To conclude, we take a general $\sigma(\mathcal{S})$ -measurable function and decompose it into the positive and negative part and treat each part as before. \square

An important property of measurable functions is that the space of measurable functions is closed under limit procedures.

Lemma 1.3.9 *Let f_n , $n \in \mathbb{N}$, be real valued random variables. Then the functions*

$$f^+ \equiv \limsup_{n \uparrow \infty} f_n, \quad \text{and} \quad f^- \equiv \liminf_{n \uparrow \infty} f_n \quad (1.30)$$

are measurable. In particular, if the $f_n \rightarrow f$ pointwise, then f is measurable.

The proof is left as an exercise.

If Ω is a topological space, we have the natural class of continuous functions from Ω to \mathbb{R} . It is easy to see that all continuous functions are measurable if Ω and \mathbb{R} are equipped with their Borel σ -algebras. Thus, all functions that are pointwise limits of continuous functions are measurable, etc..

Remark 1.3.1 Instead of introducing the Borel- σ -algebra, one could go a different path and introduce what is called the *Baire- σ -algebra*. Here one proceeds from the idea that on a topological space one naturally has the notion of continuous functions. One certainly will want all of these to be measurable functions, but certainly one will want more: any pointwise limit of a continuous function should be measurable, as well as limits of sequences of such functions. In this way one arrives at a class of functions, called *Baire-functions*, that is defined as the smallest class of functions that is closed under pointwise limits and that contains the continuous functions. One can then define the *Baire- σ -algebra* as the smallest σ -algebra that makes all Baire-functions measurable. It is in general true that the Borel- σ -algebra contains the Baire- σ -algebra, but in general they are not the same. However, on most spaces we will consider (Polish spaces), the two concepts coincide.

1.4 Integrals

We will now recall the notion of an integral of a measurable function (respectively expectation value of random variables).

To do this one first introduces the notion of *simple functions*:

Definition 1.4.1 A function, $g : \Omega \rightarrow \mathbb{R}$, is called simple if it takes only finitely many values, i.e. if there are numbers, w_1, \dots, w_k , and a partition of Ω , $A_i \in \mathcal{F}$ with $\bigcup_{i=1}^k A_i = \Omega$, such that $A_i = \{\omega \in \Omega : g(\omega) = w_i\}$. Then we can write

$$g(\omega) = \sum_{i=1}^k w_i \mathbb{1}_{A_i}(\omega).$$

The space of simple measurable functions is denoted by \mathcal{E}_+ .

It is obvious what the integral of a simple function should be.

Definition 1.4.2 Let $(\Omega, \mathcal{F}, \mu)$ be a measure space and $g = \sum_{i=1}^k w_i \mathbb{1}_{A_i}$ a simple function. Then

$$\int_{\Omega} g d\mu = \sum_{i=1}^k w_i \mu(A_i). \quad (1.31)$$

The integral of a general measurable function is defined by approximation with simple functions.

Definition 1.4.3(i) Let f be non-negative and measurable. Then

$$\int_{\Omega} f d\mu \equiv \sup_{g \leq f, g \in \mathcal{E}_+} \int_{\Omega} g d\mu \quad (1.32)$$

Note the the value of the integral is in $\mathbb{R} \cup \{+\infty\}$.

(ii) If f is measurable, set

$$f(\omega) = \mathbb{1}_{f(\omega) \geq 0} f(\omega) + \mathbb{1}_{f(\omega) < 0} f(\omega) \equiv f_+(\omega) - f_-(\omega)$$

If either $\int_{\Omega} f_+(\omega) < \infty$ or $-\int_{\Omega} f_-(\omega) d\mu < \infty$, define

$$\int_{\Omega} f d\mu \equiv \int_{\Omega} f_+(\omega) d\mu - \int_{\Omega} f_-(\omega) d\mu \quad (1.33)$$

(iii) We call a function f integrable or *absolutely integrable*, if

$$\int_{\Omega} |f| d\mu < \infty.$$

We state the key properties of the integral without proof.

The most fundamental property is the monotone convergence theorem, which to a large extent justifies the (otherwise strange) definition above.

Theorem 1.4.10 Let $(\Omega, \mathcal{F}, \mu)$ be a measure space and f a real valued non-negative measurable function. Let $f_1 \leq f_2 \leq \dots \leq f_k \leq f_{k+1} \leq \dots \leq f$ be a monotone increasing sequence of non-negative measurable functions that converge pointwise to f . Then

$$\int_{\Omega} f d\mu = \lim_{k \uparrow \infty} \int_{\Omega} f_k d\mu \quad (1.34)$$

The monotone convergence theorem allows to provide an “explicit” construction of the integral as originally used by Lebesgue as a definition.

Lemma 1.4.11 Let f be a non-negative measurable function. Then

$$\int_{\Omega} f d\mu \equiv \lim_{n \uparrow \infty} \left[\sum_{k=0}^{n2^n - 1} 2^{-n} k \mu(\omega : 2^{-n} k \leq f(\omega) < 2^{-n}(k+1)) \right. \\ \left. + n \mu(\omega : f(\omega) \geq n) \right] \quad (1.35)$$

The following lemma is known as *Fatou's lemma*:

Lemma 1.4.12 *Let f_n be a sequence of measurable non-negative functions. Then*

$$\int \liminf_n f_n d\mu \leq \liminf_n \int f_n d\mu. \quad (1.36)$$

Equally central is *Lebesgue's dominated convergence theorem*:

Theorem 1.4.13 *Let f_n be a sequence of absolutely integrable functions, and let f be a measurable function such that*

$$\lim_n f_n(\omega) = f(\omega), \text{ for } \mu\text{-almost all } \omega.$$

Let $g \geq 0$ be a positive function such that $\int g d\mu < \infty$ and

$$|f_n(\omega)| \leq g(\omega), \text{ for } \mu\text{-almost all } \omega.$$

Then f is absolutely integrable with respect to μ and

$$\lim_{n \uparrow \infty} \int f_n d\mu = \int f d\mu. \quad (1.37)$$

In the case when we are dealing with integrals with respect to a probability measure, there exists a very useful improvement of the dominated convergence theorem that leads us to the important notion of *uniform integrability*.

Let us first make the following observation.

Lemma 1.4.14 *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let X be an integrable real valued random variables on this space. Then, for any $\epsilon > 0$, there exists $K < \infty$, such that*

$$\mathbb{E} [|X| \mathbb{1}_{|X| > K}] < \epsilon. \quad (1.38)$$

Proof. This is a direct consequence from the monotone convergence theorem. We leave the details to the reader. \square

When dealing with families of random variables, one problem is that this property will in general not hold uniformly. A nice situation occurs if it does:

Definition 1.4.4 Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. A class, \mathcal{C} , of real valued random variables, X , is called *uniformly integrable*, if, for any $\epsilon > 0$, there exists $K < \infty$, such that, for all $X \in \mathcal{C}$,

$$\mathbb{E} [|X| \mathbb{1}_{|X| > K}] < \epsilon. \quad (1.39)$$

Note that, in particular, if \mathcal{C} is uniformly integrable, then there exists a constant, $C < \infty$, such that, for all $X \in \mathcal{C}$, $\mathbb{E}|X| \leq C$.

Remark 1.4.1 The simplest example of a class of random variables that is not uniformly integrable is given as follows. Take X_n such that

$$\mathbb{P}[X_n = 1] = 1 - 1/n \quad \text{and} \quad \mathbb{P}[X_n = n] = 1/n. \quad (1.40)$$

Clearly, for any K , $\lim_{n \uparrow \infty} \mathbb{E}|X| \mathbb{I}_{|x| > K} = 1$. One should always keep this example in mind when reflecting upon uniform integrability.

Note that on the other hand the class of functions, Y_n , with

$$\mathbb{P}[X_n = 1] = 1 - 1/n \quad \text{and} \quad \mathbb{P}[X_n = \sqrt{n}] = 1/n \quad (1.41)$$

is uniformly integrable.

Theorem 1.4.15 Let X_n , $n \in \mathbb{N}$ and X be integrable random variables on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then $\lim_{n \uparrow \infty} \mathbb{E}|X_n - X| = 0$, if and only if

- (i) $X_n \rightarrow X$ in probability, and
- (ii) the family X_n , $n \in \mathbb{N}$ is uniformly integrable.

Proof. We show the “if” part. Define

$$\phi_K(x) \equiv \begin{cases} K, & \text{if } x > K, \\ x, & \text{if } |x| \leq K, \\ -K, & \text{if } x < -K. \end{cases} \quad (1.42)$$

We have obviously from the uniform integrability that

$$\mathbb{E}|\phi_K(X_n) - X_n| \leq \epsilon, \quad (1.43)$$

for $n \geq 0$ (where for convenience we set $X \equiv X_0$). Moreover, since $|\phi_K(x) - \phi_K(y)| \leq |x - y|$, (i) implies that $\phi_K(X_n) \rightarrow \phi_K(X)$ in probability. Since, moreover, $\phi_K(X_n)$ is bounded, we may choose n_0 such that, for $n \geq n_0$, $\mathbb{P}[|\phi_K(X_n) - \phi_K(X)| > \delta] \leq \epsilon/K$. Then

$$\mathbb{E}|\phi_K(X_n) - \phi_K(X)| \leq \delta + 2\epsilon$$

and so $\lim_{n \uparrow \infty} \mathbb{E}|\phi_K(X_n) - \phi_K(X)| = 0$. In view of the fact that (1.43) holds for any ϵ , it follows that $\mathbb{E}|X_n - X| \rightarrow 0$.

Let us now show the converse (“only”) direction: When $\mathbb{E}|X_n - X| \rightarrow 0$, then by Chebychev’s inequality, $\mathbb{P}[|X_n - X| > \epsilon] \leq \frac{\mathbb{E}|X_n - X|}{\epsilon} \rightarrow 0$, so $X_n \rightarrow X$ in probability. Moreover, X is absolutely integrable.

Now write $X_n = X_n - X + X$ and use that, by the triangle inequality,

$$\mathbb{E}|X_n| \leq \mathbb{E}|X| + \mathbb{E}|X_n - X|. \quad (1.44)$$

Now for any $\epsilon > 0$, there exists n_0 such that, for all $n \geq n_0$, $\mathbb{E}|X_n - X| < \epsilon$. Since all X_i and X are integrable, there exists K such that, for all $n \leq n_0$, $\mathbb{E}|X_n| \mathbb{1}_{|X_n| > K} < \epsilon$. Hence

$$\mathbb{E}|X_n| \mathbb{1}_{|X_n| > 2K} \leq \begin{cases} \epsilon, & \text{if } n \leq n_0 \\ \mathbb{E}|X| \mathbb{1}_{|X_n| > 2K} + \epsilon. \end{cases} \quad (1.45)$$

Finally we use that

$$\begin{aligned} \mathbb{E}|X| \mathbb{1}_{|X_n| > 2K} &\leq \mathbb{E}|X| \mathbb{1}_{|X| > 2K - |X - X_n|} \\ &\leq \mathbb{E}|X| \mathbb{1}_{|X| > K} + \mathbb{E}|X| \mathbb{1}_{|X| \leq K} \mathbb{1}_{|X - X_n| > K} \\ &\leq \epsilon + K \mathbb{P}(|X - X_n| > K) \leq 2\epsilon. \end{aligned} \quad (1.46)$$

This concludes the proof. \square

The importance of this result lies in the fact that in probability theory, we are very often dealing with functions that are not really bounded, and where Lebesgue's theorem is not immediately applicable either. Uniform integrability is the the best possible condition for convergence of the integrals. Note that the simple example (1.41) of a uniformly integrable family given above furnishes a nice example where $\mathbb{E}|X_n - X| \rightarrow 0$, but where Lebesgue's dominated convergence theorem cannot be applied. (Exercise: show this!)

Exercise: Use the previous criterion to prove Lebesgue's dominated convergence theorem in the case of probability measures.

1.5 \mathcal{L}^p and L^p spaces

I will only rather briefly summarize some frequently used notions concerning spaces of integrable functions. Given a measure space, $(\Omega, \mathcal{F}, \mu)$, one defines, for $p \in [1, \infty]$ and measurable functions, f ,

$$\|f\|_{p, \mu} \equiv \|f\|_p \equiv (\mathbb{E}|f|^p)^{1/p} = \left(\int_{\Omega} |f|^p d\mu \right)^{1/p}. \quad (1.47)$$

The set of functions, f , such that $\|f\|_{p, \mu} < \infty$ is denoted by $\mathcal{L}^p(\Omega, \mathcal{F}, \mu) \equiv \mathcal{L}^p$.


There are two crucial inequalities. The Minkowski inequality:

Lemma 1.5.16 For $f, g \in \mathcal{L}^p$,

$$\|f + g\|_p \leq \|f\|_p + \|g\|_p, \quad (1.48)$$

and the Hölder inequality

Fig. 1.1. Convex function



Lemma 1.5.17 For measurable f, g , if $p, q \in [1, \infty]$ are such that $\frac{1}{p} + \frac{1}{q} = 1$, then

$$|\mathbb{E}fg| \leq \|f\|_p \|g\|_q, \quad (1.49)$$

Both inequalities follow from one of the most important inequalities in integration theory, *Jensen's inequality*.

Theorem 1.5.18 Let $(\Omega, \mathcal{F}, \mu)$ be a probability space, let X be an absolutely integrable random variable, and let $G : \mathbb{R} \rightarrow \mathbb{R}$ be a convex function. Then, for any $c \in \mathbb{R}$,

$$\mathbb{E}G(X - \mathbb{E}X + c) \geq G(c), \quad (1.50)$$

and in particular

$$\mathbb{E}G(X) \geq G(\mathbb{E}X). \quad (1.51)$$

Proof. If G is convex, then for any y there is a straight line below G that touches G at $(y, G(y))$, i.e. there exists $m \in \mathbb{R}$ such that $G(x) \geq G(y) + (x - y)m$. Choosing $x = X - \mathbb{E}X + c$ and $y = c$ and taking expectations on both sides yields (1.5.18). \square

Exercise: Prove the Hölder inequalities using Jensen's inequality.

Since Minkowski's inequality is really a triangle inequality and linearity is trivial, we would be inclined to think that $\|\cdot\|_p$ is a norm and \mathcal{L}^p is a normed space. In fact, the only problem is that $\|f\|_p = 0$ does not imply $f = 0$, since f maybe non-zero on sets of μ -measure zero. Therefore to define a normed space, one considers equivalence classes of functions in \mathcal{L}^p by calling two functions, f, f' equivalent, if $f - f'$ is

non-zero only on set of measure zero. The space of these equivalence classes is called $L^p \equiv L^p(\Omega, \mathcal{F}, \mu)$.

The following fact about L^p spaces will be useful to know.

Lemma 1.5.19 *The spaces $L^p(\Omega, \mathcal{F}; \mu)$ are Banach spaces (i.e. complete normed vector space).*

Proof. The by now only non-trivial fact that needs to be proven is the completeness of L^p . Let $f_i \in L^p$, $i \in \mathbb{N}$ be a Cauchy sequence. Then there are $n_k \in \mathbb{N}$, such that, for all $i, j \geq n_k$, $\|f_i - f_j\|_p \leq 2^{-k-k/p}$. Set $g_k \equiv f_{n_k}$, and

$$F \equiv \sum_{k \in \mathbb{N}} 2^{kp} |g_k - g_{k+1}|^p. \quad (1.52)$$

Then

$$\mathbb{E}F = \sum_{k \in \mathbb{N}} 2^{kp} \mathbb{E}|g_k - g_{k+1}|^p = \sum_{k \in \mathbb{N}} 2^{kp} \|g_k - g_{k+1}\|_p^p < \frac{1}{2}. \quad (1.53)$$

Therefore, F is integrable and hence finite except possibly on a set of measure zero. It follows that, for all ω s.t. $F(\omega)$ is finite, $|g_k(\omega) - g_{k+1}(\omega)| \leq 2^{-k} F(\omega)^{1/p}$. It follows further, using telescopic expansion and the triangle inequality, that $g_k(\omega)$ is a Cauchy sequence of real numbers, and hence convergent. Set $f(\omega) = \lim_{k \uparrow \infty} g_k(\omega)$. For the ω in the null-set where $F(x) = +\infty$, we set $f(\omega) = 0$. It follows readily that

$$\mathbb{E}|g_k - f|^p \rightarrow 0, \quad (1.54)$$

and using once more the Cauchy property of f_n , that

$$\mathbb{E}|f_n - f|^p \rightarrow 0. \quad (1.55)$$

□

The case $p = 2$ is particularly nice, in that the space L^2 is not only a Banach space, but a Hilbert space. The point here is that the Hölder inequality, applied for the case $p = 2$, yields

$$\mathbb{E}fg \leq \sqrt{\mathbb{E}f^2 \mathbb{E}g^2} = \|f\|_2 \|g\|_2. \quad (1.56)$$

This means that on L^2 , there exists a quadratic form,

$$(f, g)_\mu \equiv \int fg d\mu \equiv \mathbb{E}fg \quad (1.57)$$

which has the properties of a *scalar product*. The L^2 -norm being the derived norm, $\|f\|_2 = \sqrt{(f, g)_\mu}$. Although somehow L^2 spaces are not the most natural settings for probability, it is sometime quite convenient to exploit this additional structure.

1.6 Fubini's theorem

An always important tool for the computation of integral on product spaces is Fubini's theorem. We consider first the case of non-negative functions.

Theorem 1.6.20 [Fubini-Tonnelli] *Let $(\Omega_1, \mathcal{F}_1, \mu_1)$, and $(\Omega_2, \mathcal{F}_2, \mu_2)$ be two measure spaces, and let f be a real-valued, non-negative measurable function on $(\Omega_1 \times \Omega_2, \mathcal{F}_1 \otimes \mathcal{F}_2)$. Then the two functions*

$$h(x) \equiv \int_{\Omega_2} f(x, y) \mu_2(dy)$$

and

$$g(y) \equiv \int_{\Omega_1} f(x, y) \mu_1(dx)$$

are measurable with respect to \mathcal{F}_1 resp. \mathcal{F}_2 , and

$$\int_{\Omega_1 \times \Omega_2} f d(\mu_1 \otimes \mu_2) = \int_{\Omega_1} h d\mu_1 = \int_{\Omega_2} g d\mu_2 \quad (1.58)$$

Now we turn to the general case.

Theorem 1.6.21 [Fubini-Lebesgue] *Let $f : (\Omega_1 \times \Omega_2, \mathcal{F}_1 \otimes \mathcal{F}_2) \rightarrow (\mathbb{R}, \mathcal{B}(\mathbb{R}))$ be absolutely integrable with respect to the product measure $\mu_1 \otimes \mu_2$. Then*

- (i) *For μ_1 -almost all x , $f(x, y)$ is absolutely integrable with respect to μ_2 , and vice versa.*
- (ii) *The functions, $h(x) \equiv \int_{\Omega_2} f(x, y) \mu_2(dy)$ and $g(y) \equiv \int_{\Omega_1} f(x, y) \mu_1(dx)$, are well-defined except possibly on a set of measure zero with respect to the measures μ_1 , resp. μ_2 , and absolutely integrable with respect to these same measures.*
- (iii) *The equation*

$$\int_{\Omega_1 \times \Omega_2} f d(\mu_1 \otimes \mu_2) = \int_{\Omega_1} h(x) \mu_1(dx) = \int_{\Omega_2} g(y) \mu_2(dy) \quad (1.59)$$

holds.

1.7 Densities, Radon-Nikodým derivatives



In Probability 1 we have encountered the notion of a probability density. In fact, we had constructed the Lebesgue-Stieltjes measure on \mathbb{R} by prescribing a distribution function, F , (i.e. a non-decreasing, right-continuous function) in term of which any interval $(a, b]$ had measure $\mu((a, b]) = F(b) - F(a)$. In the special case when there was a positive function f , such that for all $a < b$, $F(b) - F(a) = \int_a^b f(x)dx$, where dx indicates the standard Lebesgue measure, we called f the density of μ and said that μ is *absolutely continuous* with respect to Lebesgue measure.

We now want to generalise these notions to the general context of positive measures. In particular, we want to be able to say when two measures are absolutely continuous with respect to each other, and define the corresponding relative densities.

First we notice that it is rather easy to modify a given measure μ on a measurable space (Ω, \mathcal{F}) with the help of a measurable function f . To do so, we set, for any $A \in \mathcal{F}$,

$$\mu_f(A) \equiv \int_A f d\mu. \quad (1.60)$$

Exercise: Show that if f is measurable and integrable, but not necessarily non-negative, μ_f , defined as in (1.60), defines an additive set-function. Show that, if $f \geq 0$, μ_f is indeed a measure on (Ω, \mathcal{F}) .

We see that in the case when μ is the Lebesgue measures, μ_f is the absolutely continuous measure with density f . In the general case, we have that, if $\mu(\mathcal{O}) = 0$, then it is also true that $\mu_f(\mathcal{O}) = 0$. The latter property will define the notion of absolute continuity between general measures.


Definition 1.7.1 Let μ, ν be two measures on a measurable space (Ω, \mathcal{F}) . We say that ν is *absolutely continuous* with respect to μ , or $\nu \ll \mu$, if and only if, all μ -null sets, \mathcal{O} (i.e. all sets \mathcal{O} with $\mu(\mathcal{O}) = 0$), are ν -null sets.

We say that two measures, μ, ν , are equivalent if $\mu \ll \nu$ and $\nu \ll \mu$.

We say that a measure ν is *singular* with respect to μ , if there exists a set $\mathcal{O} \in \mathcal{F}$ such that $\mu(\mathcal{O}) = 0$ and $\nu(\mathcal{O}^c) = 0$.

It is important to keep in mind that the notion of absolute continuity is not symmetric.

The following important theorem, called the *Radon-Nikodým theorem*,



asserts that relative absolute continuity is equivalent to the existence of a density.

Theorem 1.7.22 *Let μ, ν be two σ -finite measures on a measurable space (Ω, \mathcal{F}) . Then the following two statements are equivalent:*

- (i) $\nu \ll \mu$.
- (ii) *There exists a non-negative measurable function, f , such that $\nu = \mu_f$.*

Moreover, f is unique up to null sets.

Definition 1.7.2 *If $\nu \ll \mu$, then a positive measurable function f such that $\nu = \mu_f$ is called the *Radon-Nikodým derivative* of ν with respect to μ , denoted*

$$f = \frac{d\nu}{d\mu}. \quad (1.61)$$

Proof. Note that the implication (ii) \Rightarrow (i) is obvious from the definition. The other direction is more tricky.

We consider for simplicity the case when μ, ν are finite measures. The extension to σ -finite measures can then easily be carried through by using suitable partitions of Ω .

We need a few concepts and auxiliary results. The first is the notion of the *essential supremum*.

Definition 1.7.3 *Let $(\Omega, \mathcal{F}, \mu)$ be a measurable space and T an arbitrary non-empty set. The *essential supremum*, $g \equiv \text{esup}_{t \in T} g_t$, of a class, $\{g_t, t \in T\}$, of measurable functions $g_t : \Omega \rightarrow [-\infty, +\infty]$, is defined by the properties*

- (i) g is measurable;
- (ii) $g \geq g_t$, almost everywhere, for each $t \in T$;
- (iii) for any h that satisfies (i) and (ii), $h \geq g$ a.e.

Note that by definition, if there are two g that satisfy this definition, then they are a.e. equal.

The first fact we need to establish is that the essential supremum is always equal to the supremum over a countable set.

Lemma 1.7.23 *Let $(\Omega, \mathcal{F}, \mu)$ be a measure space with μ a σ -finite measure. Let $\{g_t, t \in T\}$ be a non-empty class of real measurable functions. Then there exists a countable subset $T_0 \subset T$, such that*

$$\sup_{t \in T_0} g_t = \text{esup}_{t \in T} g_t. \quad (1.62)$$

Proof. It is enough to consider the case when μ is finite. Moreover, we may restrict ourselves to the case when $|g_t| < C$, for all $t \in T$ (e.g. by passing from g_t to $\tanh^{-1}(g_t)$, which is monotone and preserves all properties of the definition). Let \mathcal{S} denote the class of all countable subsets of T . Set

$$\alpha \equiv \sup_{I \in \mathcal{S}} \mathbb{E} \left(\sup_{t \in I} g_t \right). \quad (1.63)$$

Now let $I_n \in \mathcal{S}$ be a sequence such that

$$\sup_{n \in \mathbb{N}} \mathbb{E} \left(\sup_{t \in I_n} g_t \right) = \alpha, \quad (1.64)$$

and set $T_0 = \bigcup_{n \in \mathbb{N}} I_n$. Of course, T_0 is countable, and $\alpha = \mathbb{E} \left(\sup_{t \in T_0} g_t \right)$. Then, the function $g \equiv \sup_{t \in T_0} g_t$ is measurable, since it is the supremum over a countable set of measurable functions. To see that it also satisfies (ii), assume that there exists $t \in T$, such that $g_t > g$ on a set of positive measure. Then for this t , $\mathbb{E} \max(g, g_t) > \mathbb{E} g = \alpha$. On the other hand, $T_0 \cup \{t\}$ is a countable subset of T , and so by definition of α , $\mathbb{E} \max(g, g_t) \leq \alpha$, which yields a contradiction. Thus (ii) holds. To show (iii), assume that there exists h satisfying (i) and (ii). By (ii), $h \geq g_t$, a.e., for each $t \in T$, and thus also $h \geq \sup_{t \in T_0} g_t$, a.e., since a countable union of null-sets is a null set. Thus g satisfies property (iii), too. Therefore, $g = \text{esup}_{t \in T} g_t$. \square

The notion of essential supremum is used in the next lemma, which is the major step in the proof of the Radon-Nikodým theorem.

Lemma 1.7.24 *Let $(\Omega, \mathcal{F}, \mu)$ be a measure space, with μ a σ -finite measure, and let ν be another σ -finite measure on (Ω, \mathcal{F}) . Let \mathcal{H} be the family of all measurable functions, $h \geq 0$, such that, for all $A \in \mathcal{F}$, $\int_A h d\mu \leq \nu(A)$. Then, for all $A \in \mathcal{F}$,*

$$\nu(A) = \psi(A) + \int_A g d\mu, \quad (1.65)$$

where ψ is a measure that is singular with respect to μ and

$$g = \text{esup}_{h \in \mathcal{H}} h. \quad (1.66)$$

Proof. We again assume μ, ν to be finite, and leave the extension to σ -finite measures as an easy exercise. We also exclude the trivial case of $\mu = 0$. From Lemma 1.7.23 we know that there exists a sequence of functions $h_n \in \mathcal{H}$, such that $g = \sup_{n \in \mathbb{N}} h_n$. Let us first note that if $h_1, h_2 \in \mathcal{H}$, then so is $h \equiv \max(h_1, h_2)$. To see this, note that the disjoint sets

$$A_1 \equiv \{\omega \in A : h_1(\omega) \geq h_2(\omega)\}, \quad A_2 \equiv \{\omega \in A : h_2(\omega) > h_1(\omega)\} \quad (1.67)$$

are measurable and $A_1 \cup A_2 = A$. But

$$\int_A h d\mu = \int_{A_1} h_1 d\mu + \int_{A_2} h_2 d\mu \leq \nu(A_1) + \nu(A_2) = \nu(A), \quad (1.68)$$

which implies $h \in \mathcal{H}$. We may therefore assume the sequence h_n ordered such that $h_n \leq h_{n+1}$, for all $n \geq 1$. Then $g = \lim_{n \uparrow \infty} h_n$, and by monotone convergence, for all $A \in \mathcal{F}$,

$$\int_A g d\mu = \lim_{n \uparrow \infty} \int_A h_n d\mu \leq \nu(A). \quad (1.69)$$

As a consequence, ψ defined by (1.65) satisfies $\psi(A) \geq 0$, for all $A \in \mathcal{F}$. Moreover, trivially $\psi(\emptyset) = 0$, as both ν and $g d\mu$ are measures, ψ defined as their difference is σ -additive. Thus ψ is a measure.

It remains to show that ψ is singular with respect to μ . To this end we construct a set of zero ψ -measure whose complement has zero μ -measure. Of course, this can only be done through a delicate limiting procedure. To begin we define collections of sets that whose ψ -measure is much smaller than their μ -measure. More precisely, for $n \in \mathbb{N}$ and $A \in \mathcal{F}$ with $\mu(A) > 0$, let

$$\mathcal{D}_n(A) \equiv \{B \in \mathcal{F} : B \subset A, \psi(B) < n^{-1}\mu(B)\}. \quad (1.70)$$

The key fact is that for any set A of positive μ measure contains such subsets, i.e. $\mathcal{D}_n(A) \neq \emptyset$ whenever $\mu(A) \neq 0$. This is proven by contradiction: assume that $\mathcal{D}_N(A) = \emptyset$. Then set $h_0 = n^{-1}\mathbb{I}_A$, and then, for all $B \in \mathcal{F}$,

$$\int_B h_0 d\mu = n^{-1}\mu(A \cap B) \leq \psi(A \cap B) \leq \psi(B) = \nu(B) - \int_B g d\mu. \quad (1.71)$$


But then $\int_B (h_0 + g) d\mu \leq \nu(B)$, for all $B \in \mathcal{F}$, so that $g + h_0 \in \mathcal{H}$, which contradicts the fact that $g = \text{esup}_{h \in \mathcal{H}} h$.

Since *any* set of positive μ -measure contains ψ -tiny subsets, one may expect that a set of full μ -measure is ψ -tiny. Below we show this by successively collecting all the μ mass in such sets.

We can now choose $B_{1,n} \in \mathcal{D}_n(\Omega)$ with the property that

$$\mu(B_{1,n}) \geq \frac{1}{2} \sup \{\mu(B) : B \in \mathcal{D}_n(\Omega)\} \equiv \alpha_{1,n}. \quad (1.72)$$

Morally, $B_{1,n}$ is our first attempt to pick up as much μ -mass as we can from the ψ -tiny sets. If we were lucky, and $\mu(B_{1,n}^c) = 0$, then we stop

Fig. 1.2. Construction of the sets $B_{i,n}$

the procedure. Otherwise, we continue by picking up as much mass as we can from what was left, i.e. we choose $B_{2,n} \in \mathcal{D}_n(B_{1,n}^c)$ with

$$\mu(B_{2,n}) \geq \frac{1}{2} \sup \{ \mu(B) : B \in \mathcal{D}_n(B_{1,n}^c) \} \equiv \alpha_{2,n}. \quad (1.73)$$

If $\mu((B_{2,n} \cup B_{1,n})^c) = 0$, we are happy and stop. Otherwise, we continue and choose $B_{3,n} \in \mathcal{D}_n((B_{1,n} \cup B_{2,n})^c)$ with

$$\mu(B_{3,n}) \geq \frac{1}{2} \sup \{ \mu(B) : B \in \mathcal{D}_n(B_{1,n}^c \cap B_{2,n}^c) \} \equiv \alpha_{3,n}, \quad (1.74)$$

and so on. If the process stops at some k_n -th step, set $B_{j,n} = \emptyset$ for $j > k_n$.

It is obvious from the definition that $B_{j,n} \in \mathcal{D}_n(\Omega)$, if $B_{j,n} \neq \emptyset$. Since $\mathcal{D}_n(\Omega)$ is closed under countable disjoint unions (both ψ and μ being measures), also $M_n \equiv \bigcup_{j=1}^{\infty} B_{j,n} \in \mathcal{D}_n(\Omega)$. We want to show that $\mu(M_n^c) = 0$, that is we have picked up all the mass eventually. To do this, note again that, if $\mu(M_n^c) > 0$, then there exists $D \in \mathcal{D}_n(M_n^c)$ with $\mu(D) > 0$.

On the other hand, for any $m \in \mathbb{N}$,

$$\begin{aligned} 2\alpha_{m,n} &= \sup \left\{ \mu(B) : B \in \mathcal{D}_n \left(\bigcap_{j=1}^{m-1} B_{j,n}^c \right) \right\} \\ &\geq \sup \{ \mu(B) : B \in \mathcal{D}_n(M_n^c) \} \geq \mu(D). \end{aligned} \quad (1.75)$$

Thus, if $\mu(D) > 0$, then there exists some $\alpha > 0$, such that $\alpha_{m,n} =$

$\mu(B_{m,n} \geq \alpha)$, for all m . Since all $B_{j,n}$ are disjoint, this would imply that $\mu(M_n) = \infty$, which contradicts the assumption that μ is a finite measure. Thus we conclude that $\mu(M_n^c) = 0$, and so $\psi(M_n) < n^{-1}\mu(M_n) = n^{-1}\mu(\Omega)$. Therefore,

$$\begin{aligned} \psi\left(\liminf_{n \uparrow \infty} M_n\right) &\leq \liminf_{n \uparrow \infty} \psi(M_n) = 0, \\ \mu\left(\left(\liminf_{n \uparrow \infty} M_n\right)^c\right) &= \mu\left(\limsup_{n \uparrow \infty} M_n^c\right) = \lim_{n \uparrow \infty} \mu\left(\bigcup_{k=n}^{\infty} M_k^c\right) = 0. \end{aligned} \quad (1.76)$$

This proves that ψ is singular with respect to μ . \square

As the first consequence of this lemma, we state the famous Lebesgue decomposition theorem.

Theorem 1.7.25 *If μ, ν are σ -finite measures on a measurable space (Ω, \mathcal{F}) , then there exist two uniquely determined measures, λ_c, λ_s , such that $\nu = \lambda_s + \lambda_c$, where λ_c is absolutely continuous with respect to μ and λ_s is singular with respect to μ .*

Proof. Lemma 1.7.24 provides the existence of two measures λ_s and λ_c with the desired properties. To prove the uniqueness of this decomposition, assume that there are λ'_s, λ'_c with the same properties. Due to the definition of singularity, there exists a set \mathcal{O} such that $\mu(\mathcal{O}) = 0$ and $\lambda_s(\mathcal{O}^c) = 0$. Hence λ_c is fully determined if we know it on all sets of μ -measure zero. The same holds for λ'_c . If $\lambda_s \neq \lambda'_s$, then for some $A \in \mathcal{F}$ with $\mu(A) = 0$, $\lambda_s(A) \neq \lambda'_s(A)$. But then it must also be true that $\lambda'_c(A) \neq \lambda_c(A)$. However, this is impossible since these measures are both absolutely continuous w.r.t. μ , and so $\lambda'_c(A) = \lambda_c(A) = 0$! Hence λ_s and λ'_s agree on null sets of μ and thus are identical. \square

The Radon-Nikodým theorem is now immediate: Assume that ν is absolutely continuous with respect to μ . The decomposition (1.65) applied to μ -null sets A then implies that for all these sets, $\psi(A) = 0$. But ψ is singular with respect to μ , so there should be a μ -null set, A , for which $\psi(A^c) = 0$. But since for all A , $\psi(A) = 0$, it follows that $\psi(\Omega) = \psi(A) + \psi(A^c) = 0$, and so ψ is the zero-measure.

All that remains is to assert that the Radon-Nikodým derivative is unique. To do this, assume that there exists another measurable function, g^* , such that

$$\nu(A) = \int_A g^* d\mu. \quad (1.77)$$

Now define the measurable set $A = \{\omega : C > g^* > g > -C\}$. Then, by assumption,

$$\int_A g^* d\mu = \nu(A) = \int_A g d\mu. \quad (1.78)$$

But since on A $g^* > g$, this can only hold if $\mu(A) = 0$, for all $C < \infty$. Thus, $\mu(g^* > g) = 0$. In the same way one shows that $\mu(g^* < g) = 0$, implying that g and g^* differ at most on sets of measure zero. \square

Remark 1.7.1 We have said (and seen in the proof), that the Radon-Nikodým derivative is defined modulo null-sets (w.r.t. μ). This is completely natural. Note that if μ and ν are equivalent, then $0 < \frac{d\nu}{d\mu} < \infty$ almost everywhere, and $\frac{d\nu}{d\mu} = \frac{1}{\frac{d\mu}{d\nu}}$.

The following property of the Radon-Nikodým derivative will be needed later.

Lemma 1.7.26 *Let μ, ν be σ -finite measures on (Ω, \mathcal{F}) , and let $\nu \ll \mu$. If X is \mathcal{F} -measurable and ν -integrable, then, for any $A \in \mathcal{F}$,*

$$\int_A X d\nu = \int_A X \frac{d\nu}{d\mu} d\mu. \quad (1.79)$$

Proof. We may assume that μ is finite and X non-negative. Appealing to the monotone convergence theorem, it is also enough to consider bounded X (otherwise, approximate and pass to the limit on both sides). Let \mathcal{H} be the class of all bounded non-negative \mathcal{F} -measurable functions for which (1.79) is true. Then \mathcal{H} satisfies the hypothesis of Theorem 1.3.8: for, clearly, (i) \mathcal{H} is a vector space, (ii) the function 1 is contained in \mathcal{H} by definition of the Radon-Nikodým derivative, and the property (1.79) is stable under monotone convergence by the monotone convergence theorem. Also, \mathcal{H} contains the indicator functions of all elements of \mathcal{F} . Then the assertion of Theorem 1.3.8 implies that \mathcal{H} contains all bounded \mathcal{F} -measurable function, as claimed. \square

2

Conditional expectations and conditional probabilities

In this chapter we will generalise the notion of conditional expectations and conditional probabilities from elementary probability theory considerably. In elementary probability, we could condition only on *events* of positive probability. This notion is too restrictive, as we have seen in the context of Markov processes, where this limited us to consider discrete state spaces. The new notions we will introduce is conditioning on σ -algebras.

2.1 Conditional expectations

Definition 2.1.1 Consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\mathcal{G} \subset \mathcal{F}$ be sub- σ -algebra of \mathcal{F} . Let X be a random variable, i.e. a \mathcal{F} -measurable (real-valued) function on Ω such that $|\mathbb{E}X| \leq \infty$. We say that a function Y is a conditional expectation of X given \mathcal{G} , written $Y = \mathbb{E}(X|\mathcal{G})$, if

- (i) Y is \mathcal{G} -measurable, and
- (ii) For all $A \in \mathcal{G}$,

$$\mathbb{E}\mathbb{1}_A Y = \mathbb{E}\mathbb{1}_A X. \quad (2.1)$$

Remark 2.1.1 If two functions Y, Y' both satisfy the conditions of a conditional expectation, then they can differ only on sets of probability zero, i.e. $\mathbb{P}[Y = Y'] = 1$. One calls such different realisations of a conditional expectation “versions”.

Remark 2.1.2 The condition $|\mathbb{E}X| \leq \infty$ means that $\mathbb{E}X$ is well-defined, in the sense that $\mathbb{E}X = \mathbb{E}X^+ - \mathbb{E}X^-$ and either $\mathbb{E}X^+ < \infty$ or $\mathbb{E}X^- < \infty$. It is the weakest possible under which a definition of conditional expectation can make sense. Existence of conditional expectations can be established under just this condition (see [4]), however, we will in

the sequel only treat the simple case when X is absolutely integrable, $\mathbb{E}|X| < \infty$.

Intuitively, this notion of conditional expectation can be seen as “integrating” the random variable partially, i.e. with respect to all degrees of freedom that do not affect the σ -algebra \mathcal{G} . A trivial example would be the case where $\Omega = \mathbb{R}^2$, and \mathcal{G} is the σ -algebra of events that depend only on the first coordinate, say x . Then the conditional expectation of a function $f(x, y)$ is just the integral with respect of the variables y (recall the construction of the integral in Fubini’s theorem), modulo re-normalisation. What is left is, of course, a function that depends only on x , and that also satisfies property (ii). The advantage of the notion of a conditional expectation given a σ -algebra is that it largely generalises this concept.

Before we discuss the existence of conditional expectations with respect to σ -algebras, we want to discuss the relation to the more elementary notion of conditional expectations with respect to sets. Recall that if $A \in \mathcal{F}$ has positive mass, $\mathbb{P}(A) > 0$, we can define the conditional expectation, given A , as

$$\mathbb{E}[X|A] = \frac{\int_A X d\mathbb{P}}{\mathbb{P}(A)}. \quad (2.2)$$

Recall that when we were studying Markov chains, we wanted to define conditional expectations of the form

$$\mathbb{E}[f(X_{n+1})|X_n = x]. \quad (2.3)$$

In the case of finite state spaces, we could do this using the definition (2.2), because we could without loss generality assume that $\mathbb{P}[X_n = x]$ was strictly positive. In the case of continuous state space, the canonical situation would be that $\mathbb{P}[X_n = x] = 0$, for any $x \in S$, and so the definition (2.2) is not applicable. It is to overcome this difficulty that we introduce our new notion of conditional expectation given a σ -algebra. Note that, in fact, the objects (2.3) are precisely interpretable as such: Let \mathcal{F} be the product σ -algebra generated by the two random variables X_n, X_{n+1} , and let \mathcal{G} be the sub- σ -algebra generated by X_n . Then, if we consider the expression in (2.3) as a function of x , we see that (in the finite state space case, when it is well-defined) it is a \mathcal{G} -measurable function (to be precise, we should interpret it as a function on Ω that depends on ω through the value of $X_n(\omega)$). Moreover, we see that, by the law of total probability,

$$\mathbb{E}\mathbb{E}[f(X_{n+1})|X_n = \cdot] = \sum_{x \in \mathcal{S}} \mathbb{P}[X_n = x] \mathbb{E}[f(X_{n+1})|X_n = x] = \mathbb{E}f(X_{n+1}), \quad (2.4)$$

so that property (ii) holds (with $X \equiv f(X_{n+1})$).

As we can see, the difficulty associated with constructing conditional expectations in the general case relates to making sense of expressions of the form $0/0$. However, in view of our discussion of the Radon-Nikodým derivative, we can give a precise interpretation of this ratio. Namely, in the discrete case, define the measure P_X on (Ω, \mathcal{G}) through

$$P_X(A) = \int_A X d\mathbb{P}, \quad (2.5)$$

for all $A \in \mathcal{G}$. Then it is clear that P_X is absolutely continuous with respect to the restriction of \mathbb{P} to \mathcal{G} , and the Radon-Nikodým derivative is precisely $\mathbb{E}[X|X_n = \cdot]$, since for all $A \in \mathcal{G}$,

$$P_X(A) = \sum_{x \in A} \mathbb{E}[X|X_n = x] \mathbb{P}[X_n = x]. \quad (2.6)$$

It is thus clear that the construction of the conditional expectation in the general case will be done through our general concept of the Radon-Nikodým derivative.

Theorem 2.1.1 *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, let X be a random variable such that $\mathbb{E}|X| < \infty$, and let $\mathcal{G} \subset \mathcal{F}$ be a sub- σ -algebra of \mathcal{F} . Then*

- (i) *there exists a \mathcal{G} -measurable function, $\mathbb{E}[X|\mathcal{G}]$, unique up to within sets of measure zero, the conditional expectation of X given \mathcal{G} , such that for all $A \in \mathcal{G}$,*

$$\int_A \mathbb{E}[X|\mathcal{G}] d\mathbb{P} = \int_A X d\mathbb{P}. \quad (2.7)$$

- (ii) *If X is absolutely integrable and Z is an absolutely integrable, \mathcal{G} -measurable random variable such that, for some Π -System \mathcal{D} with $\sigma(\mathcal{D}) = \mathcal{G}$,*

$$\mathbb{E}Z = \mathbb{E}X, \quad \text{and } \forall A \in \mathcal{D} \int_A Z d\mathbb{P} = \int_A X d\mathbb{P}, \quad (2.8)$$

then $Z = \mathbb{E}[X|\mathcal{G}]$, almost everywhere.

Proof. We begin by proving (i). Define the set functions $\lambda, \lambda^+, \lambda^-$ as

$$\lambda^\pm(A) \equiv \int_A X^\pm d\mathbb{P}, \quad \lambda \equiv \lambda^+ - \lambda^- \quad (2.9)$$

Now we can consider the restriction of λ to \mathcal{G} , denoted by $\lambda_{\mathcal{G}}$, and the restriction of \mathbb{P} to \mathcal{G} , $\mathbb{P}_{\mathcal{G}}$. Clearly, λ^{\pm} are absolutely continuous with respect to \mathbb{P} , and their restrictions to \mathcal{G} , $\lambda_{\mathcal{G}}^{\pm}$, are absolutely continuous with respect to the restriction of \mathbb{P} to \mathcal{G} , $\mathbb{P}_{\mathcal{G}}$. But since X is assumed to be absolutely integrable with respect to \mathbb{P} and \mathbb{P} is a probability measure, it follows that also $\lambda_{\mathcal{G}}^{\pm}$ are finite measures. Therefore, the Radon-Nikodým theorem 1.7.22 implies that there exist \mathcal{G} -measurable functions, $Y^{\pm} = \frac{d\lambda_{\mathcal{G}}^{\pm}}{d\mathbb{P}_{\mathcal{G}}}$, such that, for all $A \in \mathcal{G}$,

$$\int_A Y^{\pm} d\mathbb{P} = \lambda^{\pm}(A) = \int_A X^{\pm} d\mathbb{P}, \quad (2.10)$$

and hence $Y = \frac{d\lambda_{\mathcal{G}}}{d\mathbb{P}_{\mathcal{G}}} \equiv Y^+ - Y^-$, such that

$$\int_A Y d\mathbb{P} = \lambda(A) = \int_A X d\mathbb{P}. \quad (2.11)$$

Thus, Y has the properties of a conditional expectation and we may set $\mathbb{E}[X|\mathcal{G}] = Y = \frac{d\lambda_{\mathcal{G}}}{d\mathbb{P}_{\mathcal{G}}}$. Note that Y is unique to within sets of measure zero. Finally, to show that the conditional measure is unique in the same sense, assume that there is a function Y' satisfying the conditions of the conditional expectation that differs from Y on a set of positive measure. Then one may set $A^{\pm} = \{\omega : \pm(Y'(\omega) - Y(\omega)) > 0\}$, and at least one of these sets, say A^+ , has positive measure. Then

$$\int_{A^+} X d\mathbb{P} = \int_{A^+} Y' d\mathbb{P} > \int_{A^+} Y d\mathbb{P} = \int_{A^+} X d\mathbb{P}, \quad (2.12)$$

which is impossible. This proves uniqueness and hence (i) is established.

To prove (ii), set

$$\mathcal{A} \equiv \left\{ A \in \mathcal{F} : \int_A Z d\mathbb{P} = \int_A X d\mathbb{P} \right\}. \quad (2.13)$$

then $\Omega \in \mathcal{A}$, and $\mathcal{D} \subset \mathcal{A}$, by assumption. Also, \mathcal{A} is a λ -system, and so by Dynkin's theorem, $\mathcal{A} \supset \sigma(\mathcal{D}) = \mathcal{G}$, and so Z is the desired conditional expectation. \square

In many cases that we will encounter, the σ -algebra, \mathcal{G} , with respect to which we are conditioning is the σ -algebra, $\sigma(Y)$, generated by some other random variable, Y . In that case we will often write

$$\mathbb{E}[X|\sigma(Y)] \equiv \mathbb{E}[X|Y] \quad (2.14)$$

and call this the conditional expectation of X given Y .

2.2 Elementary properties of conditional expectations

Conditional expectations share most of the properties of ordinary expectations. The following is a list of elementary properties:

Lemma 2.2.2 *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let $\mathcal{G} \subset \mathcal{F}$ be a sub- σ -algebra. Then:*

- (i) *If X is \mathcal{G} -measurable, then $\mathbb{E}[X|\mathcal{G}] = X$, a.s.;*
- (ii) *The map $X \rightarrow \mathbb{E}[X|\mathcal{G}]$ is linear;*
- (iii) *$\mathbb{E}\mathbb{E}[X|\mathcal{G}] = \mathbb{E}X$;*
- (iv) *If $\mathcal{B} \subset \mathcal{G}$ is a σ -algebra, then $\mathbb{E}[\mathbb{E}[X|\mathcal{G}]|\mathcal{B}] = \mathbb{E}[X|\mathcal{B}]$, a.s..*
- (v) *$|\mathbb{E}[X|\mathcal{G}]| \leq \mathbb{E}[|X||\mathcal{G}]$, a.s.;*
- (vi) *If $X \leq Y$, then $\mathbb{E}[X|\mathcal{G}] \leq \mathbb{E}[Y|\mathcal{G}]$, a.s.;*

Proof. Left as an **Exercise!** □

The following theorem summarises the most important properties of conditional expectations with regard to limits.

Theorem 2.2.3 *Let $X_n, n \in \mathbb{N}$ and Y be absolutely integrable random variables on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, and let $\mathcal{G} \subset \mathcal{F}$ be a sub- σ -algebra. Then*

- (i) *If $Y \leq X_n \uparrow X$, a.s., then $\mathbb{E}[X_n|\mathcal{G}] \uparrow \mathbb{E}[X|\mathcal{G}]$, a.s..*
- (ii) *If $Y \leq X_n$ a.s., then*

$$\mathbb{E} \left[\liminf_n X_n | \mathcal{G} \right] \leq \liminf_n \mathbb{E} [X_n | \mathcal{G}]. \quad (2.15)$$

- (iii) *If $X_n \rightarrow X$, a.s., and $|X_n| \leq |Y|$, for all n , then $\mathbb{E}[X_n|\mathcal{G}] \rightarrow \mathbb{E}[X|\mathcal{G}]$, a.s..*

Of course, these are just the analogs of the three basic convergence theorems for ordinary expectations. We leave the proofs as exercises.

A useful, but not unexpected, property is the following lemma.

Lemma 2.2.4 *Let X be integrable and let Y be bounded and \mathcal{G} -measurable. Then*

$$\mathbb{E}[XY|\mathcal{G}] = Y\mathbb{E}[X|\mathcal{G}], \text{ a.s.} \quad (2.16)$$

Proof. We may assume that X, Y are non-negative; otherwise decompose them into positive and negative parts and use linearity of the conditional expectation. Moreover, it is enough to consider bounded random variables; otherwise, consider increasing sequences of bounded random

variables that converge to them and use the monotone convergence theorem.

Define, for any $A \in \mathcal{F}$,

$$\nu(A) \equiv \int_A XY d\mathbb{P}, \quad \mu(A) \equiv \int_A X d\mathbb{P}. \quad (2.17)$$

Both μ and ν are finite measures that are absolutely continuous with respect to \mathbb{P} . Then

$$\frac{d\nu_{\mathcal{G}}}{d\mathbb{P}_{\mathcal{G}}} = \mathbb{E}[XY|\mathcal{G}], \quad \frac{d\mu_{\mathcal{G}}}{d\mathbb{P}_{\mathcal{G}}} = \mathbb{E}[X|\mathcal{G}], \quad \frac{d\mu}{d\mathbb{P}} = X. \quad (2.18)$$

Then, using Lemma 1.7.26, for any $A \in \mathcal{G}$,

$$\int_A Y d\mu_{\mathcal{G}} = \int_A Y \frac{d\mu_{\mathcal{G}}}{d\mathbb{P}_{\mathcal{G}}} d\mathbb{P}_{\mathcal{G}} = \int_A Y \mathbb{E}[X|\mathcal{G}] d\mathbb{P}_{\mathcal{G}}, \quad (2.19)$$

whereas for any $A \in \mathcal{F}$,

$$\int_A Y d\mu = \int_A Y \frac{d\mu}{d\mathbb{P}} d\mathbb{P} = \int_A YX d\mathbb{P}. \quad (2.20)$$

Specialising the second equality to the case when $A \in \mathcal{G}$, we find that for those A ,

$$\int_A Y \mathbb{E}[X|\mathcal{G}] d\mathbb{P} = \int_A XY d\mathbb{P}. \quad (2.21)$$

Now $Z \equiv Y\mathbb{E}[X|\mathcal{G}]$ is \mathcal{G} -measurable, and (2.21) is precisely the defining property for Z to be the conditional expectation of XY . This concludes the proof. \square

There should be a natural connection between independence and conditional expectation, as was the case for the elementary notion of conditional expectation. Here it is.

Theorem 2.2.5 *Two σ -algebras, $\mathcal{G}_1, \mathcal{G}_2$, are independent, if and only if, for all \mathcal{G}_2 -measurable integrable random variables, X ,*

$$\mathbb{E}[X|\mathcal{G}_1] = \mathbb{E}X. \quad (2.22)$$

Note that in the theorem we can replace “for all integrable \mathcal{G}_2 measurable random variable” by “for all random variables of the form $X = \mathbb{1}_A$, $A \in \mathcal{G}_2$ ”.

Proof. Assume first that \mathcal{G}_1 and \mathcal{G}_2 are independent. Then, if X_i are \mathcal{G}_i measurable and integrable,

$$\mathbb{E}X_1X_2 = \mathbb{E}X_1\mathbb{E}X_2 = \mathbb{E}[X_1\mathbb{E}[X_2]].$$

Therefore, the constant $\mathbb{E}X_2$ satisfies the condition of the conditional expectation $\mathbb{E}[X_2|\mathcal{G}_1]$, which shows that (2.22) holds.

Now assume that (2.22) holds. Choose $X = \mathbb{1}_A$, $A \in \mathcal{G}_2$. Then

$$\mathbb{E}[\mathbb{1}_A|\mathcal{G}_1] = \mathbb{E}\mathbb{1}_A = \mathbb{P}[A].$$

Then, for all $B \in \mathcal{G}_1$,

$$\mathbb{P}[A \cap B] = \mathbb{E}\mathbb{1}_A\mathbb{1}_B = \mathbb{E}[\mathbb{E}\mathbb{1}_A|\mathcal{G}_1]\mathbb{1}_B = \mathbb{E}[\mathbb{P}[A]\mathbb{1}_B] = \mathbb{P}[A]\mathbb{P}[B].$$

Thus \mathcal{G}_1 and \mathcal{G}_2 are independent. \square

2.3 The case of random variables with absolutely continuous distributions

Let us consider some cases where conditional expectations can be computed more “explicitly”. For this, consider two random variables, X, Y , with values in \mathbb{R}^n and \mathbb{R}^m (in the sequel, nothing but notation changes if we assume $n = m = 1$, so we will do this). We will assume that the joint distribution of X and Y is absolutely continuous with respect to Lebesgue’s measure with density $p(x, y)$. That is, for any function $f : \mathbb{R}^m \times \mathbb{R}^n \rightarrow \mathbb{R}_+$,

$$\mathbb{E}f(X, Y) = \int f(x, y)p(x, y)dx dy.$$

The density of the random variable Y is then

$$q(y) = \int p(x, y)dx.$$

(where we should modify the density to be zero, when $\int p(x, y)dx = \infty$). Let now $h : \mathbb{R}^m \rightarrow \mathbb{R}_+$ be a measurable function. We want to compute $\mathbb{E}[h(X)|Y]$. To do this, take a measurable function $g : \mathbb{R}^n \rightarrow \mathbb{R}_+$. Then

$$\begin{aligned} \mathbb{E}[h(X)g(Y)] &= \int h(x)g(y)p(x, y)dx dy & (2.23) \\ &= \int \left(\int h(x)p(x, y)dx \right) g(y)dy \\ &= \int \left(\frac{\int h(x)p(x, y)dx}{q(y)} \right) g(y)q(y)\mathbb{1}_{q(y)>0}dy \\ &\equiv \int \phi(y)g(y)q(y)\mathbb{1}_{q(y)>0}dy \\ &= \mathbb{E}[\phi(Y)g(y)]. \end{aligned}$$

From this calculation we can derive the following

Proposition 2.3.6 *With the notation above, let $\nu(y, dx)$ be the measure on \mathbb{R}^m defined by*

$$\nu(y, dx) \equiv \begin{cases} \frac{p(x,y)}{q(y)} dx, & \text{if } q(y) > 0, \\ \delta_0(dx), & \text{if } q(y) = 0. \end{cases} \quad (2.24)$$

Then for any measurable function $h : \mathbb{R}^m \rightarrow \mathbb{R}_+$,

$$\mathbb{E}[h(X)|Y] = \int h(x)\nu(Y, dx). \quad (2.25)$$

The function $\frac{p(x,y)}{q(y)}$ as a function of x is called the conditional density given $Y = y$. We see that in this context, we are formally quite close to the discrete case and the intuitive notion of conditional expectations.

2.4 The special case of L^2 -random variables

Conditional expectations have a particularly nice interpretation in the case when the random variable X is square-integrable, i.e. if $X \in L^2(\Omega, \mathcal{F}, \mathbb{P})$ (since for the moment we think of conditional expectations as equivalence classes modulo sets of measure zero, we may consider X as an element of L^2 rather than \mathcal{L}^2). We will identify that space $L^2(\Omega, \mathcal{G}, \mathbb{P})$ as the subspace of $L^2(\Omega, \mathcal{F}, \mathbb{P})$ for which at least one representative of each equivalence class is \mathcal{G} -measurable.

Theorem 2.4.7 *If $X \in L^2(\Omega, \mathcal{F}, \mathbb{P})$, then $\mathbb{E}[X|\mathcal{G}]$ is the orthogonal projection of X on $L^2(\Omega, \mathcal{G}, \mathbb{P})$.*

Proof. The Jensen-inequality applied to the conditional expectation yields that $\mathbb{E}[X^2|\mathcal{G}] \geq \mathbb{E}[X|\mathcal{G}]^2$, and hence $\mathbb{E}\mathbb{E}[X|\mathcal{G}]^2 \leq \mathbb{E}\mathbb{E}[X^2|\mathcal{G}] = \mathbb{E}X^2 < \infty$, so that $\mathbb{E}[X|\mathcal{G}] \in L^2(\Omega, \mathcal{G}, \mathbb{P})$. Moreover, for any bounded, \mathcal{G} -measurable function Z ,

$$\mathbb{E}[Z(X - \mathbb{E}[X|\mathcal{G}])] = \mathbb{E}[XZ] - \mathbb{E}[Z\mathbb{E}[X|\mathcal{G}]] = 0. \quad (2.26)$$

Thus, $X - \mathbb{E}[X|\mathcal{G}]$ is orthogonal to all bounded \mathcal{G} -measurable random variables, and using that these form a dense set in $L^2(\Omega, \mathcal{G}, \mathbb{P})$, it is orthogonal to $L^2(\Omega, \mathcal{G}, \mathbb{P})$. This proves the theorem. \square

Note that this interpretation of the conditional expectation can be used to *define* the conditional expectation for L^2 -random variables.

2.5 Conditional probabilities and conditional probability measures

From conditional expectations we now want to construct conditional probability measures. These seems quite straightforward, but there are some non-trivial technicalities that arise from the version business of conditional expectations.

As before we consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and a sub- σ -algebra \mathcal{G} . For any $A \in \mathcal{F}$, we can define

$$\mathbb{P}[A|\mathcal{G}] \equiv \mathbb{E}[\mathbb{I}_A|\mathcal{G}], \quad (2.27)$$

and call it the *conditional probability of A given \mathcal{G}* . It is a \mathcal{G} -measurable function that satisfies

$$\int_G \mathbb{P}[A|\mathcal{G}] d\mathbb{P} = \int_G \mathbb{I}_A d\mathbb{P} = \mathbb{P}[A \cap G],$$

for any $G \in \mathcal{F}$.

It clearly inherits from the conditional expectation the following properties:

- (i) $0 \leq \mathbb{P}[A|\mathcal{G}] \leq 1$, a.s.;
- (ii) $\mathbb{P}[A|\mathcal{G}] = 0$, a.s., if and only if $\mathbb{P}[A] = 0$; also $\mathbb{P}[A|\mathcal{G}] = 1$, a.s., if and only if $\mathbb{P}[A] = 1$;
- (iii) If $A_n \in \mathcal{F}$, $n \in \mathbb{N}$, are disjoint sets, then

$$\mathbb{P} \left[\bigcup_{n \in \mathbb{N}} A_n | \mathcal{G} \right] = \sum_{n \in \mathbb{N}} \mathbb{P}[A_n | \mathcal{G}], \text{ a.s.}; \quad (2.28)$$

- (iv) If $A_n \in \mathcal{F}$, such that $\lim_{n \uparrow \infty} A_n = A$, then

$$\lim_{n \uparrow \infty} \mathbb{P}[A_n | \mathcal{G}] = \mathbb{P}[A | \mathcal{G}], \text{ a.s.}; \quad (2.29)$$

These observations bring us close to thinking that conditional probabilities can be thought of as \mathcal{G} -measurable functions taking values in the probability measures, at least for almost all ω . The problem, however, is that the requirement of σ -additivity which seems to be satisfied due to (iii) is in fact problematic: (iii) says, that, for any sequence A_n , there exists a set of measure one, such that, for all ω in this set,

$$\mathbb{P} \left[\bigcup_{n \in \mathbb{N}} A_n | \mathcal{G} \right] (\omega) = \sum_{n \in \mathbb{N}} \mathbb{P}[A_n | \mathcal{G}] (\omega). \quad (2.30)$$

However, this set may depend on the sequence, and since that space is not countable, it is unclear whether there exists a set of full measure on which (2.30) holds for all sequences of sets.

These considerations lead to the definition of so-called *regular conditional probabilities*.

Definition 2.5.1 Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space and let \mathcal{G} be a sub- σ -algebra. A *regular conditional probability measure* or *regular conditional probability* on \mathcal{F} given \mathcal{G} is a function, $P(\omega, A)$, defined for all $A \in \mathcal{F}$ and all $\omega \in \Omega$, such that

- (i) for each $\omega \in \Omega$, $P(\omega, \cdot)$ is a probability measure on (Ω, \mathcal{F}) ;
- (ii) for each $A \in \mathcal{F}$, $P(\cdot, A)$ is a \mathcal{G} -measurable function coinciding with the conditional probability $\mathbb{P}[A|\mathcal{G}]$ almost everywhere.

The point is that, if we have a regular conditional probability, then we can express conditional expectations as expectations with respect normal probability measures.

Theorem 2.5.8 *With the notation from above, if $P_\omega[A] \equiv P(\omega, A)$ is a regular conditional probability on \mathcal{F} given \mathcal{G} , then for a \mathcal{F} -measurable integrable random variable, X ,*

$$\mathbb{E}[X|\mathcal{G}](\omega) = \int X dP_\omega \text{ a.s.} \quad (2.31)$$

Proof. As often, we may assume X positive. The prove then goes through the monotone class theorem 1.3.8, quite similar to the proof of theorem 1.7.26. One defines the class of functions where (2.31) holds, verifies that it satisfies the hypothesis of the monotone class theorem and notices that it is true for all indicator functions of sets in \mathcal{F} . \square

The question remains whether and when regular conditional probabilities exist. A central result for us is the existence in the case when Ω is a Polish space.

Theorem 2.5.9 *Let $(\Omega, \mathcal{B}(\Omega), \mathbb{P})$ be a probability space where Ω is a Polish space and $\mathcal{B}(\Omega)$ is the Borel- σ -algebra. Let $\mathcal{G} \subset \mathcal{B}(\Omega)$ be a sub- σ -algebra. Then there exists a regular conditional probability $P(A, \omega)$ given \mathcal{G} .*

We will not give the proof of this theorem here.

3

Stochastic processes

3.1 Definition of stochastic processes

We are finally ready to come to the main topic of this course, stochastic processes.

There are various equivalent ways in which stochastic processes can be defined, and it will be useful to always keep them in mind.

The traditional definition. The standard way to define stochastic processes is as follows. We begin with an abstract probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Next we need a measurable space (S, \mathcal{B}) (which in almost all cases will be a Polish space together with its Borel σ -algebra). The space S is called the *state space*. Next, we need a set I , called the *index set*. Then a stochastic process with state space S and index set I is a collection of (S, \mathcal{B}) -valued *random variables*, $\{X_t, t \in I\}$ defined on $(\Omega, \mathcal{F}, \mathbb{P})$.

We call such a stochastic process also a stochastic process *indexed by* I . The term stochastic process is often reserved to the cases when I is either $\mathbb{N}, \mathbb{Z}, \mathbb{R}_+$, or \mathbb{R} . The index set is then interpreted as a *time parameter*. Depending on whether the index set is discrete or continuous, one refers to stochastic processes with *discrete* or *continuous* time. However, there is also an extensive theory of stochastic processes indexed by more complicated sets, such as $\mathbb{R}^d, \mathbb{Z}^d$, etc.. Often these are also referred to as *stochastic fields*. We will mostly be concerned with the standard case of one-dimensional index sets, but I will give examples of the more general case below.

From the point of view of mappings, we have the picture that for any $t \in I$, there is a measurable map,

$$X_t : \Omega \rightarrow S,$$

whose inverse maps \mathcal{B} into \mathcal{F} .

For this to work, we do want, of course, \mathcal{F} to be so rich that it makes all functions X_t measurable, that is, it should contain the σ -algebra that is generated by the entire family of r.v.'s X_t ,

$$\sigma(X_t; t \in I). \quad (3.1)$$

An example of a stochastic process with discrete time are families of independent random variables.

Sample paths. Given a stochastic process as defined above, we can take a different perspective and view, for each $\omega \in \Omega$, $X(\omega)$ as a map from I to S ,

$$\begin{aligned} X(\omega) : I &\rightarrow S \\ t &\rightarrow X_t(\omega) \end{aligned}$$

We call such a function a *sample path* of X , or a *realisation* of X . Clearly here we want to see the stochastic process as a random variable with values in the space of functions,

$$\begin{aligned} X : \Omega &\rightarrow S^I \\ \omega &\rightarrow X(\omega), \end{aligned}$$

where we view S^I as the space of functions $I \rightarrow S$. To complete this image, we need to endow S with a σ -algebra, \mathcal{B}^I . How should we choose the σ -algebra on S^I ? Our picture will be that X maps (Ω, \mathcal{F}) to (S^I, \mathcal{B}^I) . If this map is measurable, then the marginals $X_t : \Omega \rightarrow S$ should be measurable. This will be the case if the maps X_t , seen as maps from $S^I \rightarrow S$, are measurable from \mathcal{B}^I to \mathcal{B} .

Lemma 3.1.1 *Let \mathcal{B}^I be the smallest σ -algebra that contains all subsets of S^I of the form*

$$C(A, t) \equiv \{x \in S^I : x_t \in A\}. \quad (3.2)$$

with $A \in \mathcal{B}$, $t \in I$. Then \mathcal{B}^I is the smallest σ -algebra such that all the maps $X_t : S^I \rightarrow S$ that map $x \rightarrow x_t$, are measurable, i.e. $\mathcal{B}^I = \sigma(X_t, t \in I)$.

Proof. We first show that all X_t are measurable from

$$\sigma(C(A, t), A \in \mathcal{B}, t \in I) \rightarrow \mathcal{B}.$$

To do this, let $A \in \mathcal{B}$, and chose $t \in I$. Then

$$X_t^{-1}(A) = C(A, t).$$

Thus each X_t is measurable. On the other hand, assume that there is some t and some A such that $C(A, t) \notin \mathcal{B}^I$. Then clearly $X_t^{-1}(A) \notin \mathcal{B}^I$, and then X_t is not measurable! So all $C(A, t)$ must be contained, but none more have to. This proves the lemma. \square

Definition 3.1.1 If $J \subset I$ is finite, and $B \in \mathcal{B}^J$, we call a set

$$C(B, J) \equiv \{x \in S^I : x_J \equiv \{x_t\}_{t \in J} \in B\} \quad (3.3)$$

a *cylinder set* or more precisely *finite dimensional cylinder sets*. If B is of the form $B = \times_{t \in J} A_t$, $A_t \in \mathcal{B}$, we call such a set a *special cylinder*.

It is clear that \mathcal{B}^I contains all finite dimensional cylinder sets, but of course it contains much more. We call \mathcal{B}^I the product σ -algebra, or the algebra generated by the cylinder sets.

It is easy to check that the special cylinders form a π -system, and the cylinders form an algebra; both generate \mathcal{B}^I .

Lemma 3.1.2 *The map $X : \Omega \rightarrow S^I$ is measurable from $\mathcal{F} \rightarrow \mathcal{B}^I$, if and only if, for each t , X_t is measurable from $\mathcal{F} \rightarrow \mathcal{B}$.*

Proof. Since a map, X , is measurable from a σ -algebra $\mathcal{F} \rightarrow \mathcal{B}$, if $X^{-1}(C) \in \mathcal{F}$ for all C in a class that generates \mathcal{B} , to check measurability it is enough to consider C of the form $C(A, t)$. But

$$X^{-1}(C(A, t)) = \{\omega \in \Omega : X_t(\omega) \in A\},$$

which is in \mathcal{F} whenever X_t is measurable. To prove the converse implication is equally trivial. \square

Thus we see that the choice of the σ -algebra \mathcal{B}^I is just the right one to make the two points of view on stochastic processes equivalent from the point of view of measurability.

The law of a stochastic process. Once we view X as a map from Ω to the S -valued functions on I , we can define the probability distribution induced by \mathbb{P} on the space (S^I, \mathcal{B}^I) ,

$$\mu_X \equiv \mathbb{P} \circ X^{-1} \quad (3.4)$$

on (S^I, \mathcal{B}^I) as the distribution of the random variable X .

Canonical process. Given a stochastic process with law μ , one can of course realise this process on the probability space $(S^I, \mathcal{B}^I, \mu)$. In that case the random variable X is the trivial map

$$\begin{aligned} X : S^I &\rightarrow S^I \\ x &\rightarrow X(x) = x. \end{aligned}$$

The viewpoint of the canonical process is, however, not terribly helpful, since more often than not, we want to keep a much richer probability space on which many other random objects can be defined.

3.2 Construction of stochastic processes; Kolmogorov's theorem

The construction of a stochastic process may appear rather formidable, but we may draw encouragement from the fact that we have introduced a rather coarse σ -algebra on the space S^I . The most fundamental observation is that stochastic processes are determined by their observation on just finitely many points in time. We will make this precise now.

For any $J \subset I$, we will denote by π_J the canonical projection from S^I to S^J , i.e. $\pi_J X \in S^J$, such that, for all $t \in J$, $(\pi_J X)_t = X_t$. Naturally, we can define the distributions

$$\mu_X^J \equiv \mathbb{P} \circ (\pi_J X)^{-1}$$


on S^J .

Definition 3.2.1 Let $F(I)$ denote the set of all finite, non-empty subsets of I . Then the collection of probability measures

$$\{\mu_X^J : J \in F(I)\} \tag{3.5}$$

is called the collection of finite dimensional distributions¹ of X .

¹ Alternative appellations are “finite dimensional marginal distributions” or “finite dimensional marginals”.



Note that the finite dimensional distributions determine μ_X on the algebra of finite dimensional cylinder sets. Hence, by Dynkin's theorem, they determine the distribution on the σ -algebra \mathcal{B}^I . This is nice. What is nicer, is that one can also go the other way and *construct* the law of a stochastic process from specified finite dimensional distributions. This will be the content of the fundamental theorem of Daniell and Kolmogorov.

Theorem 3.2.3 *Let S be a compact metrisable space, and let $\mathcal{B} = \mathcal{B}(S)$ be its Borel- σ -algebra. Let I be a set. Suppose that, for each $J \in F(I)$, there exists a probability measure, μ^J , on (S^J, \mathcal{B}^J) , such that for any $J_1 \subset J_2 \in F(I)$,*

$$\mu^{J_1} = \mu^{J_2} \circ \pi_{J_1}^{-1}, \quad (3.6)$$

where π_{J_1} denotes the canonical projection from $S^{J_2} \rightarrow S^{J_1}$. Then there exists a unique measure, μ , on (S^I, \mathcal{B}^I) , such that, for all $J \in F(I)$,

$$\mu \circ \pi_J^{-1} = \mu^J. \quad (3.7)$$

Proof. It will not come as a surprise that we will use Carathéodory's theorem to prove our result. This means, that what we have to do is to construct a σ -additive set function on an algebra that generates the σ -algebra \mathcal{B}^I . Of course, this algebra will be the algebra of all finite-dimensional cylinder events. It is rather easy to see what this set function will have to be. Namely, if B is a finite dimensional cylinder event, then there exists $J \in F(I)$, and $A_J \in \mathcal{B}^J$, such that $B = A_J \times S^{I \setminus J}$ (we call in such a case J the *base* of the cylinder). Then we can set

$$\mu_0(B) = \mu_J(A_J). \quad (3.8)$$

Clearly $\mu_0(\emptyset) = 0$, and μ_0 is finitely additive: if B_1, B_2 are disjoint finite dimensional cylinders with basis J_i , then we can write B_i , $i = 1, 2$, in the form $A_i \times S^{I \setminus J}$, where $J = J_1 \cup J_2$, and $A_i \in \mathcal{B}^J$ are disjoint. Then it is clear that

$$\mu_0(B_1 \cup B_2) = \mu_J(A_1 \cup A_2) + \mu_J(A_1) + \mu(A_2) = \mu_0(B_1) + \mu_0(B_2). \quad (3.9)$$

The usual way to prove σ -additivity is to use the fact that an additive set-function, μ_0 , is σ -additive if and only if for any sequence $G_n \downarrow \emptyset$, $\mu(G_n) \downarrow 0$.

Therefore, the proof will be finished once we establish the following lemma.

Lemma 3.2.4 *Let B_n , $n \in \mathbb{N}$ be a sequence of cylinder sets such that $B_n \supset B_{n+1}$, for all n . If there exists $\epsilon > 0$, such that for all $n \in \mathbb{N}$, $\mu_0(B_n) \geq 2\epsilon$, then $\lim_{n \uparrow \infty} B_n \neq \emptyset$.*

Proof. If B_n satisfies the assumptions of the lemma, then there exists a sequence $J_n \in F(I)$ and $A_n \in \mathcal{B}^{J_n}$, such that $B_n = A_n \times S^{I \setminus J_n}$, and

$$\mu_0(B_n) = \mu_{J_n}(A_n).$$

Since μ_{J_n} is a probability measure on the compact metrisable space S^{J_n} , Theorem 1.2.7 implies that there exists a compact subset, $K_n \subset A_n$, such that

$$\mu_{J_n}(K_n) \geq \mu_{J_n}(A_n) - 2^{-n}\epsilon,$$

or, with $H_n \equiv K_n \times S^{I \setminus J_n}$,

$$\mu_0(H_n) \geq \mu_0(B_n) - 2^{-n}\epsilon. \quad (3.10)$$

Now, under the hypothesis of the lemma, for all $n \in \mathbb{N}$,

$$\mu(H_1 \cap \cdots \cap H_n) \geq \mu(B_1 \cap \cdots \cap B_n) - \sum_{i=1}^n \mu_0(B_n \setminus H_n) \geq 2\epsilon - \epsilon \sum_{i=1}^{\infty} 2^{-i} = \epsilon. \quad (3.11)$$

Hence, in particular, for any n , $H_1 \cap \cdots \cap H_n \neq \emptyset$. Now let $x_n \in H_1 \cap \cdots \cap H_n$, and hence $\pi_{J_k} x_n \in K_1 \cap \cdots \cap K_k$, for any $k \leq n$. By compactness of this set, there exist a subsequence, n_i , such that $\lim_{i \uparrow \infty} \pi_{J_k} x_{n_i} = x \in \bigcap_{j=1}^k K_k$.

Taking subsequently sub-subsequences, we can construct in fact such a sequence in such a way that $\pi_{J_k} x_{n_i} \rightarrow \pi_{J_k} x \in \bigcap_{j=1}^k K_k$ for all k . Then there exist an $x \in \Sigma^I$ whose projections are equal to these limits for all k and hence $x \in \bigcup_{n=1}^k B_k$ for all k , hence $x \in \bigcup_{n=1}^{\infty} B_n$ and so $\bigcap_{n \in \mathbb{N}} B_n \neq \emptyset$. But this is the claim of the lemma. \square

So we are done: μ_0 is σ -additive on the algebra of finite dimensional cylinders, and so there exists a unique probability measure on the σ -algebra \mathcal{B}^J with the advertised properties. \square

Remark 3.2.1 Note that we have used the assumption on the space \mathcal{S} only to ensure that the measures μ^J , for $J \in F(I)$, are all inner regular. Thus we can replace the assertion of the theorem by:

Theorem 3.2.5 *Let S be a topological space space, and let $\mathcal{B} = \mathcal{B}(S)$ be its Borel- σ -algebra. Let I be a set. Suppose that, for each $J \in F(I)$, there exists an inner regular probability measure, μ^J , on (S^J, \mathcal{B}^J) , such that for any $J_1 \subset J_2 \in F(I)$,*

$$\mu^{J_1} = \mu^{J_2} \circ \pi_{J_1}^{-1}, \quad (3.12)$$

where π_{J_1} denotes the canonical projection from $S^{J_2} \rightarrow S^{J_1}$. Then there exists a unique measure, μ , on (S^I, \mathcal{B}^I) , such that, for all $J \in F(I)$,

$$\mu \circ \pi_J^{-1} = \mu^J. \quad (3.13)$$

Finally, one can show that the assumption that Ω be compact and metrisable in Theorem 1.2.7 can be replaced by assuming that Ω be polish: in fact by inspecting the proof one sees that if we replace the requirement “compact” by “closed”, then the compactness requirement on Ω is no longer needed. Thus all what remains to be seen is that the closed sets that approximate B well from within can be chosen bounded on a seperable metric space. But this follows for instance since on a metric space, $\mathbb{P}(B_n(x)) \uparrow 1$, where $B_n(x)$ is the closed metric ball of radius n around x and setting $K_\epsilon = F_{\epsilon/2} \cap B_{n_\epsilon}(0)$ with n_ϵ chosen such that $\mathbb{P}(B_{n_\epsilon}) \geq 1 - \epsilon/2$.

Remark 3.2.2 Note that we have seen no need to distinguish cases according to the nature of the set I .

3.3 Examples of stochastic processes

The Kolmogorov-Daniell theorem goes a long way in helping to construct stochastic processes. However, one should not be deceived: prescribing a *consistent family of finite dimensional distributions* (i.e. distributions satisfying (3.7)) is by no means an easy task and in practise we want to have a simpler way of describing a stochastic process of our liking.

In this section I discuss some of the most important classes of examples without going into too much detail.

3.3.1 Independent random variables

We have of course already encountered independent random variables in the first course of probability. We can now formulate the existence of independent random variables in full generality and with full rigour.

Theorem 3.3.6 *Let I be a set and let, for each $t \in I$, μ_t be a probability measure on $(S, \mathcal{B}(S))$, where S is a polish space. Then there exists a unique probability measure, μ , on (S^I, \mathcal{B}^I) , such that, for $J \in F(I)$, and $A_t \in \mathcal{B}$,*

$$\mu \left(\bigcap_{t \in J} \pi_t^{-1}(A_t) \right) = \prod_{t \in J} \mu_t(A_t). \quad (3.14)$$

Proof. Under the hypothesis that S is polish, the proof is direct from the Kolmogorov-Daniell theorem. Note that this hypothesis is not, however, necessary. \square

Remark 3.3.1 Note that we don't assume I to be countable. However, in the uncountable case, the theorem is rather useless, since hardly anything can be done with a genuine uncountable collection of independent random variables. When we discuss seriously the issue of stochastic processes with continuous time, we will see that we always will want additional properties of sample paths that the theorem above does not provide.

Independent random variables are a major building block for more interesting stochastic processes. We have already encountered sums of independent random variables. Other interesting processes are e.g. maxima of independent random variables: If $X_i, i \in \mathbb{N}$ are independent random variables, define

$$M_n = \max_{1 \leq k \leq n} X_k. \quad (3.15)$$

The study of such maxima is an interesting topic in itself.

Of course one can look at many more functions of independent random variables.

3.3.2 Gaussian processes

Gaussian processes are one of the most important class of stochastic process that can be defined with the help of densities. Let us proceed in two steps.

First, we consider finite dimensional Gaussian vectors. Let $n \in \mathbb{N}$ be fixed, and let C be a positive definite $n \times n$ matrix with real entries. We denote by C^{-1} its inverse. Define the *Gaussian density*,

$$f_C(x_1, \dots, x_n) \equiv \frac{1}{(2\pi)^{n/2} \sqrt{\det C}} \exp \left(-\frac{1}{2} (x, C^{-1}x) \right). \quad (3.16)$$

You see that the necessity of having C positive derives from the fact that we want this density to be integrable with respect to the n -dimensional Lebesgue measure.

Definition 3.3.1 A family of n real random variables is called jointly Gaussian with mean zero and covariance C , if and only if their distribution is absolutely continuous w.r.t. Lebesgue measure on \mathbb{R}^n with density given by f_C .

Remark 3.3.2 In this section I will always consider only Gaussian random variables with mean zero. The corresponding expressions in the general case can be recovered by simple computations.

The definition of Gaussian vectors is no problem. The question is, can we define Gaussian processes? From what we have learned, it will be crucial to be able to define compatible families of finite dimensional distributions.

The following result will be important.

Lemma 3.3.7 Let X_1, \dots, X_n be random variables whose joint distribution is Gaussian with density covariance matrix C and mean zero. Then

(i)

$$\mathbb{E}X_k X_\ell = C_{k\ell}. \quad (3.17)$$

(ii) If $I \subset \{1, \dots, n\}$ with $|I| = m$, then the random variables $X_\ell, \ell \in I$ are jointly Gaussian with covariance given by the $m \times m$ -matrix C^I that with elements $C_{\ell k}^I = C_{\ell k}$, if $\ell, k \in I$.

Proof. For technical reasons it is very convenient to compute the moment generating function, or Laplace transform, of our jointly Gaussian vector. We define, for $u \in \mathbb{C}^n$,

$$\phi_C(u) \equiv \mathbb{E}e^{(u, X)} \equiv \mathbb{E}e^{\sum_{i=1}^n u_i X_i} = \int d^n x f_C(x_1, \dots, x_n) e^{\sum_{i=1}^n u_i x_i}. \quad (3.18)$$

It is easy to see that this integral is always finite. Its computation involves a nice trick, that is well worth remembering! To understand it, recall that a real positive matrix can always be written in the form $C = A^t A$, where A^t denotes the transpose of A , and A is itself invertible. Then likewise $C^{-1} = A^{-1}(A^t)^{-1}$. For simplicity we will write $B = (A^t)^{-1}$.

$$\begin{aligned}
\phi_C(u) &= \frac{1}{(2\pi)^{n/2}\sqrt{\det C}} \int d^n x \exp\left(-\frac{1}{2}(x, C^{-1}x) + (u, x)\right) & (3.19) \\
&= \frac{1}{(2\pi)^{n/2}\sqrt{\det C}} \int d^n x \exp\left(-\frac{1}{2}(Bx, Bx) - (u, x)\right) \\
&= \frac{1}{(2\pi)^{n/2}\sqrt{\det C}} \int d^n x \exp\left(-\frac{1}{2}(Bx - Au, Bx - Au) + \frac{1}{2}(Au, Au)\right) \\
&= \frac{\exp\left(\frac{1}{2}(u, Cu)\right)}{(2\pi)^{n/2}\sqrt{\det C}} \int d^n x \exp\left(-\frac{1}{2}(x - Cu, C^{-1}(x - Cu))\right) \\
&= \exp\left(\frac{1}{2}(u, Cu)\right),
\end{aligned}$$

where in the last line we used that the domain of integration in the integral is invariant under translation (in the case when u is real; if u is complex, we use in addition Cauchy's theorem).

We set that the computation reduces to the completion of the square.

Now it is easy to compute the correlation function. Clearly,

$$\mathbb{E}X_k X_\ell = \frac{d^2}{du_k du_\ell} \phi_C(u=0) = C_{\ell,k}.$$

This establishes (i). (ii) is now quite simple. To compute the Laplace transform of the vector $X_\ell, \ell \in I$, we just need to set $u = u^I$ where $u^I \equiv 0$, if $i \notin I$. The result is precisely the Laplace transform of a Gaussian vector with covariance C^I . Since the Laplace transform determines the distribution uniquely, (ii) follows. \square

This result is very encouraging for the prospect of defining Gaussian vectors. If we can specify an infinite dimensional positive matrix, C then all its finite dimensional sub-matrices, $C^I, I \in F(\mathbb{N})$, are positive and the ensuing family of finite dimensional distributions are Gaussian distributions that do satisfy the consistency requirements of Kolmogorov's theorem! The result is:

Theorem 3.3.8 *Let C be a positive quadratic form on $\mathbb{R}^{\mathbb{N}}$. Then there exists a unique Gaussian process with index set \mathbb{N} , state space \mathbb{R} , such that, for all any finite $J \subset \mathbb{N}$, the marginal distributions are $|J|$ -dimensional Gaussian vectors with covariance C^J .*

Thus the trick is to construct positive quadratic forms. Of course is easy to guess a few by going the other way, and using independent Gaussian random variables as building blocks:

E.g. Let X_n , $n \in \mathbb{N}$ be independent, identically distributed Gaussian random variables with mean zero and variance σ_n^2 . Set $Z_n \equiv \sum_{k=1}^n X_k$. Then

$$\mathbb{E}Z_n Z_m = \sum_{i=1}^n \mathbb{E}X_i \sum_{j=1}^m X_j = \sum_{i=1}^{m \wedge n} \mathbb{E}X_i^2 = \sum_{i=1}^{m \wedge n} \sigma_i^2 \equiv C_{n,m}.$$

Thus the quadratic form $C_{n,m} = \sum_{i=1}^{m \wedge n} \sigma_i^2$ is apparently positive. In fact, if $u_n \in \mathbb{R}^{\mathbb{N}}$,

$$\begin{aligned} (u, Cu) &= \sum_{n,m \in \mathbb{N}} u_n u_m \sum_{i=1}^{m \wedge n} \sigma_i^2 \\ &= \sum_{i \in \mathbb{N}} \sigma_i^2 \sum_{m \geq i} u_m \sum_{n \geq i} u_n \\ &= \sum_{i \in \mathbb{N}} \sigma_i^2 \left(\sum_{m \geq i} u_m \right)^2 \geq 0. \end{aligned}$$

Now we have seen that in the construction of stochastic processes, the fact to have discrete time did not appear (so far) to much of an advantage. Thus the above example may make us courageous to attempt to define a Gaussian process on \mathbb{R}_+ . To this end, define the function

$$C(t, s) \equiv t \wedge s. \quad (3.20)$$

What we have to check is that, for any $J \in F(\mathbb{R}_+)$, the restriction of C to a quadratic form on \mathbb{R}^J is positive. But indeed,

$$\sum_{t,s \in J} u_t u_s (t \wedge s) = \sum_{t,s \in J} u_t u_s \int_0^{(t \wedge s)} dr = \int_0^\infty dr \left(\sum_{t \in J, t \geq r} u_t \right)^2 \geq 0.$$

Thus all finite dimensional distributions exist as Gaussian vectors, and the compatibility conditions are trivially satisfied. Therefore there exists a Gaussian process on \mathbb{R}_+ with this covariance. This process is called “Brownian motion”. Note, however, that this constructs this process only in the product topology, which does not yet yield any nice path properties. We will later see that this process can actually be constructed on the space of continuous functions, and this object will then more properly called Brownian motion.

Exercise. Let $X_k, k \in \mathbb{N}$, be independent Gaussian random variables

with mean zero and variance $\sigma^2 = 1$. Define, for $n \in \mathbb{N}$, and $t \in [0, 1]$,

$$Z_n(t) \equiv \frac{1}{\sqrt{n}} \sum_{k=1}^{\lfloor nt \rfloor} X_k,$$

where $\lfloor \cdot \rfloor$ denotes the largest integer smaller than \cdot . Show that

- (i) $Z_n(t)$ is a stochastic process with indexset $[0, 1]$ and state space \mathbb{R} .
- (ii) Compute the covariance, C_n , of Z_n and show that for any $I \in F([0, 1])$, $C_n^I \rightarrow C^I$, where $C(s, t) = s \wedge t$.
- (iii) Show that the finite dimensional distributions of the processes Z_n converge, as $n \uparrow \infty$, to those of the “Brownian motion” defined above.
- (iv) Show that the results (i)–(iii) remain true if instead of requiring that the X_k are Gaussian we just assume that their variance equals to 1.

Note that to prove (iv), you need to prove the multi-dimensional analogue of the central limit theorem. This requires, however, little more than an adaptation of the notation from the standard CLT in dimension one.

3.3.3 Markov processes

Gaussian processes were built from independent random variables using densities. Another important way to construct non-trivial processes uses conditional probabilities. Markov processes are the most prominent examples. In the case of Markov processes we really think of the index set, \mathbb{N} or \mathbb{R}_+ , as time. The process X_t then shall have two properties: first, it should be *causal*, i.e. we want an expression for the law of X_t given the σ -algebra $\mathcal{F}_{t-1} \equiv \sigma(X_s, s < t)$; second, we want this law to be forgetful of the past: if we know the position (value; we will think mostly of a Markov process as a “particle” moving around in S) of X at some time $s < t$, then the law of X_t should be independent of the positions of $X_{s'}$ with $s' < s$. In a way, Markov processes are meant to be the stochastic analogues of deterministic evolution (differential equations).

To set such a process up, let us consider the (much simpler) case of discrete time, i.e. $I = \mathbb{N}_0$ (we always want zero in our index set as). The main building block for a Markov chain is then the so called (one-step) transition kernel, $\mathcal{P} : \mathbb{N}_0 \times S \times \mathcal{B} \rightarrow [0, 1]$, with the following properties:

- (i) For each $t \in \mathbb{N}_0$ and $x \in S$, $\mathcal{P}_t(x, \cdot)$ is a probability measure on (S, \mathcal{B}) .
- (ii) For each $A \in \mathcal{B}$, and $t \in \mathbb{N}_0$, $\mathcal{P}_t(\cdot, A)$ is a \mathcal{B} -measurable function on S .

Then, a stochastic process X with state space S and index set \mathbb{N}_0 is a discrete time Markov process with law P , if, for all $A \in \mathcal{B}$, $t \in \mathbb{N}$,

$$\mathbb{P}(X_t \in A | \mathcal{F}_{t-1})(\omega) = \mathcal{P}_{t-1}(X_{t-1}(\omega), A), \mathbb{P} - \text{a.s.} \quad (3.21)$$

The remarkable thing is that this requirement fixes the law \mathbb{P} up to one more probability measure on (S, \mathcal{B}) , the so-called *initial distribution*, p_0 .

Theorem 3.3.9 *Let (S, \mathcal{B}) be a Polish space, let \mathcal{P} be a transition kernel and p_0 an probability measure on (S, \mathcal{B}) . Then there exists a unique stochastic process satisfying (3.21) and $\mathbb{P}(X_0 \in A) = p_0(A)$, for all A .*

Proof. In view of the Kolmogorov-Daniell theorem, we have to show that our requirements fix all finite dimensional distributions, and that these satisfy the compatibility conditions. This is more a problem of notation than anything else. We will need to be able to derive formulas for

$$\mathbb{P}(X_{t_n} \in A_n, \dots, X_{t_1} \in A_1).$$

To get started, we consider

$$\mathbb{P}(X_t \in A | \mathcal{F}_s),$$

$s < t$. To do this, we use that by the elementary properties of conditional expectations (we drop the a.s.) that applies to all equations relating to conditional expectations),

$$\begin{aligned} \mathbb{P}(X_t \in A | \mathcal{F}_s) &= \mathbb{E}[\mathbb{P}(X_t \in A | \mathcal{F}_{t-1}) | \mathcal{F}_s] & (3.22) \\ &= \mathbb{E}[\mathcal{P}_{t-1}(X_{t-1}(\omega), A) | \mathcal{F}_s] \\ &= \mathbb{E}\left[\mathbb{E}\left[\int \mathcal{P}_{t-1}(x_{t-1}, A) \mathcal{P}_{t-2}(X_{t-2}(\omega), dx_{t-1}) | \mathcal{F}_{t-2}\right] | \mathcal{F}_s\right] \\ &= \int \mathcal{P}_{t-1}(x_{t-1}, A) \mathcal{P}_{t-2}(x_{t-2}, dx_{t-1}) \dots \\ &\quad \dots \mathcal{P}_{s+1}(x_{s+1}, dx_{s+2}) \mathcal{P}_s(X_s(\omega), dx_{s+1}). \end{aligned}$$

We will set

$$\begin{aligned} P_{s,t}(A, x) &\equiv \int \mathcal{P}_{t-1}(x_{t-1}, A) \mathcal{P}_{t-2}(x_{t-2}, dx_{t-1}) \dots & (3.23) \\ &\quad \dots \mathcal{P}_{s+1}(x_{s+1}, dx_{s+2}) \mathcal{P}_s(x, dx_{s+1}) \end{aligned}$$

and call $P_{s,t}$ the transition kernel from time s to time t . With this object defines, we can now proceed to more complicated expressions:

$$\begin{aligned}
& \mathbb{P}(X_{t_n} \in A_n, \dots, X_{t_1} \in A_1) \\
&= \mathbb{E} \left[\mathbb{P}(X_{t_n} \in A | \mathcal{F}_{t_{n-1}}) \mathbb{1}_{A_{n-1}}(X_{t_{n-1}}) \dots \mathbb{1}_{A_{n-1}}(X_{t_{n-1}}) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[P_{t_{n-1}, t_n}(X_{t_{n-1}}, A_n) | \mathcal{F}_{t_{n-1}} \right] \mathbb{1}_{A_{n-1}}(X_{t_{n-1}}) \dots \mathbb{1}_{A_1}(X_{t_1}) \right] \\
&= \mathbb{E} \left[\mathbb{E} \left[\int_{A_{n-1}} P_{t_{n-1}, t_n}(x_{n-1}, A_n) P_{t_{n-1}, t_{n-2}}(X_{t_{n-2}}(\omega), dx_{n-1}) | \mathcal{F}_{t_{n-2}} \right] \times \right. \\
&\quad \left. \times \mathbb{1}_{A_{n-2}}(X_{t_{n-2}}) \dots \mathbb{1}_{A_1}(X_{t_1}) \right] \\
&= \int_{A_{n-1}} P_{t_{n-1}, t_n}(x_{n-1}, A_n) \int_{A_{n-2}} P_{t_{n-1}, t_{n-2}}(x_{n-2}, dx_{n-1}) \dots \\
&\quad \dots \int_{A_1} P_{t_1, t_2}(x_1, dx_2) \int P_{t_1, 0}(x_0, dx_1) P_0(dx_0). \tag{3.24}
\end{aligned}$$

Thus, we have the desired expression of the marginal distributions in terms of the transition kernel P and the initial distribution P_0 . The compatibility relations follow from the following obvious, but important property of the transition kernels.

Lemma 3.3.10 *The transition kernels $P_{t,s}$ satisfy the Chapman-Kolmogorov equations: For any $t > r > s$,*

$$P_{s,t}(A, x) = \int P_{r,t}(y, A) P_{s,r}(x, dy). \tag{3.25}$$

Proof. This is obvious from the definition. \square

The proof of the compatibility relations is now also obvious; if some of the A_i are equal to S , we can use (3.25) and recover the expressions for the lower dimensional marginals. \square

Exercise. Consider the Brownian motion process from the last subsection. Show that this process is Markov in the sense that all finite dimensional distributions satisfy the Markov property. **Hint:** Let $J = t_n > t_{n-1} > \dots > t_1$. Show that the family of random variables $Y_n \equiv X_{t_n} - X_{t_{n-1}}, X_{t_{n-1}}, X_{t_{n-2}}, \dots, X_{t_1}$ are jointly Gaussian, and that Y_n is independent of the σ -algebra generated by $X_{t_{n-1}}, \dots, X_{t_1}$. To do this compute the covariance matrix of the new family.

3.3.4 Gibbs measures

As an aside, I will briefly explain another important way to construct stochastic processes with the help of conditional expectations and densities, that is central in *statistical mechanics*. It is particularly useful in the setting where I is not an ordered set, the most prominent example is $I = \mathbb{Z}^d$.

In order not to introduce too much notation, I will stick to a simple example, the so-called *Ising-model*. In this case, $S = \{-1, 1\}$. The principle object is family of functions, $H_\Lambda : S^{\mathbb{Z}^d} \rightarrow \mathbb{R}$, called *Hamiltonians*, that are defined for every finite $\Lambda \subset \mathbb{Z}^d$, and are given by

$$H_\Lambda(X) = -\beta \sum_{i,j:i \vee j \in \Lambda} X_i X_j J_{ij}. \quad (3.26)$$

Using this function, we will construct a family of probability kernels, μ_Λ , that have the following properties:

- (i) For each $y \in S^{\mathbb{Z}^d}$, $\mu_\Lambda(\cdot, \tau)$ is a probability measure on $S^{\mathbb{Z}^d}$;
- (ii) For each $A \in \mathcal{B}^{\mathbb{Z}^d}$, $\mu_\Lambda(A, \cdot)$ is a \mathcal{F}_{Λ^c} -measurable function, where $\mathcal{F}_{\Lambda^c} = \sigma(X_i, i \in \Lambda^c)$;
- (iii) for any pair of volumes, Λ, Λ' , with $\Lambda \subset \Lambda'$, and any $A \in \mathcal{B}^{\mathbb{Z}^d}$,

$$\int \mu_\Lambda(z, A) \mu_{\Lambda'}(x, dz) = \mu_{\Lambda'}(x, A). \quad (3.27)$$

We will indeed give an explicit formula for μ_Λ :

$$\mu_\Lambda(A, y) = \frac{\sum_{x_i, i \in \Lambda} \mathbb{1}_{(x_\Lambda, y_{\Lambda^c}) \in A} e^{-H_\Lambda((x_\Lambda, y_{\Lambda^c}))}}{\sum_{x_i, i \in \Lambda} e^{-H_\Lambda((x_\Lambda, y_{\Lambda^c}))}}. \quad (3.28)$$

It is easily checked that this expression indeed defines a kernel with properties (i) and (ii). An expression of this type is called a *local Gibbs specification*.

Now we see that the properties of these kernels are reminiscent of those of regular conditional probabilities.

One defines the notion of a Gibbs measure as follows:

Definition 3.3.2 A probability measure on $S^{\mathbb{Z}^d}$ is called a Gibbs measure, if and only if, for any finite $\Lambda \subset \mathbb{Z}^d$, the kernel μ_Λ is a regular conditional probability for μ given \mathcal{F}_{Λ^c} .

More specifically, if the kernel is the Gibbs specification (3.28), it will be called a Gibbs measure for the d -dimensional Ising model at temperature β^{-1} .

One can prove that such Gibbs measures exist; for this one shows that any accumulation point of a sequence $\mu_{\Lambda_n}(\cdot, x)$, where $\Lambda_n \uparrow \mathbb{Z}^d$ is any

increasing sequence of volumes that converges to \mathbb{Z}^d (in the sense, that, for any finite Λ , there exists n_0 , such that, for all $n \geq n_0$, $\Lambda \subset \Lambda_n$), will be a Gibbs measure. This is relatively straightforward, by writing equation (3.27) for a sequence of volumes $\Lambda_n \uparrow \mathbb{Z}^d$:

$$\int \mu_\Lambda(z, A) \mu_{\Lambda_n}(x, dz) = \mu_{\Lambda_n}(x, A).$$

If μ_{Λ_n} converges weakly to some measure μ , then the right-hand side converges to $\mu(A)$. The left-hand side will converge to $\int \mu_\Lambda(z, A) \mu(x, dz)$, since one can easily see that $\mu_\Lambda(z, A)$ is a continuous function, if A is a cylinder event (in fact, in our example, it is a local function on a discrete space). But then μ satisfies the desired properties of a Gibbs measure.

The existence of accumulation points is then guaranteed by the fact that $S^{\mathbb{Z}^d}$ is compact (Tychonov, since $S = \{-1, 1\}$ is compact), and that the set of probability measures over a compact space is compact. What makes this setting interesting is that there is no general uniqueness result. In fact, if $d \geq 2$, and $\beta > \beta_c$, for a certain β_c , then it is known that there exists more than one Gibbs measure. This mathematical fact is connected to the physical phenomenon of a so-called *phase transition*, and this is what makes the study of Gibbs measures so interesting. For deeper material on Gibbs measures see [3, 12, 6].

Martingales

In this chapter we introduce the fundamental concept of *martingales*, which will keep playing a central rôle in our investigation of stochastic processes. Martingales are “truly random” stochastic processes, in the sense that their observation in the past does not allow for useful prediction of the future. By useful we mean here that no gambling strategies can be devised that would allow for systematic gains. In this chapter we will always assume that random variables take values in \mathbb{R} , unless specified otherwise. The treatment of martingales follows largely the book of Rogers and Williams [11], with the exception of a section on the central limit theorem, which is inspired by [2].



4.1 Definitions

We begin by formally introducing the notion of a filtration of a σ -algebra that we have already briefly encountered in the context of Markov processes. We remain in the context of discrete index sets.

Definition 4.1.1 Let (Ω, \mathcal{F}) be a measurable space. A family of sub- σ -algebras, $\{\mathcal{F}_n, n \in \mathbb{N}_0\}$ of \mathcal{F} that satisfies

$$\mathcal{F}_0 \subset \mathcal{F}_1 \subset \cdots \subset \cdots \mathcal{F}_\infty \equiv \sigma \left(\bigcup_{n \in \mathbb{N}_0} \mathcal{F}_n \right) \subset \mathcal{F}, \quad (4.1)$$

is called a *filtration* of the σ -algebra \mathcal{F} . We call a quadruple $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_n, n \in \mathbb{N}_0\})$ a filtered (probability) space.

In this chapter we will henceforth always assume that we are given a filtered space.

Filtrations and stochastic processes are closely linked. We will see that this goes in two ways.

Definition 4.1.2 A stochastic process, $\{X_n, n \in \mathbb{N}_0\}$, is called *adapted* to the filtration $\{\mathcal{F}_n, n \in \mathbb{N}_0\}$, if, for every n , X_n is \mathcal{F}_n -measurable.

Now the other direction:

Definition 4.1.3 Let $\{X_n, n \in \mathbb{N}_0\}$ be a stochastic process on $(\Omega, \mathcal{F}, \mathbb{P})$. The natural filtration, $\{\mathcal{W}_n, n \in \mathbb{N}_0\}$ with respect to X is the smallest filtration such that X is adapted to it, that is,

$$\mathcal{W}_n = \sigma(X_0, \dots, X_n). \quad (4.2)$$

We see that the basic idea of the natural filtration is that functions of the process that are measurable with respect to \mathcal{W}_n depend only on the observations of the process up to time n .

We now define martingales.

Definition 4.1.4 A stochastic process, X , on a filtered space is called a *martingale*, if and only if the following hold:

- (i) The process X is adapted to the filtration $\{\mathcal{F}_n, n \in \mathbb{N}_0\}$;
- (ii) For all $n \in \mathbb{N}_0$, $\mathbb{E}|X_n| < \infty$;
- (iii) For all $n \in \mathbb{N}$,

$$\mathbb{E}(X_n | \mathcal{F}_{n-1}) = X_{n-1}, \text{ a.s.} \quad (4.3)$$

If (i) and (ii) hold, but instead to (iii), it is only true that $\mathbb{E}(X_n | \mathcal{F}_{n-1}) \geq X_{n-1}$, respectively $\mathbb{E}(X_n | \mathcal{F}_{n-1}) \leq X_{n-1}$, then the process X is called a *sub-martingale*, respectively a *super-martingale*.

It is clear that the property (iii) is what makes martingales special: intuitively, it means that the best guess for what X_n could be, knowing what happened up to time $n - 1$ is simply X_{n-1} . No prediction on the direction of change is possible.

We will now head for the fundamental theorem concerning the impossibility of winning systems in games build on martingales.

To put us into the gambling mood, we think of the increments of the process, $Y_n \equiv X_n - X_{n-1}$, as the result of (not necessarily independent games (Examples: (i) Coin tosses, or (ii) the daily increase of the price of a stock). We are allowed to bet on the outcome in the following way: at each moment in time, $n - 1$, we choose a number $C_n \in \mathbb{R}$. Then our wealth will increase by the amount $C_n Y_n$, i.e. the wealth process, W_n is given by $W_n = \sum_{k=1}^n C_k Y_k$ (Example: (i) in the coin toss, case, chose

$C_n > 0$ means to bet on head ($= \{Y_n = +1\}$) an amount C_n , and $C_n < 0$ means to bet on the outcome tails ($= \{Y_n = -1\}$) the amount $-C_n$; (ii) in the stock case, C_n represents the amount of stock an investor decides to hold at time $n - 1$ up to time n (here negative values can be realised by short-selling).

The choice of the C_n is done knowing the process up to time $n - 1$. This justifies the following definition.

Definition 4.1.5 A stochastic process $\{C_n, n \in \mathbb{N}\}$ is called *previsible*, if, for all $n \in \mathbb{N}$, C_n is \mathcal{F}_{n-1} -measurable.

Given an adapted stochastic process, X and a previsible process C , we can define the wealth process

$$W_n \equiv \sum_{k=1}^n C_k (X_k - X_{k-1}) \equiv (C \bullet X)_n. \quad (4.4)$$

Definition 4.1.6 The process $(C \bullet X)$ is called the *martingale transform* of X by C or the *discrete stochastic integral* of C with respect to X .

Now we can formulate the general “no-system” theorem for martingales:

Theorem 4.1.1(i) Let C be a bounded non-negative previsible process such that there exists $K < \infty$, such that, for all n , and all $\omega \in \Omega$, $|C_n(\omega)| \leq K$. Let X be a supermartingale. Then $C \bullet X$ is a supermartingale that vanishes for $n = 0$.

- (ii) Let C be a bounded previsible process (boundedness as above) and X be a martingale. Then $C \bullet X$ is a martingale that vanishes at zero.
 (iii) Both in (i) and (ii), the condition of boundedness can be replaced by $C_n \in \mathcal{L}^2$, if also $X_n \in \mathcal{L}^2$.

Remark 4.1.1 In terms of gambling, (i) says that, if the underlying process has a tendency to fall, then playing against the trend (“investing in a falling stock”) leads to a wealth process that tends to fall. On the other hand, (ii) says that, if the underlying process X is a martingale, then no matter what strategy you use, the wealth process has mean zero.

Proof. The proofs are excessively simple. We have that $W_n - W_{n-1} = C_n(X_n - X_{n-1})$. Then

$$\mathbb{E}(W_n - W_{n-1} | \mathcal{F}_{n-1}) = C_n \mathbb{E}(X_n - X_{n-1} | \mathcal{F}_{n-1}), \quad (4.5)$$

by Lemma 2.2.4. If X is a martingale, the conditional expectation on the right is zero, so $\mathbb{E}(W_n - W_{n-1} | \mathcal{F}_{n-1}) = 0$, and W is a martingale. If

X is a supermartingale, the conditional expectation is non-positive and this remains true for the product, if C_n is non-negative. This proves (i) and (ii). We leave (iii) as an exercise. \square

The quantities $Y_n = X_n - X_{n-1}$ are called *martingale differences*. A sequence $S_n \equiv \sum_{k=1}^n Y_k$ where $\mathbb{E}(Y_n | \mathcal{F}_{n-1}) = 0$ is called a *martingale difference sequence*. If Y_n are square integrable, then the variance of a martingale difference sequence satisfies

$$\mathbb{E}S_n^2 = \sum_{k=1}^n \mathbb{E}Y_k^2. \quad (4.6)$$

Some examples. A canonical way to construct a martingale is to take any random variable, X , on a filtered probability space, $(\Omega, \mathcal{F}, \mathcal{P})$, and to define

$$X_n \equiv \mathbb{E}[X | \mathcal{F}_n].$$

Then, by the properties of conditional expectation,

$$\mathbb{E}[X_n | \mathcal{F}_{n-1}] = \mathbb{E}[\mathbb{E}[X | \mathcal{F}_n] | \mathcal{F}_{n-1}] = \mathbb{E}[X | \mathcal{F}_{n-1}] = X_{n-1}, \text{ a.s.}$$

In this case, we should expect that $\lim_{n \uparrow \infty} X_n = X$, a.s..

Another example is a Markov chain whose transition kernel has the property that

$$\int xP(y, dx) = y.$$

In particular, sums of iid random variables with mean zero are martingales.

4.2 Upcrossings and convergence

Consider an interval $[a, b]$. We want to count the number of times a process crosses this interval from below. We will say that an upcrossing of $[a, b]$ occurs between times s and t , if $X_s < a, X_t > b$, and $\forall s < r < t, X_r \in [a, b]$. Given N and X , we denote by $U_N(X, [a, b])(\omega)$ the number of upcrossings in the time interval $[0, N]$.

We will now consider a (obviously) previsible process constructed as follows:

$$C_1 = \mathbb{1}_{X_0 < a}; \quad C_n = \mathbb{1}_{C_{n-1}=1} \mathbb{1}_{X_{n-1} \leq b} + \mathbb{1}_{C_{n-1}=0} \mathbb{1}_{X_{n-1} < a}, \text{ for } n \geq 2. \quad (4.7)$$

This process represents a “winning” strategy: wait until the process

(say, price of ...) drops below a . Buy the stock, and hold it until its price exceeds b ; sell, wait until the price drops below a , and so on. Our wealth process is $W = C \bullet X$.

Now each time there is an upcrossing of $[a, b]$ we win at least $(b - a)$. Thus, at time N , we have

$$W_N \geq (b - a)U_N(X, [a, b]) - |X_N - a|\mathbb{1}_{X_N < a}, \quad (4.8)$$

where the last term count is the maximum loss that we could have incurred if we are invested at time N and the price is below a .

Naive intuition would suggest that in the long run, the first term must win. Our theorem above says that this is false, if we are in a fair or disadvantageous game (that is, in practice, always).

Theorem 4.2.2 [Doob's upcrossing lemma] *Let X be a supermartingale and U_N as above. Then*

$$(b - a)\mathbb{E}U_N(X, [a, b]) \leq \mathbb{E}(|X_N - a|\mathbb{1}_{X_N < a}). \quad (4.9)$$

Proof. C is a bounded, non-negative previsible process, so (i) of Theorem 4.1.1 implies that W is supermartingale with $W_0 = 0$. This implies immediately the claim. \square

The result has the following, quite remarkable consequence:

Corollary 4.2.3 *Let X_n be a \mathcal{L}^1 -bounded supermartingale, i.e. such that $\sup_n \mathbb{E}|X_n| < \infty$. Then, for any interval $[a, b]$, define $U_\infty(X, [a, b]) = \lim_{n \uparrow \infty} U_n(X, [a, b])$. Then*

$$(b - a)\mathbb{E}U_\infty(X, [a, b]) \leq a + \sup_n \mathbb{E}|X_n| < \infty. \quad (4.10)$$

In particular, $\mathbb{P}(U_\infty(X, [a, b]) = \infty) = 0$.

Proof. Exercise! \square

This is quite impressive: a (super) martingale that is integrable cannot cross any interval infinitely often. The next result is even more striking, and in fact one of the most important results about martingales.

Theorem 4.2.4 [Doob's supermartingale convergence theorem] *Let X_n be a \mathcal{L}^1 -bounded supermartingale (i.e. $\sup_n \mathbb{E}|X_n| < \infty$). Then, almost surely, $X_\infty \equiv \lim_{n \uparrow \infty} X_n$ exists and is a finite random variable.*

Proof. Define

$$\begin{aligned} \Lambda &\equiv \{\omega : X_n(\omega) \text{ does not converge to a limit in } [-\infty, +\infty]\} & (4.11) \\ &= \{\omega : \limsup X_n(\omega) > \liminf X_n(\omega)\} \\ &= \bigcup_{a < b \in \mathbb{Q}} \{\omega : \limsup X_n(\omega) > b > a \geq \liminf X_n(\omega)\} \equiv \bigcup_{a < b \in \mathbb{Q}} \Lambda_{a,b}. \end{aligned}$$

But:

$$\Lambda_{a,b} \subset \{\omega : U_\infty(X, [a, b])(\omega) = \infty\}. \quad (4.12)$$

Therefore, by Corollary 4.2.3, $\mathbb{P}(\Lambda_{a,b}) = 0$, and thus also $\mathbb{P}(\bigcup_{a < b \in \mathbb{Q}} \Lambda_{a,b}) = 0$, since countable unions of null-sets are null-sets.

Thus the limit of X_n exists in $[-\infty, \infty]$ with probability one. It remains to show that it is finite. To do this, we use Fatou's lemma:

$$\mathbb{E}|X_\infty| = \mathbb{E}[\liminf_n |X_n|] \leq \liminf_n \mathbb{E}|X_n| \leq \sup_n \mathbb{E}|X_n| < \infty. \quad (4.13)$$

So X_∞ is almost surely finite. \square

Doob's convergence theorem implies that positive supermartingale always converge a.s.. This is because the supermartingale property ensures in this case that $\mathbb{E}|X_n| = \mathbb{E}X_n \leq \mathbb{E}X_0$, so the uniform boundedness in \mathcal{L}^1 is always guaranteed.

Our next result gives a sharp criterion for convergence that brings to light the importance of the notion of the uniform integrability.

Theorem 4.2.5 *Let X be a \mathcal{L}^1 -bounded supermartingale, so that, by Theorem 4.2.4 $\lim_n X_n \equiv X_\infty$ exists a.s.. Then $X_n \rightarrow X_\infty$ in \mathcal{L}^1 , if and only if the sequence $\{X_n, n \in \mathbb{N}_0\}$ is uniformly integrable. Then, for $n \in \mathbb{N}_0$,*

$$\mathbb{E}(X_\infty | \mathcal{F}_n) \leq X_n, \text{ a.s.} \quad (4.14)$$

with equality holding if X is a martingale.

Proof. The first statement follows from Theorem 1.4.15. Now for $m \geq n$, and all $F \in \mathcal{F}_n$, $\mathbb{E}(X_m | F) \leq \mathbb{E}(X_n | F)$. Letting m tend to infinity, and using L^1 -convergence, $\lim_{m \uparrow \infty} \mathbb{E}(X_m | F) \geq \mathbb{E}(\lim_{m \uparrow \infty} X_m | F) = \mathbb{E}(X_\infty | F)$, we obtain (4.14). \square

A martingale with the property that there exists integrable X_∞ such that $X_n = \mathbb{E}(X_\infty | \mathcal{F}_n)$ is called a closed martingale. The same applies to super (sub) martingales upon appropriate modification of the equality

relation. The preceding theorem thus says in particular that (sup,super) martingales that are uniformly integrable and converge a.s. are closed.

The martingales of our first example are by definition closed. The next result implies that such martingales converge almost surely and in \mathcal{L}^1 . To show this, we need yet another result of Doob that implies the uniform integrability of conditional expectations.

Theorem 4.2.6 *Let X be an absolutely integrable random variable on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then the family*

$$\{\mathbb{E}(X|\mathcal{G}) : \mathcal{G} \text{ is a sub-}\sigma\text{-algebra of } \mathcal{F}\} \quad (4.15)$$

is uniformly integrable.

Proof. Since X is absolutely integrable, for any $\epsilon > 0$, we can find $\delta > 0$ such that, if $F \in \mathcal{F}$ with $\mathbb{P}(F) < \delta$, then $\mathbb{E}(X\mathbb{1}_F) < \epsilon$. Let such ϵ and δ be given. Choose K such that $K^{-1}\mathbb{E}|X| < \delta$. Let now $\mathcal{G} \subset \mathcal{F}$ be a σ -algebra, and let Y be a version of $\mathbb{E}(X|\mathcal{G})$. Then Jensen's inequality for conditional expectations implies that

$$|Y| \leq \mathbb{E}(|X||\mathcal{G}), \text{ a.s.}$$

Now by Chebeychev, $K\mathbb{P}(|Y| > K) \leq \mathbb{E}|Y| \leq \mathbb{E}|X|$. Thus $\mathbb{P}(|Y| > K) \leq \delta$. Moreover, since the event $\{|Y| > K\} \in \mathcal{G}$, we can argue that

$$\begin{aligned} \mathbb{E}(|Y|\mathbb{1}_{|Y|>K}) &\leq \mathbb{E}(\mathbb{E}(|X||\mathcal{G})\mathbb{1}_{|Y|>K}) = \mathbb{E}(\mathbb{E}(|X|\mathbb{1}_{|Y|>K})|\mathcal{G}) \\ &= \mathbb{E}(|X|\mathbb{1}_{|Y|>K}) < \epsilon. \end{aligned}$$

This is the uniform integrability property we want to prove. \square

Theorem 4.2.7 *Let ξ be an absolutely integrable random variable on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P}, \{\mathcal{F}_n, n \in \mathbb{N}_0\})$. Define $X_n \equiv \mathbb{E}(\xi|\mathcal{F}_n)$, a.s.. Then X_n is a uniformly integrable martingale and*

$$X_n \rightarrow X_\infty = \mathbb{E}(\xi|\mathcal{F}_\infty), \quad (4.16)$$

almost surely and in \mathcal{L}^1 .

Proof. X_n is a \mathcal{L}^1 -bounded martingale by the properties of conditional expectations. The preceding Theorem 4.2.6 implies that X_n is uniformly integrable. Thus X_n converges almost surely and in \mathcal{L}^1 . We have to show the last equality in (4.16). For any n , and any $F \in \mathcal{F}_n$,

$$\mathbb{E}(\mathbb{E}(\xi|\mathcal{F}_\infty)\mathbb{1}_F) = \mathbb{E}(X_n\mathbb{1}_F) = \mathbb{E}(X_\infty\mathbb{1}_F),$$

and so $\mathbb{E}(\mathbb{E}(\xi|\mathcal{F}_\infty)\mathbb{1}_F) = \mathbb{E}(X_\infty\mathbb{1}_F)$ for any F in the π -system $\bigcup_{n \in \mathbb{N}_0} \mathcal{F}_n$

that generates the σ -algebra \mathcal{F}_∞ . The claim follows from the definition of conditional expectation. \square

Note that, when $\mathcal{F} = \mathcal{F}_\infty$, the theorem says that $\mathbb{E}(X|\mathcal{F}_n) \rightarrow \xi$.

An application of this result is Kolmogorov's 0 – 1 law.

Theorem 4.2.8 *Let $X_n, n \in \mathbb{N}$ be a sequence of independent random variables. Define $\mathcal{T}_n = \sigma(X_{n+1}, X_{n+2}, \dots)$ and $\mathcal{T} \equiv \bigcap_{n \in \mathbb{N}} \mathcal{T}_n$. Then, if $F \in \mathcal{T}$, $\mathbb{P}(F) \in \{0, 1\}$.*

Proof. Let $\mathcal{F}_n \equiv \sigma(X_1, \dots, X_n)$. Let $F \in \mathcal{T}$, and $\eta \equiv \mathbb{1}_F$. Since η is bounded and \mathcal{F}_∞ -measurable, the preceding theorem tells us that

$$\eta = \mathbb{E}(\eta|\mathcal{F}_\infty) = \lim_n \mathbb{E}(\eta|\mathcal{F}_n), \text{ a.s.}$$

Now η is \mathcal{T}_n -measurable for each n and hence independent of \mathcal{F}_n . Thus, for any n

$$\mathbb{E}(\eta|\mathcal{F}_n) = \mathbb{E}\eta = \mathbb{P}(F), \text{ a.s.}$$

and so $\eta = \mathbb{P}(F)$, a.s.. But η takes only the values 0 and 1, being an indicator function. Thus $\mathbb{P}(F) \in \{0, 1\}$, proving the theorem. \square

The next theorem relates to filtrations to the infinite past. It is called the Lévy-Doob downward theorem. It is somehow an inverted version of the upward theorem.

Theorem 4.2.9 *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, and let $\{\mathcal{G}_{-n}, n \in \mathbb{N}\}$ be a collection of sub- σ -algebras of \mathcal{F} such that, for all $k, n \in \mathbb{N}$,*

$$\mathcal{G}_{-\infty} \equiv \bigcap_{k \in \mathbb{N}} \mathcal{G}_{-k} \subset \dots \subset \mathcal{G}_{-n-1} \subset \mathcal{G}_{-n} \subset \dots \subset \mathcal{G}_{-1}. \quad (4.17)$$

Let $\{X_{-n}, n \in \mathbb{N}\}$ be a supermartingale relative to $\{\mathcal{G}_{-n}, n \in \mathbb{N}\}$, i.e.

$$\mathbb{E}(X_{-n}|\mathcal{G}_{-m}) \leq X_{-m}, \text{ a.s.}$$

for $m \geq n$. Assume that $\sup_{n \geq 1} \mathbb{E}X_{-n} < \infty$. Then the process X is uniformly integrable and the limit

$$X_{-\infty} = \lim_{n \uparrow \infty} X_{-n}$$

exists a.s. and in \mathcal{L}^1 . Moreover,

$$\mathbb{E}(X_{-n}|\mathcal{G}_{-\infty}) \leq X_{-\infty}, \text{ a.s.}$$

with equality in the martingale case.

Remark 4.2.1 Note that the limit we are considering here is really quite different from the one in the previous convergence theorems. We are really looking backward in time: as n tends to infinity, X_{-n} is measurable with respect to smaller and smaller σ -algebras, contrary to the usual X_n , that depend on more information. Therefore, while a convergent martingale X_n can converge to a constant only if the entire sequence is a constant, but usually is a random variable, a convergent X_{-n} has a much better chance to converge to a real constant. We will see shortly why this can be used to prove things like the strong law of large numbers.

Proof. The nice thing about the upcrossing theorem is that it also provides a proof of the convergence of X_{-n} . In fact, just as before, if $\mathbb{E}|X_{-1}|$ is bounded, it follows that the number of upcrossings of any $[a, b]$ by the process X_{-n} is a.s. finite. Therefore the limit exists in $[-\infty, \infty]$. The finiteness then follows since the condition $\sup_n \mathbb{E}X_{-n} < \infty$, and the supermartingale property imply that $\infty > \mathbb{E}X_{-\infty} \geq \mathbb{E}X_{-1} > -\infty$. This implies that $X_{-\infty}$ is finite almost surely.

Thus we just need to prove uniform integrability to obtain convergence in \mathcal{L}^1 . Now we know that $\mathbb{E}X_{-n}$ is monotone increasing, and $\lim_{n \uparrow \infty} \mathbb{E}X_{-n} < \infty$. Thus, for any $\epsilon > 0$, there is $k \in \mathbb{N}$, such that

$$0 \leq \mathbb{E}X_{-n} - \mathbb{E}X_{-k} \leq \epsilon/2, \quad (4.18)$$

for all $n \geq k$. Now, for such n, k , and $\lambda > 0$,

$$\begin{aligned} \mathbb{E}(|X_{-n}| \mathbb{1}_{|X_{-n}| > \lambda}) &= -\mathbb{E}(X_{-n} \mathbb{1}_{X_{-n} < -\lambda}) + \mathbb{E}X_{-n} - \mathbb{E}(X_{-n} \mathbb{1}_{X_{-n} \leq \lambda}) \\ &\leq -\mathbb{E}(X_{-k} \mathbb{1}_{X_{-n} < -\lambda}) + \mathbb{E}X_{-n} - \mathbb{E}(X_{-k} \mathbb{1}_{X_{-n} \leq \lambda}) \end{aligned}$$

where we used the supermartingale property to replace n by k . Next we can replace $\mathbb{E}X_{-n}$ by $\mathbb{E}X_{-k}$ with an error of at most $\epsilon/2$, after which the right-hand side reproduces the left hand one with n replaced by k in the first place, i.e.

$$\mathbb{E}(|X_{-n}| \mathbb{1}_{|X_{-n}| > \lambda}) \leq \mathbb{E}(|X_{-k}| \mathbb{1}_{|X_{-n}| > \lambda}) + \epsilon/2. \quad (4.19)$$

Since X_{-k} is absolutely integrable, there exists $\delta > 0$ such that for all F with

$$\mathbb{P}(F) < \delta \Rightarrow \mathbb{E}(|X_{-k}| \mathbb{1}_F) < \epsilon/2. \quad (4.20)$$

But $\mathbb{P}(|X_{-n}| > \lambda) \leq \lambda^{-1} \mathbb{E}|X_{-n}|$. To control $\mathbb{E}|X_{-n}|$, set $X^- \equiv \max(-X, 0)$, and write

$$\mathbb{E}|X_{-n}| = \mathbb{E}X_{-n} + 2\mathbb{E}X_{-n}^-.$$

But X^- is a submartingale, and so

$$\mathbb{E}|X_{-n}| \leq \sup_n \mathbb{E}X_{-n} + 2\mathbb{E}X_{-1}. \quad (4.21)$$

Thus we can choose $K < \infty$ such that

$$\begin{aligned} \mathbb{P}(|X_{-n}| > K) &\leq \delta, \text{ if } n \geq k, \\ \mathbb{E}(|X_{-j}| \mathbb{1}_{|X_{-j}| > K}) &< \epsilon, \text{ if } j < k. \end{aligned} \quad (4.22)$$

(for the second we just use the integrability for the finitely many values of i ; for the first we use the uniform bound (4.21)). Then the first inequalities imply that $\mathbb{E}(|X_{-n}| \mathbb{1}_{|X_{-n}| > K}) \leq \epsilon$ for $n \geq k$ via (4.19) and the implication (4.20). This proves the uniform integrability. \square

As an application we give a new proof of Kolmogorov's Law of Large Numbers.

Theorem 4.2.10 *Let $X_n, n \in \mathbb{N}$ be iid random variables with $\mathbb{E}|X_n| < \infty$. Let $\mu = \mathbb{E}X_n$. Set $S_n \equiv \sum_{i=1}^n X_n$. Then*

$$n^{-1}S_n \rightarrow \mu, \quad (4.23)$$

a.s. and in \mathcal{L}^1 .

Proof. Define $\mathcal{G}_{-n} = \sigma(S_n, S_{n+1}, \dots)$. Then, for $n \geq 1$,

$$\mathbb{E}(X_1 | \mathcal{G}_{-n}) = \mathbb{E}(X_2 | \mathcal{G}_{-n}) = \dots = \mathbb{E}(X_n | \mathcal{G}_{-n}) = n^{-1} \mathbb{E}(S_n | \mathcal{G}_{-n}) = n^{-1} S_n, \text{ a.s.} \quad (4.24)$$

The reason for these equalities is simply that knowing something about the sums S_n, S_{n+1} , etc. effects the expectation of the $X_k, k \leq n$ all in the same way: we could simply re-label the first indices without changing anything. Thus $L_{-n} \equiv n^{-1}S_n$ is a martingale as in the preceding theorem, and thus $L \equiv \lim_{n \uparrow \infty} L_{-n}$ exists a.s.

But clearly we also have, for any finite k , that $L = \lim_{n \uparrow \infty} n^{-1}(X_{k+1} + \dots + X_{n+k})$, which means that L is measurable with respect to \mathcal{T}_k , for any k . Now Kolmogorov's zero-one law implies that, for any c , $\mathbb{P}(L \leq c) \in \{0, 1\}$. Since as a function of c this is monotone and right-continuous, there must be exactly one c , such that $\mathbb{P}(L = c) = 1$. Then $\mathbb{E}L = c$. But $\mathbb{E}L_{-n} = \mu$, for all n , so $c = \mu$. \square

The proof above shows some of the power of martingales!

4.3 Inequalities

In this section we derive some fundamental inequalities for martingales. One of the most useful ones is the following *maximum inequality*.

Theorem 4.3.11 *Let Z be a non-negative submartingale. Then, for $c > 0$, and $n \in \mathbb{N}$,*

$$c\mathbb{P}\left(\max_{k \leq n} Z_k \geq c\right) \leq \mathbb{E}\left(Z_n \mathbb{1}_{\max_{k \leq n} Z_k \geq c}\right) \leq \mathbb{E}Z_n. \quad (4.25)$$

Remark 4.3.1 You may recall a similar result for sums of iid random variables as Kolmogorov's inequality. The estimate is extremely powerful, since it gives the same estimate for the probability of the maximum to exceed c as Chebyshev's inequality would give for just the endpoint!

Proof. Define the sequence of disjoint events $F_0 \equiv \{Z_0 \geq c\}$,

$$F_k \equiv \bigcap_{\ell < k} \{Z_\ell < c\} \cap \{Z_k \geq c\} = \{\omega : \min(\ell \leq n : Z_\ell \geq c) = k\}. \quad (4.26)$$

Then

$$F \equiv \{\sup_{k \leq n} Z_k \geq c\} = \bigcup_{k=0}^n F_k. \quad (4.27)$$

Clearly, the events $F_k \in \mathcal{F}_k$. Moreover, on F_k ; we know that $Z_k > c$. Thus

$$\mathbb{E}(Z_n \mathbb{1}_{F_k}) \geq \mathbb{E}(Z_k \mathbb{1}_{F_k}) \geq c\mathbb{P}(F_k). \quad (4.28)$$

Here the first inequality used of course the submartingale property of Z . Thus

$$\mathbb{E}(Z_n \mathbb{1}_F) = \sum_{k=0}^n \mathbb{E}(Z_n \mathbb{1}_{F_k}) \geq c \sum_{k=0}^n \mathbb{P}(F_k) = c\mathbb{P}(F). \quad (4.29)$$

This implies the assertion of the theorem. \square

Note that if M_n is a martingale and f a convex function, then $f(M_n)$ is a submartingale. This allows to obtain many useful inequalities from the one of Theorem 4.3.11! In particular, Kolmogorov's inequality follows by choosing $f(X) = X^2$. Other useful choices are the exponential function.

Our next target is Doob's \mathcal{L}^p inequality. The next lemma is a first step in this direction.

Lemma 4.3.12 *Let X and Y be non-negative random variables such that, for all $c > 0$,*

$$c\mathbb{P}(X \geq c) \leq \mathbb{E}(Y \mathbb{1}_{X \geq c}). \quad (4.30)$$

Then, for $p > 1$ and $q^{-1} = 1 - p^{-1}$,

$$\|X\|_p \leq q\|Y\|_p. \quad (4.31)$$

Proof. By our hypothesis, it holds that

$$L \equiv \int_0^\infty pc^{p-1}\mathbb{P}(X \geq c)dc \leq \int_0^\infty pc^{p-2}\mathbb{E}(Y \mathbb{1}_{X \geq c})dc \equiv R.$$

Using Fubini's theorem for non-negative integrands, we can write

$$\begin{aligned} L &= \int_0^\infty pc^{p-1} \left(\int_\Omega \mathbb{1}_{X \geq c}(\omega) \mathbb{P}(d\omega) \right) dc \\ &= \int_\Omega \left(\int_0^{X(\omega)} pc^{p-1} dc \right) \mathbb{P}(d\omega) = \int_\Omega X(\omega)^p \mathbb{P}(d\omega) = \mathbb{E}X^p. \end{aligned}$$

Starting from the right-hand side, we can perform the same calculation, and derive that

$$R = q\mathbb{E}(X^{p-1}Y) \leq q\|Y\|_p\|X^{p-1}\|_q,$$

where the second inequality is just Hölder's inequality. Then

$$\mathbb{E}X^p \leq q\|Y\|_p\|X^{p-1}\|_q. \quad (4.32)$$

Assume that $\|Y\|_p$ and $\|X\|_q$ are finite. Clearly, $(p-1)q = p$, and so

$$\|X^{p-1}\|_q = \left(\mathbb{E}X^{q(p-1)} \right)^{1/q} = (\mathbb{E}X^p)^{1/q}.$$

Therefore (4.32) reads

$$\|X\|_p^p \leq q\|Y\|_p\|X\|_p^{p/q},$$

or $\|X\|_p \leq q\|Y\|_p$, as claimed. If $\|X\|_p = \infty$, one derives the inequality first for $X \wedge n$, and then uses monotone convergence. This proves the lemma. \square

We can now formulate Doob's \mathcal{L}^p -inequality.

Theorem 4.3.13 *Let $p > 1$ and $q^{-1} = 1 - p^{-1}$. Let Z be a non-negative submartingale bounded in \mathcal{L}^p , and define*

$$Z^* \equiv \sup_{k \in \mathbb{N}_0} Z_k. \quad (4.33)$$

Then $Z^ \in \mathcal{L}^p$, and*

$$\|Z^*\|_p \leq q \sup_{n \in \mathbb{N}_0} \|Z_n\|_p. \quad (4.34)$$

The limit, $Z_\infty \equiv \lim_{n \uparrow \infty} Z_n$, exists a.s. and in \mathcal{L}^p , and

$$\|Z_\infty\|_p \leq \sup_{n \in \mathbb{N}_0} \|Z_n\|_p = \lim_{n \uparrow \infty} \|Z_n\|_p. \quad (4.35)$$

If Z is of the form $Z = |M|$, where M is a martingale bounded in \mathcal{L}^p , then $M_\infty \equiv \lim_{n \uparrow \infty} M_n$ exists a.s. and in \mathcal{L}^p , and $Z_\infty = |M_\infty|$, a.s..

Proof. Define $Z_n^* \equiv \sup_{k \leq n} Z_k$. Theorem 4.3.11 and Lemma 4.3.12 imply that

$$\|Z_n^*\|_p \leq q \|Z_n\|_p \leq q \sup_{k \leq n} \|Z_k\|_p.$$

Using the monotone convergence theorem, we get (4.34). Now $-Z$ is a supermartingale bounded in \mathcal{L}^p , and hence in \mathcal{L}^1 , it follows that Z_∞ exists a.s.. But

$$|Z_\infty - Z_n|^p \leq (\max(Z_\infty, Z_n))^p \leq (Z_n^*)^p \in \mathcal{L}^1,$$

so that, by Lebesgue's dominated convergence theorem, $\mathbb{E}|Z_\infty - Z_n|^p \rightarrow 0$, i.e. $Z_n \rightarrow Z_\infty$ in \mathcal{L}^p .

The last assertion in (4.35) follows since by Jensen's inequality and the submartingale property

$$\mathbb{E}Z_n^p \geq \mathbb{E}(\mathbb{E}(Z_n^p | \mathcal{F}_{n-1})) \geq \mathbb{E}(\mathbb{E}(Z_n | \mathcal{F}_{n-1})^p) \geq \mathbb{E}Z_{n-1}^p,$$

and so $\|Z_n\|_p$ is a non-decreasing sequence. The remaining assertions are straightforward. \square

4.4 Doob decomposition

One of the games when dealing with stochastic processes is to “extract the martingale part”. There are several such decompositions, but the following *Doob decomposition* is very important and its continuous time analogue will be fundamental for the theory of stochastic integration.

Theorem 4.4.14(i) *Let $\{X_n, n \in \mathbb{N}_0\}$ be an adapted process on a filtered space $(\Omega, \mathcal{F}, \mathbb{P}; \{\mathcal{F}_n, n \in \mathbb{N}_0\})$ with $X_n \in \mathcal{L}^1$ for all n . Then X can be written in the form¹*

¹ To make sure that there is no confusion about notation: the following equation is to be understood in the sense that $X_0 = X_0$, and for $n \geq 1$, $X_n = X_0 + M_n + A_n$.

$$X = X_0 + M + A, \quad (4.36)$$

where M is a martingale with $M_0 = 0$ and A is a previsible process with $A_0 = 0$. This decomposition is unique modulo indistinguishability, i.e. if for some other M', A' , $X = X_0 + M' + A'$, then

$$\mathbb{P}(M_n = M'_n, A_n = A'_n, \forall n \in \mathbb{N}) = 1.$$

(ii) The process X is a submartingale, if and only if A is an increasing process in the sense that

$$\mathbb{P}(A_n \leq A_{n+1}, \forall n \in \mathbb{N}) = 1.$$

Proof. The proof is unsurprisingly very easy. All we need to do is to derive explicit formulae for M and A . Now assume that a decomposition of the claimed form exists. Then

$$\begin{aligned} \mathbb{E}((X_n - X_{n-1})|\mathcal{F}_{n-1}) &= \mathbb{E}((M_n - M_{n-1})|\mathcal{F}_{n-1}) + \mathbb{E}((A_n - A_{n-1})|\mathcal{F}_{n-1}) \\ &= 0 + A_n - A_{n-1} \end{aligned} \quad (4.37)$$

by the martingale and predictability properties. Therefore

$$A_n = \sum_{k=1}^n \mathbb{E}((X_k - X_{k-1})|\mathcal{F}_{k-1}), \text{ a.s.} \quad (4.38)$$

So now just *define* A_n by (4.38), and M_n by $M_n \equiv X_n - X_0 - A_n$. Clearly M is then a martingale, and A is by construction predictable. This ends the proof of (i). The assertion of (ii) is obvious from (4.37). \square

An immediate application of the decomposition theorem is a maximum inequality without positivity assumption.

Lemma 4.4.15 *If X is either a submartingale or a supermartingale then, for $n \in \mathbb{N}$ and $c > 0$,*

$$c\mathbb{P}\left(\sup_{k \leq n} |X_k| \geq 3c\right) \leq 4\mathbb{E}|X_0| + 3\mathbb{E}|X_n|. \quad (4.39)$$

Proof. We consider the case when X is a submartingale, the case of the supermartingale is identical by passing to $-X$. Then there is a Doob decomposition

$$X = X_0 + M + A$$

with A an increasing process. Thus

$$\sup_{k \leq n} |X_k| \leq |X_0| + \sup_{k \leq n} |M_k| + \sup_{k \leq n} |A_k| = |X_0| + \sup_{k \leq n} |M_k| + A_n.$$

Note that $|M|$ is a non-negative submartingale, so for the supremum of $|M_k|$ we can use Theorem (4.3.11). We use the simple observation that, if $x + y + z > 3c$, then at least one of the x, y, z must exceed c . Thus,

$$\begin{aligned} c\mathbb{P}\left(\sup_{k \leq n} |X_k| \geq 3c\right) &\leq c\mathbb{P}(|X_0| \geq c) + c\mathbb{P}\left(\sup_{k \leq n} |M_k| \geq c\right) + c\mathbb{P}(A_n \geq c) \\ &\leq \mathbb{E}|X_0| + \mathbb{E}|M_n| + \mathbb{E}A_n \end{aligned} \quad (4.40)$$

Now

$$\mathbb{E}|M_n| = \mathbb{E}|X_n - X_0 - A_n| \leq \mathbb{E}|X_n| + \mathbb{E}|X_0| + \mathbb{E}A_n$$

and

$$\mathbb{E}A_n = \mathbb{E}(X_n - X_0 - M_n) = \mathbb{E}(X_n - X_0) \leq \mathbb{E}|X_n| + \mathbb{E}|X_0|.$$

Inserting these two bounds into (4.40) gives the claimed inequality. \square

The Doob decomposition gives rise to two important derived processes associated to a martingale, M , the *bracket*, $\langle M \rangle$, and $[M]$.

Definition 4.4.1 Let M be a martingale in \mathcal{L}^2 with $M_0 = 0$. Then M^2 is a submartingale with Doob decomposition

$$M^2 = N + \langle M \rangle,$$

where N is a martingale that vanishes at zero and $\langle M \rangle$ is a previsible process that vanishes at zero. The process $\langle M \rangle$ is called the *bracket* of M .

Note that boundedness in \mathcal{L}^1 of $\langle M \rangle$ is equivalent to boundedness in \mathcal{L}^2 of M .

From the formulas associated with the Doob decomposition, we derive that

$$\langle M \rangle_n - \langle M \rangle_{n-1} = \mathbb{E}((M_n^2 - M_{n-1}^2) | \mathcal{F}_{n-1}) = \mathbb{E}((M_n - M_{n-1})^2 | \mathcal{F}_{n-1}). \quad (4.41)$$

Definition 4.4.2 Let M be as before. We define

$$[M]_n \equiv \sum_{k=1}^n (M_k - M_{k-1})^2. \quad (4.42)$$

Lemma 4.4.16 *If M is as before,*

$$M^2 - [M] \equiv V = (C \bullet M), \quad (4.43)$$

where $C_n \equiv 2M_{n-1}$. V is a martingale. If M is bounded in \mathcal{L}^2 , then V is bounded in \mathcal{L}^1 .

Proof. Exercise! □

4.5 A discrete time Itô formula.

We will now give in some way a justification of the name “discrete stochastic integral” for the martingale transform. We consider a martingale M zero in zero and a function $F : \mathbb{R} \rightarrow \mathbb{R}$. We want to consider the process $F(M_T)$ and ask whether we can represent $F(M_t) - F(0)$ as a “stochastic integral. Since we have called $C \bullet M$ a stochastic integral, we might expect that this formula could simply read $F(M_T) = (F' \bullet X)_T + F(M_0)$, as in the usual fundamental theorem of calculus, but this will not turn out to be the case in general.

Let us consider the situation when the increments of M_t are getting very small; the idea here is that the spacings between consecutive times are really small. So we introduce parameter $\epsilon > 0$ that will later tend to zero, while we think that $T = \epsilon^{-1}C$. We will also assume that $\mathbb{E}(M_t - M_{t-1})^2 = O(\epsilon)$. In this situation we can expand in Taylor series, assuming that F will be a smooth function:

$$\begin{aligned} F(M_t) &= F(M_{t-1}) + (M_t - M_{t-1})F'(M_{t-1}) \\ &\quad + \frac{1}{2}F''(M_{t-1})(M_t - M_{t-1})^2 + O((M_t - M_{t-1})^3) \end{aligned} \quad (4.44)$$

where we assume that

$$\mathbb{E}O((M_T - M_{T-1})^3) \leq K\epsilon^{3/2},$$

and therefore $T\mathbb{E}O((M_T - M_{T-1})^3) \leq K\epsilon^{1/2} \downarrow 0$, so that as $\epsilon \downarrow 0$ these error terms will be negligible (and will henceforth not be written anymore). Now we may iterate this procedure to obtain

$$\begin{aligned} F(M_T) &= F(M_0) + \sum_{t=1}^T F'(M_{t-1})(M_t - M_{t-1}) \\ &\quad + \frac{1}{2} \sum_{t=1}^T F''(M_{t-1})(M_t - M_{t-1})^2 + O(\epsilon^{1/2}). \end{aligned} \quad (4.45)$$

This expression looks almost like the Doob decomposition of the process

$F(M_t)$, except that the last term is not exactly predictable. In fact, from the Doob decomposition, we would instead expect a predictable term of the form

$$\sum_{t=1}^T F''(M_{t-1}) \mathbb{E}[(M_t - M_{t-1})^2 | \mathcal{F}_{t-1}]. \quad (4.46)$$

However, under reasonable assumptions (on F and on the behaviour of the increments of the martingale M), the martingale

$$\Delta_T \equiv \sum_{t=1}^T F''(M_{t-1}) ((M_t - M_{t-1})^2 - \mathbb{E}[(M_t - M_{t-1})^2 | \mathcal{F}_{t-1}])$$

satisfies $\mathbb{E}\Delta_T^2 = O(\epsilon)$, and is therefore negligible in our approximation. This implies the discrete version of *Itô's formula*:

$$\begin{aligned} F(M_T) &= F(M_0) + \sum_{t=1}^T F'(M_{t-1})(M_t - M_{t-1}) \\ &\quad + \frac{1}{2} \sum_{t=1}^T F''(M_{t-1}) \mathbb{E}[(M_t - M_{t-1})^2 | \mathcal{F}_{t-1}] + O(\epsilon^{1/2}). \end{aligned} \quad (4.47)$$

4.6 Central limit theorem for martingales

One important further result for martingales concerns *central limit theorems*. There are various different formulations of such theorems. We will present one which emphasises the rôle of the bracket.

Theorem 4.6.17 *Let M be a martingale with $M_0 = 0$. Set $s_n^2 \equiv \mathbb{E}[M]_n \equiv \sum_{i=1}^n \mathbb{E}(M_i - M_{i-1})^2$. Assume that $s_n^{-2} \max_{k \leq n} \mathbb{E}(M_k - M_{k-1})^2 \downarrow 0$, and that, for all $\epsilon > 0$,*

$$s_n^{-2} \sum_{k=1}^n \mathbb{E}[(M_k - M_{k-1})^2 \mathbb{I}_{|M_k - M_{k-1}| > \epsilon s_n} | \mathcal{F}_{k-1}] \downarrow 0, \text{ a.s.} \quad (4.48)$$

Then, if $\langle M \rangle_n / s_n^2 \rightarrow 1$ in probability, then

$$s_n^{-1} M_n \rightarrow \mathcal{N}(0, 1) \quad (4.49)$$

in distribution.

Remark 4.6.1 Condition (4.48) is called the *conditional Lindeberg condition*. In the case when $M_n = S_n = \sum_{i=1}^n X_i$ with independent centered random variables X_i , (4.48) reduces to the usual Lindeberg condition

$$s_n^{-2} \sum_{k=1}^n \mathbb{E} [X_k^2 \mathbb{1}_{|X_k| > \epsilon s_n}] \downarrow 0. \quad (4.50)$$

Moreover, in that case $\mathbb{E}[M]_n = \langle M \rangle_n$, and so condition (4.49) is trivially verified. Thus the above theorem implies the usual CLT for sums of independent random variables under the weakest possible conditions.

Proof. The proof will use characteristic functions. We know that (4.49) will hold, if and only if, for all $u \in \mathbb{R}$,

$$\lim_{n \uparrow \infty} \mathbb{E} e^{ius_n^{-1} M_n} = e^{-u^2/2}. \quad (4.51)$$

Let us set $Y_k \equiv M_k - M_{k-1}$. Things are a little tricky, and the following decomposition is quite helpful:

$$\begin{aligned} & \left| \mathbb{E} \left[e^{ius_n^{-1} M_n} - e^{-\frac{u^2}{2}} \right] \right| \quad (4.52) \\ &= \left| \mathbb{E} \left[e^{ius_n^{-1} M_n} \left(1 - e^{\frac{u^2}{2s_n^2} \langle M \rangle_n} e^{-\frac{u^2}{2}} \right) + e^{-\frac{u^2}{2}} \left(e^{ius_n^{-1} M_n} e^{\frac{u^2}{2s_n^2} \langle M \rangle_n} - 1 \right) \right] \right| \\ &\leq \mathbb{E} \left[\left| 1 - e^{\frac{u^2}{2s_n^2} \langle M \rangle_n} e^{-\frac{u^2}{2}} \right| \right] + \left| \mathbb{E} \left[e^{ius_n^{-1} M_n} e^{\frac{u^2}{2s_n^2} \langle M \rangle_n} - 1 \right] \right| \end{aligned}$$

We want to show that both terms converge to zero as $n \uparrow \infty$. For the first one, this is quite clear from the fact that $e^{\frac{u^2}{2s_n^2} \langle M \rangle_n} \rightarrow e^{\frac{u^2}{2}}$ in probability, provided that

$$\langle M \rangle_n / s_n^2 \leq C < \infty \quad (4.53)$$

We will deal with this condition later. Let us now deal with the second term.

Then we write

$$\begin{aligned} \mathbb{E} e^{ius_n^{-1} M_n + \frac{u^2}{2s_n^2} \langle M \rangle_n} &= \mathbb{E} \left[e^{\sum_{k=1}^n \left(ius_n^{-1} Y_k + \frac{u^2}{2s_n^2} \mathbb{E}[Y_k^2 | \mathcal{F}_{k-1}] \right)} \right] \quad (4.54) \\ &= \mathbb{E} \left[\mathbb{E} \left[\prod_{k=1}^n e^{\left(ius_n^{-1} Y_k + \frac{u^2}{2s_n^2} \mathbb{E}[Y_k^2 | \mathcal{F}_{k-1}] \right)} \middle| \mathcal{F}_{n-1} \right] \right] \\ &= \mathbb{E} \left[\mathbb{E} \left[e^{ius_n^{-1} Y_n + \frac{u^2}{2s_n^2} \mathbb{E}[Y_n^2 | \mathcal{F}_{n-1}]} \middle| \mathcal{F}_{n-1} \right] \prod_{k=1}^{n-1} e^{\left(ius_n^{-1} Y_k + \frac{u^2}{2s_n^2} \mathbb{E}[Y_k^2 | \mathcal{F}_{k-1}] \right)} \right] \end{aligned}$$

Note that in the last lines we have inserted the conditional expectation with respect to \mathcal{F}_{n-1} and pulled the \mathcal{F}_{n-1} -measurable functions as far as possible.

Let us now work these conditional expectations up one after the other. We use that

$$|e^{iux} - 1 - iux + u^2x^2/2| \equiv |R_1(ux)| \leq \min(u^2x^2, |u|^3|x|^3),$$

and

$$|e^{u^2x^2/2} - 1 - u^2x^2/2| \equiv |R_2(u^2x^2)| \leq u^4\sigma^4e^{u^2x^2/2}.$$

To keep notation tractable, define $\sigma_n^2 \equiv \mathbb{E}(Y_n^2|\mathcal{F}_{n-1})$. Then

$$\begin{aligned} \mathbb{E}\left[e^{ius_n^{-1}Y_n + \frac{u^2}{2s_n^2}\mathbb{E}[Y_n^2|\mathcal{F}_{n-1}]}\right] &= \mathbb{E}\left[1 + ius_n^{-1}Y_n - \frac{u^2}{2s_n^2}Y_n^2 + R_1(s_n^{-1}uY_n)\right. \\ &+ \frac{u^2}{2s_n^2}\sigma_n^2 + iu^3s_n^{-3} + Y_n\sigma_n^2 - \frac{u^4}{4s_n^4}Y_n^2\sigma_n^2 + \frac{u^2}{2s_n^2}\sigma_n^2R_1(us_n^{-1}Y_n) \\ &\left.+ R_2(u^2\sigma_n^2/s_n^2)(1 + ius_n^{-1}Y_n - \frac{u^2}{2s_n^2}Y_n^2 + R_1(us_n^{-1}Y_n))\right|\mathcal{F}_{n-1}] \\ &= 1 + \mathbb{E}\left[R_1(s_n^{-1}uY_n) - \frac{u^4}{4s_n^4}Y_n^2\sigma_n^2 + \frac{u^2}{2s_n^2}\sigma_n^2R_1(us_n^{-1}Y_n)\right. \\ &\left.+ R_2(u^2\sigma_n^2/s_n^2)(1 - \frac{u^2}{2s_n^2}Y_n^2 + R_1(us_n^{-1}Y_n))\right|\mathcal{F}_{n-1}] \\ &= 1 + \mathbb{E}\left[R_1(s_n^{-1}uY_n)|\mathcal{F}_{n-1}\right]\left(1 + \frac{u^2\sigma_n^2}{2s_n^2} + R_2(u^2\sigma_n^2/s_n^2)\right) \\ &\quad - \frac{u^2\sigma_n^2}{2s_n^2}\left(\frac{u^2\sigma_n^2}{2s_n^2} + R_2(u^2\sigma_n^2/s_n^2)\right). \end{aligned} \tag{4.55}$$

Here we used that σ_n^2 is \mathcal{F}_{n-1} -measurable and can be treated like a constant in the conditional expectation. The term 1 is just what we want. All other terms are really small, even then summed up in the course of the iteration later on. To see how to deal with them we use the following properties:

- (i) $\sigma_n^2 \leq \epsilon^2 s_n^2 + \mathbb{E}[Y_n^2 \mathbb{1}_{|Y_n| > \epsilon \sigma_n} | \mathcal{F}_{n-1}]$. Both these terms will be nice: the first is order ϵ^2 , and the second is controlled by the Lindeberg condition (4.48).
- (ii) σ_k^2/s_n^2 is sumable over k . Thus this bound will be kept together with ϵ -terms produced using (i).
- (iii) $\sigma_n^2/s_n^2 \leq C$. This bound is applied where $\exp(u^2\sigma_n^2/2s_n^2)$ appears in R_2 , as well as to replace any useless powers of σ_n^2 that may be around.

(iv) Finally, by a similar decomposition

$$\begin{aligned}
& \min(|Y_n^3||u|^3/s_n^3, u^2 Y_n^2/s_n^2) & (4.56) \\
& = \min(|Y_n^3||u|^3/s_n^3, u^2 Y_n^2/s_n^2) (\mathbb{1}_{|Y_n| \leq \epsilon s_n} + \mathbb{1}_{|Y_n| > \epsilon s_n}) \\
& \leq |Y_n^3||u|^3/s_n^3 \mathbb{1}_{|Y_n| \leq \epsilon s_n} + u^2 Y_n^2/s_n^2 \mathbb{1}_{|Y_n| > \epsilon s_n} \\
& \leq \epsilon |u|^3 \frac{Y_n^2}{s_n^2} + u^2 Y_n^2/s_n^2 \mathbb{1}_{|Y_n| > \epsilon s_n}
\end{aligned}$$

and so

$$|\mathbb{E}[R_1(uY_n/s_n)|\mathcal{F}_{n-1}]| \leq \epsilon \frac{\sigma_n^2}{s_n^2} |u|^3 + u^2 s_n^{-2} \mathbb{E}[Y_n^2 \mathbb{1}_{|Y_n| \geq \epsilon s_n} | \mathcal{F}_{n-1}]. \quad (4.57)$$

Using these estimates, we arrive at the following result.

Lemma 4.6.18

$$\begin{aligned}
& \left| \mathbb{E} e^{ius_n^{-1} M_n + \frac{u^2}{2s_n^2} \langle M \rangle_n} - \mathbb{E} e^{ius_n^{-1} M_{n-1} + \frac{u^2}{2s_n^2} \langle M \rangle_{n-1}} \right| & (4.58) \\
& \leq u^2 \left(\epsilon \sigma_n^2 / s_n^2 ((1 + |u|\epsilon) + \mathbb{E}[Y_n^2 \mathbb{1}_{|Y_n| \geq \epsilon s_n} | \mathcal{F}_{n-1}]) \right. \\
& \quad \left. \times (1 + 2u^2 C + 2u^4 C^2 e^{u^2 C/2}) e^{Cu^2/2} \right)
\end{aligned}$$

Proof. Insert the estimate (4.55) into (4.54) The term “1” produces the “ $n - 1$ ”-term in (4.58). Use the estimates from the list above in the error terms, and use the fact that

$$\left| e^{ius_n^{-1} M_{n-1} + \frac{u^2}{2s_n^2} \langle M \rangle_{n-1}} \right| \leq e^{u^2 C/2}.$$

□

To conclude the proof under the hypothesis (4.53), we now have to just iterate. This yields

$$\begin{aligned}
& \left| \mathbb{E} e^{ius_n^{-1} M_n + \frac{u^2}{2s_n^2} \langle M \rangle_n} - 1 \right| & (4.59) \\
& \leq \sum_{k=1}^n \mathbb{E} \left(\epsilon |u| \sigma_k^2 / s_n^2 + \mathbb{E}[Y_k^2 \mathbb{1}_{|Y_k| \geq \epsilon s_n} | \mathcal{F}_{k-1}] \right) |u|^2 e^{u^2 C} C' \\
& \leq \left(\epsilon C |u| + \sum_{k=1}^n \mathbb{E}[Y_k^2 \mathbb{1}_{|Y_k| \geq \epsilon s_n} | \mathcal{F}_{k-1}] \right) |u|^2 C' e^{u^2 C}.
\end{aligned}$$

The Lindeberg type term tends to zero, as n goes to infinity, and the other term can be made arbitrarily small by choosing ϵ as small as desired. This shows proves the CLT under the hypothesis (4.53).

Finally we deal with the boundedness assumption (4.53). Define $A_{n,m} \equiv \{\omega \in \Omega : s_n^{-2} \sum_{k=1}^m \mathbb{E}[Y_k | \mathcal{F}_{k-1}] \leq C\}$. Of course, $\mathbb{P}(A_{n,n})$ bounds the probabilities $\mathbb{P}(A_{n,m})$, $m \leq n$ from below, and $\lim_{n \uparrow \infty} \mathbb{P}(A_{n,n}) = 1$. Notice that $\sum_{k=1}^m \mathbb{E}[Y_k^2 | \mathcal{F}_{k-1}]$ is \mathcal{F}_{m-1} measurable, and hence so is $\mathbb{1}_{A_{n,m}}$. Thus, if we set $Z_{n,m} \equiv Y_m \mathbb{1}_{A_{n,m}}$, it holds that $\mathbb{E}[Z_{n,m} | \mathcal{F}_{m-1}] = 0$ for all $m \leq n$. Thus the variables $\{Z_{n,m}, m \leq n\}$, for fixed n , form a martingale difference sequence (which can be extended to an infinite sequence by setting $Z_{n,m} = 0, m > n$). Since $|Z_{n,m}| \leq |Y_m|$, all the properties used in the calculations above carry over to the $Z_{n,m}$. Therefore, repeating the calculations above with M_n replaced by $\widehat{M}_n \equiv \sum_{m=1}^n Z_{n,m}$, we find that

$$\lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu \widehat{M}_n} \right] = e^{-u^2/2}. \quad (4.60)$$

Finally,

$$\begin{aligned} \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu M_n} \right] &= \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu M_n} \mathbb{1}_{A_{n,n}} \right] + \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu M_n} \mathbb{1}_{A_{n,n}^c} \right] \quad (4.61) \\ &= \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu \widehat{M}_n} \mathbb{1}_{A_{n,n}} \right] + 0 \\ &= \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu \widehat{M}_n} \right] - \lim_{n \uparrow \infty} \mathbb{E} \left[e^{iu \widehat{M}_n} \mathbb{1}_{A_{n,n}^c} \right] \\ &= e^{-u^2/2} - 0. \end{aligned}$$

This concludes the proof of the theorem. \square

Remark 4.6.2 It is obvious from the proof that the result has an immediate extension to triangular arrays of martingale difference sequences, like \widehat{M} .

Very similar computations like those presented above play an important rôle in what is called the *concentration of measure phenomenon*. Without going into too many details, let me briefly describe this. The setting one is considering is the following. We have n independent, identically distributed random variables, X_1, \dots, X_n , assumed to have mean zero, variance one, and to satisfy, e.g. $\mathbb{E}e^{uX_i} < \infty$, for all $u \in \mathbb{R}$. Let $f : \mathbb{R}^n \rightarrow \mathbb{R}$ be a differentiable function that satisfies

$$\sup_{k=1}^n \sup_{x \in \mathbb{R}^n} \left| \frac{\partial f}{\partial x_k}(x_1, \dots, x_n) \right| \leq 1.$$

Set $F \equiv f(X_1, \dots, X_n)$. Then one can show that for some constant, $C > 0$,

$$\mathbb{P}(|F - \mathbb{E}F| > \rho n) \leq 2e^{-\frac{\rho^2}{2C}}. \quad (4.62)$$

The proof relies on the exponential Markov inequality, that states that

$$\mathbb{P}(F - \mathbb{E}F > n\rho) \leq \inf_{t \geq 0} e^{-tn\rho} \mathbb{E}e^{t(F - \mathbb{E}F)}.$$

The trick is to compute the Laplace transform, i.e. to show that

$$\mathbb{E}e^{t(F - \mathbb{E}F)} \leq e^{t^2 nC/2}.$$

(and not, as one might worry, of order $\exp(n^2)$!!).

To do this, one writes $F - \mathbb{E}F$ as a martingale difference sequence with respect to the filtration generated by the random variables X_i :

$$F - \mathbb{E}F = \sum_{k=1}^n (\mathbb{E}(F|\mathcal{F}_k) - \mathbb{E}(F|\mathcal{F}_{k-1})). \quad (4.63)$$

The computations one has to do are quite similar to those we have performed in the proof of the central limit theorem. There is one small trick that is useful to use: Set $F^u \equiv f(X_1, \dots, uX_k, X_{k+1}, \dots, X_n)$. Then

$$F - F^0 = \int_0^1 du \frac{d}{du} F^u = \int_0^1 du X_k \frac{\partial}{\partial x_k} f(X_1, \dots, uX_k, X_{k+1}, \dots, X_n)$$

and

$$\begin{aligned} \mathbb{E}[F|\mathcal{F}_k] - \mathbb{E}[F|\mathcal{F}_{k-1}] &= \int_0^1 du \left(\mathbb{E} \left[\frac{d}{du} F^u | \mathcal{F}_k \right] - \mathbb{E} \left[\frac{d}{du} F^u | \mathcal{F}_{k-1} \right] \right) \\ &\equiv \mathbb{E}[Z_k | \mathcal{F}_k] - \mathbb{E}[Z_k | \mathcal{F}_{k-1}], \end{aligned} \quad (4.64)$$

where $|Z_k| \leq |X_k|$. Hence

$$\begin{aligned} &\mathbb{E} \left(e^{\lambda(\mathbb{E}(F|\mathcal{F}_k) - \mathbb{E}(F|\mathcal{F}_{k-1}))} - 1 - \lambda (\mathbb{E}(F|\mathcal{F}_k) - \mathbb{E}(F|\mathcal{F}_{k-1})) \middle| \mathcal{F}_{k-1} \right) \\ &\leq \lambda^2 \mathbb{E} \left[(\mathbb{E}(F|\mathcal{F}_k) - \mathbb{E}(F|\mathcal{F}_{k-1}))^2 e^{\lambda|\mathbb{E}(Z_k|\mathcal{F}_k) - \mathbb{E}(Z_k|\mathcal{F}_{k-1})|} \middle| \mathcal{F}_{k-1} \right] \\ &\leq \lambda^2 C \end{aligned} \quad (4.65)$$

by the assumption on the law of X_k . We leave the remaining details of the calculation as an exercise. For more on concentration of measure, see e.g. [8, 9].

4.7 Stopping times, optional stopping

In a stochastic process we often want to consider random times that are determined by the occurrence of a particular event. If this event depends only on what happens “in the past”, we call it a *stopping time*. Stopping times are nice, since we can determine their occurrence as we observe

the process; hence, if we are only interested in them, we can stop the process at this moment, hence the name.

Definition 4.7.1 A map $\tau : \Omega \rightarrow \mathbb{N}_0 \cup \{+\infty\}$ is called a stopping time (with respect to a filtration $\{\mathcal{F}_n, n \in \mathbb{N}_0\}$), if, for all $n \in \mathbb{N}_0 \cup \{+\infty\}$,

$$\{\tau = n\} \in \mathcal{F}_n. \quad (4.66)$$

Example. The most important examples of stopping times are hitting time. Let X be an adapted process, and let $B \in \mathcal{B}$. Define

$$\tau_B \equiv \inf\{t > 0 : X_t \in B\}.$$

Then τ_B is a stopping time. To see this, note that, if $n \in \mathbb{N}$.

$$\{\tau_B = n\} = \{\omega : X_n(\omega) \in B, X_k(\omega) \notin B, \forall 0 < k < n\}.$$

This event is manifestly in \mathcal{F}_n . The event $\{\tau_B = \infty\}$ occurs if $\{X_n \notin B, \forall n \in \mathbb{N}\} \subset \mathcal{F}_\infty$.

In principle all stopping times can be realised as first hitting times of some process. To do so, define

$$I_{[T, \infty)}(n, \omega) = \begin{cases} 1, & \text{if } n \geq T(\omega), \\ 0, & \text{otherwise.} \end{cases}$$

This process is adapted, and $T = \tau_1$.

It is sometimes very convenient to have the notion of a σ -algebra of events that take place before a stopping time.

Definition 4.7.2 The pre- T - σ -algebra, \mathcal{F}_T , is the set of events $F \subset \Omega$, such that, for all $n \in \mathbb{N}_0 \cup \{+\infty\}$,

$$F \cap \{T \leq n\} \in \mathcal{F}_n. \quad (4.67)$$

Pre- T - σ -algebras will play an important rôle in formulation the strong Markov property.

There are some useful elementary facts associated with this concept.

Lemma 4.7.19 *Let S, T be stopping times. Then:*

- (i) *If X is an adapted process, then X_T is \mathcal{F}_T -measurable.*
- (ii) *If $S < T$, then $\mathcal{F}_S \subset \mathcal{F}_T$.*
- (iii) *$\mathcal{F}_{T \wedge S} = \mathcal{F}_T \cap \mathcal{F}_S$.*
- (iv) *If $F \in \mathcal{F}_{S \vee T}$, then $F \cap \{S \leq T\} \in \mathcal{F}_T$.*
- (v) *$\mathcal{F}_{S \vee T} = \sigma(\mathcal{F}_T, \mathcal{F}_S)$.*

Proof. Exercise. □

We now return to our gambling mode. We consider a supermartingale X and we want to play a strategy, C , that depends of a stopping time, T : say, we keep one unit of stock until the random time T . Then

$$C_n \equiv C_n^T \equiv \mathbb{1}_{n \leq T}.$$

Note that C^T is a previsible process. Namely,

$$\{C_n^T = 0\} = \{T \leq n - 1\} \in \mathcal{F}_{n-1},$$

and since C_n^T only takes the two values 0, 1, this suffices to show that $C_n^T \in \mathcal{F}_{n-1}$. The wealth process associated to this strategy is then

$$(C^T \bullet X)_n = X_{T \wedge n} - X_0.$$

If we define the stopped process X^T , via

$$X_n^T(\omega) \equiv X_{T(\omega) \wedge n}(\omega),$$

we have alternatively

$$C^T \bullet X = X^T - X_0.$$

Since C^T is positive and bounded, Theorem 4.1.1, (i), implies the following:

Theorem 4.7.20(i) *If X is a supermartingale and T is a stopping time, then the stopped process, $X^T \equiv \{X_{T \wedge n}, n \in \mathbb{N}_0\}$, is a supermartingale. In particular, for all $n \in \mathbb{N}$,*

$$\mathbb{E}X_{T \wedge n} \leq \mathbb{E}X_0. \quad (4.68)$$

(ii) If X is a martingale and T is a stopping time, then X^T is a martingale. In particular

$$\mathbb{E}X_{T \wedge n} = \mathbb{E}X_0. \quad (4.69)$$

Proof. Immediate from Theorem 4.1.1, (i). □

This theorem is disappointing news for those who might have hoped to reach a certain gain by playing until they have won a preset sum of money, and stopping then. In a martingale setting, the sure gain that will occur *if* this stopping time is reached before time n is offset by the expected loss, if the target has not yet been reached.

Note, however, that the theorem does not assert that $\mathbb{E}X_T \leq \mathbb{E}X_0$. The following theorem, called Doob's Optional Stopping Theorem, gives conditions under which even that holds.

Theorem 4.7.21(i) *Let T be a stopping time, and let X be a supermartingale. Then, X_T is integrable and*

$$\mathbb{E}X_T \leq \mathbb{E}X_0, \quad (4.70)$$

if one of the following conditions holds:

- (a) T is bounded (i.e. there exists $N \in \mathbb{N}$, s.t. $\forall \omega \in \Omega T(\omega) \leq N$);
- (b) X is bounded, and T is a.s. finite;
- (c) $\mathbb{E}T < \infty$, and, for some $K < \infty$,

$$|X_n(\omega) - X_{n-1}(\omega)| \leq K, \quad (4.71)$$

for all $n \in \mathbb{N}, \omega \in \Omega$.

(ii) *If X is a martingale, then in any of the situations above $\mathbb{E}X_T = \mathbb{E}X_0$.*

Remark 4.7.1 This theorem may look strange, and contradict the “no strategy” idea: take a simple random walk, S_n , (i.e. a series of fair games, and define a stopping time $T = \inf\{n : S_n = 10\}$. Then clearly $X_T = 10 \neq \mathbb{E}X_0 = 0$! So we conclude, using (c), that $\mathbb{E}T = +\infty$. In fact, the “sure” gain if we achieve our goal is offset by the fact that on average, it takes infinitely long to reach it (of course, most games will end quickly, but chances are that some may take very very long!

Proof. We already know that $\mathbb{E}X_{n \wedge T} - \mathbb{E}X_0 \leq 0$ for all $n \in \mathbb{N}$. Consider case (a). Then we know that $T \wedge N = T$, and so $\mathbb{E}X_T = \mathbb{E}X_{T \wedge N} \leq \mathbb{E}X_0$, as claimed.

In case (b), we start from $\mathbb{E}X_{n \wedge T} - \mathbb{E}X_0 \leq 0$ and let $n \uparrow \infty$. Since T is almost surely finite, $\lim_{n \uparrow \infty} X_{n \wedge T} = X_T$, a.s., and since X_n is uniformly bounded,

$$\lim_{n \uparrow \infty} \mathbb{E}X_{T \wedge n} = \mathbb{E} \lim_{n \uparrow \infty} X_{T \wedge n} = \mathbb{E}X_T,$$

which implies the result.

In the last case, (c), we observe that

$$|X_{T \wedge n} - X_0| = \left| \sum_{k=1}^{T \wedge n} (X_k - X_{k-1}) \right| \leq KT,$$

and by assumption $\mathbb{E}KT < \infty$. Thus, we can again take the limit $n \uparrow \infty$ and use Lebesgue’s dominated convergence theorem to justify that the inequality survives.

Finally, to justify (ii), use that if X is a martingale, then both X and $-X$ are supermartingales. The ensuing two inequalities imply the desired equality. \square

Case (c) in the above theorem is certainly the most frequent situation one may hope to be in. For this it is good to know how to show that $\mathbb{E}T < \infty$, if that is the case. The following lemma states that this is always the case, whenever, eventually, the probability that the event leading to T is reasonably big.

Lemma 4.7.22 *Suppose that T is a stopping time and that there exists $N \in \mathbb{N}$ and $\epsilon > 0$, such that, for all $n \in \mathbb{N}$,*

$$\mathbb{P}(T \leq n + N | \mathcal{F}_n) > \epsilon, \text{ a.s.} \quad (4.72)$$

Then $\mathbb{E}T < \infty$.

Proof. Consider $\mathbb{P}(T > kN)$. Clearly we can write

$$\begin{aligned} \mathbb{P}(T > kN) &= \mathbb{E}(\mathbb{1}_{T > (k-1)N} \mathbb{1}_{T > kN}) & (4.73) \\ &= \mathbb{E}(\mathbb{E}(\mathbb{1}_{T > (k-1)N} \mathbb{1}_{T > kN} | \mathcal{F}_{(k-1)N})) \\ &= \mathbb{E}(\mathbb{1}_{T > (k-1)N} \mathbb{E}(\mathbb{1}_{T > kN} | \mathcal{F}_{(k-1)N})) \\ &\leq (1 - \epsilon) \mathbb{E}(\mathbb{1}_{T > (k-1)N}) \\ &\leq (1 - \epsilon)^k, \end{aligned}$$

by iteration. The exponential decay of the probability implies the finiteness of the expectation of T immediately. \square

Finally we state Doob's supermartingale inequalities for non-negative supermartingales.

Theorem 4.7.23 *Let X be non-negative supermartingale and T a stopping time. Then*

$$\mathbb{E}X_T \leq \mathbb{E}X_0. \quad (4.74)$$

Moreover, for any $c > 0$,

$$c\mathbb{P}\left(\sup_k X_k > c\right) \leq \mathbb{E}X_0. \quad (4.75)$$

Proof. We know that $\mathbb{E}X_{T \wedge n} \leq \mathbb{E}X_0$. Using Fatou's lemma allows to pass to the limit $n \uparrow \infty$. For (4.74), set $T = \inf\{n : X_n > c\}$. Clearly $X_T \geq c$, if $\sup_k X_k > c$, and zero else. Thus $\mathbb{E}X_T \geq c\mathbb{P}(\sup_k X_k > c)$, and (4.75) follows from (4.74). \square

Markov processes

We have seen the definition and construction of discrete time Markov chains already in Chapter 3. Markov chains are among the most important stochastic processes that are used to model real live phenomena that involve disorder. This is because the construction of these processes is very much adapted to our thinking about such processes. Moreover, Markov processes can be very easily implemented in numerical algorithms. This allows to numerically simulate even very complicated systems. We will always imagine a Markov process as a “particle” moving around in state space; mind, however, that these “particles” can represent all kinds of very complicated things, once we allow the state space to be sufficiently general. In this section, S will always be a complete separable metric space.

5.1 Markov processes with stationary transition probabilities

In general, we call a stochastic process whose index set supports the action of a group (or semi-group) *stationary* (with respect to the action of this (semi) group, if all finite dimensional distributions are invariant under the simultaneous shift of all time-indices. Specifically, if our index sets, I , are \mathbb{R}_+ or \mathbb{Z} , resp. \mathbb{N} , then a stochastic process is stationary if for all $\ell \in \mathbb{N}$, $s_1, \dots, s_\ell \in I$, all $A_1, \dots, A_\ell \in \mathcal{B}$, and all $t \in I$,

$$\mathbb{P}[X_{s_1} \in A_1, \dots, X_{s_\ell} \in A_\ell] = \mathbb{P}[X_{s_1+t} \in A_1, \dots, X_{s_\ell+t} \in A_\ell]. \quad (5.1)$$

We can express this also as follows: Define the shift θ , for any $t \in I$, as $(X \circ \theta_t)_s \equiv X_{t+s}$. Then X is stationary, if and only if, for all $t \in I$, the processes X and $X \circ \theta_t$ have the same finite dimensional distributions.

In the case of Markov processes, a necessary (but not sufficient) condition for stationarity is the stationarity of the transitions kernels. Recall

that we have defined the one-step transition kernel \mathcal{P}_t of a Markov process in Section 3.3.

Definition 5.1.1 A Markov process with discrete time \mathbb{N}_0 and state space S is said to have *stationary transition probabilities (kernels)*, if its one step transition kernel, \mathcal{P}_t , is independent of t , i.e. there exists a probability kernel, $P(x, A)$, s.t.

$$\mathcal{P}_t(x, A) = P(x, A), \quad (5.2)$$

for all $t \in \mathbb{N}$, $x \in S$, and $A \in \mathcal{B}$.

Remark 5.1.1 With the notation $P_{t,s}$ for the transition kernel from time s to time t , we could alternatively state that a Markov process has *stationary transition probabilities (kernels)*, if there exists a family of transition kernels $P_t(x, A)$, s.t.

$$P_{s,t}(x, A) = P_{t-s}(x, A), \quad (5.3)$$

for all $s < t \in \mathbb{N}$, $x \in S$, and $A \in \mathcal{B}$. Note that there is a potential conflict of notation between \mathcal{P}_t and P_t which should not be confused.

A key concept for Markov chains with stationary transition kernels is the notion of an *invariant* distribution.

Definition 5.1.2 Let P be the transition kernel of a Markov chain with stationary transition kernels. Then a probability measure, π , on (S, \mathcal{B}) is called an invariant (probability) distribution, if

$$\int \pi(dx) P(x, A) = \pi(A), \quad (5.4)$$

for all $A \in \mathcal{B}$. More generally, a positive, Σ -finite measure, π , satisfying (5.4), is called an *invariant measure*.

Lemma 5.1.1 *A Markov chain with stationary probability kernels and initial distribution $P_0 = \pi$ is a stationary stochastic process, if and only if π is an invariant probability distribution.*

Proof. Exercise. □

In the case when the state space, S , is finite, we have seen that there is always at least one invariant measure, which then can be chosen to be a probability measure. In the case of general state spaces, while there still will always be an invariant measure (through a generalisation of the Perron-Frobenius theorem to the operator setting), there appears a new issue, namely whether there is an invariant measure that is finite, viz. whether there exists a invariant probability distribution.

5.2 The strong Markov property

The setting of Markov processes is very much suitable for the application of the notions of stopping times introduced in the last section. In fact, one of the very important properties of Markov processes is the fact that we can split expectations between past and future also at random times.

Theorem 5.2.2 *Let X be a Markov process with stationary transition kernels. Let \mathcal{F}_n be a filtration such that X is adapted, and let T be a stopping time. Let F and G be \mathcal{F} -measurable functions, and let F in addition be measurable with respect to the pre- T - σ -algebra \mathcal{F}_T . Then*

$$\mathbb{E}[\mathbb{1}_{T < \infty} FG \circ \theta_T | \mathcal{F}_0] = \mathbb{E}[\mathbb{1}_{T < \infty} F \mathbb{E}'[G | \mathcal{F}'_0](X_T) | \mathcal{F}_0] \quad (5.5)$$

where \mathbb{E}' and refers to an independent copy, X' , of the Markov chain X .

Proof. We have

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{T < \infty} FG \circ \theta_T | \mathcal{F}_0] & (5.6) \\ &= \mathbb{E}[\mathbb{E}[\mathbb{1}_{T < \infty} FG \circ \theta_T | \mathcal{F}_T] | \mathcal{F}_0] \\ &= \mathbb{E}[\mathbb{1}_{T < \infty} F \mathbb{E}[G \circ \theta_T | \mathcal{F}_T] | \mathcal{F}_0]. \end{aligned}$$

Now $\mathbb{E}[G \circ \theta_T | \mathcal{F}_T]$ depends only on X_T and by stationarity is equal to $\mathbb{E}'[G | \mathcal{F}'_0](X_T)$, which yields the claim of the theorem. \square

5.3 Markov processes and martingales

We now want to develop some theory that will be more important and more difficult in the continuous time case. First we want to see how the transition kernels can be seen as operators acting on spaces of measures respectively spaces of function.

If μ is a σ -finite measure on S , and P is a Markov transition kernel, we define the measure μP as

$$\mu P(A) \equiv \int_S P(x, A) d\mu(x), \quad (5.7)$$

and similarly, for the t -step transition kernel, P_t ,

$$\mu P_t(A) \equiv \int_S P_t(x, A) d\mu(x). \quad (5.8)$$

By the Markov property, we have of course the

$$\mu P_t(A) = \mu P^t(A). \quad (5.9)$$

The action on measures has of course the following natural interpretation in terms of the process: if $\mathbb{P}(X_0 \in A) = \mu(A)$, then

$$\mu(X_t \in A) = \mu P_t(A). \quad (5.10)$$

Alternatively, if f is a bounded, measurable function on S , we define

$$(Pf)(x) \equiv \int_S f(y)P(x, dy), \quad (5.11)$$

and

$$(P_t f)(x) \equiv \int_S f(y)P_t(x, dy), \quad (5.12)$$

where again

$$P_t f = P^t f. \quad (5.13)$$

We say that P_t is a semi-group acting on the space of measures, respectively on the space of bounded measurable functions. The interpretation of the action on functions is given as follows.

Lemma 5.3.3 *Let P_t be a Markov semi-group acting on bounded measurable functions f . Then*

$$(P_t f)(x) = \mathbb{E}(f(X_t) | \mathcal{F}_0)(x) \equiv \mathbb{E}_x f(X_t). \quad (5.14)$$

Proof. We only need to show this for $t = 1$. Then, by definition,

$$\mathbb{E}_x f(X_1) = \int_S f(y) \mathbb{P}[X_1 \in dy | \mathcal{F}_0](x) = \int_S f(y) P(x, dy).$$

□

Notice that, by telescopic expansion, we have the elementary formula

$$P_t f - f = \sum_{s=0}^{t-1} P_s (P - \mathbb{I}) f \equiv \sum_{s=0}^{t-1} P_s L f, \quad (5.15)$$

where we call $L \equiv P - \mathbb{I}$ the (discrete) generator of our Markov process (this formula will have a complete analogon in the continuous-time case).

An interesting consequence is the following observation:

Lemma 5.3.4 *[Discrete time martingale problem]. Let L be the generator of a Markov process, X_t , and let f be a bounded measurable function. Then*

$$M_t \equiv f(X_t) - f(X_0) - \sum_{s=0}^{t-1} Lf(X_s) \quad (5.16)$$

is a martingale.

Proof. Let $t, r \geq 0$. Then

$$\begin{aligned}
\mathbb{E}(M_{t+r}|\mathcal{F}_t) &= \mathbb{E}(f(X_{t+r})|\mathcal{F}_t) - \mathbb{E}(f(X_0)|\mathcal{F}_t) - \sum_{s=0}^{t+r-1} \mathbb{E}(Lf(X_s)|\mathcal{F}_t) \\
&= P^r f(X_t) - f(X_t) + f(X_t) - f(X_0) \\
&\quad - \sum_{s=t}^{t+r-1} \mathbb{E}(Lf(X_s)|\mathcal{F}_t) - \sum_{s=0}^{t-1} \mathbb{E}(Lf(X_s)|\mathcal{F}_t) \\
&= f(X_t) - f(X_0) - \sum_{s=0}^{t-1} (Lf(X_s)) \\
&\quad + P^r f(X_t) - f(X_t) - \sum_{s=0}^{r-1} P^r (Lf(X_t)) \\
&= M_t + 0.
\end{aligned} \tag{5.17}$$

This proves the lemma. \square

Remark 5.3.1 (5.16) is of course the Doob decomposition of the process $f(X_t)$, since $\sum_{s=0}^{t-1} Lf(X_s)$ is a previsible process. One may check that this can be obtained directly using the formula (4.38) [Exercise!].

What is important about this observation is that it gives rise to a characterisation of the generator that will be extremely useful in the general continuous time setting.

Namely, one can ask whether the requirement that M_t be a martingale given a family of pairs (f, Lf) characterises fully a Markov process.

Theorem 5.3.5 *Let X be a discrete time stochastic process on a filtered space such that X is adapted. Then X is a Markov process with transition kernel $P \equiv \mathbb{I} + L$, if and only if, for all bounded measurable functions, f , the expression on the right-hand side of (5.16) is a martingale.*

Proof. Lemma 5.3.4 already provides the “only if” part, so it remains to show the “if” part.

First, if we assume that X is a Markov process, setting $r = 1$ above and taking conditional expectations given \mathcal{F}_0 , we see that $\mathbb{E}f(X_1) - f(X_0) = (Lf)(X_0)$, implying that the transition kernel must be $\mathbb{I} + L$.

It remains to show that X is indeed a Markov process. To see this, we just use the above calculation to see that

$$\begin{aligned}
\mathbb{E}(f(X_{t+r})|\mathcal{F}_t) &= \mathbb{E}(M_{t+r}|\mathcal{F}_t) + f(X_0) \\
&\quad + \sum_{s=0}^{t-1} (Lf)(X_s) + \sum_{s=t}^{t+r-1} \mathbb{E}((Lf)(X_s)|\mathcal{F}_t) \\
&= M_t + f(X_0) + \sum_{s=0}^{t-1} (Lf)(X_s) + \sum_{s=t}^{t+r-1} \mathbb{E}((Lf)(X_s)|\mathcal{F}_t) \\
&= f(X_t) + \sum_{s=0}^{r-1} \mathbb{E}((Lf)(X_{t+s})|\mathcal{F}_t) \tag{5.18}
\end{aligned}$$

Now let again $r = 1$. Then

$$\mathbb{E}(f(X_{t+1})|\mathcal{F}_t) = f(X_t) + (Lf)(X_t) = ((\mathbb{I} + L)f)(X_t) \equiv Pf(X_t), \tag{5.19}$$

In view of the definition of discrete time Markov processes, choosing $f = \mathbb{1}_A$, for $A \in \mathcal{B}(S)$, this gives (3.21), and hence X is a Markov process. Thus the theorem is proven. \square

In view of continuous time Markov processes it is, however, instructive to see that we can also derive easily the more general formula

$$\mathbb{E}(f(X_{t+s})|\mathcal{F}_t) = (\mathbb{I} + L)^s f(X_t) \equiv P^s f(X_t), \tag{5.20}$$

from the martingale problem. We have seen that it holds for $s = 1$; Now proceed by induction: assume that it holds for all bounded measurable functions for $s \leq r - 1$. We must show that it then also holds for $s = r$. To do this, we use (5.18) and use the induction hypothesis for the terms in the sum (where $s \leq r - 1$) with f replaced by Lf . This gives

$$\begin{aligned}
\mathbb{E}(f(X_{t+r})|\mathcal{F}_t) &= f(X_t) + \sum_{s=0}^{r-1} ((\mathbb{I} + L)^s Lf)(X_t) \tag{5.21} \\
&= f(X_t) + \sum_{s=0}^{r-1} (((\mathbb{I} + L)^s (L + 1)f)(X_t) - ((\mathbb{I} + L)^s f)(X_t)) \\
&= ((\mathbb{I} + L)^r f)(X_t),
\end{aligned}$$

as claimed. Hence (5.20) holds for all s , by induction.

Remark 5.3.2 The full strength of this theorem will come out in the continuous time case, where it remains valid. A crucial point is that it will not be necessary to even consider all bounded functions, but just sufficiently rich classes. This allows to formulate martingale problems

even then one cannot write down the generator in an explicit form. The idea of characterising Markov processes by the associated martingale problem goes back to Stroock and Varadhan, see [13].

5.4 Harmonic functions and martingales

We have seen that measures that satisfy $\mu L = 0$ are of special importance in the theory of Markov processes. Also of central importance are functions that satisfy $Lf = 0$. In this section we will assume that the transition kernels of our Markov chains have bounded support, so that for some $K < \infty$, $|X_{t+1} - X_t| \leq K < \infty$ for all t .

Definition 5.4.1 Let L be the generator of a Markov process. A measurable function that satisfies

$$Lf(x) = 0, \forall x \in S, \quad (5.22)$$

is called a *harmonic function*. A function is called *subharmonic* (resp. *superharmonic*), if $Lf \geq 0$, (resp. $Lf \leq 0$).

Theorem 5.4.6 Let X_t be a Markov process with generator L . Then, a non-negative function f is

- (i) *harmonic, if and only if $f(X_t)$ is a martingale;*
- (ii) *subharmonic, if and only if $f(X_t)$ is a submartingale;*
- (iii) *superharmonic, if and only if $f(X_t)$ is a supermartingale;*

Proof. Simply use Lemma 5.3.4. □

Remark 5.4.1 Theorem 5.4.6 establishes a profound relationship between potential theory and martingales. It also explains, the strange choice of super and sub in martingale theory.

A nice application of the preceding result is the maximum principle.

Theorem 5.4.7 Let X be a Markov process and let D be a bounded open domain such that $\mathbb{E}\tau_{D^c} < \infty$. Assume that f is a non-negative subharmonic function on D . Then

$$\sup_{x \in D} f(x) \leq \sup_{x \in D^c} f(x). \quad (5.23)$$

Proof. Let us define $T \equiv \tau_{D^c}$. Then $f(X^T)$ is a submartingale, and thus

$$\mathbb{E}(f(X_T)|\mathcal{F}_0)(x) \geq f(x). \quad (5.24)$$

Since $X_T \in D^c$, it must be true that

$$\sup_{y \in D^c} f(y) \geq \mathbb{E}(f(X_T)|\mathcal{F}_0)(x) \geq f(x), \quad (5.25)$$

for all $x \in D$, hence the claim of the theorem. Of course we used again the Doob's optional stopping theorem in case (i,c). \square

The theorem can be phrased as saying that (sub) harmonic functions take on their maximum on the boundary, since of course the set D^c in (5.23) can be replaced by a subset, $\partial D \subset D^c$ such that $\mathbb{P}_x(X_T \in \partial D) = 1$. The above proof is an example of how intrinsically analytic results can be proven with probabilistic means. The next section will further develop this theme.

5.5 Dirichlet problems

Let us now consider a connected bounded open subset of S . We define the stopping time $T \equiv \tau_{D^c} \equiv \inf\{t > 0 : X_t \in D^c\}$.

If g is a measurable function on D , we consider the Dirichlet problem associated to a generator, L , of a Markov process, X :

$$\begin{aligned} -(Lf)(x) &= g(x), & x \in D, \\ f(x) &= 0, & x \in D^c. \end{aligned} \quad (5.26)$$

Theorem 5.5.8 *Assume the $\mathbb{E}T < \infty$. Then (5.26) has a unique solution given by*

$$f(x) = \mathbb{E} \left(\sum_{t=0}^{T-1} g(X_t) | \mathcal{F}_0 \right) (x) \quad (5.27)$$

Proof. Consider the martingale M_t from Lemma 5.3.4. We know from Theorem 4.7.20 that M^T is also a martingale. Moreover,

$$M_T = f(X_T) - f(X_0) - \sum_{t=0}^{T-1} (Lf)(X_t) = 0 - f(X_0) - \sum_{t=0}^{T-1} (Lf)(X_t). \quad (5.28)$$

But we want f such that $-Lf = g$ on D . Thus, (5.28) seen as a problem for f , reads

$$M_T = -f(X_0) + \sum_{t=0}^{T-1} g(X_t). \quad (5.29)$$

Taking expectations conditioned on \mathcal{F}_0 , yields

$$0 = -f(X_0) + \mathbb{E} \left(\sum_{t=0}^{T-1} g(X_t) \middle| \mathcal{F}_0 \right), \quad (5.30)$$

or

$$f(x) = \mathbb{E}_x \left(\sum_{t=0}^{T-1} g(X_t) \right) \quad (5.31)$$

Here we relied of course on Doob's optimal stopping theorem for $\mathbb{E}M_T = 0$.

Thus any solution of the Dirichlet problem is given by (5.31). To verify existence, we just need to check that (5.31) solves $-Lf = g$ on D . To do this we use the Markov property "backwards", to see that

$$\begin{aligned} Pf(x) &= P\mathbb{E}_x \left(\sum_{t=0}^{T-1} g(X_t) \right) = \mathbb{E}_x \left[\sum_{t=1}^{T-1} g(X_t) \right] \\ &= \mathbb{E}_x \left[\sum_{t=0}^{T-1} g(X_t) \right] - g(x) = f(x) - g(x). \end{aligned} \quad (5.32)$$

□

We see that the Markov process produces a solution of the Dirichlet problem. We can express the solution in terms of an integral kernel, called the Green's kernel, $G_D(x, dy)$, as

$$f(x) = \int G_D(x, dy)g(y) \equiv \mathbb{E}_x \left(\sum_{t=0}^{T-1} g(X_t) \right), \quad (5.33)$$

or, in more explicit terms,

$$G_D(x, dy) = \sum_{t=0}^{\infty} \mathbb{P}_x(T > t) P_D^t(x, dy), \quad (5.34)$$

where

$$P_D^t(x, dy) = \int_D P(x, dz_1) \int_D P(z_1, dz_2) \int_D \dots \int_D P(z_{t-1}, dy). \quad (5.35)$$

The preceding theorem has an obvious extension to more complicated boundary value problems:

Theorem 5.5.9 *Let D be as above, and let h be a bounded function on D^c . Assume the $\mathbb{E}T < \infty$. Then*

$$f(x) \equiv \mathbb{E}_x \left(\sum_{t=0}^{T-1} g(X_t) \right) + \mathbb{E}_x (h(X_T)). \quad (5.36)$$

Then f is the unique solution of the Dirichlet problem

$$\begin{aligned} -(Lf)(x) &= g(x), & x \in D, \\ f(x) &= h(x), & x \in D^c. \end{aligned} \tag{5.37}$$

Proof. Identical to the previous one. \square

Theorem 5.5.9 is a two way game: it allows to produce solutions of analytic problems in terms of stochastic processes, and it allows to compute interesting probabilistic problems analytically. As an example, assume that $D^c = A \cup B$ with $A \cap B = \emptyset$. Set $h = \mathbb{1}_A$. Then, clearly, for $x \in D$,

$$\mathbb{E}_x h(X_T) = \mathbb{P}_x(X_T \in A) \equiv \mathbb{P}_x(\tau_A < \tau_B), \tag{5.38}$$

and so $\mathbb{P}_x(X_T \in A)$ can be represented as the solution of the boundary value problem

$$\begin{aligned} (Lf)(x) &= 0, & x \in D, \\ f(x) &= 1, & x \in A, \\ f(x) &= 0, & x \in B. \end{aligned} \tag{5.39}$$

This is a generalisation of the *ruin* problem for the random walk that we discussed in Probability 1.

Exercise. Derive the formula for $\mathbb{P}_x(\tau_A < \tau_B)$ directly from the Markov property without using Lemma 5.3.4.

5.6 Doob's h -transform

Let us consider a discrete time Markov process, X , with generator $P - 1$ given. We may want to consider modification of the process. One important type of conditioning is that to reach some set in particular places (e.g. consider a random walk in a finite interval; we may be interested to consider this walk conditioned on the fact that it exits on a specific side of the interval; this may correspond to consider a sequence of games conditioned on the player to win).

How and when can we do this, and what is the nature of the resulting process? In particular, is the resulting process again a Markov process, and if so, what is its generator?

As an example, let us try to condition a Markov process to hit a domain B for the first time in a subset $A \subset B$. We may assume that $\mathbb{E}\tau_B < \infty$. Define $h(x) \equiv \mathbb{P}_x[\tau_A = \tau_B]$, if $x \notin B$. Let \mathbb{P} be the law of X .

Let us define a new measure, \mathbb{P}^h , on the space of paths as follows: If Y is a \mathcal{F}_t -measurable random variable, then

$$\mathbb{E}^h[Y|\mathcal{F}_0] = \frac{1}{h(X_0)}\mathbb{E}[h(X_t)Y|\mathcal{F}_0]. \quad (5.40)$$

Lemma 5.6.10 *With the notation above, if Y is a \mathcal{F}_{τ_B-1} -measurable function,*

$$\mathbb{E}_x^h[Y] = \mathbb{E}_x[Y|\tau_A = \tau_B]. \quad (5.41)$$

Proof. This is an application of the strong Markov property. We have by definition

$$\begin{aligned} \mathbb{E}_x^h[Y] &= \frac{1}{h(x)}\mathbb{E}_x[Yh(X_{\tau_B-1})] & (5.42) \\ &= \frac{1}{h(x)}\mathbb{E}_x[Y\mathbb{E}_x[\mathbb{1}_{\tau_A=\tau_B}|\mathcal{F}'_0](X_{\tau_B-1})] \\ &= \frac{1}{h(x)}\mathbb{E}_x[Y\mathbb{1}_{\tau_A=\tau_B}\theta_{\tau_B-1}] \\ &= \frac{1}{\mathbb{P}_x[\tau_A = \tau_B]}\mathbb{E}_x[Y\mathbb{1}_{\tau_A=\tau_B}] \\ &= \mathbb{E}_x[Y|\tau_A = \tau_B]. \end{aligned}$$

Here the first equality is just the definition of h and reproduces the form of the right-hand side of the strong Markov property; the second equality is the strong Markov property; the last equality uses that fact that the event $\{\tau_A = \tau_B\}$ depends only on what happens after $\tau_B - 1$, and so $\mathbb{1}_{\tau_A=\tau_B}\theta_{\tau_B-1} = \mathbb{1}_{\tau_A=\tau_B}$. \square

Let us now look at the transformed measure \mathbb{P}^h in the general case. The first thing to check is of course whether this defines in a consistent way a probability measure. Some thought shows that all that is to show for this is the following lemma.

Lemma 5.6.11 *Let Y be \mathcal{F}_s -measurable. Then, for any $t \geq s$,*

$$\mathbb{E}^h[Y|\mathcal{F}_0] \equiv \frac{1}{h(X_0)}\mathbb{E}[h(X_s)Y|\mathcal{F}_0] = \frac{1}{h(X_0)}\mathbb{E}[h(X_t)Y|\mathcal{F}_0]. \quad (5.43)$$

In particular, $\mathbb{P}^h[\Omega|\mathcal{F}_0] = 1$.

Proof. Just introduce a conditional expectation:

$$\mathbb{E}[h(X_t)Y|\mathcal{F}_0] = \mathbb{E}[\mathbb{E}[h(X_t)Y|\mathcal{F}_s]|\mathcal{F}_0] = \mathbb{E}[Y\mathbb{E}[h(X_t)|\mathcal{F}_s]|\mathcal{F}_0], \quad (5.44)$$

and use that $h(X_t)$ is a martingale

$$= \mathbb{E}[Yh(X_s)|\mathcal{F}_0],$$

from which the result follows. \square

This lemma shows in particular, why it is important that h be a harmonic function.

Now we turn to the question of whether the law \mathbb{P}^h is a Markov chain. To this end we turn to the martingale problem. We will show that there exists a generator, L^h , such that

$$M_t^h \equiv f(X_t) - f(X_0) - \sum_{s=0}^{t-1} (L^h f)(X_s) \quad (5.45)$$

is a martingale under the law \mathbb{E}^h , i.e. that, for $t > t'$,

$$\mathbb{E}^h[M_t^h | \mathcal{F}_{t'}] = M_{t'}^h. \quad (5.46)$$

Note first that, by definition

$$\begin{aligned} \mathbb{E}^h[M_t^h | \mathcal{F}_{t'}] &= \frac{1}{h(X_{t'})} \mathbb{E}[h(X_t)f(X_t) | \mathcal{F}_{t'}] - f(X_0) - \sum_{s=0}^{t'-1} (L^h f)(X_s) \\ &\quad - \sum_{s=t'}^{t-1} \frac{1}{h(X_{t'})} \mathbb{E}[h(X_s)L^h f(X_s) | \mathcal{F}_{t'}]. \end{aligned} \quad (5.47)$$

The middle terms are part of $M_{t'}^h$ and we must consider $\mathbb{E}[f(X_t)h(X_t) | \mathcal{F}_{t'}]$. This is done by applying the martingale problem for \mathbb{P} and the function fh . This yields

$$\mathbb{E}[f(X_t)h(X_t) | \mathcal{F}_{t'}] = f(X_{t'})h(X_{t'}) + \sum_{s=t'}^{t-1} \mathbb{E}[(L(fh))(X_s) | \mathcal{F}_{t'}]$$

Inserting this in (5.47) gives

$$\begin{aligned} \mathbb{E}^h[M_t^h | \mathcal{F}_{t'}] &= f(X_{t'}) - f(X_0) - \sum_{s=0}^{t'-1} (L^h f)(X_s) \\ &\quad + \frac{1}{h(X_{t'})} \sum_{s=t'}^{t-1} [\mathbb{E}[(L(fh))(X_s) | \mathcal{F}_{t'}] - \mathbb{E}[h(X_s)L^h f(X_s) | \mathcal{F}_{t'}]] \\ &= M_{t'}^h \\ &\quad + \frac{1}{h(X_{t'})} \sum_{s=t'}^{t-1} [\mathbb{E}[(L(fh))(X_s) | \mathcal{F}_{t'}] - \mathbb{E}[h(X_s)L^h f(X_s) | \mathcal{F}_{t'}]]. \end{aligned}$$

The second term will vanish if we choose L^h defined through $Lf(x) = h(x)^{-1}(L(hf))(x)$, i.e.

$$L^h f(x) \equiv \frac{1}{h(x)} \int P(x, dy) h(y) f(y) - f(x). \quad (5.48)$$

Hence we see that under \mathbb{P}^h , X solves the martingale problem corresponding to the generator L^h , and so is a Markov chain with transition kernel $P^h = L^h + 1$. The process X under \mathbb{P}^h is called the (Doob) h -transform of the original Markov process.

Exercise. As a simple example, consider a simple random walk on $\{-N, -N + 1, \dots, N\}$. Assume we want to condition this process on hitting $+N$ before $-N$. Then let

$$h(x) = \mathbb{P}_x[\tau_N = \tau_{\{N\} \cup \{-N\}}] = \mathbb{P}_x[\tau_N < \tau_{-N}].$$

Compute $h(x)$ and use this to compute the transition rates of the h -transformed walk? Plot the probabilities to jump down in the new chain!

5.7 Markov chains with countable state space

The setting of discrete time Markov chains does in some sense not go too well with general state spaces. In fact, in these cases, it is usually more appropriate to consider continuous time. In this chapter we will do three things. Here we provide some results on Markov chains with countable state space, in particular introduce the notions of recurrence and transience and discuss the existence and uniqueness of invariant distributions.

Much of the theory of Markov chains with countable state space is similar to the case of finite state space. In particular, the notions of communicating classes, irreducibility, and periodicity carry over. There are, however, important new concepts in the case when the state space is infinite. These are the notions of *recurrence* and *transience*. It will be useful to use a notation close to the matrix notation of finite chains. Thus we set

$$P(i, \{j\}) = p(i, j) \quad (5.49)$$

We will place us in the setting of an irreducible Markov chain, i.e. the all states in S communicate (i.e. for any $i, j \in S$, $\mathbb{P}_j(\tau_j < \infty) > 0$). We may also for simplicity assume that our chain is aperiodic. In the case of finite state space, we have seen that such chains are ergodic in the sense that there exists a unique invariant probability distribution,

and the marginal distributions at time t , converge to this distribution independently of the starting measure. Essentially this is true because the chain is trapped on the finite set. If S is infinite, a new phenomenon is possible: the chain may run “to infinity”.

Definition 5.7.1 Let X be an irreducible aperiodic Markov chain with countable state space S . Then:

- (i) X is called *transient*, if for any $i \in S$,

$$\mathbb{P}_i[\tau_i < \infty] < 1; \quad (5.50)$$

- (ii) X is called *recurrent*, if it is not transient.
 (iii) X is called *positive recurrent* or *ergodic*, if, for all $i \in S$,

$$\mathbb{E}_i T_i < \infty. \quad (5.51)$$

Remark 5.7.1 The notion of recurrence and transience can be defined for states rather than for the entire chain. In the case of irreducible and aperiodic chains, all states have the same characteristics.

Some simple consequences of the definition are the following.

Lemma 5.7.12 Let X be a Markov chain with countable state space, irreducible and aperiodic. Then X is transient, iff

$$\mathbb{P}_\ell [X_t = \ell, \text{i.o.}] = 0. \quad (5.52)$$

Positive recurrent chains are called ergodic, because they are ergodic in the same sense as finite Markov chains.

Lemma 5.7.13 Let X positive recurrent Markov chain with countable state space, S . Then, for any $j, \ell \in S$,

$$\mu(j) = \frac{\mathbb{E}_\ell [\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_t=j}]}{\mathbb{E}_\ell \tau_\ell}. \quad (5.53)$$

is the unique invariant probability distribution of X .

Proof. Define $\nu_\ell(j) = \mathbb{E}_\ell [\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_t=j}]$. We show first that ν is an invariant measure. Obviously, $1 = \sum_{m \in S} \mathbb{1}_{X_{\ell-1}=m}$, and hence, using the strong Markov property,

$$\begin{aligned}
\nu_\ell(j) &\equiv \mathbb{E}_\ell \left[\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_t=j} \right] = \mathbb{E}_\ell \left[\sum_{m \in S} \sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_t=j} \mathbb{1}_{X_{t-1}=m} \right] \\
&= \sum_{m \in S} \mathbb{E}_\ell \left[\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_{t-1}=m} \mathbb{P}[X_t = j | \mathcal{F}_{t-1}](m) \right] \\
&= \sum_{m \in S} \mathbb{E}_\ell \left[\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_{t-1}=m} \right] p(m, j) \\
&= \sum_{m \in S} \mathbb{E}_\ell \left[\sum_{t=1}^{\tau_\ell} \mathbb{1}_{X_t=m} \right] P(m, j) \\
&= \sum_{m \in S} \nu_\ell(m) \cdot P(m, j)
\end{aligned}$$

Thus μ_ℓ solves the invariance equation and thus is an invariant measure. It remains to show that ν_ℓ is normalizable. But

$$\sum_{j \in \Sigma} \nu_\ell(j) = \mathbb{E}_\ell t_\ell < \infty,$$

by assumption. Thus $\nu_\ell(j) / \sum_{i \in S} \nu_\ell(i) = \mu(j)$ is an invariant probability distribution.

Next we want to show uniqueness. Note first that for any irreducible Markov chain (with discrete state space) it holds that, if μ is an invariant measure and $\mu(i) = 0$, for some $i \in S$, then $\mu \equiv 0$. Namely, if for some $j, \mu(j) > 0$, then there exists t finite such that $P_{ji}^t > 0$, and $\mu(i) \geq \mu(j) P_{ji}^t > 0$, in contradiction to the hypothesis.

We will now actually show that ν_ℓ is the only invariant measure such that $\nu_\ell(\ell) = 1$ (which implies the desired uniqueness result immediately). To do so, we will show that for any other invariant measure, ν , such that $\nu(\ell) = 1$, we have that $\nu(j) \geq \nu_\ell(j)$ for all j . For then, $\nu - \nu_\ell$ is a positive invariant measure as well, and being zero in ℓ , must vanish identically. Hence $\nu = \nu_\ell$.

Now we clearly have that

$$\nu(i) = \sum_{j_1 \neq \ell} (p(j_1, i) \nu(j_1) + p(\ell, i)), \quad (5.54)$$

since $\nu(\ell) = 1$, by hypothesis. We want to think of $p(\ell, i)$ as

$$p(\ell, i) = \mathbb{E}_\ell (\mathbb{1}_{\tau_\ell \geq 1} \mathbb{1}_{X_s=i}).$$

Now iterate the same relation in the first term in (5.54). Thus

$$\begin{aligned}
\nu(i) &= \sum_{j_1, j_2 \neq \ell} p(j_2, j_1) p(j_1, i) \nu(j_2) + \sum_{j_1 \neq \ell} p(\ell, j_1) p(j_1, i) + \mathbb{E}_\ell (\mathbb{1}_{\tau_\ell \geq 1} \mathbb{1}_{X_1=i}) \\
&= \sum_{j_1, j_2 \neq \ell} p(j_2, j_1) p(j_1, i) \nu(j_2) + \mathbb{E}_\ell \left(\sum_{s=1}^{2 \wedge \tau_\ell} \mathbb{1}_{X_s=i} \right). \tag{5.55}
\end{aligned}$$

Further iteration yields for any $n \in \mathbb{N}$

$$\begin{aligned}
\nu(i) &= \sum_{j_1, j_2, \dots, j_n \neq \ell} p(j_n, j_{n-1}) \cdots p(j_2, j_1) p(j_1, i) \nu(j_n) + \mathbb{E}_\ell \left(\sum_{s=1}^{n \wedge \tau_\ell} \mathbb{1}_{X_s=i} \right) \\
&\geq \mathbb{E}_\ell \left(\sum_{s=1}^{n \wedge \tau_\ell} \mathbb{1}_{X_s=i} \right). \tag{5.56}
\end{aligned}$$

This implies $\nu(i) \geq \nu_\ell(i)$, as desired, and the proof is complete. \square

Corollary 5.7.14 *An ergodic Markov chain satisfies*

$$\mu(j) = \frac{1}{\mathbb{E}_j \tau_j}. \tag{5.57}$$

Proof. Just set $\ell = j$ in the definition of $\mu(j)$, and note that $\nu_j(j) = \mathbb{E}_j \sum_{t=1}^{\tau_j} \mathbb{1}_{X_t=j} = 1$. \square

6

Random walks and Brownian motion

The goal of this chapter is to introduce Brownian motion as a continuous time stochastic process with continuous paths and to explain its connection to random walks through Donsker's invariance principle.

6.1 Random walks

The innocent looking stochastic processes

$$S_n \equiv \sum_{i=1}^n X_i, \quad (6.1)$$

with $X_i, i \in \mathbb{N}$ iid random variables are generally called *random walks* and receive a considerable attention in probability theory. A special case is the so-called *simple random walk on \mathbb{Z}^d* , characterized by the fact that the random variables X_i take values in the set of \pm unit vectors in the lattice \mathbb{Z}^d . Consequently, $S_n \in \mathbb{Z}^d$, is a stochastic process with discrete state space. Obviously, S_n is a Markov chain, and, moreover, the coordinate processes, $S_n^\mu, \mu = 1, \dots, d$, are sub-, super-, or martingales, depending on whether $\mathbb{E}X_0^\mu$ is positive, negative, or zero.

Let us focus on the centered case, $\mathbb{E}X_1 = 0$. In this case we have seen that $Z_n \equiv n^{-1/2}S_n$ converges in distribution to a Gaussian random variable. By considering the process coordinate wise, it will also be enough to think about $d = 1$. We now want to extend this result to a convergence result on the level of stochastic process. That is, rather saying something about the position of the random walk at a time n , we want to trace the entire trajectories of the process and try give a description of their statistical properties in terms of some limiting stochastic process.

It is rather clear from the central limit theorem that we must consider a rescaling somewhat like

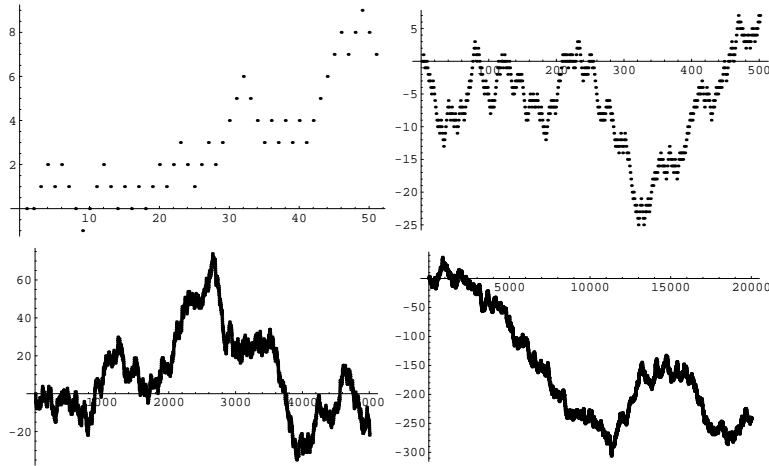


Fig. 6.1. Paths of $\sqrt{n}Z_n(t)$ for various values of n .

$$Z_n(t) \equiv n^{-1/2} \sum_{k=1}^{[tk]} X_k. \quad (6.2)$$

In that case we have from the central limit theorem, that for any $t \in (0, 1]$,

$$Z_n(t) \xrightarrow{\mathcal{D}} B_t,$$

($[x]$ denotes the lower integer part of x) where B_t is a centered Gaussian random variable with variance t . Moreover, for any finite collection of indices t_1, \dots, t_ℓ , define $Y_n(i) \equiv Z_n(t_i) - Z_n(t_{i-1})$. Then the random variables $Y_n(i)$ are independent and it is easy to see that they converge, as $n \uparrow \infty$, jointly to a family of independent centered Gaussian variables with variances $t_i - t_{i-1}$. This implies that the finite dimensional distributions of the processes $Z_n(t), t \in (0, 1]$, converge to the finite dimensional distributions of the Gaussian process with covariance $C(t, s) \equiv t \wedge s$, that we introduced in Section 3.3. and that we have preliminarily called Brownian motion.

We now want to go a step further and discuss the properties of the paths of our processes.

From looking at pictures, it is clear that the limiting process B_t should have rather continuous looking sample paths. In the way we defined $Z_n(t)$, for each n fixed these processes have discontinuous paths with




Fig. 6.2. Paths of $\sqrt{n}S_n(t)$ for various values of n .

jumps of size $n^{-1/2}$. This can easily be mended by considering the piecewise linear continuous interpolation of Z_n , to be written formally as

$$S_n(t) \equiv n^{-1/2} \sum_{k=0}^n \mathbb{I}_{k/n < t \leq (k+1)/n} \left(\left(t - \frac{k}{n} \right) S_{k+1} + \left(\frac{k+1}{n} - t \right) S_k \right), \quad (6.3)$$

We have now a sequence of processes with continuous sample paths, and we should hope that the limit shares this property. On the other hand, it is simple to see that the finite dimensional distributions of $S_n(t)$ have the same limit as those of $Z_n(t)$.

6.2 Construction of Brownian motion

Before stating the desired convergence result, we have to define and construct the limiting object, the Brownian motion.

Definition 6.2.1 A stochastic process $\{B_t \in \mathbb{R}^d, t \in \mathbb{R}_+\}$, defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$, is called a d -dimensional Brownian motion starting in 0, iff

- (o) $B_0 = 0$, a.s..
- (i) For any $p \in \mathbb{N}$, and any $t_0 < t_1 < \dots < t_p$, the random variables

$B_{t_1}, B_{t_2} - B_{t_1}, \dots, B_{t_p} - B_{t_{p-1}}$, are independent and each $B_{t_i} - B_{t_{i-1}}$ is a centered Gaussian r.v. with variance $t_i - t_{i-1}$.

(ii) For any $\omega \in \Omega$, the map $t \rightarrow B_t(\omega)$ is continuous.

We now must proceed to the construction of B . The first property can, as we have seen, be established with the help of Kolmogorov's theorem. The problem with this is that it constructs the process on the space $((\mathbb{R}^d)^{\mathbb{R}_+}, \mathcal{B}^{\mathbb{R}_+}(\mathbb{R}^d))$; but the second requirement, the continuity of the sample paths, is not a measurable property with respect to the product σ -algebra. Therefore, we have to proceed differently. In fact, we want to construct Brownian motion as a random variable with values in the space $C(\mathbb{R}_+, \mathbb{R}^d)$.

Theorem 6.2.1 *Brownian motion exists.*

Proof. We consider the case $d = 1$, the extension to higher dimensions will be trivial. We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which an infinite family of independent standard Gaussian random variables is defined. We now define the so-called Haar-functions, h_n^k on $[0, 1]$ via

$$h_0^0(t) \equiv 1 \tag{6.4}$$

$$h_n^k(t) \equiv 2^{n/2} [\mathbb{I}_{[(2k)2^{-n-1}, (2k+1)2^{-n-1})}(t) - \mathbb{I}_{[(2k+1)2^{-n-1}, (2k+2)2^{-n-1})}(t)]$$

for $k \in \{0, \dots, 2^n - 1\}$. The functions h_n^k , $n \in \mathbb{N}_0$, $k \in \{0, \dots, 2^n - 1\}$ form a complete orthonormal system of functions in $L^2([0, 1])$, as one may easily check. Now set

$$f_n^k(t) = \int_0^t h_n^k(u) du, \tag{6.5}$$

and set

$$B_t^{(n)} \equiv \sum_{m=0}^n \sum_{k=0}^{2^m-1} f_m^k(t) X_{m,k} \tag{6.6}$$

for $t \in [0, 1]$, where $X_{m,k}$ are our independent standard normal random variables. We will show that (i) the continuous functions $B^{(n)}(\omega)$ converge uniformly, almost surely, and hence to continuous functions, and (ii) that the covariances of $B^{(n)}$ converge to the correct limit. The limit, modified to be $B_t(\omega) \equiv 0$ when $B_t^{(n)}(\omega)$ does not converge to a continuous function, will then be BM on $[0, 1]$.

Let us now prove (i). The point here is that, of course, that the functions $f_n^k(t)$ are very small, namely,

$$|f_n^k(t)| \leq 2^{-(n+1)/2}.$$

Therefore,

$$\begin{aligned}
& \mathbb{P} \left[\sup_{0 \leq t \leq 1} |B_t^{(n)} - B_t^{(n-1)}| > a_n \right] \\
&= \mathbb{P} \left[\sup_{k=0}^{2^n-1} |X_{n,k}| > 2^{(n+1)/2} a_n \right] \\
&\leq 2^n \mathbb{P} \left[|X_{n,1}| > 2^{(n+1)/2} a_n \right] \\
&\leq 2^n \frac{e^{-a_n^2 2^n}}{\sqrt{2\pi a_n 2^{(n+1)/2}}} = \frac{2^{n/2} e^{-a_n^2 2^n}}{\sqrt{4\pi a_n}},
\end{aligned} \tag{6.7}$$

where we used the very useful bounds for Gaussian probabilities

$$\frac{1}{u\sqrt{2\pi}} e^{-u^2/2} (1 - 2u^{-2}) \leq \mathbb{P}[X > u] \leq \frac{1}{u\sqrt{2\pi}} e^{-u^2/2} \tag{6.8}$$

Now we are close to being done: Chose a sequence a_n such that $\sum_{n=0}^{\infty} a_n < \infty$ and

$$\sum_{n=1}^{\infty} \mathbb{P} \left[\sup_{0 \leq t \leq 1} |B_t^{(n)} - B_t^{(n-1)}| > a_n \right] < \infty.$$

Clearly, the choice $a_n = 2^{-n/4}$ will do. Then, by the Borel-Cantelli lemma,

$$\mathbb{P} \left[\sup_{0 \leq t \leq 1} |B_t^{(n)} - B_t^{(n-1)}| > a_n \text{ i.o.} \right] = 0,$$

which implies that the sequence $B^{(n)}$ converges uniformly in the interval $[0, 1]$ and since uniformly convergent sequences of continuous functions converge to continuous functions, $\lim_{n \uparrow \infty} B_t^{(n)} \equiv B_t(\omega)$ in $C([0, 1], \mathbb{R})$, for almost all ω .

To check (ii), we compute the covariances:

$$\begin{aligned}
\mathbb{E}B_t^{(n)}B_s^{(n)} &= \sum_{m=0}^n \sum_{k=0}^{2^m-1} \sum_{m'=0}^n \sum_{k'=0}^{2^{m'}-1} f_m^k(t)f_{m'}^{k'}(s)\mathbb{E}(X_{m,k}X_{m',k'}) \quad (6.9) \\
&= \sum_{m=0}^n \sum_{k=0}^{2^m-1} f_m^k(t)f_m^k(s) \\
&= \int_0^1 dudv \mathbb{I}_{[0,t]}(u)\mathbb{I}_{[0,s]}(v) \sum_{m=0}^n \sum_{k=0}^{2^m-1} h_m^k(u)h_m^k(v) \\
&\rightarrow \int_0^1 dudv \mathbb{I}_{[0,t]}(u)\mathbb{I}_{[0,s]}(v) \sum_{m=0}^{\infty} \sum_{k=0}^{2^m-1} h_m^k(u)h_m^k(v) \\
&= \int_0^1 du \mathbb{I}_{[0,t]}(u)\mathbb{I}_{[0,s]}(u) = t \wedge s.
\end{aligned}$$

due to the ortho-normality of the system h_n^k .

This provides B_t on $[0, 1]$. To construct B_t for $t \in (k, k+1]$, just take $k+1$ independent copies of the B we just constructed, via

$$B_t = \sum_{i=1}^k B_{i,1} + B_{k+1,t-k}.$$

Finally, to construct d -dimensional Brownian motion, take d independent copies of B_t , say $B_{t,1}, \dots, B_{t,d}$ and let e^μ , $\mu = 1, \dots, d$, be a orthonormal basis of \mathbb{R}^d . Then set

$$\widehat{B}_t \equiv \sum_{\mu=1}^d e^\mu B_{t,\mu}. \quad (6.10)$$

It is easily checked that this process is a Brownian motion in \mathbb{R}^d . This concludes the existence proof. \square

Having constructed the random variable B_t in $C(\mathbb{R}_+, \mathbb{R}^d)$, we can now define its distribution, the so-called *Wiener measure*.

For this is it useful to observe that

Lemma 6.2.2 *The smallest σ -algebra, \mathcal{C} , on $C(\mathbb{R}_+, \mathbb{R}^d)$ that makes all coordinate functions, $t \rightarrow w(t)$, measurable coincides with the Borel- σ -algebra, $\mathcal{B} \equiv \mathcal{B}(C(\mathbb{R}_+, \mathbb{R}^d))$, of the metrisable space $C(\mathbb{R}_+, \mathbb{R}^d)$ equipped with the topology of uniform convergence on compact sets.*

Proof. First, $\mathcal{C} \subset \mathcal{B}$ since all functions $t \rightarrow w(t)$ are continuous and hence measurable with respect to the Borel- σ -algebra \mathcal{B} . To prove that

$\mathcal{B} \subset \mathcal{C}$, we note that the topology of uniform convergence is equivalent to the metric topology relative to the metric

$$d(w, w') \equiv \sum_{n \in \mathbb{N}} 2^{-n} \sup_{0 \leq t \leq n} (|w(t) - w'(t)| \wedge 1). \quad (6.11)$$

We thus have to show that any ball with respect to this distance is measurable with respect to \mathcal{C} . But since w are continuous functions,

$$\sup_{t \in [0, n]} (|w(t) - w'(t)| \wedge 1) = \sup_{t \in [0, n] \cap \mathbb{Q}} (|w(t) - w'(t)| \wedge 1),$$

we see that e.g. the set $d(w, 0) < \rho$ is in fact in \mathcal{C} . \square

Note that by construction, the map $\omega \rightarrow B(\omega)$ is measurable, since the maps $\omega \rightarrow B_t(\omega)$ are measurable for all t , and by definition of \mathcal{C} , all coordinate maps $B \rightarrow B_t$ are measurable. Thus the following definition makes sense.

Definition 6.2.2 Let B_t a Brownian motion in \mathbb{R}^d defined on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. The d -dimensional Wiener measure is the probability measure on $(C(\mathbb{R}_+, \mathbb{R}^d), \mathcal{B}(C(\mathbb{R}_+, \mathcal{R}^d)))$ that is the image of \mathbb{P} under the map $\omega \rightarrow \{B_t, t \in \mathbb{R}^t\}$.

Note that uniqueness of the Wiener measure is a consequence of the Kolmogorov-Daniell theorem, since we have already seen that the finite-dimensional distributions are fixed by the prescription of the covariances, (i).

6.3 Donsker's invariance principle

We are now in the position to prove *Donsker's theorem*.

Theorem 6.3.3 Let X_i be independent, identically distributed random variables with mean zero and variance one. Let $S_n(t)$ be as defined in (6.3). Then

$$\{S_n(t), t \in [0, 1]\} \xrightarrow{\mathcal{D}} \{B_t, t \in [0, 1]\} \quad (6.12)$$

Proof. We will give an interesting proof of this theorem which will not use what we already know about finite dimensional distributions. For simplicity we consider the case $d = 1$ only. It will be based on the famous *Skorokhod embedding*. What this will do is to construct any desired random walk from a Brownian motion. This goes as follows: we assume that F is the common distribution function of our random

variables X_i , assumed to have finite second moments σ^2 . We now want to construct stopping times, T , for the Brownian motion, B , such that (i) the law of B_T is F , and (ii) $\mathbb{E}T = \sigma^2$. This is a little tricky. First, we construct a probability measure on $(-\mathbb{R}_+) \times \mathbb{R}_+$, from the restrictions, F_{\pm} , of F to the positive and negative axis:

$$\mu(da, db) \equiv \gamma(b-a)dF_-(a)dF_+(b). \quad (6.13)$$

where γ provides the normalization, i.e.

$$\gamma^{-1} = \int_0^{\infty} bdF_+(b) = - \int_{-\infty}^0 adF_-(a). \quad (6.14)$$

We need some elementary facts that follow easily once we know that B_t is a Markov chain with continuous time and generator $\Delta/2$:

Lemma 6.3.4 *Let $a < 0 < b$, and $\tau \equiv \inf\{t > 0 : B_t \notin (a, b)\}$. Then*

- (i) $\mathbb{P}(B_{\tau} = a) = \frac{b}{b-a}$;
- (ii) $\mathbb{E}\tau = |ab|$.

Proof. As we will discuss shortly, B_t is a martingale and let us anticipate that we Doob's optional stopping theorem also holds for Brownian motion. Then $0 = \mathbb{E}B_{\tau} = b\mathbb{P}[B_{\tau} = b] + a\mathbb{P}[B_{\tau} = a] = b + (a-b)\mathbb{P}[B_{\tau} = a]$, which gives (i). To prove (ii) consider

$$M_t = (B_t - a)(b - B_t) + t,$$

which is a martingale with $M_0 = -ba$. On the other hand (again assuming that we can use the optional stopping theorem,

$$\mathbb{E}M_0 = \mathbb{E}M_{\tau} = \mathbb{E}\tau + \mathbb{E}(B_{\tau} - a)(b - B_{\tau}) = \mathbb{E}\tau + 0.$$

which gives the claimed result. \square

The Skorokhod embedding is now constructed by choosing $\alpha < 0 < \beta$ at random from μ , and $T = \inf\{t > 0 : B_t \notin (\alpha, \beta)\}$. Then:

Theorem 6.3.5 *The law of B_T is F and $\mathbb{E}T = \sigma^2$.*

Proof. Let $b > 0$. Then

$$\mathbb{P}(B_T \in db) = \int_{-\infty}^0 \frac{-a}{b-a} \gamma(b-a)dF_+(b)dF_-(a) = dF_+(b).$$

Analogously, for $a > 0$, $\mathbb{P}(B_T \in da) = dF_-(a)$. This proves the first

assertion. Finally, by a simple computation,

$$\mathbb{E}T = \int_0^\infty \int_{-\infty}^0 \mu(da, db)|ab| = \int_{-\infty}^\infty x^2 F(dx) = \sigma^2.$$

This proves (ii). \square

Exercise. Construct the Skorokhod embedding for the simple random walk on \mathbb{Z} .

We can now define a sequence of stopping times $T_1 = T$, $T_2 = T_1 + T'_2, \dots$, where T'_i are independent and constructed in the same way as T on the Brownian motions $B_{T_{i-1}+t} - B_{T_{i-1}}$. Then it follows immediately from the preceding theorem that

Theorem 6.3.6 *The process $\tilde{S}_n, n \in \mathbb{N}$ where $\tilde{S}_n \equiv B_{T_n}$, for all $n \in \mathbb{N}$, has the same distribution as the process $S_n \equiv \sum_{i=1}^n X_i$, where X_i are iid with distribution F .*

The Skorokhod embedding now provides the means to prove Donsker's theorem. Namely, we will show that the process $\tilde{S}_n(t)$ converges *uniformly* to B_t in probability. This is possible, since it is coupled to B_t realisationwise, unlike the original $S_n(t)$ which would not know which particular $B_t(\omega)$ it should stick with. We will set $\sigma^2 = 1$.

Note first that by the continuity of Brownian motion, for any $\epsilon > 0$, we can find $\delta > 0$ such that

$$\mathbb{P}(\exists_{s,t \in [0,1]; |s-t| \leq \delta} |B_s - B_t| > \epsilon) \leq \epsilon/2.$$

Next, by the independence of the T'_i , and the law of large numbers,

$$\lim_{n \uparrow \infty} \frac{T_n}{n} = \mathbb{E}T = 1, \text{ a.s..}$$

Thus

$$\lim_{n \uparrow \infty} n^{-1} \sup_{k \leq n} |T_k - k| = 0 \text{ a.s..}$$

Therefore, there exists n_1 such that for all $n \geq n_1$,

$$\mathbb{P} \left[n^{-1} \sup_{k \leq n} |T_k - k| \geq \delta/3 \right] \leq \epsilon/2. \quad (6.15)$$

Finally, the interpolated process $\tilde{S}_n(t)$ will coincide for any $t \in [k/n, (k+1)/n]$ with $n^{-1/2}B_{nu}$ for some $nu \in [T_k, T_{k+1}]$. Hence, if n is large enough such that $n \geq 3/\delta$, for any $t \in [0, 1]$, there exists u such that $|u - t| \leq \delta$, such that $\tilde{S}_n(u) = n^{-1/2}B_{nu}$, on the event estimated in (6.15). Therefore,

$$\begin{aligned}
& \mathbb{P} \left[\sup_{0 \leq t \leq 1} \left| \tilde{S}_n(t) - n^{-1/2} B_{nt} \right| \geq \epsilon \right] \\
& \leq \mathbb{P} \left[\sup_{0 \leq t \leq 1} \left| \tilde{S}_n(t) - n^{-1/2} B_{nt} \right| \geq \epsilon, |T_k - k| \leq n\delta/3, \forall k \leq n \right] \\
& \quad + \mathbb{P} \left[n^{-1} \sup_{k \leq n} |T_k - k| \geq \delta/3 \right] \\
& \leq \mathbb{P} \left[\exists u, t \in [0, 1], |u-t| \leq \delta : \left| n^{-1/2} B_{un} - n^{-1/2} B_{nt} \right| \geq \epsilon \right] + \epsilon/2 \\
& = \mathbb{P} \left[\exists u, t \in [0, 1], |u-t| \leq \delta : |B_u - B_t| \geq \epsilon \right] + \epsilon/2 \leq \epsilon.
\end{aligned}$$

Here we used that $n^{-1/2}B_{nt}$ and B_t have the same law (this is an important fact, called Brownian scaling, which follows immediately from the fact that both processes have the same covariance).

This implies that the difference between $\tilde{S}_n(t)$ and $n^{-1/2}B_{nt}$ converges uniformly in $t \in [0, 1]$ to zero in probability. On the other hand, $\tilde{S}_n(t)$ has the same law as $S_n(t)$, and $n^{-1/2}B_{nt}$ has the same law as B_t . This implies weak convergence as claimed. \square

6.4 Martingale and Markov properties

Although we have not studied with full rigor the concepts of martingales and Markov processes in continuous time, Brownian motion is a good example to get provisionally acquainted with them. The nice thing here is that we know already that it has continuous paths, so that we need not worry about discontinuities; moreover, a path is determined by knowing it on a dense set of times, say the rationals, so we also need not worry about uncountabilities.

Proposition 6.4.7 *Brownian motion is a continuous time martingale, in the sense that, if \mathcal{F}_t is a filtration such that B_t is adapted, for any $s < t$,*

$$\mathbb{E}[B_t | \mathcal{F}_s] = B_s. \quad (6.16)$$

Proof. Of course we have not defined what a continuous time filtration is, but we will not worry at this moment, and just take \mathcal{F}_t as the σ -algebra generated by $\{B_s\}_{s \leq t}$. Now we know that $B_t = B_t - B_s + B_s$, where $B_t - B_s$ and B_s are independent. Thus

$$\mathbb{E}[B_t | \mathcal{F}_s] = \mathbb{E}[B_t - B_s + B_s | \mathcal{F}_s] = \mathbb{E}[B_t - B_s | \mathcal{F}_s] + \mathbb{E}[B_s | \mathcal{F}_s] = 0 + B_s,$$

as claimed. \square

Next we show that Brownian motion is also a Markov-process. As a definition of a continuous time Markov process, we adopt the obvious generalization of (3.21): A stochastic process with state space S and index set \mathbb{R}_+ is called a continuous time Markov process, if there exists a two-parameter family of probability kernels, $P_{s,t}$, satisfying the Chapman-Kolmogorov equations (3.25), such that for all $A \in \mathcal{B}$, $s < t \in \mathbb{R}_+$,

$$\mathbb{P}[B_t \in A | \mathcal{F}_s](\omega) = P_{s,t}(B_s(\omega), A), \text{ a.s.} \quad (6.17)$$

This definition may not sound abstract enough, because it stipulates that we search for the kernels $P_{s,t}$; one may replace this by saying that

$$\mathbb{P}[B_t \in A | \mathcal{F}_s] \quad (6.18)$$

is independent of the σ -algebras \mathcal{F}_r , for all $r < s$; or in other words, that $\mathbb{P}[B_t \in A | \mathcal{F}_s](\omega)$ is a function of $B_s(\omega)$, a.s.. You can see that we will have to worry a little bit about these definitions in general, but by the continuity of Brownian motion, we may just look at rational times and then no problem arises. We come to these things in the next course. We see that the two definitions are really the same, using the existence of regular conditional probabilities: namely, $P_{s,t}$ will be just the regular version of $\mathbb{P}[B_t \in A | \mathcal{F}_s]$.

Proposition 6.4.8 *Brownian motion in dimension d is a continuous time Markov process with transition kernel*

$$P_{s,t}(x, A) = \frac{1}{(2\pi(t-s))^{d/2}} \int_A \exp\left(-\frac{|y-x|^2}{2(t-s)}\right) dy. \quad (6.19)$$

Proof. The proof is next to trivial from the defining property (i) of Brownian motion and left as an exercise. \square

We now come, again somewhat informally, to the martingale problem associated with Brownian motion.

Theorem 6.4.9 *Let f be a bounded twice differentiable function, and let B_t be Brownian motion. Then*

$$M_t \equiv f(B_t) - f(B_0) - \frac{1}{2} \int_0^t \Delta f(B_s) ds \quad (6.20)$$

is a martingale.

Proof. We consider for simplicity only the case $d = 1$; the general case works the same way. We proceed as in the discrete time case.

$$\begin{aligned}
\mathbb{E}[M_{t+r}|\mathcal{F}_t] &= f(B_t) - f(B_0) - \frac{1}{2} \int_0^t f''(B_s) ds & (6.21) \\
&\quad + \mathbb{E}[f(B_{t+r}) - f(B_t)|\mathcal{F}_t] - \frac{1}{2} \int_0^r \mathbb{E}[f''(B_{t+s})|\mathcal{F}_t] ds \\
&= M_t + \frac{1}{\sqrt{2\pi r}} \int f(y) \exp\left(-\frac{(y-B_t)^2}{2r}\right) dy - f(B_t) \\
&\quad - \frac{1}{2} \int_0^r \frac{1}{\sqrt{2\pi s}} \int f''(y) \exp\left(-\frac{(y-B_t)^2}{2s}\right) dy ds \\
&= M_t
\end{aligned}$$

The last inequality holds since, using integration by parts

$$\begin{aligned}
&\frac{1}{\sqrt{2\pi s}} \int f''(y) \exp\left(-\frac{(y-x)^2}{2s}\right) dy & (6.22) \\
&= \int f(y) \frac{d^2}{dy^2} \frac{1}{\sqrt{2\pi s}} \exp\left(-\frac{(y-x)^2}{2s}\right) dy \\
&= \frac{1}{\sqrt{2\pi}} \int f(y) \left[-s^{-3/2} + (y-x)^2 s^{-5/2}\right] \exp\left(-\frac{(y-x)^2}{2s}\right) dy \\
&= 2 \int f(y) \frac{d}{ds} \frac{1}{\sqrt{2\pi s}} \exp\left(-\frac{(y-x)^2}{2s}\right) dy
\end{aligned}$$

Integrating the last expression over s yields

$$\frac{2}{\sqrt{2\pi r}} \int f(y) \exp\left(-\frac{(x-y)^2}{2r}\right) dy - f(x),$$

where we used that

$$\lim_{h \downarrow 0} \frac{2}{\sqrt{2\pi h}} \int f(y) \exp\left(-\frac{(x-y)^2}{2h}\right) dy = f(x).$$

Inserting this into (6.4.8) concludes the proof. \square

Note that we really used that the function

$$e(t, x) \equiv \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{\|x\|^2}{2t}\right) \quad (6.23)$$

satisfies the (parabolic) partial differential equation

$$\frac{\partial}{\partial t} e(x, t) = \frac{1}{2} \Delta e(x, t), \quad (6.24)$$

with the (singular) initial condition

$$e(x, t) = \delta(x), \quad (6.25)$$

(where δ here denotes the Dirac-delta function, i.e. for any bounded integrable function $\int \delta(x)f(x)dx = f(0)$). $e(t, x)$ is called the *heat kernel* associated to (one-dimensional) Brownian motion.

Remark 6.4.1 Let us note that if we rewrite (6.20) in the form

$$f(B_t) = f(B_0) + M_t + \frac{1}{2} \int_0^t \Delta f(B_s) ds, \quad (6.26)$$

it formally resembles the Itô-formula, or the Doob decomposition. The martingale M_t should then play the rôle of the stochastic integral, i.e. we would like to think of

$$M_t = \int_0^t \nabla f(B_s) \cdot dB_s.$$

It will turn out that this is indeed a correct interpretation if and that (6.26) is the Itô-formula for Brownian motion.

The preceding theorem justifies to call $L = \frac{\Delta}{2}$ the generator of Brownian motion, and to think of (6.21) as the associated martingale problem. The connection between Markov processes and potential theory, established for discrete time Markov processes, also carries over to Brownian motion; in this case, this links to the classical potential theory associated to the Laplace operator Δ .

6.5 Sample path properties

We have constructed Brownian motion on a space of continuous paths. What else can we say about the properties of these paths? The striking feature is that Brownian paths are almost surely *nowhere* differentiable!

The following theorem shows that it is not even Lipschitz continuous anywhere:

Theorem 6.5.10 *For almost all ω , $B(\omega)$ is nowhere Lipschitz continuous.*

Proof. Let $K > 0$ and define

$$A_n \equiv \{\omega \in \Omega : \exists_{s \in [0,1]} \forall_{|t-s| \leq 2/n} |B_t - B_s| \leq K|t-s|\}. \quad (6.27)$$

Clearly

$$A_n \subset \cup_{k=2}^n \{|B_{j/n} - B_{(j-1)/n}| \leq 4K/n, \text{ for } j \in \{k-1, k, k+1\}\}. \quad (6.28)$$

Now

$$\begin{aligned}\mathbb{P}[A_n] &\leq (n-2) (\mathbb{P}[|B_{1/n} - B_{(0)/n}| \leq 4K/n])^3 \\ &\leq (n-2) (\mathbb{P}[|B_1| \leq 4K/n])^3 \leq Cn^{-1/2}.\end{aligned}\quad (6.29)$$

Now $A_n \subset A_{n+1}$, and so for all n and all K ,

$$\mathbb{P}[A_n] \leq \lim_{k \uparrow \infty} \mathbb{P}[A_k] = 0. \quad (6.30)$$

Finally, by monotonicity of the Lipschitz property, it follows that

$$\mathbb{P}[\exists K < \infty A_n] \leq \sum_{K \in \mathbb{N}} \mathbb{P}[A_n(K)] = 0.$$

□

An important notion is that of the *quadratic variation*. Let $t_k^n \equiv (k2^{-n}) \wedge t$ and set

$$[B]_t^n \equiv \sum_{k=1}^{\infty} [B_{t_k^n} - B_{t_{k-1}^n}]^2. \quad (6.31)$$

Lemma 6.5.11 *With probability one, as $n \uparrow \infty$, $[B]_t^n \rightarrow t$, uniformly on compact intervals.*

Proof. Note that all sums over k contain only finitely many non-zero terms, and that all the summands in (6.36) are independent random variables, satisfying

$$\mathbb{E} \left(B_{t_k^n} - B_{t_{k-1}^n} \right)^2 = 2^{-n}, \quad (6.32)$$

$$\text{var} \left(\left(B_{t_k^n} - B_{t_{k-1}^n} \right)^2 \right) = 22^{-2n}. \quad (6.33)$$

Thus

$$\mathbb{E}[B]_t^n = t, \quad \text{var}([B]_t^n) = 2^{-n}t, \quad (6.34)$$

and thus

$$\lim_{n \uparrow \infty} [B]_t^n = t, \quad \text{a.s.} \quad (6.35)$$

By telescopic expansion,

$$\begin{aligned}
B_t^2 - B_0^2 &= \sum_{k=1}^{\infty} \left(B_{t_k^n}^2 - B_{t_{k-1}^n}^2 \right) \\
&= \sum_{k=1}^{\infty} \left(B_{t_k^n} - B_{t_{k-1}^n} \right) \left(B_{t_k^n} + B_{t_{k-1}^n} \right) \\
&= \sum_{k=1}^{\infty} 2B_{t_{k-1}^n} \left(B_{t_k^n} - B_{t_{k-1}^n} \right) + [B]_t^n.
\end{aligned} \tag{6.36}$$

Now set

$$V_t^n \equiv B_t^2 - [B]_t^n = \sum_{k=1}^{\infty} 2B_{t_{k-1}^n} \left(B_{t_k^n} - B_{t_{k-1}^n} \right). \tag{6.37}$$

One can check easily that for any n , V^n is a martingale. Then also

$$V_t^n - V_t^{n+1} = [B]_t^{n+1} - [B]_t^n \tag{6.38}$$

is a martingale. If we accept that Doob's \mathcal{L}^2 -inequality (Theorem 4.3.13) applies in the continuous martingale case as well, we get that, for any $T < \infty$,

$$\left\| \sup_{0 \leq t \leq T} ([B]_t^{n+1} - [B]_t^n) \right\|_2 \leq 2 \sup_{0 \leq t \leq T} \|[B]_t^{n+1} - [B]_t^n\|_2 = 2\sqrt{T2^{-n-1}}, \tag{6.39}$$

where the last inequality is obtained by explicit computation. This implies that $[B]_t^n$ converges uniformly on compact intervals. \square

Remark 6.5.1 Lemma 6.5.11 plays a crucial rôle in stochastic calculus. It justifies the claim that $d[B]_t = dt$. If we go with this into our “discrete Itô formula (Section 4.6), this means this justifies in a more precise way the step from Eq. (4.45) to Eq. (4.47).

Remark 6.5.2 The definition of the quadratic variation we adopt here via di-adic partitions is different from the “true” quadratic variation that would be

$$\sup \left\{ \sum_{k=1}^n [B_{t_k} - B_{t_{k-1}}]^2, n \in \mathbb{N}, 0 = t_0 < t_1 < \dots < t_n = 1 \right\},$$

which can be shown to be infinite almost surely (note that the choices of the t_i can be adapted to the specific realisation of the BM). The diadic version above is, however, important in the construction of stochastic integrals.

Remark 6.5.3 The fact that the quadratic variation of BM converges to t implies that the linear variation,

$$\sum_{k=1}^{\infty} |B_{t_k^n} - B_{t_{k-1}^n}|$$

is infinite on every interval. This means in particular that the length of a Brownian path between any times t, t' is infinite.

6.6 The law of the iterated logarithm

How precisely random phenomena can be controlled is witnessed by the so-called *law of the iterated logarithm (LIL)*. It states (not in its most general form) that

Theorem 6.6.12 Let $S_n = \sum_{i=1}^n X_i$, where X_i are independent identically distributed random variables with mean zero and variance σ^2 . Then

$$\mathbb{P} \left[\limsup_{n \uparrow \infty} \frac{S_n}{\sigma \sqrt{2n \ln \ln n}} = 1 \right] = 1. \quad (6.40)$$

Remark 6.6.1 Just as the CLT, the LIL has extensions to the case of non-identically distributed random variables. For a host of results, see [4], Chapter 10. Furthermore, there are extensions to the case of martingales, under similar conditions as for the CLT.

The nicest proof of this result passes through the analogous result for Brownian motion and then uses the Skorokhod embedding theorem. The proof below follows [11].

Thus we want to first prove:

Theorem 6.6.13 Let B_t be a one-dimensional Brownian motion. Then

$$\mathbb{P} \left[\limsup_{t \uparrow \infty} \frac{B_t}{\sqrt{2t \ln \ln t}} = 1 \right] = 1, \quad (6.41)$$

and

$$\mathbb{P} \left[\limsup_{t \downarrow 0} \frac{B_t}{\sqrt{2t \ln \ln(1/t)}} = 1 \right] = 1. \quad (6.42)$$

Proof. Note first that the two statements are equivalent since the two processes B_t and $tB_{1/t}$ have the same law (Exercise!).

We concentrate on (6.42). Set $h(t) = \sqrt{2t \ln \ln(1/t)}$. Basically, the idea is to use exponentially shrinking subsequences $t_n \equiv \theta^n$ in such a way

that the variables B_{t_n} are essentially independent. Then, for the lower bound, it is enough to show that along such a subsequence, the $h(t_n)$ is reached infinitely often: this will prove that the \limsup is as large as claimed. For the lower bound, one shows that along such subsequences, the threshold $h(t_n)$ is not exceeded, and then uses a maximum inequality for martingales to control the intermediate values of t .

We first show that $\limsup_{t \downarrow 0} B_t \leq 1$. For this we will assume that we can use Doob's submartingale inequality, Theorem (4.3.11) also in the continuous time case. Define

$$Z_t \equiv \exp\left(\alpha B_t - \frac{1}{2}\alpha^2 t\right). \quad (6.43)$$

A simple calculation shows that Z_t is a martingale (with $\mathbb{E}Z_t = 1$), and so

$$\mathbb{P}\left[\sup_{s \leq t} B_s - \alpha s/2 > \beta\right] = \mathbb{P}\left[\sup_{s \leq t} e^{\alpha B_s - \alpha^2 s/2} > e^{-\alpha\beta}\right] \leq \mathbb{E}e^{-\alpha\beta} Z_t = e^{\alpha\beta}.$$

Let $\theta, \delta \in (0, 1)$, and chose $t_n = \theta^n$, $\alpha_n = \theta^{-n}(1 + \delta)h(\theta^n)$, and $\beta_n = \frac{1}{2}h(\theta^n)$. Then

$$\mathbb{P}\left[\sup_{s \leq \theta^n} B_s - \alpha_n s/2 > \beta\right] \leq n^{-(1+\delta)}(\ln 1/\theta)^{-(1+\delta)},$$

since $\alpha_n \beta_n = (1 + \delta) \ln \ln \theta^n = (1 + \delta)(\ln n + \ln \ln \theta^{-1})$. Therefore, the Borel-Cantelli lemma implies that, almost surely, for all but finitely many values of n ,

$$\sup_{s \leq \theta^n} B_s - \frac{s}{2}(1 + \delta)\theta^{-n}h(\theta^n) \leq \frac{1}{2}h(\theta^n).$$

It follows that

$$\sup_{s \leq \theta^n} B_s \leq \frac{\theta^n}{2}(1 + \delta)\theta^{-n}h(\theta^n) + \frac{1}{2}(2 + \delta)h(\theta^n) \quad (6.44)$$

and so for any $\theta^{n+1} \leq t \leq \theta^n$,

$$B_t \leq \sup_{s \leq \theta^n} B_s \leq \frac{1}{2}(2 + \delta)\theta^{-1/2}h(t), \quad (6.45)$$

hence, almost surely,

$$\limsup_{t \downarrow 0} B_t/h(t) \leq \frac{1}{2}\theta^{-1/2}(2 + \delta). \quad (6.46)$$

Since this holds for any $\delta > 0$ and $\theta < 1$ almost surely, it holds along any countable subsequence $\delta_k \downarrow 0$, $\theta_k \uparrow 1$, almost surely, and

$$\limsup_{t \downarrow 0} B_t/h(t) \leq 1, \text{ a.s.} \quad (6.47)$$

To prove the converse inequality, consider the event

$$A_n \equiv \{B_{\theta^n} - B_{\theta^{n+1}} > (1 - \theta)^{1/2} h(\theta^n)\}.$$

The events are independent, and their probability can be bounded easily using (6.8):

$$\begin{aligned} \mathbb{P}[A_n] &= \frac{1}{\sqrt{2\pi(\theta^n(1-\theta))}} \int_{(1-\theta)^{1/2} h(\theta^n)}^{\infty} \exp\left(-\frac{x^2}{2\theta^n(1-\theta)}\right) dx \quad (6.48) \\ &= \frac{1}{\sqrt{2\pi}} \int_{\theta^{-n/2} h(\theta^n)}^{\infty} \exp\left(-\frac{x^2}{2}\right) dx \\ &\geq \frac{\exp(-\theta^{-n} h(\theta^n)^2/2)}{\sqrt{2\pi\theta^{-n/2} h(\theta^n)}} (1 - \theta^n h(\theta^n)^{-2}) \equiv \gamma_n. \end{aligned}$$

Now $\theta^{-n} h(\theta^n)^2 = 2 \ln n + 2 \ln \ln(1/\theta)$, and so

$$\gamma_n \geq C \frac{1}{n\sqrt{\ln n}},$$

so that $\sum_n \gamma_n = +\infty$; hence, the second Borel-Cantelli lemma implies that, with probability one, A_n happens infinitely often, i.e. for infinitely many n ,

$$B_{\theta^n} \geq (1 - \theta)^{1/2} h(\theta^n) + B_{\theta^{n+1}}.$$

Now, the upper bound (6.47) also holds for $-B_t$, so that, almost surely, for all but finitely many n ,

$$B_{\theta^{n+1}} \geq -h(\theta^{n+1}).$$

But by some simple estimates,

$$h(\theta^{n+1}) = \theta^{1/2} h(\theta^n) \sqrt{\frac{\ln \ln(\theta^{-n}\theta^{-1})}{\ln \ln(\theta^{-n})}} \leq \theta^{-1/2} h(\theta^n) (1 + O(\ln \theta^{-1}/n)),$$

so that, for infinitely many n ,

$$B_{\theta^n} \geq \left((1 - \theta)^{1/2} - 2\theta^{1/2}\right) h(\theta^n).$$

This implies that

$$\limsup_n B_{\theta^n}/h(\theta^n) \geq \left((1 - \theta)^{1/2} - 2\theta^{1/2}\right), \quad (6.49)$$

for all $\theta > 0$; hence,

$$\limsup_t B_t/h(t) \geq 1, \quad (6.50)$$

which completes the proof. \square

From the LIL for Brownian motion one can prove the LIL for random walk using the Skorokhod embedding.

Proof. (of Theorem 6.6.12) From the construction of the Skorokhod embedding, we know that we may choose $S_n(\omega) = B_{T_n}(\omega)$. The strong law of large numbers implies that $T_n/n \rightarrow 1$, a.s., and so also $h(T_n)/h(n) \rightarrow 1$, a.s.. Thus the upper bound follows trivially:

$$\limsup_n \frac{S_n}{h(n)} = \limsup_n \frac{B_{T_n}}{h(T_n)} \leq \limsup_t \frac{B_t}{h(t)} = 1. \quad (6.51)$$

To prove the complementing lower bound, note that by Kolmogorov's 0-1-law, $\rho \equiv \limsup_n \frac{S_n}{h(n)}$ is almost surely a constant (since the limsup is measurable with respect to the tail- σ -algebra. Assume $\rho < 1$; then, there exists $n_0 < \infty$, such that for all $n \geq n_0$, $\frac{B_{T_n}}{h(T_n)} < \rho$. We will show that this leads to a contradiction with (6.41) of Theorem 6.6.13. To show this, we must show that the Brownian motion cannot rise too far in the intervals $[T_n, T_{n+1}]$. But recall that T_{n+1} is defined as the stopping time at the random interval $[\alpha, \beta]$ of the Brownian motion B_t . We will want to show that in no such interval can the BM climb by more than $\epsilon\sqrt{2n \ln \ln n}$. An explicit computation shows that

$$\phi(x) \equiv \mathbb{P} \left[\sup_{t \leq T_1} B_t > x \right] = \gamma \int_{-\infty}^0 dF_-(a) \int_x^\infty dF_+(b) (b-a) \frac{-a}{x-a}, \quad (6.52)$$

where the ratio $\frac{-a}{x-a}$ is the probability that the BM reaches x before a (i.e. before T_1) (the logic of the formula is that for B_t to exceed x before T_1 , the random variable β must be larger than x , and then B_t may not hit the lower boundary before reaching x). Now we will be done by Borel-Cantelli, if

$$\sum_n \phi(\epsilon\sqrt{2n \ln \ln n}) < \infty,$$

or in fact the stronger but simpler condition

$$\sum_n \phi(\epsilon\sqrt{n}) < \infty \quad (6.53)$$

holds for all $\epsilon > 0$. For than, except finitely often,

$$\sup_{T_n < t < T_{n+1}} B_t \leq h(n)(\rho + \epsilon),$$

which implies

$$\limsup_t \frac{B_t}{h(t)} < \rho + \epsilon,$$

which can be made smaller than 1, thus contradicting the result for BM.

We are left we checking (6.53). We may decompose ϕ as

$$\begin{aligned} \Phi(x) &= \gamma \int_{-\infty}^0 dF_-(a) \int_x^\infty dF_+(b)(b-x) \frac{-a}{x-a} \\ &\quad + \gamma \int_{-\infty}^0 dF_-(a) \int_x^\infty dF_+(b)|a| \equiv \phi_1(x) + \phi_2(x). \end{aligned} \quad (6.54)$$

Now $\sum_n \phi_2(\epsilon\sqrt{n}) < \infty$ if $\int_0^\infty \phi_2(\epsilon\sqrt{x}) < \infty$. Recalling the formula for γ , (6.14), we see that

$$\int_0^\infty \phi_2(\epsilon\sqrt{x}) dx = \int_0^\infty (1-F(\epsilon\sqrt{x})) dx = \epsilon^{-2} \int_0^\infty (1-F(t)) t dt < \mathbb{E}X^2 < \infty.$$

To deal with ϕ_1 , use that $x-a > x$, and then as before

$$\phi_1(x) \leq x^{-1} \int_x^\infty (b-x) dF_+(x)$$

Comparing the sum to an integral, we must check the finiteness of

$$\int dx \frac{1}{\epsilon\sqrt{x}} \int_{\epsilon\sqrt{x}}^\infty dF_+(b)(b-\epsilon\sqrt{x}) = 2\epsilon^{-2} \int dt \int_t^\infty dF_+(b)(b-t),$$

which again hold since F has finite second moment. This concludes the proof. \square

On can show more than what we did. For one thing, not only is $\limsup_t B_t/h(t) = +1$ (and hence by symmetry $\liminf_t B_t/h(t) = -1$, a.s., it is also true that the set of limit points of the process $B_t/h(t)$ is the entire interval $[-1, 1]$; i.e., for any $a \in [-1, 1]$, there exist subsequences t_n , such that $\lim_n B_{t_n}/h(t_n) = a$.

Bibliography

- [1] Heinz Bauer. *Probability theory and elements of measure theory*. Academic Press Inc. [Harcourt Brace Jovanovich Publishers], London, 1981. Second edition of the translation by R. B. Burckel from the third German edition, Probability and Mathematical Statistics.
- [2] P. Billingsley. *Probability and measure*. Wiley Series in Probability and Mathematical Statistics. John Wiley & Sons Inc., New York, 1995.
- [3] Anton Bovier. *Statistical mechanics of disordered systems*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge, 2006.
- [4] Y.S. Chow and H. Teicher. *Probability theory*. Springer Texts in Statistics. Springer-Verlag, New York, 1997.
- [5] J. L. Doob. *Measure theory*, volume 143 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 1994.
- [6] H.-O. Georgii. *Gibbs measures and phase transitions*, volume 9 of *de Gruyter Studies in Mathematics*. Walter de Gruyter & Co., Berlin, 1988.
- [7] O. Kallenberg. *Random measures*. Akademie-Verlag, Berlin, 1983.
- [8] M. Ledoux and M. Talagrand. *Probability in Banach spaces*, volume 23 of *Ergebnisse der Mathematik und ihrer Grenzgebiete (3) [Results in Mathematics and Related Areas (3)]*. Springer-Verlag, Berlin, 1991.
- [9] Michel Ledoux. *The concentration of measure phenomenon*, volume 89 of *Mathematical Surveys and Monographs*. American Mathematical Society, Providence, RI, 2001.
- [10] M. M. Rao. *Measure theory and integration*, volume 265 of *Monographs and Textbooks in Pure and Applied Mathematics*. Marcel Dekker Inc., New York, second edition, 2004.
- [11] L. C. G. Rogers and David Williams. *Diffusions, Markov processes, and martingales. Vol. 1*. Cambridge Mathematical Library. Cambridge University Press, Cambridge, 2000. Foundations, Reprint of the second (1994) edition.
- [12] B. Simon. *The statistical mechanics of lattice gases. Vol. I*. Princeton Series in Physics. Princeton University Press, Princeton, NJ, 1993.

- [13] Daniel W. Stroock and S. R. Srinivasa Varadhan. *Multidimensional diffusion processes*, volume 233 of *Grundlehren der Mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1979.
- [14] Michael E. Taylor. *Measure theory and integration*, volume 76 of *Graduate Studies in Mathematics*. American Mathematical Society, Providence, RI, 2006.

Index

- 0 – 1-law
 - Kolmogorov's, 64
- L^p -space, 20
- \mathcal{L}^p -space, 20
- Π -system, 5
- λ -system, 5
- σ -additive, 6
- σ -algebra, 2
 - Borel, 3
 - generated, 3

- absolute continuity, 24
- absolutely integrable, 17
- adapted process, 58
- algebra, 1

- Baire σ -algebra, 16
- Baire function, 16
- Borel measure, 11, 12
- Borel- σ -algebra, 3
- Brownian motion, 51, 100
 - construction, 101

- Carathéodory's theorem, 7
- Cauchy sequence, 4
- central limit theorem, 73
- Chapman-Kolmogorov equations, 54
- class, 1
- closed, 3
- concentration of measure, 77
- conditional expectation, 31
- conditional probability, 39
- cylinder set, 43

- density, 24
- Dirichlet problem, 90
- Donsker's theorem, 105
- Doob decomposition, 69, 87
- Doob's supermartingale inequality, 82
- Dynkin's theorem, 5

- equivalence (of measures), 24
- ergodic, 96
- essential supremum, 25

- Fatou's lemma, 18
- filtered space, 57
- filtration, 57
- filtrations
 - natural, 58
- Fubini-Lebesgue theorem, 23
- Fubini-Tonnelli theorem, 23

- Gaussian density, 49
- Gaussian process, 48
- generator, 86
- Gibbs measure, 55

- Hölder inequality, 20
- Haar functions, 102
- harmonic function, 89
- heat kernel, 111

- independent random variables, 47
- index set, 41
- indicator function, 1
- induced measure, 14
- inequality
 - Hölder, 20
 - inequality, 20
 - Jensen, 21
 - maximum, 70
 - upcrossing, 61
- initial distribution, 53
- inner regular, 12
- integrable, 17
- invariance principle, 105
- invariant
 - distribution, 84, 96
 - measure, 84
- invariant measure, 96

- Ising model, 55
- Itô formula, 111
- Jensen inequality, 21
- Kolmogorov's 0 – 1-law, 64
- Kolmogorov's LLN, 66
- Kolmogorov-Daniell theorem, 45
- Laplace transform, 50
- law of large numbers, 66
- law of the iterated logarithm, 114
- Lebesgue decomposition theorem, 29
- Lebesgue integral, 17
- Lebesgue measure, 12
- Lebesgue's dominated convergence theorem, 18
- Levy's downward theorem, 64
- Lindeberg condition, 73, 75
- local specification, 55
- Lousin space, 4
- marginals, 44
 - finite dimensional, 44
- Markov chain, 52
- Markov inequality
 - exponential, 78
- Markov process, 52, 83
 - stationary, 83
- Markov property, 54
 - strong, 85
- martingale, 57
 - convergence theorem, 61
 - problem, 87
 - sub, 58
 - super, 58
 - transform, 59
- martingale difference sequence, 60
- maximum inequality, 67, 70
- maximum principle, 89
- measurable space, 2
- measure, 2
 - σ -finite, 2
 - finite, 2
 - invariant, 96
 - probability, 2
 - Wiener, 104
- measure space, 2
- metric, 4
- metric space, 4
- Minlowski inequality, 20
- monotone convergence theorem, 17
- open, 3
- outer measure, 7
- outer regular, 12
- Polish space, 4
- positive recurrent, 96
- pre- T - σ -algebra, 79
- previsible process, 59
- probability
 - regular conditional, 40
- probability measure, 2
- process
 - adapted, 58
- product space, 5
- product topology, 5
- quadratic variation, 112
- Radon measure, 13
- Radon-Nikodým derivative, 25
- Radon-Nikodým theorem, 25
- random variable, 14
- random walk, 99
- recurrence, 95
- recurrent
 - positive, 96
- regular conditional probability, 40, 55
- sample path, 42
- set function, 6
- simple function, 16
- Skorokhod embedding, 106
- space, 1
 - complete, 4
 - filtered, 57
 - Hausdorff, 3
 - Lousin, 4
 - measurable, 2
 - metric, 4
 - Polish, 4
 - topological, 2
- special cylinder, 43
- state space, 41
- stationary process, 83
- statistical mechanics, 55
- stochastic integral, 59
- stochastic process, 41
- strong Markov property, 85
- stopping time, 78
- time
 - continuous, 41
 - discrete, 41
- topological space, 2
- topology, 2
- transience, 95
- transition kernel, 52
 - stationary, 83
- uniform integrability, 18, 62
- upcrossing, 60

- inequality, 61
- version (of conditional expectation), 31
- Wiener measure, 104