

Continuous martingales and stochastic calculus

Alison Etheridge

March 11, 2018

Contents

1	Introduction	3
2	An overview of Gaussian variables and processes	5
2.1	Gaussian variables	5
2.2	Gaussian vectors and spaces	7
2.3	Gaussian spaces	9
2.4	Gaussian processes	10
2.5	Constructing distributions on $(\mathbb{R}^{[0,\infty)}, \mathcal{B}(\mathbb{R}^{[0,\infty)}))$	11
3	Brownian Motion and First Properties	12
3.1	Definition of Brownian motion	12
3.2	Wiener Measure	18
3.3	Extensions and first properties	18
3.4	Basic properties of Brownian sample paths	19
4	Filtrations and stopping times	23
4.1	Stopping times	24
5	(Sub/super-)Martingales in continuous time	27
5.1	Definitions	27
5.2	Doob's maximal inequalities	28
5.3	Convergence and regularisation theorems	30
5.4	Martingale convergence and optional stopping	33
6	Continuous semimartingales	36
6.1	Functions of finite variation	36
6.2	Processes of finite variation	38
6.3	Continuous local martingales	40
6.4	Quadratic variation of a continuous local martingale	41
6.5	Continuous semimartingales	47

7	Stochastic Integration	48
7.1	Stochastic integral w.r.t. L^2 -bounded martingales	48
7.2	Intrinsic characterisation of stochastic integrals using the quadratic co-variation	51
7.3	Extensions: stochastic integration with respect to continuous semimartingales	54
7.4	Itô's formula and its applications	57
A	The Monotone Class Lemma/ Dynkin's $\pi - \lambda$ Theorem	62
B	Review of some basic probability	63
B.1	Convergence of random variables.	63
B.2	Convergence in probability.	63
B.3	Uniform Integrability	64
B.4	The Convergence Theorems	64
B.5	Information and independence.	66
B.6	Conditional expectation.	66
C	Convergence of backwards discrete time supermartingales	67
D	A very short primer in functional analysis	67
E	Normed vector spaces	67
F	Banach spaces	69

These notes are based heavily on notes by Jan Obłój from last year's course, and the book by Jean-François Le Gall, *Brownian motion, martingales, and stochastic calculus*, Springer 2016. The first five chapters of that book cover everything in the course (and more). Other useful references (in no particular order) include:

1. D. Revuz and M. Yor, *Continuous martingales and Brownian motion*, Springer (Revised 3rd ed.), 2001, Chapters 0-4.
2. I. Karatzas and S. Shreve, *Brownian motion and stochastic calculus*, Springer (2nd ed.), 1991, Chapters 1-3.
3. R. Durrett, *Stochastic Calculus: A practical introduction*, CRC Press, 1996. Sections 1.1 - 2.10.
4. F. Klebaner, *Introduction to Stochastic Calculus with Applications*, 3rd edition, Imperial College Press, 2012. Chapters 1, 2, 3.1–3.11, 4.1-4.5, 7.1-7.8, 8.1-8.7.

5. J. M. Steele, *Stochastic Calculus and Financial Applications*, Springer, 2010. Chapters 3 - 8.
6. B. Oksendal, *Stochastic Differential Equations: An introduction with applications*, 6th edition, Springer (Universitext), 2007. Chapters 1 - 3.
7. S. Shreve, *Stochastic calculus for finance, Vol 2: Continuous-time models*, Springer Finance, Springer-Verlag, New York, 2004. Chapters 3 - 4.

The appendices gather together some useful results that we take as known.

1 Introduction

Our topic is part of the huge field devoted to the study of *stochastic processes*. Since first year, you've had the notion of a *random variable*, X say, on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and taking values in a state space (E, \mathcal{E}) . $X : \Omega \rightarrow E$ is just a measurable mapping, so that for each $e \in \mathcal{E}$, $X^{-1}(e) \in \mathcal{F}$ and so, in particular, we can assign a probability to the event that $X \in e$. Often (E, \mathcal{E}) is just $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ (where $\mathcal{B}(\mathbb{R})$ is the Borel sets on \mathbb{R}) and this just says that for each $x \in \mathbb{R}$ we can assign a probability to the event $\{X \leq x\}$.

Definition 1.1 A stochastic process, indexed by some set \mathcal{T} , is a collection of random variables $\{X_t\}_{t \in \mathcal{T}}$, defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and taking values in a common state space (E, \mathcal{E}) .

For us, \mathcal{T} will generally be either $[0, \infty)$ or $[0, T]$ and we think of X_t as a random quantity that evolves with time.

When we model deterministic quantities that evolve with (continuous) time, we often appeal to ordinary differential equations as models. In this course we develop the 'calculus' necessary to develop an analogous theory of *stochastic (ordinary) differential equations*.

An ordinary differential equation might take the form

$$dX(t) = a(t, X(t))dt,$$

for a suitably nice function a . A stochastic equation is often formally written as

$$dX(t) = a(t, X(t))dt + b(t, X(t))dB_t,$$

where the second term on the right models 'noise' or fluctuations. Here $(B_t)_{t \geq 0}$ is an object that we call Brownian motion. We shall consider what appear to be more general driving noises, but the punchline of the course is that under rather general conditions they can all be built from Brownian motion. Indeed if we added possible (random) 'jumps' in $X(t)$ at times given by a Poisson process, we'd capture essentially the most general theory. We are not going to allow jumps, so we'll be thinking of settings in which our stochastic equation has a continuous solution $t \mapsto X_t$, and Brownian motion will be a fundamental object.

It requires some measure-theoretic niceties to make sense of all this.

Definition 1.2 The mapping $t \mapsto X_t(\omega)$ for a fixed $\omega \in \Omega$, represents a realisation of our stochastic process, called a sample path or trajectory. We shall assume that

$$(t, \omega) \mapsto X_t(\omega) : ([0, \infty) \times \Omega, \mathcal{B}([0, \infty)) \otimes \mathcal{F}) \mapsto (\mathbb{R}, \mathcal{B}(\mathbb{R}))$$

is measurable (i.e. $\forall A \in \mathcal{B}(\mathbb{R}), \{(t, \omega) : X_t \in A\}$ is in the product σ -algebra $\mathcal{B}([0, \infty)) \otimes \mathcal{F}$). Our stochastic process is then said to be measurable.

Definition 1.3 Let X, Y be two stochastic processes defined on a common probability space $(\Omega, \mathcal{F}, \mathbb{P})$.

- i. We say that X is a modification of Y if, for all $t \geq 0$, we have $X_t = Y_t$ a.s.;
- ii. We say that X and Y are indistinguishable if $\mathbb{P}[X_t = Y_t, \text{ for all } 0 \leq t < \infty] = 1$.

If X and Y are modifications of one another then, in particular, they have the same finite dimensional distributions,

$$\mathbb{P}[(X_{t_1}, \dots, X_{t_n}) \in A] = \mathbb{P}[(Y_{t_1}, \dots, Y_{t_n}) \in A]$$

for all measurable sets A , but indistinguishability is a much stronger property.

Indistinguishability takes the *sample path* as the basic object of study, so that we could think of $(X_t(\omega), t \geq 0)$ (the path) as a random variable taking values in the space $E^{[0, \infty)}$ (of all possible paths). We are abusing notation a little by identifying the sample space Ω with the state space $E^{[0, \infty)}$. This state space then has to be endowed with a σ -algebra of measurable sets. For definiteness, we take real-valued processes, so $E = \mathbb{R}$.

Definition 1.4 An n -dimensional cylinder set in $\mathbb{R}^{[0, \infty)}$ is a set of the form

$$C = \{\omega \in \mathbb{R}^{[0, \infty)} : (\omega(t_1), \dots, \omega(t_n)) \in A\}$$

for some $0 \leq t_1 < t_2 \dots < t_n$ and $A \in \mathcal{B}(\mathbb{R}^n)$.

Let \mathcal{C} be the family of all finite-dimensional cylinder sets and $\mathcal{B}(\mathbb{R}^{[0, \infty)})$ the σ -algebra it generates. This is small enough to be able to build probability measures on $\mathcal{B}(\mathbb{R}^{[0, \infty)})$ using Carathéodory's Theorem (see B8.1). On the other hand $\mathcal{B}(\mathbb{R}^{[0, \infty)})$ only contains events which can be defined using at most countably many coordinates. In particular, the event

$$\{\omega \in \mathcal{B}(\mathbb{R}^{[0, \infty)}) : \omega(t) \text{ is continuous}\}$$

is *not* $\mathcal{B}(\mathbb{R}^{[0, \infty)})$ -measurable.

We will have to do some work to show that many processes can be assumed to be continuous, or right continuous. The sample paths are then fully described by their values at times $t \in \mathbb{Q}$, which will greatly simplify the study of quantities of interest such as $\sup_{0 \leq s \leq t} |X_s|$ or $\tau_0(\omega) = \inf\{t \geq 0 : X_t(\omega) > 0\}$.

A monotone class argument (see Appendix A) will tell us that a probability measure on $\mathcal{B}(\mathbb{R}^{[0, \infty)})$ is characterised by its finite-dimensional distributions - so if we can take continuous paths, then we only need to find the probabilities of cylinder sets to characterise the distribution of the process.

2 An overview of Gaussian variables and processes

Brownian motion is a special example of a Gaussian process - or at least a version of one that is assumed to have continuous sample paths. In this section we give an overview of Gaussian variables and processes, but in the next we shall give a direct construction of Brownian motion, due to Lévy, from which continuity of sample paths is an immediate consequence.

2.1 Gaussian variables

Definition 2.1 A random variable X is called a centred standard Gaussian, or standard normal, if its distribution has density

$$p_X(x) = \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} \quad x \in \mathbb{R}$$

with respect to Lebesgue measure. We write $X \sim \mathcal{N}(0, 1)$.

It is elementary to calculate its Laplace transform:

$$\mathbb{E}[e^{\lambda X}] = e^{\frac{\lambda^2}{2}}, \quad \lambda \in \mathbb{R},$$

or extending to complex values the characteristic function

$$\mathbb{E}[e^{i\xi X}] = e^{-\frac{\xi^2}{2}}, \quad \xi \in \mathbb{R}.$$

We say Y has Gaussian (or normal) distribution with mean m and variance σ^2 , written $Y \sim \mathcal{N}(m, \sigma^2)$, if $Y = \sigma X + m$ where $X \sim \mathcal{N}(0, 1)$. Then

$$\mathbb{E}[e^{i\xi Y}] = \exp\left(im\xi - \frac{\sigma^2 \xi^2}{2}\right), \quad \xi \in \mathbb{R},$$

and if $\sigma > 0$, the density on \mathbb{R} is

$$p_Y(x) = \frac{1}{\sqrt{2\pi\sigma}} e^{-\frac{(x-m)^2}{2\sigma^2}}, \quad x \in \mathbb{R}.$$

We think of a *constant* ‘random’ variable as being a *degenerate Gaussian*. Then the space of Gaussian variables (resp. distributions) is closed under convergence in probability (resp. distribution).

Proposition 2.2 Let (X_n) be a sequence of Gaussian random variables with $X_n \sim \mathcal{N}(m_n, \sigma_n^2)$, which converges in distribution to a random variable X . Then

- (i) X is also Gaussian, $X \sim \mathcal{N}(m, \sigma^2)$ with $m = \lim_{n \rightarrow \infty} m_n$, $\sigma^2 = \lim_{n \rightarrow \infty} \sigma_n^2$; and
- (ii) if $(X_n)_{n \geq 1}$ converges to X in probability, then the convergence is also in L^p for all $1 \leq p < \infty$.

Proof. Convergence in distribution is equivalent to saying that the characteristic functions converge:

$$\mathbb{E} \left[e^{i\xi X_n} \right] = \exp(im_n \xi - \sigma_n^2 \xi^2 / 2) \longrightarrow \mathbb{E} \left[e^{i\xi X} \right], \quad \xi \in \mathbb{R}. \quad (1)$$

Taking absolute values we see that the sequence $\exp(-\sigma_n^2 \xi^2 / 2)$ converges, which in turn implies that $\sigma_n^2 \rightarrow \sigma^2 \in [0, \infty)$ (where we ruled out the case $\sigma_n \rightarrow \infty$ since the limit has to be the absolute value of a characteristic function and so, in particular, has to be continuous). We deduce that

$$e^{im_n \xi} \longrightarrow e^{\frac{1}{2} \sigma^2 \xi^2} \mathbb{E} \left[e^{i\xi X} \right], \quad \xi \in \mathbb{R}.$$

We now argue that this implies that the sequence m_n converges to some finite m . Suppose first that the sequence $\{m_n\}_{n \geq 1}$ is bounded, and consider any two convergent subsequences, converging to m and m' say. Then rearranging (1) yields $m = m'$ and so the sequence converges.

Now suppose that $\limsup_{n \rightarrow \infty} m_n = \infty$ (the case $\liminf_{n \rightarrow \infty} m_n = -\infty$ is similar). There exists a subsequence $\{m_{n_k}\}_{k \geq 1}$ which tends to infinity. Given M , for k large enough that $m_{n_k} > M$,

$$\mathbb{P}[X_{n_k} \geq M] \geq \mathbb{P}[X_{n_k} \geq m_{n_k}] = \frac{1}{2},$$

and so using convergence in distribution

$$\mathbb{P}[X \geq M] \geq \limsup_{k \rightarrow \infty} \mathbb{P}[X_{n_k} \geq M] \geq \frac{1}{2}, \quad \text{for any } M > 0.$$

This is clearly impossible for any fixed random variable X and gives us the desired contradiction. This completes the proof of (i).

To show (ii), observe that the convergence of σ_n and m_n implies, in particular, that

$$\sup_n \mathbb{E} \left[e^{\theta X_n} \right] = \sup_n e^{\theta m_n + \theta^2 \sigma_n^2 / 2} < \infty, \quad \text{for any } \theta \in \mathbb{R}.$$

Since $\exp(|x|) \leq \exp(x) + \exp(-x)$, this remains finite if we take $|X_n|$ instead of X_n . This implies that $\sup_n \mathbb{E}[|X_n|^p] < \infty$ for any $p \geq 1$ and hence also

$$\sup_n \mathbb{E}[|X_n - X|^p] < \infty, \quad \forall p \geq 1. \quad (2)$$

Fix $p \geq 1$. Then the sequence $|X_n - X|^p$ converges to zero in probability (by assumption) and is uniformly integrable, since, for any $q > p$,

$$\mathbb{E}[|X_n - X|^p : |X_n - X|^p > r] \leq \frac{1}{r^{(q-p)/p}} \mathbb{E}[|X_n - X|^q] \rightarrow 0 \text{ as } r \rightarrow \infty$$

(by equation (2) above). It follows that we also have convergence of X_n to X in L^p .

□

2.2 Gaussian vectors and spaces

So far we've considered only real-valued Gaussian variables.

Definition 2.3 A random vector taking values in \mathbb{R}^d is called Gaussian if and only if

$$\langle u, X \rangle := u^T X = \sum_{i=1}^d u_i X_i \quad \text{is Gaussian for all } u \in \mathbb{R}^d.$$

It follows immediately that the image of a Gaussian vector under a linear transformation is also Gaussian: if $X \in \mathbb{R}^d$ is Gaussian and A is an $m \times d$ matrix, then AX is Gaussian in \mathbb{R}^m .

Lemma 2.4 Let X be a Gaussian vector and define $m_X := (\mathbb{E}[X_1], \dots, \mathbb{E}[X_d])$ and $\Gamma_X := (\text{cov}(X_i, X_j)_{1 \leq i, j \leq d})$, the mean vector and the covariance matrix respectively. Then $q_X(u) := u^T \Gamma_X u$ is a non-negative quadratic form and

$$u^T X \sim \mathcal{N}(u^T m_X, q_X(u)), \quad u \in \mathbb{R}^d. \quad (3)$$

Proof. Clearly $\mathbb{E}[u^T X] = u^T m_X$ and

$$\text{var}(u^T X) = \mathbb{E} \left[\left(\sum_{i=1}^d u_i (X_i - \mathbb{E}[X_i]) \right)^2 \right] = \sum_{1 \leq i, j \leq d} u_i u_j \text{cov}(X_i, X_j) = u^T \Gamma_X u = q_X(u),$$

which also shows that $q_X(u) \geq 0$. □

The identification in (3) is equivalent to

$$\mathbb{E} [e^{i \langle u, X \rangle}] = e^{i \langle u, m_X \rangle - \frac{1}{2} q_X(u)}, \quad u \in \mathbb{R}^d.$$

From this we derive easily the following important fact:

Proposition 2.5 Let X be a Gaussian vector and Γ_X its covariance matrix. Then X_1, \dots, X_d are independent if and only if Γ_X is a diagonal matrix (i.e. the variables are pairwise uncorrelated).

Warning: Note It is crucial to assume that the vector X is Gaussian and not just that X_1, \dots, X_d are Gaussian. For example, consider $X_1 \sim \mathcal{N}(0, 1)$ and ε an independent random variable with $\mathbb{P}[\varepsilon = 1] = 1/2 = \mathbb{P}[\varepsilon = -1]$. Let $X_2 := \varepsilon X_1$. Then $X_2 \sim \mathcal{N}(0, 1)$ and $\text{cov}(X_1, X_2) = 0$, while clearly X_1, X_2 are *not* independent.

By definition, a Gaussian vector X remains Gaussian if we add to it a deterministic vector $m \in \mathbb{R}^d$. Hence, without loss of generality, by considering $X - m_X$, it suffices to consider centred Gaussian vectors. The variance-covariance matrix Γ_X is symmetric and non-negative definite (as observed above). Conversely, for any such matrix Γ , there exists a Gaussian vector X with $\Gamma_X = \Gamma$, and, indeed, we can construct it as a linear transformation of a Gaussian vector with i.i.d. coordinates.

Theorem 2.6 Let Γ be a symmetric non-negative definite $d \times d$ matrix. Let $(\varepsilon_1, \dots, \varepsilon_d)$ be an orthonormal basis in \mathbb{R}^d which diagonalises Γ , i.e. $\Gamma \varepsilon_i = \lambda_i \varepsilon_i$ for some $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_r > 0 = \lambda_{r+1} = \dots = \lambda_d$, where $1 \leq r \leq d$ is the rank of Γ .

(i) A centred Gaussian vector X with covariance matrix $\Gamma_X = \Gamma$ exists.

(ii) Further, any such vector can be represented as

$$X = \sum_{i=1}^r Y_i \varepsilon_i, \quad (4)$$

where Y_1, \dots, Y_r are independent Gaussian random variables with $Y_i \sim \mathcal{N}(0, \lambda_i)$.

(iii) If $r = d$, then X admits a density given by

$$p_X(x) = \frac{1}{(2\pi)^{d/2}} \frac{1}{\sqrt{\det(\Gamma)}} \exp\left(-\frac{1}{2}x^T \Gamma^{-1}x\right), \quad x \in \mathbb{R}^d.$$

Proof. Let A be a matrix whose columns are ε_i so that $\Gamma_X = A\Lambda A^T$ where Λ is the diagonal matrix with entries λ_i on the diagonal. Let Z_1, \dots, Z_n be i.i.d. standard centred Gaussian variables and $Y_i = \sqrt{\lambda_i}Z_i$. Let X be given by (4), i.e. $X = AY$. Then

$$\langle u, X \rangle = \langle u, AY \rangle = \langle A^T u, Y \rangle = \sum_{i=1}^d \sqrt{\lambda_i} \langle u, \varepsilon_i \rangle Z_i$$

is Gaussian, centred. Its variance is given by

$$\text{var}(\langle u, X \rangle) = \sum_{i=1}^d \lambda_i (u^T \varepsilon_i)^2 = \sum_{i=1}^d u^T \varepsilon_i \lambda_i \varepsilon_i^T u = (A^T u)^T \Lambda A^T u = u^T A \Lambda A^T u = u^T \Gamma u,$$

and (i) is proved.

Conversely, suppose X is a centred Gaussian vector with covariance matrix Γ and let $Y = A^T X$. For $u \in \mathbb{R}^d$, $\langle u, Y \rangle = \langle Au, X \rangle$ is centred Gaussian with variance

$$(Au)^T \Gamma Au = u^T A^T \Gamma Au = u^T \Lambda u$$

and we conclude that Y is also a centred Gaussian vector with covariance matrix Λ . Independence between Y_1, \dots, Y_d then follows from Proposition 2.5. It follows that when $r = d$, Y admits a density on \mathbb{R}^d given by

$$p_Y(y) = \frac{1}{(2\pi)^{d/2}} \frac{1}{\sqrt{\det(\Lambda)}} \exp\left(-\frac{1}{2}y^T \Lambda^{-1}y\right), \quad y \in \mathbb{R}^d.$$

Change of variables, together with $\det(\Lambda) = \det(\Gamma)$ and $|\det(A)| = 1$ gives the desired density for x . \square

Once we know that we can write $X = \sum_{i=1}^r Y_i \varepsilon_i$ in this way, we have an easy way to compute conditional expectations within the family of random variables

which are linear transformations of a Gaussian vector X . To see how it works, suppose that X is a Gaussian vector in \mathbb{R}^d and define $Z := (X_1 - \sum_{i=2}^d a_i X_i)$ with the coefficients a_i chosen in such a way that Z and X_i are uncorrelated for $i = 2, \dots, d$; that is

$$\text{cov}(X_1, X_i) - \sum_{j=2}^d a_j \text{cov}(X_j, X_i) = 0, \quad 2 \leq i \leq d.$$

Evidently Z is Gaussian (it is a linear combination of Gaussians) and since it is uncorrelated with X_2, \dots, X_d , by Proposition 2.5, it is independent of them. Then

$$\begin{aligned} \mathbb{E}[X_1 | \sigma(X_2, \dots, X_d)] &= \mathbb{E}[Z + \sum_{i=2}^d a_i X_i | \sigma(X_2, \dots, X_d)] \\ &= \mathbb{E}[Z | \sigma(X_2, \dots, X_d)] + \mathbb{E}[\sum_{i=2}^d a_i X_i | \sigma(X_2, \dots, X_d)] = \sum_{i=2}^d a_i X_i, \end{aligned}$$

where we have used independence to see that $\mathbb{E}[Z | \sigma(X_2, \dots, X_d)] = \mathbb{E}[Z] = 0$.

The most striking feature is that $\mathbb{E}[X | \sigma(K)]$ is an element of the vector space K itself and not a general $\sigma(K)$ -measurable random variable. In particular, if (X_1, X_2, X_3) is a Gaussian vector then the best (in the L^2 sense, see Appendix B.6) approximation of X_1 in terms of X_2, X_3 is in fact a *linear* function of X_2 and X_3 . This extends to the more general setting of Gaussian spaces to which we now turn.

2.3 Gaussian spaces

Note that to a Gaussian vector X in \mathbb{R}^d , we can associate the vector space spanned by its coordinates:

$$\left\{ \sum_{i=1}^d u_i X_i : u_i \in \mathbb{R} \right\},$$

and by definition all elements of this space are Gaussian random variables. This is a simple example of a *Gaussian space* and it is useful to think of such spaces in much greater generality.

Definition 2.7 A closed linear subspace $H \subset L^2(\Omega, \mathcal{F}, \mathbb{P})$ is called a Gaussian space if all of its elements are centred Gaussian random variables.

In analogy to Proposition 2.5, two elements of a Gaussian space are independent if and only if they are uncorrelated, which in turn is equivalent to being orthogonal in L^2 . More generally we have the following result.

Theorem 2.8 Let H_1, H_2 be two Gaussian subspaces of a Gaussian space H . Then H_1, H_2 are orthogonal if and only if $\sigma(H_1)$ and $\sigma(H_2)$ are independent.

The theorem follows from monotone class arguments, which (see Appendix A) reduce it to checking that it holds true for any finite subcollection of random variables - which is Proposition 2.5.

Corollary 2.9 *Let H be a Gaussian space and K a closed subspace. Let p_K denote the orthogonal projection onto K . Then for $X \in H$*

$$\mathbb{E}[X|\sigma(K)] = p_K(X). \quad (5)$$

Proof. Let $Y = X - p_K(X)$ which, by Theorem 2.8 is independent of $\sigma(K)$. Hence

$$\mathbb{E}[X|\sigma(K)] = \mathbb{E}[p_K(X)|\sigma(K)] + \mathbb{E}[Y|\sigma(K)] = p_K(X) + \mathbb{E}[Y] = p_K(X),$$

where we have used that Y is a centred Gaussian and so has zero mean. \square

Warning: For an arbitrary $X \in L^2$ we would have

$$\mathbb{E}[X|\sigma(K)] = p_{L^2(\Omega, \sigma(K), \mathbb{P})}(X).$$

It is a special property of Gaussian random variables X that it is enough to consider the projection onto the much smaller space K .

2.4 Gaussian processes

Definition 2.10 *A stochastic process $(X_t : t \geq 0)$ is called a (centred) Gaussian process if any finite linear combination of its coordinates is a (centred) Gaussian variable.*

Equivalently, X is a centred Gaussian process if for any $n \in \mathbb{N}$ and $0 \leq t_1 < t_2 < \dots < t_n$, $(X_{t_1}, X_{t_2}, \dots, X_{t_n})$ is a (centred) Gaussian vector. It follows that the distribution of a centred Gaussian process on $\mathcal{B}(\mathbb{R}^{[0, \infty)})$, is characterised by the covariance function $\Gamma : [0, \infty)^2 \rightarrow \mathbb{R}$, i.e.

$$\Gamma(s, t) := \text{cov}(X_s, X_t).$$

For any fixed n -tuple $(X_{t_1}, \dots, X_{t_n})$ the covariance matrix $(\Gamma(t_i, t_j))$ has to be symmetric and positive semi-definite. As the following result shows, the converse also holds - for any such function, Γ , we may construct an associated Gaussian process.

Theorem 2.11 *Let $\Gamma : [0, \infty)^2 \rightarrow \mathbb{R}$ be symmetric and such that for any $n \in \mathbb{N}$ and $0 \leq t_1 < t_2 < \dots < t_n$,*

$$\sum_{1 \leq i, j \leq n} u_i u_j \Gamma(t_i, t_j) \geq 0, \quad u \in \mathbb{R}^d.$$

Then there exists a centred Gaussian process with covariance function Γ .

This result will follow from the (more general) Daniell-Kolmogorov Theorem 2.13 below.

Recalling from Proposition 2.2 that an L^2 -limit of Gaussian variables is also Gaussian, we observe that the closed linear subspace of L^2 spanned by the variables $(X_t : t \geq 0)$ is a Gaussian space.

2.5 Constructing distributions on $(\mathbb{R}^{[0,\infty)}, \mathcal{B}(\mathbb{R}^{[0,\infty)}))$

In this section, we're going to provide a very general result about constructing continuous time stochastic processes and a criterion due to Kolmogorov which gives conditions under which there will be a version of the process with continuous paths.

Let \mathbb{T} be the set of finite increasing sequences of non-negative numbers, i.e. $\mathbf{t} \in \mathbb{T}$ if and only if $\mathbf{t} = (t_1, t_2, \dots, t_n)$ for some n and $0 \leq t_1 < t_2 < \dots < t_n$.

Suppose that for each $\mathbf{t} \in \mathbb{T}$ of length n we have a probability measure $\mathbb{Q}_{\mathbf{t}}$ on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$. The collection $(\mathbb{Q}_{\mathbf{t}} : \mathbf{t} \in \mathbb{T})$ is called a family of finite-dimensional (marginal) distributions.

Definition 2.12 *A family of finite dimensional distributions is called consistent if for any $\mathbf{t} = (t_1, t_2, \dots, t_n) \in \mathbb{T}$ and $1 \leq j \leq n$*

$$\begin{aligned} \mathbb{Q}_{\mathbf{t}}(A_1 \times A_2 \times \dots \times A_{j-1} \times \mathbb{R} \times A_{j+1} \times \dots \times A_n) \\ = \mathbb{Q}_{\mathbf{s}}(A_1 \times A_2 \times \dots \times A_{j-1} \times A_{j+1} \times \dots \times A_n) \end{aligned}$$

where $A_i \in \mathcal{B}(\mathbb{R})$ and $\mathbf{s} := (t_1, t_2, \dots, t_{j-1}, t_{j+1}, \dots, t_n)$.

(In other words, if we integrate out over the distribution at the j th time point then we recover the corresponding marginal for the remaining lower dimensional vector.)

If we have a probability measure \mathbb{Q} on $(\mathbb{R}^{[0,\infty)}, \mathcal{B}(\mathbb{R}^{[0,\infty)}))$ then it defines a consistent family of marginals via

$$\mathbb{Q}_{\mathbf{t}}(A) = \mathbb{Q}(\{\omega \in \mathbb{R}^{[0,\infty)} : (\omega(t_1), \dots, \omega(t_n)) \in A\})$$

where $\mathbf{t} = (t_1, t_2, \dots, t_n)$, $A \in \mathcal{B}(\mathbb{R}^n)$, and we note that the set in question is in $\mathcal{B}(\mathbb{R}^{[0,\infty)})$ as it depends on finitely many coordinates. But we'd like a converse - if I give you $\mathbb{Q}_{\mathbf{t}}$, when does there exist a corresponding measure \mathbb{Q} .

Theorem 2.13 (Daniell-Kolmogorov) *Let $\{\mathbb{Q}_{\mathbf{t}} : \mathbf{t} \in \mathbb{T}\}$ be a consistent family of finite-dimensional distributions. Then there exists a probability measure \mathbb{P} on $(\mathbb{R}^{[0,\infty)}, \mathcal{B}(\mathbb{R}^{[0,\infty)}))$ such that for any n , $\mathbf{t} = (t_1, \dots, t_n) \in \mathbb{T}$ and $A \in \mathcal{B}(\mathbb{R}^n)$,*

$$\mathbb{Q}_{\mathbf{t}}(A) = \mathbb{P}[\{\omega \in \mathbb{R}^{[0,\infty)} : (\omega(t_1), \dots, \omega(t_n)) \in A\}]. \quad (6)$$

We won't prove this, but notice that (6) defines \mathbb{P} on the cylinder sets and so if we have countable additivity then the proof reduces to an application of Carathéodory's extension theorem. Uniqueness is a consequence of the Monotone Class Lemma.

This is a remarkably general result, but it doesn't allow us to say anything meaningful about the paths of the process. For that we appeal to Kolmogorov's criterion.

Theorem 2.14 (Kolmogorov's continuity criterion) *Suppose that a stochastic process $(X_t : t \leq T)$ defined on $(\Omega, \mathcal{F}, \mathbb{P})$ satisfies*

$$\mathbb{E}[|X_t - X_s|^\alpha] \leq C|t - s|^{1+\beta}, \quad 0 \leq s, t \leq T \quad (7)$$

for some strictly positive constants α , β and C .

Then there exists \tilde{X} , a modification of X , whose paths are γ -locally Hölder continuous $\forall \gamma \in (0, \beta/\alpha)$ a.s., i.e. there are constants $\delta > 0$ and $\varepsilon > 0$ s.t.

$$\sup_{0 \leq t-s \leq \delta, s, t \in [0, T]} \frac{|\tilde{X}_t - \tilde{X}_s|}{|t-s|^\gamma} \leq \varepsilon \quad \text{a.s.} \quad (8)$$

In particular, the sample paths of \tilde{X} are a.s. continuous.

3 Brownian Motion and First Properties

3.1 Definition of Brownian motion

Our fundamental building block will be Brownian motion. It is a centred Gaussian process, so we can characterise it in terms of its covariance structure. It is often described as an ‘infinitesimal random walk’, so to motivate the definition, we take a quick look at simple random walk.

Definition 3.1 *The discrete time stochastic process $\{S_n\}_{n \geq 0}$ is a symmetric simple random walk under the measure \mathbb{P} if $S_n = \sum_{i=1}^n \xi_i$, where the ξ_i can take only the values ± 1 , and are i.i.d. under \mathbb{P} with $\mathbb{P}[\xi_i = -1] = 1/2 = \mathbb{P}[\xi_i = 1]$.*

Lemma 3.2 *$\{S_n\}_{n \geq 0}$ is a \mathbb{P} -martingale (with respect to the natural filtration) and*

$$\text{cov}(S_n, S_m) = n \wedge m.$$

To obtain a ‘continuous’ version of simple random walk, we appeal to the Central Limit Theorem. Since $\mathbb{E}[\xi_i] = 0$ and $\text{var}(\xi_i) = 1$, we have

$$\mathbb{P} \left[\frac{S_n}{\sqrt{n}} \leq x \right] \rightarrow \int_{-\infty}^x \frac{1}{\sqrt{2\pi}} e^{-y^2/2} dy \text{ as } n \rightarrow \infty.$$

More generally,

$$\mathbb{P} \left[\frac{S_{[nt]}}{\sqrt{n}} \leq x \right] \rightarrow \int_{-\infty}^x \frac{1}{\sqrt{2\pi t}} e^{-y^2/2t} dy \text{ as } n \rightarrow \infty,$$

where $[nt]$ denotes the integer part of nt .

Heuristically at least, passage to the limit from simple random walk suggests the following definition of Brownian motion.

Definition 3.3 (Brownian motion) *A real-valued stochastic process $\{B_t\}_{t \geq 0}$ is a \mathbb{P} -Brownian motion (or a \mathbb{P} -Wiener process) if for some real constant σ , under \mathbb{P} ,*

- i. for each $s \geq 0$ and $t > 0$ the random variable $B_{t+s} - B_s$ has the normal distribution with mean zero and variance $\sigma^2 t$,
- ii. for each $n \geq 1$ and any times $0 \leq t_0 \leq t_1 \leq \dots \leq t_n$, the random variables $\{B_{t_r} - B_{t_{r-1}}\}$ are independent,
- iii. $B_0 = 0$,
- iv. B_t is continuous in $t \geq 0$.

When $\sigma^2 = 1$, we say that we have a standard Brownian motion.

Notice in particular that for $s < t$,

$$\text{cov}(B_s, B_t) = \mathbb{E}[B_s B_t] = \mathbb{E}[B_s^2 + B_s(B_t - B_s)] = \mathbb{E}[B_s^2] = s \quad (= s \wedge t).$$

We can write down the finite dimensional distributions using the independence of increments. They admit a density with respect to Lebesgue measure. We write $p(t, x, y)$ for the transition density

$$p(t, x, y) = \frac{1}{\sqrt{2\pi t}} \exp\left(-\frac{(x-y)^2}{2t}\right).$$

For $0 = t_0 \leq t_1 \leq t_2 \leq \dots \leq t_n$, writing $x_0 = 0$, the joint probability density function of B_{t_1}, \dots, B_{t_n} is

$$f(x_1, \dots, x_n) = \prod_1^n p(t_j - t_{j-1}, x_{j-1}, x_j).$$

Although the sample paths of Brownian motion are continuous, it does not mean that they are nice in any other sense. In fact the behaviour of Brownian motion is distinctly odd. Here are just a few of its strange behavioural traits.

- i. Although $\{B_t\}_{t \geq 0}$ is continuous everywhere, it is (with probability one) differentiable nowhere.
- ii. Brownian motion will eventually hit any and every real value no matter how large, or how negative. No matter how far above the axis, it will (with probability one) be back down to zero at some later time.
- iii. Once Brownian motion hits a value, it immediately hits it again *infinitely* often, and then again from time to time in the future.
- iv. It doesn't matter what scale you examine Brownian motion on, it looks just the same. Brownian motion is a fractal.

The last property is really a consequence of the construction of the process. We'll formulate the second and third more carefully later.

We could recover the existence of Brownian motion from the general principles outlined so far (Daniell-Kolmogorov Theorem and the Kolmogorov continuity criterion plus what we know about Gaussian processes), but we are now going to take a short digression to describe a beautiful (and useful) construction due to Lévy.

The idea is that we can simply produce a path of Brownian motion by direct polygonal interpolation. We require just one calculation.

Lemma 3.4 *Suppose that $\{B_t\}_{t \geq 0}$ is standard Brownian motion. Conditional on $B_{t_1} = x_1$, the probability density function of $B_{t_1/2}$ is*

$$p_{t_1/2}(x) \triangleq \sqrt{\frac{2}{\pi t_1}} \exp\left(-\frac{1}{2} \left(\frac{(x - \frac{1}{2}x_1)^2}{t_1/4}\right)\right).$$

In other words, the conditional distribution is a normally distributed random variable with mean $x_1/2$ and variance $t_1/4$. The proof is an exercise.

The construction: Without loss of generality we take the range of t to be $[0, 1]$. Lévy's construction builds (inductively) a polygonal approximation to the Brownian motion from a countable collection of *independent* normally distributed random variables with mean zero and variance one. We index them by the dyadic points of $[0, 1]$, a generic variable being denoted $\xi(k2^{-n})$ where $n \in \mathbb{N}$ and $k \in \{0, 1, \dots, 2^n\}$.

The induction begins with

$$X_1(t) = t\xi(1).$$

Thus X_1 is a linear function on $[0, 1]$.

The n th process, X_n , is linear in each interval $[(k-1)2^{-n}, k2^{-n}]$ is continuous in t and satisfies $X_n(0) = 0$. It is thus determined by the values $\{X_n(k2^{-n}), k = 1, \dots, 2^n\}$.

The inductive step: We take

$$X_{n+1}(2k2^{-(n+1)}) = X_n(2k2^{-(n+1)}) = X_n(k2^{-n}).$$

We now determine the appropriate value for $X_{n+1}((2k-1)2^{-(n+1)})$. Conditional on $X_{n+1}(2k2^{-(n+1)}) - X_{n+1}(2(k-1)2^{-(n+1)})$, Lemma 3.4 tells us that

$$X_{n+1}\left((2k-1)2^{-(n+1)}\right) - X_{n+1}\left(2(k-1)2^{-(n+1)}\right)$$

should be normally distributed with mean

$$\frac{1}{2} \left(X_{n+1}\left(2k2^{-(n+1)}\right) - X_{n+1}\left(2(k-1)2^{-(n+1)}\right) \right)$$

and variance $2^{-(n+2)}$.

Now if $X \sim \mathcal{N}(0, 1)$, then $aX + b \sim \mathcal{N}(b, a^2)$ and so we take

$$X_{n+1} \left((2k-1)2^{-(n+1)} \right) - X_{n+1} \left(2(k-1)2^{-(n+1)} \right) = 2^{-(n/2+1)} \xi \left((2k-1)2^{-(n+1)} \right) \\ + \frac{1}{2} \left(X_{n+1} \left(2k2^{-(n+1)} \right) - X_{n+1} \left(2(k-1)2^{-(n+1)} \right) \right).$$

In other words

$$X_{n+1} \left((2k-1)2^{-(n+1)} \right) = \frac{1}{2} X_n \left((k-1)2^{-n} \right) \\ + \frac{1}{2} X_n \left(k2^{-n} \right) + 2^{-(n/2+1)} \xi \left((2k-1)2^{-(n+1)} \right) \\ = X_n \left((2k-1)2^{-(n+1)} \right) \\ + 2^{-(n/2+1)} \xi \left((2k-1)2^{-(n+1)} \right), \quad (9)$$

where the last equality follows by linearity of X_n on $[(k-1)2^{-n}, k2^{-n}]$.

The construction is illustrated in Figure 1

Brownian motion will be the process constructed by letting n increase to infinity. To check that it exists we need some technical lemmas. The proofs are adapted from Knight (1981).

Lemma 3.5

$$\mathbb{P} \left[\lim_{n \rightarrow \infty} X_n(t) \text{ exists for } 0 \leq t \leq 1 \text{ uniformly in } t \right] = 1.$$

Proof: Notice that $\max_t |X_{n+1}(t) - X_n(t)|$ will be attained at a vertex, that is for $t \in \{(2k-1)2^{-(n+1)}, k = 1, 2, \dots, 2^n\}$ and using (9),

$$\mathbb{P} \left[\max_t |X_{n+1}(t) - X_n(t)| \geq 2^{-n/4} \right] \\ = \mathbb{P} \left[\max_{1 \leq k \leq 2^n} \left| \xi \left((2k-1)2^{-(n+1)} \right) \right| \geq 2^{n/4+1} \right] \\ \leq 2^n \mathbb{P} \left[|\xi(1)| \geq 2^{n/4+1} \right].$$

Now

$$\mathbb{P}[\xi(1) \geq x] \leq \frac{1}{x\sqrt{2\pi}} e^{-x^2/2},$$

(exercise), and, combining this with the fact that

$$\exp\left(-2^{(n/2+1)}\right) < 2^{-2n+2},$$

we obtain that for $n \geq 4$

$$2^n \mathbb{P} \left[\xi(1) \geq 2^{n/4+1} \right] \leq \frac{2^n}{2^{n/4+1}} \frac{1}{\sqrt{2\pi}} \exp\left(-2^{(n/2+1)}\right) \leq \frac{2^n}{2^{n/4+1}} 2^{-2n+2} < 2^{-n}.$$

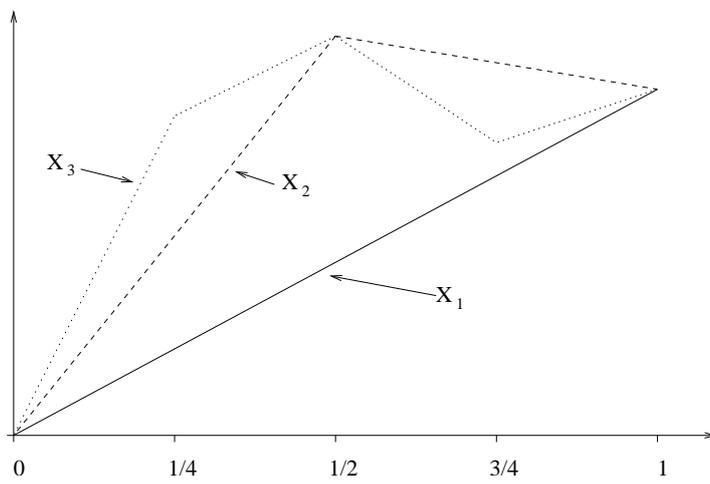


Figure 1: Lévy's sequence of polygonal approximations to Brownian motion.

Consider now for $k > n$,

$$\mathbb{P} \left[\max_t |X_k(t) - X_n(t)| \geq 2^{-n/4+3} \right] = 1 - \mathbb{P} \left[\max_t |X_k(t) - X_n(t)| \leq 2^{-n/4+3} \right]$$

and

$$\begin{aligned} \mathbb{P} \left[\max_t |X_k(t) - X_n(t)| \leq 2^{-n/4+3} \right] & \geq \mathbb{P} \left[\sum_{j=n}^{k-1} \max_t |X_{j+1}(t) - X_j(t)| \leq 2^{-n/4+3} \right] \\ & \geq \mathbb{P} \left[\max_t |X_{j+1}(t) - X_j(t)| \leq 2^{-j/4}, \right. \\ & \qquad \qquad \qquad \left. j = n, \dots, k-1 \right] \\ & \geq 1 - \sum_{j=n}^{k-1} 2^{-j} \geq 1 - 2^{-n+1}. \end{aligned}$$

Finally we have that

$$\mathbb{P} \left[\max_t |X_k(t) - X_n(t)| \geq 2^{-n/4+3} \right] \leq 2^{-n+1},$$

for all $k \geq n$. The events on the left are increasing (since the maximum can only increase by the addition of a new vertex) so

$$\mathbb{P} \left[\max_t |X_k(t) - X_n(t)| \geq 2^{-n/4+3} \text{ for some } k > n \right] \leq 2^{-n+1}.$$

In particular, for $\varepsilon > 0$,

$$\lim_{n \rightarrow \infty} \mathbb{P} [\text{For some } k > n \text{ and } t \leq 1, |X_k(t) - X_n(t)| \geq \varepsilon] = 0,$$

which proves the lemma. \square

To complete the proof of existence of the Brownian motion, we must check the following.

Lemma 3.6 *Let $X(t) = \lim_{n \rightarrow \infty} X_n(t)$ if the limit exists uniformly and zero otherwise. Then $X(t)$ satisfies the conditions of Definition 3.3 (for t restricted to $[0, 1]$).*

Proof: By construction, the properties 1–3 of Definition 3.3 hold for the approximation $X_n(t)$ restricted to $T_n = \{k2^{-n}, k = 0, 1, \dots, 2^n\}$. Since we don't change X_k on T_n for $k > n$, the same must be true for X on $\cup_{n=1}^{\infty} T_n$. A uniform limit of continuous functions is continuous, so condition 4 holds and now by approximation of any $0 \leq t_1 \leq t_2 \leq \dots \leq t_n \leq 1$ from within the dense set $\cup_{n=1}^{\infty} T_n$ we see that in fact all four properties hold without restriction for $t \in [0, 1]$. \square

3.2 Wiener Measure

Let $C(\mathbb{R}_+, \mathbb{R})$ be the space of continuous functions from $[0, \infty)$ to \mathbb{R} . Given a Brownian motion $(B_t : t \geq 0)$ on $(\Omega, \mathcal{F}, \mathbb{P})$, consider the map

$$\Omega \rightarrow C(\mathbb{R}_+, \mathbb{R}), \quad \text{given by } \omega \mapsto (B_t(\omega) : t \geq 0) \quad (10)$$

which is measurable w.r.t. $\mathcal{B}(C(\mathbb{R}_+, \mathbb{R}))$ - the smallest σ -algebra such that the coordinate mappings (i.e. $(\omega_t : t \geq 0) \mapsto \omega(t_0)$ for a fixed t_0) are measurable. (In fact $\mathcal{B}(C(\mathbb{R}_+, \mathbb{R}))$ is also the Borel σ -algebra generated by the topology of uniform convergence on compacts).

Definition 3.7 *The Wiener measure \mathbf{W} is the image of \mathbb{P} under the mapping in (10); it is the probability measure on the space of continuous functions such that the canonical process, i.e. $(B_t(\omega) = \omega(t), t \geq 0)$, is a Brownian motion.*

In other words, \mathbf{W} is the unique probability measure on $(C(\mathbb{R}_+, \mathbb{R}), \mathcal{B}(C(\mathbb{R}_+, \mathbb{R})))$ such that

- i. $\mathbf{W}(\{\omega \in C(\mathbb{R}_+, \mathbb{R}), \omega(0) = 0\}) = 1$;
- ii. for any $n \geq 1, \forall 0 = t_0 < t_1 < \dots < t_n, A \in \mathcal{B}(\mathbb{R}^n)$

$$\begin{aligned} & \mathbf{W}(\{\omega \in C(\mathbb{R}_+, \mathbb{R}) : (\omega(t_1), \dots, \omega(t_n)) \in A\}) \\ &= \int_A \frac{1}{(2\pi)^{\frac{n}{2}}} \frac{dy_1 \cdots dy_n}{\sqrt{t_1(t_2 - t_1) \cdots (t_n - t_{n-1})}} \exp\left(-\sum_{i=1}^n \frac{(y_i - y_{i-1})^2}{2(t_i - t_{i-1})}\right), \end{aligned}$$

where $y_0 := 0$.

(Uniqueness follows from the Monotone Class Lemma, since $\mathcal{B}(C(\mathbb{R}_+, \mathbb{R}))$ is generated by finitely dimensional projections.)

3.3 Extensions and first properties

Definition 3.8 *Let μ be a probability measure on \mathbb{R}^d . A d -dimensional stochastic process $(B_t : t \geq 0)$ on $(\Omega, \mathcal{F}, \mathbb{P})$ is called a d -dimensional Brownian motion with initial distribution μ if*

- i. $\mathbb{P}[B_0 \in A] = \mu(A), \quad A \in \mathcal{B}(\mathbb{R}^d)$;
- ii. $\forall 0 \leq s \leq t$ the increment $(B_t - B_s)$ is independent of $\sigma(B_u : u \leq s)$ and is normally distributed with mean 0 and covariance matrix $(t - s) \times \mathbf{I}_d$;
- iii. B has a.s. continuous paths.

Writing the d -dimensional Brownian motion as $B_t = (B_t^{(1)}, \dots, B_t^{(d)})$, the coordinate processes $(B_t^{(i)}), 1 \leq i \leq d$, are independent one-dimensional Brownian motions. If $\mu(\{x\}) = 1$ for some $x \in \mathbb{R}^d$, we say that B starts at x .

Proposition 3.9 *Let B be a standard real-valued Brownian motion. Then*

- i. $-B_t$ is also a Brownian motion, (symmetry)
- ii. $\forall c \geq 0$, cB_{t/c^2} is a Brownian motion, (scaling)
- iii. $X_0 = 0$, $X_t := tB_{\frac{1}{t}}$ is a Brownian motion, (time reversal)
- iv. $\forall s \geq 0$, $\tilde{B}_t = B_{t+s} - B_s$ is a Brownian motion independent of $\sigma(B_u : u \leq s)$, (simple Markov property).

The proof is an exercise.

3.4 Basic properties of Brownian sample paths

From now on, when we say ‘‘Brownian motion’’, we mean a standard real-valued Brownian motion.

We know that $t \mapsto B_t(\omega)$ is almost surely continuous.

Exercise: Use the Kolmogorov continuity criterion to show that Brownian motion admits a modification which is locally Holder continuous of order γ for any $0 < \gamma < 1/2$.

On the other hand, as we have already remarked, the path is actually rather ‘rough’. We’d like to have a way to quantify this roughness.

Definition 3.10 *Let π be a partition of $[0, T]$, $N(\pi)$ the number of intervals that make up π and $\delta(\pi)$ be the mesh of π (that is the length of the longest interval in the partition). Write $0 = t_0 < t_1 < \dots < t_{N(\pi)} = T$ for the endpoints of the intervals of the partition. Then the variation of a function $f : [0, T] \rightarrow \mathbb{R}$ is*

$$\lim_{\delta \rightarrow 0} \left\{ \sup_{\pi: \delta(\pi) = \delta} \sum_{j=1}^{N(\pi)} |f(t_j) - f(t_{j-1})| \right\}.$$

If the function is ‘nice’, for example differentiable, then it has bounded variation. Our ‘rough’ paths will have *unbounded* variation. To quantify roughness we can extend the idea of variation to that of p -variation.

Definition 3.11 *In the notation of Definition 3.10, the p -variation of a function $f : [0, T] \rightarrow \mathbb{R}$ is defined as*

$$\lim_{\delta \rightarrow 0} \left\{ \sup_{\pi: \delta(\pi) = \delta} \sum_{j=1}^{N(\pi)} |f(t_j) - f(t_{j-1})|^p \right\}.$$

Notice that for $p > 1$, the p -variation will be finite for functions that are much rougher than those for which the variation is bounded. For example, roughly speaking, finite 2-variation will follow if the fluctuation of the function over an interval of order δ is order $\sqrt{\delta}$.

For a typical Brownian path, the 2-variation will be infinite. However, a slightly weaker analogue of the 2-variation *does* exist.

Theorem 3.12 Let B_t denote Brownian motion under \mathbb{P} and for a partition π of $[0, T]$ define

$$S(\pi) = \sum_{j=1}^{N(\pi)} |B_{t_j} - B_{t_{j-1}}|^2.$$

Let π_n be a sequence of partitions with $\delta(\pi_n) \rightarrow 0$. Then

$$\mathbb{E} \left[|S(\pi_n) - T|^2 \right] \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (11)$$

We say that the quadratic variation process of Brownian motion, which we denote by $\{\langle B \rangle_t\}_{t \geq 0}$ is $\langle B \rangle_t = t$. More generally, we can define the quadratic variation process associated with any bounded continuous martingale.

Definition 3.13 Suppose that $\{M_t\}_{t \geq 0}$ is a bounded continuous \mathbb{P} -martingale. The quadratic variation process associated with $\{M_t\}_{t \geq 0}$ is the process $\{\langle M \rangle_t\}_{t \geq 0}$ such that for any sequence of partitions π_n of $[0, T]$ with $\delta(\pi_n) \rightarrow 0$,

$$\mathbb{E} \left[\left| \sum_{j=1}^{N(\pi_n)} |M_{t_j} - M_{t_{j-1}}|^2 - \langle M \rangle_T \right|^2 \right] \rightarrow 0 \quad \text{as } n \rightarrow \infty. \quad (12)$$

Remark: We don't prove it here, but the limit in (12) will be independent of the sequence of partitions. \square

Proof of Theorem 3.12: We expand the expression inside the expectation in (11) and make use of our knowledge of the normal distribution. Let $\{t_{n,j}\}_{j=0}^{N(\pi_n)}$ denote the endpoints of the intervals that make up the partition π_n . First observe that

$$|S(\pi_n) - T|^2 = \left| \sum_{j=1}^{N(\pi_n)} \left\{ |B_{t_{n,j}} - B_{t_{n,j-1}}|^2 - (t_{n,j} - t_{n,j-1}) \right\} \right|^2.$$

It is convenient to write $\delta_{n,j}$ for $|B_{t_{n,j}} - B_{t_{n,j-1}}|^2 - (t_{n,j} - t_{n,j-1})$. Then

$$|S(\pi_n) - T|^2 = \sum_{j=1}^{N(\pi_n)} \delta_{n,j}^2 + 2 \sum_{j < k} \delta_{n,j} \delta_{n,k}.$$

Note that since Brownian motion has independent increments,

$$\mathbb{E}[\delta_{n,j} \delta_{n,k}] = \mathbb{E}[\delta_{n,j}] \mathbb{E}[\delta_{n,k}] = 0 \quad \text{if } j \neq k.$$

Also

$$\mathbb{E}[\delta_{n,j}^2] = \mathbb{E} \left[|B_{t_{n,j}} - B_{t_{n,j-1}}|^4 - 2 |B_{t_{n,j}} - B_{t_{n,j-1}}|^2 (t_{n,j} - t_{n,j-1}) + (t_{n,j} - t_{n,j-1})^2 \right].$$

For a normally distributed random variable, X , with mean zero and variance λ , $\mathbb{E}[|X|^4] = 3\lambda^2$, so we have

$$\begin{aligned} \mathbb{E}[\delta_{n,j}^2] &= 3(t_{n,j} - t_{n,j-1})^2 - 2(t_{n,j} - t_{n,j-1})^2 + (t_{n,j} - t_{n,j-1})^2 \\ &= 2(t_{n,j} - t_{n,j-1})^2 \\ &\leq 2\delta(\pi_n)(t_{n,j} - t_{n,j-1}). \end{aligned}$$

Summing over j

$$\begin{aligned}\mathbb{E}\left[|S(\pi_n) - T|^2\right] &\leq 2 \sum_{j=1}^{N(\pi_n)} \delta(\pi_n)(t_{n,j} - t_{n,j-1}) \\ &= 2\delta(\pi_n)T \\ &\rightarrow 0 \quad \text{as } n \rightarrow \infty.\end{aligned}$$

□

Corollary 3.14 *Brownian sample paths are of infinite variation on any interval almost surely.*

Corollary 3.15 *Brownian sample paths are almost surely nowhere locally Hölder continuous of order $\gamma > \frac{1}{2}$.*

(The proofs are exercises.)

To study the very small time behaviour of Brownian motion, it is useful to establish the following 0–1 law.

Fix a Brownian motion $(B_t)_{t \geq 0}$ on $(\Omega, \mathcal{F}, \mathbb{P})$. For every $t \geq 0$ we set $\mathcal{F}_t := \sigma(B_u : u \leq t)$. Note that $\mathcal{F}_s \subset \mathcal{F}_t$ if $s \leq t$. We also set $\mathcal{F}_{0+} := \bigcap_{s>0} \mathcal{F}_s$.

Theorem 3.16 (Blumenthal’s 0-1 law)

The σ -field \mathcal{F}_{0+} is trivial in the sense that $\mathbb{P}[A] = 0$ or 1 for every $A \in \mathcal{F}_{0+}$.

Proof. Let $0 < t_1 < t_2 \cdots < t_k$ and let $g : \mathbb{R}^k \rightarrow \mathbb{R}$ be a bounded continuous function. Also, fix $A \in \mathcal{F}_{0+}$. Then by a continuity argument,

$$\mathbb{E}[\mathbf{1}_A g(B_{t_1}, \dots, B_{t_k})] = \lim_{\varepsilon \downarrow 0} \mathbb{E}[\mathbf{1}_A g(B_{t_1} - B_\varepsilon, \dots, B_{t_k} - B_\varepsilon)].$$

If $0 < \varepsilon < t_1$, the variables $B_{t_1} - B_\varepsilon, \dots, B_{t_k} - B_\varepsilon$ are independent of \mathcal{F}_ε (by the Markov property) and thus also of \mathcal{F}_{0+} . It follows that

$$\begin{aligned}\mathbb{E}[\mathbf{1}_A g(B_{t_1}, \dots, B_{t_k})] &= \lim_{\varepsilon \downarrow 0} \mathbb{E}[\mathbf{1}_A g(B_{t_1} - B_\varepsilon, \dots, B_{t_k} - B_\varepsilon)] \\ &= \mathbb{P}[A] \mathbb{E}[g(B_{t_1}, \dots, B_{t_k})].\end{aligned}$$

We have thus obtained that \mathcal{F}_{0+} is independent of $\sigma(B_{t_1}, \dots, B_{t_k})$. Since this holds for any finite collection $\{t_1, \dots, t_k\}$ of (strictly) positive reals, \mathcal{F}_{0+} is independent of $\sigma(B_t, t > 0)$. However, $\sigma(B_t, t > 0) = \sigma(B_t, t \geq 0)$, since B_0 is the pointwise limit of B_t when $t \rightarrow 0$. Since $\mathcal{F}_{0+} \subset \sigma(B_t, t \geq 0)$, we conclude that \mathcal{F}_{0+} is independent of itself and so must be trivial. □

Proposition 3.17 *Let B be a standard real-valued Brownian motion, as above.*

i. Then, a.s., for every $\varepsilon > 0$,

$$\sup_{0 \leq s \leq \varepsilon} B_s > 0 \quad \text{and} \quad \inf_{0 \leq s \leq \varepsilon} B_s < 0.$$

In particular, $\inf\{t > 0 : B_t = 0\} = 0$ a.s.

ii. For every $a \in \mathbb{R}$, let $T_a := \inf\{t \geq 0 : B_t = a\}$ (with the convention that $\inf \emptyset = \infty$). Then a.s. for each $a \in \mathbb{R}$, $T_a < \infty$. Consequently, we have a.s.

$$\limsup_{t \rightarrow \infty} B_t = +\infty, \quad \liminf_{t \rightarrow \infty} B_t = -\infty.$$

Remark 3.18 It is not a priori obvious that $\sup_{0 \leq s \leq \varepsilon} B_s$ is even measurable, since this is an uncountable supremum of random variables, but since sample paths are continuous, we can restrict to rational values of $s \in [0, \varepsilon]$ so that we are taking the supremum over a countable set. We implicitly use this observation in what follows.

Proof. (i) Let ε_p be a sequence of strictly positive reals decreasing to zero and set $A := \bigcap_{p \geq 0} \{\sup_{0 \leq s \leq \varepsilon_p} B_s > 0\}$. Since this is a monotone decreasing intersection, $A \in \mathcal{F}_{0+}$. On the other hand, by monotonicity, $\mathbb{P}[A]$ is the decreasing limit

$$\mathbb{P}[A] = \lim_{p \rightarrow \infty} \downarrow \mathbb{P}\left[\sup_{0 \leq s \leq \varepsilon_p} B_s > 0\right]$$

and

$$\mathbb{P}\left[\sup_{0 \leq s \leq \varepsilon_p} B_s > 0\right] \geq \mathbb{P}[B_{\varepsilon_p} > 0] = \frac{1}{2}.$$

So $\mathbb{P}[A] \geq 1/2$ and by Blumenthal's 0 – 1 law $\mathbb{P}[A] = 1$. Hence a.s. for all $\varepsilon > 0$, $\sup_{0 \leq s \leq \varepsilon} B_s > 0$. Replacing B by $-B$ we obtain $\mathbb{P}[\inf_{0 \leq s \leq \varepsilon} B_s < 0] = 1$.

(ii) Write

$$1 = \mathbb{P}\left[\sup_{0 \leq s \leq 1} B_s > 0\right] = \lim_{\delta \downarrow 0} \uparrow \mathbb{P}\left[\sup_{0 \leq s \leq 1} B_s > \delta\right],$$

where \uparrow indicates an increasing limit. Now use the scale invariance property, $B_t^\lambda = B_{\lambda^2 t}/\lambda$ is a Brownian motion, with $\lambda = 1/\delta$ to see that for any $\delta > 0$,

$$\mathbb{P}\left[\sup_{0 \leq s \leq 1} B_s > \delta\right] = \mathbb{P}\left[\sup_{0 \leq s \leq 1/\delta^2} B_s^\delta > 1\right] = \mathbb{P}\left[\sup_{0 \leq s \leq 1/\delta^2} B_s > 1\right]. \quad (13)$$

If we let $\delta \downarrow 0$, we find

$$\mathbb{P}\left[\sup_{s \geq 0} B_s > 1\right] = \lim_{\delta \downarrow 0} \uparrow \mathbb{P}\left[\sup_{0 \leq s \leq 1/\delta^2} B_s > 1\right] = 1$$

(since $\lim_{\delta \downarrow 0} \uparrow \mathbb{P}[\sup_{0 \leq B_s \leq 1} B_s > \delta] = 1$).

Another scaling argument shows that, for every $M > 0$,

$$\mathbb{P}\left[\sup_{s \geq 0} B_s > M\right] = 1$$

and replacing B with $-B$,

$$\mathbb{P}[\inf_{s \geq 0} B_s < -M] = 1.$$

Continuity of sample paths completes the proof of (ii). \square

Corollary 3.19 *a.s. $t \mapsto B_t$ is not monotone on any non-trivial interval.*

4 Filtrations and stopping times

These are concepts that you already know about in the context of discrete parameter martingales. Our definitions here mirror what you already know, but in the continuous setting one has to be slightly more careful. In the end, we'll make enough assumptions to guarantee that everything goes through nicely.

Definition 4.1 *A collection $\{\mathcal{F}_t, t \in [0, \infty)\}$ of σ -algebras of sets in \mathcal{F} is a filtration if $\mathcal{F}_t \subseteq \mathcal{F}_{t+s}$ for $t, s \in [0, \infty)$. (Intuitively, \mathcal{F}_t corresponds to the information known to an observer at time t .)*

In particular, for a process X we define $\mathcal{F}_t^X = \sigma(\{X(s) : s \leq t\})$ (that is \mathcal{F}_t^X is the information obtained by observing X up to time t) to be the natural filtration associated with the process X .

We say that \mathcal{F}_t is right continuous if for each $t \geq 0$,

$$\mathcal{F}_t = \mathcal{F}_{t+} \equiv \bigcap_{\varepsilon > 0} \mathcal{F}_{t+\varepsilon}.$$

We say that $\{\mathcal{F}_t\}$ is complete if $(\Omega, \mathcal{F}, \mathbb{P})$ is complete (contains all subsets of the \mathbb{P} -null sets) and $\{A \in \mathcal{F} : \mathbb{P}[A] = 0\} \subset \mathcal{F}_0$ (and hence $\subset \mathcal{F}_t$ for all t).

A process X is adapted to a filtration $\{\mathcal{F}_t\}$ if $X(t)$ is \mathcal{F}_t -measurable for each $t \geq 0$ (if and only if $\mathcal{F}_t^X \subseteq \mathcal{F}_t$ for all t).

A process X is \mathcal{F}_t -progressive if for each $t \geq 0$, the mapping $(s, \omega) \mapsto X_s(\omega)$ is measurable on $([0, t] \times \Omega, \mathcal{B}([0, t]) \otimes \mathcal{F}_t)$.

If X is \mathcal{F}_t -progressive, then it is \mathcal{F}_t -adapted, but the converse is not necessarily true. However, every right continuous \mathcal{F}_t -adapted process is \mathcal{F}_t -progressive and since we shall only be interested in continuous processes, we won't need to dwell on these details.

Proposition 4.2 *An adapted process (X_t) whose paths are all right-continuous (or are all left-continuous) is progressively measurable.*

Proof. We present the argument for a right-continuous X . For $t > 0$, $n \geq 1$, $k = 0, 1, 2, \dots, 2^n - 1$ let $X_0^{(n)}(\omega) = X_0(\omega)$ and $X_s^{(n)}(\omega) := X_{\frac{k+1}{2^n}t}(\omega)$ for $\frac{kt}{2^n} < s \leq \frac{k+1}{2^n}t$.

Clearly $(X_s^{(n)} : s \leq t)$ takes finitely many values and is $\mathcal{B}([0, t]) \otimes \mathcal{F}_t$ -measurable. Further, by right continuity, $X_s(\omega) = \lim_{n \rightarrow \infty} X_s^{(n)}(\omega)$, and hence is also measurable (as a limit of measurable mappings). \square

Definition 4.3 A filtration $(\mathcal{F}_t)_{t \geq 0}$ (or the filtered space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$) is said to satisfy the usual conditions if it is right-continuous and complete.

Given a filtered probability space, we can always consider a natural augmentation, replacing the filtration with $\sigma(\mathcal{F}_{t+}, \mathcal{N})$, where $\mathcal{N} = \mathcal{N}(\mathbb{P}) := \{A \in \Omega : \exists B \in \mathcal{F} \text{ such that } A \subseteq B \text{ and } \mathbb{P}[B] = 0\}$. The augmented filtration satisfies the usual conditions. In Section 5.3 we'll see that if we have a martingale with respect to a filtration that satisfies the usual conditions, then it has a right continuous version.

Usually we shall consider the natural filtration associated with a process and don't specify it explicitly. Sometimes we suppose that we are given a filtration \mathcal{F}_t and then we say that a process is an \mathcal{F}_t -Brownian motion if it is adapted to \mathcal{F}_t .

4.1 Stopping times

Again the definition mirrors what you know from the discrete setting.

Definition 4.4 Let $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ be a filtered space. A random variable $\tau : \Omega \mapsto [0, +\infty]$ is called a stopping time (relative to (\mathcal{F}_t)) if $\{\tau \leq t\} \in \mathcal{F}_t, \forall t \geq 0$.

The 'first time a certain phenomenon occurs' will be a stopping time. Our fundamental examples will be first hitting times of sets. If (X_t) is a stochastic process and $\Gamma \in \mathcal{B}(\mathbb{R})$ we set

$$H_\Gamma(\omega) = H_\Gamma(X(\omega)) := \inf\{t \geq 0 : X(\omega) \in \Gamma\}. \quad (14)$$

Exercise 4.5 Show that

- i. if (X_t) is adapted to (\mathcal{F}_t) and has right-continuous paths then H_Γ , for Γ an open set, is a stopping time relative to (\mathcal{F}_{t+}) .
- ii. if (X_t) has continuous paths, then H_Γ , for Γ a closed set, is a stopping time relative to (\mathcal{F}_t) .

With a stopping time we can associate 'the information known at time τ ':

Definition 4.6 Given a stopping time τ relative to (\mathcal{F}_t) we define

$$\begin{aligned} \mathcal{F}_\tau &:= \{A \in \mathcal{F} : A \cap \{\tau \leq t\} \in \mathcal{F}_t \forall t \geq 0\}, \\ \mathcal{F}_{\tau+} &:= \{A \in \mathcal{F} : A \cap \{\tau < t\} \in \mathcal{F}_t\}, \\ \mathcal{F}_{\tau-} &:= \sigma(\{A \cap \{\tau > t\} : t \geq 0, A \in \mathcal{F}_t\}) \end{aligned}$$

(which satisfy all the natural properties).

Proposition 4.7 Let τ be a stopping time. Then

- (i) \mathcal{F}_τ is a σ -algebra and τ is \mathcal{F}_τ -measurable.
- (ii) $\mathcal{F}_{\tau-} \subseteq \mathcal{F}_\tau \subseteq \mathcal{F}_{\tau+}$ and $\mathcal{F}_\tau = \mathcal{F}_{\tau+}$ if (\mathcal{F}_t) is right-continuous.

(iii) If $\tau = t$ then $\mathcal{F}_\tau = \mathcal{F}_t$, $\mathcal{F}_{\tau+} = \mathcal{F}_{t+}$.

(iv) If τ and ρ are stopping times then so are $\tau \wedge \rho$, $\tau \vee \rho$ and $\tau + \rho$ and $\{\tau \leq \rho\} \in \mathcal{F}_{\tau \wedge \rho}$. Further if $\tau \leq \rho$ then $\mathcal{F}_\tau \subseteq \mathcal{F}_\rho$.

(v) If τ is a stopping time and ρ is a $[0, \infty]$ -valued random variable which is \mathcal{F}_τ -measurable and $\rho \geq \tau$, then ρ is a stopping time. In particular,

$$\tau_n := \sum_{k=0}^{\infty} \frac{k+1}{2^n} \mathbf{1}_{\{\frac{k}{2^n} < \tau \leq \frac{k+1}{2^n}\}} + \infty \mathbf{1}_{\{\tau = \infty\}} \quad (15)$$

is a sequence of stopping times with $\tau_n \downarrow \tau$ as $n \rightarrow \infty$.

Proof.

We give a proof of (v).

Note that $\{\rho \leq t\} = \{\rho \leq t\} \cap \{\tau \leq t\} \in \mathcal{F}_t$ since ρ is \mathcal{F}_τ -measurable. Hence ρ is a stopping time. We have $\tau_n \downarrow \tau$ by definition, and clearly τ_n is \mathcal{F}_τ -measurable since τ is \mathcal{F}_τ -measurable. \square

It is often useful to be able to ‘stop’ a process at a stopping time and know that the result still has nice measurability properties. If $(X_t)_{t \geq 0}$ is progressively measurable and τ is a stopping time, then $X^\tau := (X_{t \wedge \tau} : t \geq 0)$ is progressively measurable.

We’re going to use the sequence τ_n of stopping times in (15) to prove an important generalisation of the Markov property for Brownian motion called the *strong* Markov property. Recall that the Markov property says that Brownian motion has ‘no memory’ - we can start it again from B_s and $B_{t+s} - B_s$ is just a Brownian motion, independent of the path followed by B up to time s . The strong Markov property says that the same is true if we replace s by a stopping time.

Theorem 4.8 *Let $B = (B_t : t \geq 0)$ be a standard Brownian motion on the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ and let τ be a stopping time with respect to $(\mathcal{F}_t)_{t \geq 0}$. Then, conditional on $\{\tau < \infty\}$, the process*

$$B_t^{(\tau)} := B_{\tau+t} - B_\tau \quad (16)$$

is a standard Brownian motion independent of \mathcal{F}_τ .

Proof. Assume that $\tau < \infty$ a.s..

We will show that $\forall A \in \mathcal{F}_\tau, 0 \leq t_1 < \dots < t_p$ and continuous and bounded functions F on \mathbb{R}^n we have

$$\mathbb{E}[\mathbf{1}_A F(B_{t_1}^{(\tau)}, \dots, B_{t_p}^{(\tau)})] = \mathbb{P}(A) \mathbb{E}[F(B_{t_1}, \dots, B_{t_p})]. \quad (17)$$

Granted (17), taking $A = \Omega$, we find that B and $B^{(\tau)}$ have the same finite dimensional distributions, and since $B^{(\tau)}$ has continuous paths, it must be a Brownian

motion. On the other hand (as usual using a monotone class argument), (17) says that $(B_{t_1}^{(\tau)}, \dots, B_{t_p}^{(\tau)})$ is independent of \mathcal{F}_τ , and so $B^{(\tau)}$ is independent of \mathcal{F}_τ .

To establish (17), first observe that by continuity of B and F ,

$$F(B_{t_1}^{(\tau)}, \dots, B_{t_p}^{(\tau)}) = \lim_{n \rightarrow \infty} \sum_{k=0}^{\infty} \mathbf{1}_{\frac{k-1}{2^n} < \tau \leq \frac{k}{2^n}} F(B_{\frac{k}{2^n} + t_1} - B_{\frac{k}{2^n}}, \dots, B_{\frac{k}{2^n} + t_p} - B_{\frac{k}{2^n}}) \quad \text{a.s.},$$

and by the Dominated Convergence Theorem

$$\mathbb{E}[\mathbf{1}_A F(B_{t_1}^{(\tau)}, \dots, B_{t_p}^{(\tau)})] = \lim_{n \rightarrow \infty} \sum_{k=0}^{\infty} \mathbb{E} \left[\mathbf{1}_A \mathbf{1}_{\frac{k-1}{2^n} < \tau \leq \frac{k}{2^n}} F(B_{\frac{k}{2^n} + t_1} - B_{\frac{k}{2^n}}, \dots, B_{\frac{k}{2^n} + t_p} - B_{\frac{k}{2^n}}) \right].$$

For $A \in \mathcal{F}_\tau$, the event $A \cap \{\frac{k-1}{2^n} < \tau \leq \frac{k}{2^n}\} \in \mathcal{F}_{\frac{k}{2^n}}$, so using the simple Markov property at $k/2n$,

$$\begin{aligned} \mathbb{E} \left[\mathbf{1}_{A \cap \{\frac{k-1}{2^n} < \tau \leq \frac{k}{2^n}\}} F(B_{\frac{k}{2^n} + t_1} - B_{\frac{k}{2^n}}, \dots, B_{\frac{k}{2^n} + t_p} - B_{\frac{k}{2^n}}) \right] \\ = \mathbb{P} \left[A \cap \left\{ \frac{k-1}{2^n} < \tau \leq \frac{k}{2^n} \right\} \right] \mathbb{E} [F(B_{t_1}, \dots, B_{t_p})]. \end{aligned}$$

Sum over k to recover the desired result.

If $\mathbb{P}(\tau = \infty) > 0$, the same argument gives instead

$$\mathbb{E} \left[\mathbf{1}_{A \cap \{\tau < \infty\}} F(B_{t_1}^{(\tau)}, \dots, B_{t_p}^{(\tau)}) \right] = \mathbb{P}[A \cap \{\tau < \infty\}] \mathbb{E} [F(B_{t_1}, \dots, B_{t_p})].$$

□

It was not until the 1940's that Doob properly formulated the strong Markov property and it was 1956 before Hunt proved it for Brownian motion.

The following result, known as the reflection principle, was known at the end of the 19th Century for random walk and appears in the famous 1900 thesis of Bachelier, which introduced the idea of modelling stock prices using Brownian motion (although since he had no formulation of the strong Markov property, his proof is not rigorous).

Theorem 4.9 (The reflection principle) *Let $S_t := \sup_{u \leq t} B_u$. For $a \geq 0$ and $b \leq a$ we have*

$$\mathbb{P}[S_t \geq a, B_t \leq b] = \mathbb{P}[B_t \geq 2a - b] \quad \forall t \geq 0.$$

In particular S_t and $|B_t|$ have the same distribution.

Proof. We apply the strong Markov property to the stopping time $T_a = \inf\{t > 0 : B_t = a\}$. We have already seen that $T_a < \infty$ a.s. and so in the notation of Theorem 4.8,

$$\mathbb{P}[S_t \geq a, B_t \leq b] = \mathbb{P}[T_a \leq t, B_t \leq b] = \mathbb{P}[T_a \leq t, B_{t-T_a}^{(T_a)} \leq b - a] \quad (18)$$

(since $B_{t-T_a}^{(T_a)} = B_t - B_{T_a} = B_t - a$).

Now $B^{(T_a)}$ is a Brownian motion, independent of \mathcal{F}_{T_a} and hence of T_a . Since $B^{(T_a)}$ has the same distribution as $-B^{(T_a)}$, $(T_a, B^{(T_a)})$ has the same distribution as $(T_a, -B^{(T_a)})$.

So

$$\begin{aligned}\mathbb{P}[T_a \leq t, B^{(T_a)} \leq b - a] &= \mathbb{P}[T_a \leq t, -B_{t-T_a}^{(T_a)} \leq b - a] \\ &= \mathbb{P}[T_a \leq t, B_t \geq 2a - b] = \mathbb{P}[B_t \geq 2a - b],\end{aligned}$$

since $2a - b \geq a$ and so $\{B_t \geq 2a - b\} \subseteq \{T_a \leq t\}$.

We have proved that $\mathbb{P}[S_t \geq a, B_t \leq b] = \mathbb{P}[B_t \geq 2a - b]$. For the last assertion of the theorem, taking $a = b$ in (18), observe that

$$\begin{aligned}\mathbb{P}[S_t \geq a] &= \mathbb{P}[S_t \geq a, B_t \geq a] + \mathbb{P}[S_t \geq a, B_t \leq a] \\ &= 2\mathbb{P}[B_t \geq a] = \mathbb{P}[B_t \geq a] + \mathbb{P}[B_t \leq -a] \quad (\text{symmetry}) \\ &= \mathbb{P}[|B_t| \geq a].\end{aligned}$$

□

5 (Sub/super-)Martingales in continuous time

The results in this section will to a large extent mirror what you proved last term for discrete parameter martingales (and we use those results repeatedly in our proofs). We assume throughout that a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ is given.

5.1 Definitions

Definition 5.1 An adapted stochastic process $(X_t)_{t \geq 0}$ such that $X_t \in L^1(\mathbb{P})$ (i.e. $\mathbb{E}[|X_t|] < \infty$) for any $t \geq 0$, is called

- i. a martingale if $\mathbb{E}[X_t | \mathcal{F}_s] = X_s$ for all $0 \leq s \leq t$,
- ii. a super-martingale if $\mathbb{E}[X_t | \mathcal{F}_s] \leq X_s$ for all $0 \leq s \leq t$,
- iii. a sub-martingale if $\mathbb{E}[X_t | \mathcal{F}_s] \geq X_s$ for all $0 \leq s \leq t$.

Exercises: Suppose $(Z_t : t \geq 0)$ is an adapted process with independent increments, i.e. for all $0 \leq s < t$, $Z_t - Z_s$ is independent of \mathcal{F}_s . The following give us examples of martingales:

- i. if $\forall t \geq 0, Z_t \in L^1$, then $\tilde{Z}_t := Z_t - \mathbb{E}[Z_t]$ is a martingale,
- ii. if $\forall t \geq 0, Z_t \in L^2$, then $\tilde{Z}_t^2 - \mathbb{E}[\tilde{Z}_t^2]$ is a martingale,
- iii. if for some $\theta \in \mathbb{R}$, and $\forall t \geq 0, \mathbb{E}[e^{\theta Z_t}] < \infty$, then $\frac{e^{\theta Z_t}}{\mathbb{E}[e^{\theta Z_t}]}$ is a martingale.

In particular, B_t , $B_t^2 - t$ and $e^{\theta B_t - \theta^2 t/2}$ are all martingales with respect to a filtration $(\mathcal{F}_t)_{t \geq 0}$ for which $(B_t)_{t \geq 0}$ is a Brownian motion.

Warning: It is important to remember that a process is a martingale *with respect to a filtration* - giving yourself more information (enlarging the filtration) may destroy the martingale property. For us, even when we don't explicitly mention it, there is a filtration implicitly assumed (usually the natural filtration associated with the process, augmented to satisfy the usual conditions).

Given a martingale (or submartingale) it is easy to generate many more.

Proposition 5.2 *Let $(X_t)_{t \geq 0}$ be a martingale (respectively sub-martingale) and $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ be a convex (respectively convex and increasing) such that $\mathbb{E}[|\varphi(X_t)|] < \infty$ for any $t \geq 0$. Then $(\varphi(X_t))_{t \geq 0}$ is a sub-martingale.*

Proof. Conditional Jensen inequality. □

In particular, if $(X_t)_{t \geq 0}$ is martingale with $\mathbb{E}[|X_t|^p] < \infty$, for some $p \geq 1$ and all $t \geq 0$, then $|X_t|^p$ is a sub-martingale (and consequently, $t \mapsto \mathbb{E}[|X_t|^p]$ is non-decreasing).

5.2 Doob's maximal inequalities

Doob was the person who placed martingales on a firm mathematical foundation (beginning in the 1940's). He initially called them 'processes with the property E' ', but reverted to the term martingale in his monumental book.

Doob's inequalities are fundamental to proving convergence theorems for martingales. You already encountered them in the discrete setting and we shall recall those results that underpin our proofs in the continuous world here. They allow us to control the running maximum of a martingale.

Theorem 5.3 *If $(X_n)_{n \geq 0}$ is a discrete martingale (or a positive submartingale) w.r.t. some filtration (\mathcal{F}_n) , then for any $N \in \mathbb{N}$, $p \geq 1$ and $\lambda > 0$,*

$$\lambda^p \mathbb{P} \left[\sup_{n \leq N} |X_n| \geq \lambda \right] \leq \mathbb{E}[|X_N|^p]$$

and for any $p > 1$

$$\mathbb{E}[|X_N|^p] \leq \mathbb{E} \left[\sup_{n \leq N} |X_n|^p \right] \leq \left(\frac{p}{p-1} \right)^p \mathbb{E}[|X_N|^p].$$

We'd now like to extend this to continuous time.

Suppose that X is indexed by $t \in [0, \infty)$. Take a countable dense set D in $[0, T]$, e.g. $D = \mathbb{Q} \cap [0, T]$, and an increasing sequence of finite subsets $D_n \subseteq D$ such that $\bigcup_{n=1}^{\infty} D_n = D$.

The above inequalities hold for X indexed by $t \in D_n \cup \{T\}$. Monotone convergence then yields the result for $t \in D$. If X has regular sample paths (e.g. right continuous) then the supremum over a countable dense set in $[0, T]$ is the same as over the *whole* of $[0, T]$ and so:

Theorem 5.4 (Doob's maximal and L^p inequalities)

If $(X_t)_{t \geq 0}$ is a right continuous martingale or positive sub-martingale, then for any $T \geq 0, \lambda > 0$,

$$\begin{aligned} \lambda^p \mathbb{P} \left[\sup_{t \leq T} |X_t| \geq \lambda \right] &\leq \mathbb{E}[|X_T|^p], \quad p \geq 1 \\ E \left[\sup_{t \leq T} |X_t|^p \right] &\leq \left(\frac{p}{p-1} \right)^p \mathbb{E}[|X_T|^p], \quad p > 1. \end{aligned} \tag{19}$$

As an application of Doob's maximal inequality, we derive a useful bound for Brownian motion.

Proposition 5.5 Let $(B_t)_{t \geq 0}$ be Brownian motion and $S_t = \sup_{u \leq t} B_u$. For any $\lambda > 0$ we have

$$\mathbb{P}[S_t \geq \lambda t] \leq e^{-\frac{\lambda^2 t}{2}}.$$

Proof. Recall that $e^{\alpha B_t - \alpha^2 t/2}, t \geq 0$, is a non-negative martingale. It follows that, for $\alpha \geq 0$,

$$\begin{aligned} \mathbb{P}[S_t \geq \lambda t] &\leq \mathbb{P} \left[\sup_{u \leq t} (e^{\alpha B_u - \alpha^2 u/2}) \geq e^{\alpha \lambda t - \alpha^2 t/2} \right] \\ &\leq \mathbb{P} \left[\sup_{u \leq t} (e^{\alpha B_u - \alpha^2 u/2}) \geq e^{\alpha \lambda t - \alpha^2 t/2} \right] \\ &\leq e^{-\alpha \lambda t + \alpha^2 t/2} \underbrace{\mathbb{E} \left[e^{\alpha B_t - \alpha^2 t/2} \right]}_{=1}. \end{aligned}$$

The bound now follows since $\min_{\alpha \geq 0} e^{-\alpha \lambda t + \alpha^2 t/2} = e^{-\lambda^2 t/2}$ (with the minimum achieved when $\alpha = \lambda$). \square

In the next subsection, we are going to show that even if a supermartingale is not right continuous, it has a right continuous version (this is Doob's Regularisation Theorem). To prove this, we need a slight variant of the maximal inequality - this time for a supermartingale - which in turn relies on Doob's Optional Stopping Theorem for discrete supermartingales.

Theorem 5.6 (Doob's Optional Stopping Theorem for discrete supermartingales) (bounded case)

If $(Y_n)_{n \geq 1}$ is a supermartingale, then for any choice of bounded stopping times S and T such that $S \leq T$, we have

$$Y_S \geq \mathbb{E}[Y_T | \mathcal{F}_S].$$

Here's the version of the maximal inequality that we shall need.

Proposition 5.7 Let $(X_t : t \geq 0)$ be a supermartingale. Then

$$\mathbb{P} \left[\sup_{t \in [0, T] \cap \mathbb{Q}} |X_t| \geq \lambda \right] \leq \frac{1}{\lambda} (2\mathbb{E}[|X_T|] + \mathbb{E}[|X_0|]), \quad \forall \lambda, T > 0. \tag{20}$$

In particular, $\sup_{t \in [0, T] \cap \mathbb{Q}} |X_t| < \infty$ a.s.

Proof. Take a sequence of rational numbers $0 = t_0 < t_1 < \dots < t_n = T$. Applying Theorem 5.6 with $S = \min\{t_i : X_{t_i} \geq \lambda\} \wedge T$, we obtain

$$\mathbb{E}[X_0] \geq \mathbb{E}[X_S] \geq \lambda \mathbb{P}\left[\sup_{1 \leq i \leq n} X_{t_i} \geq \lambda\right] + \mathbb{E}[X_T \mathbf{1}_{\sup_{1 \leq i \leq n} X_{t_i} < \lambda}].$$

Rearranging,

$$\lambda \mathbb{P}\left(\sup_{1 \leq i \leq n} X_{t_i} \geq \lambda\right) \leq \mathbb{E}[X_0] + \mathbb{E}[X_T^-]$$

(where $X_T^- = -\min(X_T, 0)$). Now X_T^- is a non-negative submartingale and so we can apply Doob's inequality directly to it, from which

$$\lambda \mathbb{P}\left(\sup_{1 \leq i \leq n} X_{t_i}^- \geq \lambda\right) \leq \mathbb{E}[X_T^-],$$

and, since $\mathbb{E}[X_T^-] \leq \mathbb{E}[|X_T|]$, taking the (monotone) limit in nested sequences in $[0, T] \cap \mathbb{Q}$, gives the result. \square

5.3 Convergence and regularisation theorems

As advertised, our aim in this section is to prove that, provided the filtration satisfies 'the usual conditions', any martingale has a version with right continuous paths.

First we recall the notion of upcrossing numbers.

Definition 5.8 Let $f : I \rightarrow \mathbb{R}$ be a function defined on a subset I of $[0, \infty)$. If $a < b$, the upcrossing number of f along $[a, b]$, which we shall denote $U([a, b], (f_t)_{t \in I})$ is the maximal integer $k \geq 1$ such that there exists a sequence $s_1 < t_1 < \dots < s_k < t_k$ of elements of I such that $f(s_i) \leq a$ and $f(t_i) \geq b$ for every $i = 1, \dots, k$.

If even for $k = 1$ there is no such sequence, we take $U([a, b], (f_t)_{t \in I}) = 0$. If such a sequence exists for every $k \geq 1$, we set $U([a, b], (f_t)_{t \in I}) = \infty$.

Upcrossing numbers are a convenient tool for studying the regularity of functions. We omit the proof of the following analytic lemma.

Lemma 5.9 Let D be a countable dense set in $[0, \infty)$ and let f be a real function defined on D . Assume that for every $T \in D$

- i. f is bounded on $D \cap [0, T]$;
- ii. for all rationals a and b such that $a < b$

$$U([a, b], (f_t)_{t \in D \cap [0, T]}) < \infty.$$

Then the right limit

$$f(t+) = \lim_{s \downarrow t, s \in D} f(s)$$

exists for every real $t \geq 0$, and similarly the left limit

$$f(t-) = \lim_{s \uparrow t, s \in D} f(s)$$

exists for any real $t > 0$.

Furthermore, the function $g : \mathbb{R}_+ \rightarrow \mathbb{R}$ defined by $g(t) = f(t+)$ is càdlàg ('continue à droite avec des limites à gauche'; i.e. right continuous with left limits) at every $t > 0$.

Lemma 5.10 (Doob's upcrossing lemma in discrete time) *Let $(X_t)_{t \geq 0}$ be a supermartingale and F a finite subset of $[0, T]$. If $a < b$ then*

$$\mathbb{E} \left[U([a, b], (X_n : n \in F)) \right] \leq \sup_{n \in F} \frac{\mathbb{E}[(X_n - a)^-]}{b - a} \leq \frac{\mathbb{E}[(X_T - a)^-]}{b - a}.$$

The last inequality follows since $(X_t - a)^-$ is a submartingale.

Taking an increasing sequence F_n and setting $\cup_n F_n = F$, this immediately extends to a countable $F \subset [0, T]$. From this we deduce:

Theorem 5.11 *If (X_t) is a right-continuous super-martingale and $\sup_t \mathbb{E}[X_t^-] < \infty$ then $\lim_{t \rightarrow \infty} X_t$ exists a.s. In particular, a non-negative right-continuous supermartingale converges a.s. as $t \rightarrow \infty$.*

Proof. This is immediate since, by right continuity, upcrossings of $[a, b]$ over $t \in [0, \infty)$ are the same as upcrossings of (a, b) over $t \in [0, \infty) \cap \mathbb{Q}$ and a sequence $(x_n)_{n \geq 1}$ converges if and only if $U([a, b], (x_n)_{n \geq 1}) < \infty$ for all $a < b$ with $a, b \in \mathbb{Q}$. \square

The next result says that we can talk about left and right limits for a general supermartingale and then our analytic lemma will tell us how to find a càdlàg version.

Theorem 5.12 *If $(X_t : t \geq 0)$ is a supermartingale then for \mathbb{P} -almost every $\omega \in \Omega$,*

$$\forall t \in (0, \infty) \quad \lim_{r \uparrow t, r \in \mathbb{Q}} X_r(\omega) \text{ and } \lim_{r \downarrow t, r \in \mathbb{Q}} X_r(\omega) \text{ exist and are finite.} \quad (21)$$

Proof. Fix $T > 0$. From Lemmas 5.7 and 5.10, there exists $\Omega^T \subseteq \Omega$, with $\mathbb{P}(\Omega^T) = 1$, such that for any $\omega \in \Omega^T$

$$\forall a, b \in \mathbb{Q} \text{ with } a < b, \quad U([a, b], (X_t(\omega) : t \in [0, T] \cap \mathbb{Q})) < \infty,$$

and

$$\sup_{t \in [0, T] \cap \mathbb{Q}} |X_t(\omega)| < \infty.$$

It follows that the limits in (21) are well defined and finite for all $t \leq T$ and $\omega \in \Omega^T$. To complete the proof, take $\Omega := \Omega^1 \cap \Omega^2 \cap \Omega^3 \cap \dots$ \square

Using this, even if X is not right-continuous, its right-continuous version is a.s. well defined. The following fundamental regularisation result is again due to Doob.

Theorem 5.13 *Let X be a supermartingale with respect to a right-continuous and complete filtration (\mathcal{F}_t) . If $t \mapsto \mathbb{E}[X_t]$ is right continuous (e.g. if X is a martingale) then X admits a modification with càdlàg paths, which is also an (\mathcal{F}_t) -supermartingale.*

Remark 5.14 *If X is a martingale then its càdlàg modification is also a martingale.*

Proof. By Theorem 5.12, there exists $\Omega_0 \subseteq \Omega$, with $\mathbb{P}[\Omega_0] = 1$, such that the process

$$X_{t+}(\omega) = \begin{cases} \lim_{r \downarrow t, r \in \mathbb{Q}} X_r(\omega) & \omega \in \Omega_0 \\ 0 & \omega \notin \Omega_0 \end{cases}$$

is well defined and adapted to (\mathcal{F}_t) . By Lemma 5.9, it has càdlàg paths.

To check that we really have only produced a modification of X_t , that is $X_t = X_{t+}$ almost surely, let $t_n \downarrow t$ be a sequence of rationals. Set $Y_k = X_{t-k}$ for every integer $k \leq 0$. Then Y is a backwards supermartingale with respect to the (backward) discrete filtration $\mathcal{H}_k = \mathcal{F}_{t-k}$ and $\sup_{k \leq 0} \mathbb{E}[|Y_k|] < \infty$. The convergence theorem for backwards supermartingales (see Appendix C) then implies that X_{t_n} converges to X_{t+} in L^1 . In particular, $X_{t+} \in L^1$ and thanks to L^1 convergence, we can pass to the limit $n \rightarrow \infty$ in the inequality $X_t \geq \mathbb{E}[X_{t_n} | \mathcal{F}_t]$ to obtain $X_t \geq \mathbb{E}[X_{t+} | \mathcal{F}_t]$.

Right continuity of $t \mapsto \mathbb{E}[X_t]$ implies $\mathbb{E}[X_{t+} - X_t] = 0$, so that $X_t = X_{t+}$ almost surely.

Now, to check the supermartingale property, let $s < t$ and let $(s_n)_{n \geq 0}$ be a sequence of rationals decreasing to s . Assume that $s_n < t$ for all n . Then, as above, $X_{s_n} \rightarrow X_{s+} \in L^1$, so if $A \in \mathcal{F}_{s+}$, which implies $A \in \mathcal{F}_{s_n}$ for every n , with t_n as above,

$$\begin{aligned} \mathbb{E}[X_{s+} \mathbf{1}_A] &= \lim_{n \rightarrow \infty} \mathbb{E}[X_{s_n} \mathbf{1}_A] \geq \lim_{n \rightarrow \infty} \mathbb{E}[X_{t_n} \mathbf{1}_A] \\ &= \mathbb{E}[X_{t+} \mathbf{1}_A] = \mathbb{E}[\mathbb{E}[X_{t+} | \mathcal{F}_{s+}] \mathbf{1}_A]. \end{aligned}$$

Since this holds for all $A \in \mathcal{F}_{s+}$ and since X_{s+} and $\mathbb{E}[X_{t+} | \mathcal{F}_{s+}]$ are both \mathcal{F}_{s+} -measurable, this shows

$$X_{s+} \geq \mathbb{E}[X_{t+} | \mathcal{F}_{s+}].$$

For martingales all the inequalities can be replaced by equalities and for submartingales we use that $-X$ is a supermartingale. \square

Remark 5.15 *Let's make some comments on the assumptions of the theorem.*

- i. The assumption that the filtration is right continuous is necessary. For example, let $\Omega = \{-1, +1\}$, $\mathbb{P}[\{1\}] = 1/2 = \mathbb{P}[\{-1\}]$. We set $\varepsilon(\omega) = \omega$ and $X_t = 0$ if $0 \leq t \leq 1$, $X_t = \varepsilon$ if $t > 1$. Then X is a martingale with respect to the canonical filtration (which is complete since there are no nonempty negligible sets), but no modification of X can be right continuous at $t = 1$.*

- ii. Similarly, take $X_t = f(t)$, where $f(t)$ is deterministic, non-increasing and not right continuous. Then no modification can have right continuous sample paths.

5.4 Martingale convergence and optional stopping

Theorem 5.16 *Let X be a supermartingale with right continuous sample paths. Assume that $(X_t)_{t \geq 0}$ is bounded in L^1 . Then there exists $X_\infty \in L^1$ such that $\lim_{t \rightarrow \infty} X_t = X_\infty$ almost surely.*

Proof. Let D be a countable dense subset of \mathbb{R}_+ . We showed before that for $a < b$

$$\mathbb{E}[U([a, b], (X_t)_{t \in D \cap [0, T]})] \leq \frac{1}{b-a} \mathbb{E}[(X_T - a)^-].$$

By the Monotone Convergence Theorem,

$$\mathbb{E}[U([a, b], (X_t)_{t \in D})] \leq \frac{1}{b-a} \sup_{t \geq 0} \mathbb{E}[(X_t - a)^-] < \infty$$

(since X is bounded in L^1). This implies that $X_\infty = \lim_{D \ni t \rightarrow \infty} X_t$ exists a.s. in $[-\infty, \infty]$.

We can exclude the values $\pm\infty$ since Fatou's Lemma gives

$$\mathbb{E}[|X_\infty|] \leq \liminf_{D \ni t \rightarrow \infty} \mathbb{E}[|X_t|] < \infty,$$

so $X_\infty \in L^1$. Right continuity of sample paths allows us to remove the restriction that $t \in D$. □

Under the assumptions of this theorem, X_t may not converge to X_∞ in L^1 . The next result gives, for martingales, necessary and sufficient conditions for L^1 -convergence.

Definition 5.17 *A martingale is said to be closed if there exists a random variable $Z \in L^1$ such that for every $t \geq 0$, $X_t = \mathbb{E}[Z | \mathcal{F}_t]$.*

Theorem 5.18 *Let $(X_t : t \geq 0)$ be a martingale with right continuous sample paths. Then TFAE:*

- i. X is closed;
- ii. the collection $(X_t)_{t \geq 0}$ is uniformly integrable;
- iii. X_t converges almost surely and in L^1 as $t \rightarrow \infty$.

Moreover, if these properties hold, $X_t = \mathbb{E}[X_\infty | \mathcal{F}_t]$ for every $t \geq 0$, where $X_\infty \in L^1$ is the almost sure limit of X_t as $t \rightarrow \infty$.

Proof. That the first condition implies the second is easy. If $Z \in L^1$, then $\mathbb{E}[Z|\mathcal{G}]$, where \mathcal{G} varies over sub σ -fields of \mathcal{F} is uniformly integrable.

If the second condition holds, then, in particular, $(X_t)_{t \geq 0}$ is bounded in L^1 and, by Theorem 5.16, $X_t \rightarrow X_\infty$ almost surely. By the uniform integrability, we also have convergence in L^1 , so the third condition is satisfied.

Finally, if the third condition holds, for every $s \geq 0$, pass to the limit as $t \rightarrow \infty$ in the equality $X_s = \mathbb{E}[X_t|\mathcal{F}_s]$ (using the fact that conditional expectation is continuous for the L^1 -norm) and obtain $X_s = \mathbb{E}[X_\infty|\mathcal{F}_s]$. \square

We should now like to establish conditions under which we have an optional stopping theorem for continuous martingales. As usual, our starting point will be the corresponding discrete time result and we shall pass to a suitable limit.

Theorem 5.19 (Optional stopping for uniformly integrable discrete martingales)

Let $(Y_n)_{n \in \mathbb{N}}$ be a uniformly integrable martingale with respect to the filtration $(\mathcal{G}_n)_{n \in \mathbb{N}}$, and let Y_∞ be the a.s. limit of Y_n when $n \rightarrow \infty$. Then, for every choice of the stopping times S and T such that $S \leq T$, we have $Y_T \in L^1$ and

$$Y_S = \mathbb{E}[Y_T|\mathcal{G}_S],$$

where

$$\mathcal{G}_S = \{A \in \mathcal{G}_\infty : A \cap \{S = n\} \in \mathcal{G}_n \text{ for every } n \in \mathbb{N}\},$$

with the convention that $Y_T = Y_\infty$ on the event $\{T = \infty\}$, and similarly for Y_S .

Let $(X_t)_{t \geq 0}$ be a right continuous martingale or supermartingale such that X_t converges almost surely as $t \rightarrow \infty$ to a limit X_∞ . Then for every stopping time T , we define

$$X_T(\omega) = \mathbf{1}_{\{T(\omega) < \infty\}} X_{T(\omega)}(\omega) + \mathbf{1}_{\{T(\omega) = \infty\}} X_\infty(\omega).$$

Theorem 5.20 Let $(X_t)_{t \geq 0}$ be a uniformly integrable martingale with right continuous sample paths. Let S and T be two stopping times with $S \leq T$. Then X_S and X_T are in L^1 and $X_S = \mathbb{E}[X_T|\mathcal{F}_S]$.

In particular, for every stopping time S we have $X_S = \mathbb{E}[X_\infty|\mathcal{F}_S]$ and $\mathbb{E}[X_S] = \mathbb{E}[X_\infty] = \mathbb{E}[X_0]$.

Proof. For any integer $n \geq 0$ set

$$T_n = \sum_{k=0}^{\infty} \frac{k+1}{2^n} \mathbf{1}_{\{k2^{-n} < T \leq (k+1)2^{-n}\}} + \infty \mathbf{1}_{\{T = \infty\}},$$

$$S_n = \sum_{k=0}^{\infty} \frac{k+1}{2^n} \mathbf{1}_{\{k2^{-n} < S \leq (k+1)2^{-n}\}} + \infty \mathbf{1}_{\{S = \infty\}}.$$

Then T_n and S_n are sequences of stopping times that decrease respectively to T and S . Moreover, $S_n \leq T_n$ for every $n \geq 0$.

For each fixed n , $2^n S_n$ and $2^n T_n$ are stopping times of the discrete filtration $\mathcal{H}_n = \mathcal{F}_{k/2^n}$ and $Y_k^{(n)} = X_{k/2^n}$ is a discrete martingale with respect to this filtration.

From Theorem 5.19, $Y_{2^n S_n}^{(n)}$ and $Y_{2^n T_n}^{(n)}$ are in L^1 and

$$X_{S_n} = Y_{2^n S_n}^{(n)} = \mathbb{E}[Y_{2^n T_n}^{(n)} | \mathcal{H}_{2^n S_n}] = \mathbb{E}[X_{T_n} | \mathcal{F}_{S_n}].$$

Let $A \in \mathcal{F}_S$. Since $\mathcal{F}_S \subseteq \mathcal{F}_{S_n}$ we have $A \in \mathcal{F}_{S_n}$ and so $\mathbb{E}[\mathbf{1}_A X_{S_n}] = \mathbb{E}[\mathbf{1}_A X_{T_n}]$. By right continuity, $X_S = \lim_{n \rightarrow \infty} X_{S_n}$ and $X_T = \lim_{n \rightarrow \infty} X_{T_n}$. The limits also hold in L^1 (in fact, by Theorem 5.19, $X_{S_n} = \mathbb{E}[X_\infty | \mathcal{F}_{S_n}]$ for every n and so $(X_{S_n})_{n \geq 1}$ and $(X_{T_n})_{n \geq 1}$ are uniformly integrable). L^1 convergence implies that the limits X_S and X_T are in L^1 and allows us to pass to a limit, $\mathbb{E}[\mathbf{1}_A X_S] = \mathbb{E}[\mathbf{1}_A X_T]$. This holds for all $A \in \mathcal{F}_S$ and so since X_S is \mathcal{F}_S -measurable we conclude that $X_S = \mathbb{E}[X_T | \mathcal{F}_S]$, as required. \square

Corollary 5.21 *In particular, for any martingale with right continuous paths and two bounded stopping times, $S \leq T$, we have $X_S, X_T \in L^1$ and $X_S = \mathbb{E}[X_T | \mathcal{F}_S]$.*

Proof. Let a be such that $S \leq T \leq a$. The martingale $(X_{t \wedge a})_{t \geq 0}$ is closed by X_a and so we may apply our previous results. \square

Corollary 5.22 *Suppose that $(X_t)_{t \geq 0}$ is a martingale with right continuous paths and T is a stopping time.*

- i. $(X_{t \wedge T})_{t \geq 0}$ is a martingale;
- ii. if, in addition, $(X_t)_{t \geq 0}$ is uniformly integrable, then $(X_{t \wedge T})_{t \geq 0}$ is uniformly integrable and for every $t \geq 0$, $X_{t \wedge T} = \mathbb{E}[X_T | \mathcal{F}_t]$.

We omit the proof which can be found in Le Gall.

Above all, optional stopping is a powerful tool for explicit calculations.

Example 5.23 *Fix $a > 0$ and let T_a be the first hitting time of a by standard Brownian motion. Then for each $\lambda > 0$,*

$$\mathbb{E}[e^{-\lambda T_a}] = e^{-a\sqrt{2\lambda}}.$$

Recall that $N_t^\lambda = \exp(\lambda B_t - \frac{\lambda^2}{2}t)$ is a martingale. So $N_{t \wedge T_a}^\lambda$ is still a martingale and it is bounded above by $e^{\lambda a}$ and hence is uniformly integrable, so $\mathbb{E}[N_{T_a}^\lambda] = \mathbb{E}[N_0^\lambda]$. That is,

$$e^{a\lambda} \mathbb{E}[e^{-\lambda^2 T_a / 2}] = \mathbb{E}[N_0^\lambda] = 1.$$

Replace λ by $\sqrt{2\lambda}$ and rearrange.

Warning: This argument fails if $\lambda < 0$ - the reason being that we lose the uniform integrability.

6 Continuous semimartingales

Recall that our original goal was to make sense of differential equations driven by ‘rough’ inputs. In fact, we’ll recast our differential equations as integral equations and so we must develop a theory that allows us to integrate with respect to ‘rough’ driving processes. The class of processes with which we work are called *semimartingales*, and we shall specialise to the continuous ones.

We’re going to start with functions for which the integration theory that you already know is adequate - these are called functions of finite variation.

Throughout, we assume that a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ satisfying the usual conditions is given.

6.1 Functions of finite variation

Throughout this section we only consider real-valued right-continuous functions on $[0, \infty)$. Our arguments will be shift invariant so, without loss of generality, we assume that any such function a satisfies $a(0) = 0$. Recall the following definition.

Definition 6.1 *The (total) variation of a function a over $[0, T]$ is defined as*

$$V(a)_T = \sup_{\pi} \sum_{i=0}^{n_{\pi}-1} |a_{t_{i+1}} - a_{t_i}|,$$

where the supremum is over partitions $\pi = \{0 = t_0 < t_1 < \dots < t_{n_{\pi}} = T\}$ of $[0, T]$. We say that a is of finite variation on $[0, T]$ if $V(a)_T < \infty$. The function a is of finite variation if $V(a)_T < \infty$ for all $T \geq 0$ and of bounded variation if $\lim_{T \rightarrow \infty} V(a)_T < \infty$.

Remark 6.2 *Note that $t \rightarrow V(a)_t$ is non-negative, right-continuous and non-decreasing in t . This follows since any partition of $[0, s]$ may be included in a partition of $[0, t]$, $t \geq s$.*

Proposition 6.3 *The function a is of finite variation if and only if it is equal to the difference of two non-decreasing functions, a_1 and a_2 .*

Moreover, if a is of finite variation, then a_1 and a_2 can be chosen so that $V(a)_t = a_1(t) + a_2(t)$. If a is càdlàg then $V(a)_t$ is also càdlàg.

Proof.

$$V(a)_t - a(t) = \sup_{\pi} \sum_{i=0}^{n(\pi)-1} (|a(t_{i+1}) - a(t_i)| - (a(t_{i+1}) - a(t_i)))$$

is an non-decreasing function of t , as is $V(a)_t + a(t)$. □

If we define measures μ_+ , μ_- by

$$\mu_+((0, t]) = \frac{V(a)_t + a(t)}{2}, \quad \mu_-((0, t]) = \frac{V(a)_t - a(t)}{2},$$

then we can develop a theory of integration with respect to a by declaring that

$$\int_0^t f(s) da(s) = \int_0^t f(s) \mu_+(ds) - \int_0^t f(s) \mu_-(ds),$$

provided that

$$\int_0^t |f(s)| |\mu|(ds) = \int_0^t |f(s)| (\mu_+(ds) + \mu_-(ds)) < \infty.$$

We say that $\mu = \mu_+ - \mu_-$ is the *signed measure* associated with a , μ_+ , μ_- is its *Jordan decomposition* and $\int_0^t f(s) da(s)$ is the *Lebesgue-Stieltjes* integral of f with respect to a .

We sometimes use the notation

$$(f \cdot a)(t) = \int_0^t f(s) da(s).$$

The function $(f \cdot a)$ will be right continuous and of finite variation whenever a is finite variation and f is a -integrable (exercise).

Proposition 6.4 (Associativity) *Let a be of finite variation as above and f, g measurable functions, f is a -integrable and g is $(f \cdot a)$ -integrable. Then gf is a -integrable and*

$$\int_0^t g(s) d(f \cdot a)(s) = \int_0^t g(s) f(s) da(s).$$

In our ‘dot’-notation:

$$g \cdot (f \cdot a) = (gf) \cdot a. \quad (22)$$

Proposition 6.5 (Stopping) *Let a be of finite variation as above and fix $t \geq 0$. Set $a^t(s) = a(t \wedge s)$. Then a^t is of finite variation and for any measurable a -integrable function f*

$$\int_0^{u \wedge t} f(s) da(s) = \int_0^u f(s) da^t(s) = \int_0^u f(s) \mathbf{1}_{[0,t]}(s) da(s), \quad u \in [0, \infty].$$

Proposition 6.6 (Integration by parts) *Let a and b be two continuous functions of finite variation with $a(0) = b(0) = 0$. Then for any t*

$$a(t)b(t) = \int_0^t a(s) db(s) + \int_0^t b(s) da(s).$$

Proposition 6.7 (Chain-rule) *If F is a C^1 function and a is continuous of finite variation, then $F(a(t))$ is also of finite variation and*

$$F(a(t)) = F(a(0)) + \int_0^t F'(a(s)) da(s).$$

Proof. The statement is trivially true for $F(x) = x$. Now by Proposition 6.6, it is straightforward to check that if the statement is true for F , then it is also true for $xF(x)$. Hence, by induction, the statement holds for all polynomials. To complete the proof, approximate $F \in C^1$ by a sequence of polynomials. \square

Proposition 6.8 (Change of variables) *If a is non-decreasing and right-continuous then so is*

$$c(s) := \inf\{t \geq 0 : a(t) > s\},$$

where $\inf \emptyset = +\infty$. Let $a(0) = 0$. Then, for any Borel measurable function $f \geq 0$ on \mathbb{R}_+ , we have

$$\int_0^\infty f(u) da(u) = \int_0^{a(\infty)} f(c(s)) ds.$$

Proof. If $f(u) = \mathbf{1}_{[0, v]}(u)$, then the claim becomes

$$a(v) = \int_0^\infty \mathbf{1}_{\{c(s) \leq v\}} ds = \inf\{s : c(s) > v\},$$

and equality holds by definition of c . Take differences to get indicators of sets $(u, v]$. The Monotone Class Theorem allows us to extend to functions of compact support and then take increasing limits to obtain the formula in general. \square

6.2 Processes of finite variation

Recall that a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t), \mathbb{P})$ satisfying the usual conditions is given.

Definition 6.9 *An adapted right-continuous process $A = (A_t : t \geq 0)$ is called a finite variation process (or a process of finite variation) if $A_0 = 0$ and $t \mapsto A_t$ is (a function) of finite variation a.s..*

Proposition 6.10 *Let A be a finite variation process and K a progressively measurable process s.t.*

$$\forall t \geq 0, \forall \omega \in \Omega, \int_0^t |K_s(\omega)| |dA_s(\omega)| < \infty.$$

Then $((K \cdot A)_t : t \geq 0)$, defined as $(K \cdot A)_t(\omega) := \int_0^t K_s(\omega) dA_s(\omega)$, is a finite variation process.

Proof. The right continuity is immediate from the deterministic theory, but we need to check that $(K \cdot A)_t$ is adapted. For this we check that if $t > 0$ is fixed and $h : [0, t] \times \Omega \rightarrow \mathbb{R}$ is measurable with respect to $\mathcal{B}([0, t]) \otimes \mathcal{F}_t$, and if

$$\int_0^t |h(s, \omega)| |dA_s(\omega)| < \infty$$

for every $\omega \in \Omega$, then

$$\int_0^t h(s, \omega) dA_s(\omega)$$

is \mathcal{F}_t -measurable.

Fix $t > 0$. Consider first h defined by $h(s, \omega) = \mathbf{1}_{(u,v]}(s) \mathbf{1}_\Gamma(\omega)$ for $(u, v] \subseteq [0, t]$ and $\Gamma \in \mathcal{F}_t$. Then

$$(h \cdot A)_t = \mathbf{1}_\Gamma(A_v - A_u)$$

is \mathcal{F}_t -measurable. By the Monotone Class Theorem, $(h \cdot A)_t$ is \mathcal{F}_t -measurable for any $h = \mathbf{1}_G$ with $G \in \mathcal{B}([0, t]) \otimes \mathcal{F}_t$, or, more generally, any bounded $\mathcal{B}([0, t]) \otimes \mathcal{F}_t$ -measurable function h . If h is a general $\mathcal{B}([0, t]) \otimes \mathcal{F}_t$ -measurable function satisfying

$$\int_0^t |h(s, \omega)| |dA_s(\omega)| < \infty \quad \forall \omega \in \Omega,$$

then h is a pointwise limit, $h = \lim_{n \rightarrow \infty} h_n$, of simple functions with $|h| \geq |h_n|$. The integrals $\int h_n(s, \omega) dA_s(\omega)$ converge by the Dominated Convergence Theorem, and hence $\int_0^t h(s, \omega) dA_s(\omega)$ is also \mathcal{F}_t -measurable (as a limit of \mathcal{F}_t -measurable functions). In particular, $(K \cdot A)_t(\omega)$ is \mathcal{F}_t -measurable since by progressive measurability, $(s, \omega) \mapsto K_s(\omega)$ on $[0, t]$ is $\mathcal{B}([0, t]) \otimes \mathcal{F}_t$ -measurable. \square

It is worth recording that our integral can be obtained through the limiting procedure that one might expect. Let $f : [0, T] \rightarrow \mathbb{R}$ be continuous and $0 = t_0^n < t_1^n < \dots < t_{p_n}^n = T$ be a sequence of partitions of $[0, T]$ with mesh tending to zero. Then

$$\int_0^T f(s) da(s) = \lim_{n \rightarrow \infty} \sum_{i=1}^{p_n} f(t_{i-1}^n) (a(t_i^n) - a(t_{i-1}^n)).$$

The proof is easy: let $f_n : [0, T] \rightarrow \mathbb{R}$ be defined by $f_n(s) = f(t_{i-1}^n)$ if $s \in (t_{i-1}^n, t_i^n]$, $1 \leq i \leq p_n$, and $f_n(0) = 0$. Then

$$\sum_{i=1}^{p_n} f(t_{i-1}^n) (a(t_i^n) - a(t_{i-1}^n)) = \int_{[0, T]} f_n(s) \mu(ds),$$

where μ is the signed measure associated with a . The desired result now follows by the Dominated Convergence Theorem.

In the argument above, f_n took the value of f at the *left* endpoint of each interval. In the finite variation case, we could equally have approximated by f_n taking the value of f at the midpoint of the interval, or the right hand endpoint, or any other point in between. However, we are going to extend our theory to processes that do *not* have finite variation and then, even if f is continuous, it matters whether f_n takes the value of f at the left or right endpoint of the intervals.

The processes that make our theory work are slight generalisations of martingales.

6.3 Continuous local martingales

Definition 6.11 An adapted process $(M_t : t \geq 0)$ is called a continuous local martingale if $M_0 = 0$, it has continuous trajectories a.s. and if there exists a non-decreasing sequence of stopping times $(\tau_n)_{n \geq 1}$ such that $\tau_n \uparrow \infty$ a.s. and for each n , $M^{\tau_n} = (M_{t \wedge \tau_n} : t \geq 0)$ is a (uniformly integrable) martingale. We say (τ_n) reduces M .

More generally, when we do not assume that $M_0 = 0$, we say that M is a continuous local martingale if $N_t = M_t - M_0$ is a continuous local martingale.

Any martingale is a local martingale, but the converse is false.

Proposition 6.12 i. A non-negative continuous local martingale such that $M_0 \in L^1$ is a supermartingale.

ii. A continuous local martingale M such that there exists a random variable $Z \in L^1$ with $|M_t| \leq Z$ for every $t \geq 0$ is a uniformly integrable martingale.

iii. If M is a continuous local martingale and $M_0 = 0$ (or more generally $M_0 \in L^1$), the sequence of stopping times

$$T_n = \inf\{t \geq 0 : |M_t| \geq n\}$$

reduces M .

iv. If M is a continuous local martingale, then for any stopping time ρ , the stopped process M^ρ is also a continuous local martingale.

Proof. (i) Write $M_t = M_0 + N_t$. By definition, there exists a sequence T_n of stopping times that reduces N . Thus, if $s \leq t$, for every n ,

$$N_{s \wedge T_n} = \mathbb{E}[N_{t \wedge T_n} | \mathcal{F}_s].$$

We can add M_0 to both sides (M_0 is \mathcal{F}_s -measurable and in L^1) and we find

$$M_{s \wedge T_n} = \mathbb{E}[M_{t \wedge T_n} | \mathcal{F}_s].$$

Since M takes non-negative values, let $n \rightarrow \infty$ and apply Fatou's lemma for conditional expectations to find

$$M_s \geq \mathbb{E}[M_t | \mathcal{F}_s]. \quad (23)$$

Taking $s = 0$, $\mathbb{E}[M_t] \leq \mathbb{E}[M_0] < \infty$. So $M_t \in L^1$ for every $t \geq 0$, and (23) says that M is a supermartingale.

(ii) By the same argument,

$$M_{s \wedge T_n} = \mathbb{E}[M_{t \wedge T_n} | \mathcal{F}_s].$$

Since $|M_{t \wedge T_n}| \leq Z$, this time apply the Dominated Convergence Theorem to see that $M_{t \wedge T_n}$ converges in L^1 (to M_t) and $M_s = \mathbb{E}[M_t | \mathcal{F}_s]$.

The other two statements are immediate. \square

Theorem 6.13 *A continuous local martingale M with $M_0 = 0$ a.s., is a process of finite variation if and only if M is indistinguishable from zero.*

Proof. Assume M is a continuous local martingale and of finite variation. Let

$$\tau_n = \inf\{t \geq 0 : \int_0^t |dM_s| \geq n\} = \inf\{t \geq 0 : V(M)_t \geq n\},$$

which are stopping times since $V(M)_t = \int_0^t |dM_s|$ is continuous and adapted. Let $N = M^{\tau_n}$, which is bounded since

$$|N_t| = |M_{t \wedge \tau_n}| \leq \left| \int_0^{t \wedge \tau_n} dM_u \right| \leq \int_0^{t \wedge \tau_n} |dM_u| \leq n,$$

and hence (N_t) is a martingale.

Let $t > 0$ and $\pi = \{0 = t_0 < t_1 < t_2 < \dots < t_{m(\pi)} = t\}$ be a partition of $[0, t]$. Then

$$\begin{aligned} \mathbb{E}[N_t^2] &= \sum_{i=1}^{m(\pi)} \mathbb{E}[N_{t_i}^2 - N_{t_{i-1}}^2] = \sum_{i=1}^{m(\pi)} \mathbb{E}[(N_{t_i} - N_{t_{i-1}})^2] \\ &\leq \mathbb{E} \left[\left(\sup_{1 \leq i \leq m(\pi)} |N_{t_i} - N_{t_{i-1}}| \right) \cdot \underbrace{\sum_{i=1}^{m(\pi)} |N_{t_i} - N_{t_{i-1}}|}_{\leq V(N)_t = V(M)_{t \wedge \tau_n} \leq n} \right] \\ &\leq n \mathbb{E} \left[\sup_{1 \leq i \leq m(\pi)} |N_{t_i} - N_{t_{i-1}}| \right] \rightarrow 0 \quad \text{as } \delta(\pi) \rightarrow 0 \end{aligned}$$

(where $\delta(\pi)$ is the mesh of π), by the Dominated Convergence Theorem (since $|N_{t_i} - N_{t_{i-1}}| \leq V(N)_t \leq n$ and so n is a dominating function).

It then follows by Fatou's Lemma that

$$\mathbb{E}[M_t^2] = \mathbb{E}[\lim_{n \rightarrow \infty} M_{t \wedge \tau_n}^2] \leq \lim_{n \rightarrow \infty} \mathbb{E}[M_{t \wedge \tau_n}^2] = 0$$

which implies that $M_t = 0$ a.s., and so by continuity of paths, $\mathbb{P}[M_t = 0 \forall t \geq 0] = 1$. \square

6.4 Quadratic variation of a continuous local martingale

If our martingales are going to be interesting, then they're going to have unbounded variation. But remember that we said that we'd use Brownian motion as a basic building block, and that while Brownian motion has infinite variation, it has bounded *quadratic variation*, defined over $[0, T]$ by

$$\lim_{\delta(\pi_n) \rightarrow 0} \sum_{j=1}^{N(\pi_n)} (B_{t_j} - B_{t_{j-1}})^2 = T.$$

We are now going to see that the analogue of this process exists for any continuous local martingale. Ultimately, we shall see that the quadratic variation is in some sense a 'clock' for a local martingale, but that will be made more precise in the very last result of the course.

Theorem 6.14 *Let M be a continuous local martingale. There exists a unique (up to indistinguishability) non-decreasing, continuous adapted finite variation process $(\langle M, M \rangle_t : t \geq 0)$, starting in zero, such that $(M_t^2 - \langle M, M \rangle_t : t \geq 0)$ is a continuous local martingale.*

Furthermore, for any $T > 0$ and any sequence of partitions $\pi_n = \{0 = t_0^n < t_1^n < \dots < t_{n(\pi_n)}^n = T\}$ with $\delta(\pi_n) = \sup_{1 \leq i \leq n(\pi_n)} (t_i^n - t_{i-1}^n) \rightarrow 0$ as $n \rightarrow \infty$

$$\langle M, M \rangle_T = \lim_{n \rightarrow \infty} \sum_{i=1}^{n(\pi_n)} (M_{t_i^n} - M_{t_{i-1}^n})^2, \quad (24)$$

where the limit is in probability.

The process $\langle M, M \rangle$ is called the quadratic variation of M , or simply the increasing process of M , and is often denoted $\langle M, M \rangle_t = \langle M \rangle_t$.

Proof.[Sketch of the Proof (NOT EXAMINABLE)]

Uniqueness is a direct consequence of Theorem 6.13 since if A, A' are two such processes then $(M^2 - A - (M^2 - A')) = A - A'$ is a local martingale starting in zero and of finite variation, which implies $A = A'$ by Theorem 6.13.

The idea of existence is as follows. First suppose that M is *bounded*. Take a sequence of partitions $0 = t_0^n < \dots < t_{p_n}^n = T$ with mesh tending to zero. Then check that

$$X_t^n := \sum_{i=1}^{n_p} M_{t_{i-1}^n} (M_{t_i^n \wedge t} - M_{t_{i-1}^n \wedge t})$$

is a (bounded) martingale. Now observe that

$$M_{t_j^n}^2 - 2X_{t_j^n}^n = \sum_{i=1}^j (M_{t_i^n} - M_{t_{i-1}^n})^2.$$

A direct computation gives

$$\lim_{n, m \rightarrow \infty} \mathbb{E}[(X_t^n - X_t^m)^2] = 0,$$

and by Doob's L^2 -inequality

$$\lim_{n, m \rightarrow \infty} \mathbb{E}[\sup_{t \leq T} (X_t^n - X_t^m)^2] = 0.$$

By passing to a subsequence, $X^{n_k} \rightarrow Y$ almost surely on $[0, T]$ where $(Y_t)_{t \leq T}$ is a continuous process which inherits the martingale property from X .

$$M_{t_j^n}^2 - 2X_{t_j^n}^n = \sum_{i=1}^j (M_{t_i^n} - M_{t_{i-1}^n})^2 =: QV_{t_j^n}^{\pi_n}(M)$$

is non-decreasing along $t_j^n : j \leq n(\pi_n)$. Letting $n \rightarrow \infty$, $M_t^2 - 2Y_t$ is almost surely non-decreasing and we set $\langle M, M \rangle_t = M_t^2 - 2Y_t$.

To move to a general continuous local martingale, we consider a sequence of stopped processes.

Details are in, for example, Le Gall's book. \square

Our theory of integration is going to be an ' L^2 -theory'. Let us introduce the martingales with which we are going to work. We are going to think of them as being defined up to indistinguishability - nothing changes if we change the process on a null set. Think of this as analogous to considering Lebesgue integrable functions as being defined 'almost everywhere'.

Definition 6.15 Let \mathbb{H}^2 be the space of L^2 -bounded càdlàg martingales, i.e.

$$((\mathcal{F}_t), \mathbb{P})\text{-martingales } M \text{ s.t. } \sup_{t \geq 0} \mathbb{E}[M_t^2] < \infty,$$

and \mathbb{H}^2 the subspace consisting of continuous L^2 -bounded martingales. Finally, let $\mathbb{H}_0^2 = \{M \in \mathbb{H}^2 : M_0 = 0 \text{ a.s.}\}$.

We note that the space \mathbb{H}^2 is also sometimes denoted \mathcal{M}^2 .

It follows from Doob's L^2 -inequality that

$$\mathbb{E} \left[\sup_{t \geq 0} M_t^2 \right] \leq 4 \sup_{t \geq 0} \mathbb{E}[M_t^2] < +\infty, \quad M \in \mathbb{H}^2.$$

Consequently, $\{M_t : t \geq 0\}$ is bounded by a square integrable random variable ($\sup_{t \geq 0} |M_t|$) and in particular is uniformly integrable. It follows from the Optional Stopping Theorem that $M_t = \mathbb{E}[M_\infty | \mathcal{F}_t]$ for some square integrable random variable M_∞ .

Conversely, we can start with a random variable $Y \in L^2(\Omega, \mathcal{F}_\infty, \mathbb{P})$ and define a martingale $M_t := \mathbb{E}[Y | \mathcal{F}_t] \in \mathbb{H}^2$ (and $M_\infty = Y$).

Two L^2 -bounded martingales M, M' are indistinguishable if and only if $M_\infty = M'_\infty$ a.s. and so if we endow \mathbb{H} with the norm

$$\|M\|_{\mathbb{H}^2} := \sqrt{\mathbb{E}[M_\infty^2]} = \|M_\infty\|_{L^2(\Omega, \mathcal{F}_\infty, \mathbb{P})}, \quad M \in \mathbb{H}^2, \quad (25)$$

then \mathbb{H}^2 can be identified with the familiar $L^2(\Omega, \mathcal{F}_\infty, \mathbb{P})$ space.

Theorem 6.16 \mathbb{H}^2 is a closed subspace of \mathbb{H}^2 .

Proof. This is almost a matter of writing down definitions. Suppose that the sequence $M^n \in \mathbb{H}^2$ converges in $\|\cdot\|_{\mathbb{H}^2}$ to some $M \in \mathbb{H}^2$. By Doob's L^2 -inequality

$$\mathbb{E} \left[\sup_{t \geq 0} |M_t^n - M_t|^2 \right] \leq 4 \|M^n - M\|_{\mathbb{H}^2}^2 \longrightarrow 0, \quad \text{as } n \rightarrow \infty.$$

Passing to a subsequence, we have $\sup_{t \geq 0} |M_t^{n_k} - M_t| \rightarrow 0$ a.s. and hence M has continuous paths a.s., which completes the proof. \square

For continuous local martingales, the norm in (25) can be re-expressed in terms of the quadratic variation:

Theorem 6.17 *Let M be a continuous local martingale with $M_0 \in L^2$.*

i. TFAE

- (a) M is a martingale, bounded in L^2 ;
- (b) $\mathbb{E}[\langle M, M \rangle_\infty] < \infty$.

Furthermore, if these properties hold, $M_t^2 - \langle M, M \rangle_t$ is a uniformly integrable martingale and, in particular, $\mathbb{E}[M_\infty^2] = \mathbb{E}[M_0^2] + \mathbb{E}[\langle M, M \rangle_\infty]$.

ii. TFAE

- (a) M is a martingale and $M_t \in L^2$ for every $t \geq 0$;
- (b) $\mathbb{E}[\langle M, M \rangle_t] < \infty$ for every $t \geq 0$.

Furthermore, if these properties hold, $M_t^2 - \langle M, M \rangle_t$ is a martingale.

Proof. The second statement will follow from the first on applying it to $M_{t \wedge a}$ for every choice of $a \geq 0$.

To prove the first set of equivalences, without loss of generality, suppose that $M_0 = 0$ (or replace M by $M - M_0$).

Suppose that M is a martingale, bounded in L^2 . Doob's L^2 -inequality implies that for every $T > 0$,

$$\mathbb{E}[\sup_{0 \leq t \leq T} M_t^2] \leq 4\mathbb{E}[M_T^2],$$

and so, letting $T \rightarrow \infty$,

$$\mathbb{E}[\sup_{t \geq 0} M_t^2] \leq 4 \sup_{t \geq 0} \mathbb{E}[M_t^2] = C < \infty.$$

Let $S_n = \inf\{t \geq 0 : \langle M, M \rangle_t \geq n\}$. Then the continuous local martingale $M_{t \wedge S_n}^2 - \langle M, M \rangle_{t \wedge S_n}$ is dominated by $\sup_{s \geq 0} M_s^2 + n$, which is integrable. By Proposition 6.12 this continuous local martingale is a uniformly integrable martingale and hence

$$\mathbb{E}[\langle M, M \rangle_{t \wedge S_n}] = \mathbb{E}[M_{t \wedge S_n}^2] \leq \mathbb{E}[\sup_{s \geq 0} M_s^2] \leq C < \infty.$$

Let n and then t tend to infinity and use the Monotone Convergence Theorem to obtain $\mathbb{E}[\langle M, M \rangle_\infty] < \infty$.

Conversely, assume that $\mathbb{E}[\langle M, M \rangle_\infty] < \infty$. Set $T_n = \inf\{t \geq 0 : |M_t| \geq n\}$. Then the continuous local martingale $M_{t \wedge T_n}^2 - \langle M, M \rangle_{t \wedge T_n}$ is dominated by $n^2 + \langle M, M \rangle_\infty$ which is integrable. From Proposition 6.12 again, this continuous local martingale is a uniformly integrable martingale and hence for every $t \geq 0$,

$$\mathbb{E}[M_{t \wedge T_n}^2] = \mathbb{E}[\langle M, M \rangle_{t \wedge T_n}] \leq \mathbb{E}[\langle M, M \rangle_\infty] = C' < \infty. \quad (26)$$

Let $n \rightarrow \infty$ and use Fatou's lemma to see that $\mathbb{E}[M_t^2] \leq C' < \infty$, so $(M_t)_{t \geq 0}$ is bounded in L^2 .

We still have to check that $(M_t)_{t \geq 0}$ is a martingale. However, (26) shows that $(M_{t \wedge T_n})_{n \geq 1}$ is uniformly integrable and so converges both almost surely and in L^1 to M_t for every $t \geq 0$. Recalling that M^{T_n} is a martingale, L^1 convergence allows us to pass to the limit as $n \rightarrow \infty$ in the martingale property, so M is a martingale.

Finally, if the two properties hold, then $M^2 - \langle M, M \rangle$ is dominated by $\sup_{t \geq 0} M_t^2 + \langle M, M \rangle_\infty$, which is integrable, and so Proposition 6.12 again says that $M^2 - \langle M, M \rangle$ is a uniformly integrable martingale. \square

Our previous theorem immediately yields that for a martingale M with $M_0 = 0$ we have

$$\|M\|_{H^2}^2 = \mathbb{E}[M_\infty^2] = \mathbb{E}[\langle M \rangle_\infty].$$

We can also deduce a complement to Theorem 6.13.

Corollary 6.18 *Let M be a continuous local martingale with $M_0 = 0$. Then the following are equivalent:*

- i. M is indistinguishable from zero,
- ii. $\langle M \rangle_t = 0$ for all $t \geq 0$ a.s.,
- iii. M is a process of finite variation.

In other words, there is nothing ‘in between’ finite variation and finite quadratic variation for this class of processes.

Proof. We already know that the first and third statements are equivalent. That the first implies the second is trivial, so we must just show that the second implies the first. We have $\langle M \rangle_\infty = \lim_{t \rightarrow \infty} \langle M \rangle_t = 0$. From Theorem 6.17, $M \in H^2$ and $\mathbb{E}[M_\infty^2] = \mathbb{E}[\langle M \rangle_\infty] = 0$ and so $M_t = \mathbb{E}[M_\infty | \mathcal{F}_t] = 0$ almost surely. \square

We can see that the quadratic variation of a martingale is telling us something about how its variance increases with time. We also need an analogous quantity for the ‘covariance’ between two martingales. This is most easily defined through *polarisation*.

Definition 6.19 *The quadratic co-variation between two continuous local martingales M, N is defined by*

$$\langle M, N \rangle := \frac{1}{2} (\langle M + N, M + N \rangle - \langle M, M \rangle - \langle N, N \rangle). \quad (27)$$

It is often called the (angle) bracket process of M and N .

Proposition 6.20 *For two continuous local martingales M, N*

- i. *the process $\langle M, N \rangle$ is the unique finite variation process, zero in zero, such that $(M_t N_t - \langle M, N \rangle_t : t \geq 0)$ is a continuous local martingale;*
- ii. *the mapping $M, N \mapsto \langle M, N \rangle$ is bilinear and symmetric;*

iii. for any stopping time τ ,

$$\langle M^\tau, N^\tau \rangle_t = \langle M^\tau, N \rangle_t = \langle M, N^\tau \rangle_t = \langle M, N \rangle_{\tau \wedge t}, \quad t \geq 0, \text{ a.s.}; \quad (28)$$

iv. for any $t > 0$ and a sequence of partitions π_n of $[0, t]$ with mesh converging to zero

$$\sum_{t_i \in \pi_n} (M_{t_{i+1}} - M_{t_i})(N_{t_{i+1}} - N_{t_i}) \rightarrow \langle M, N \rangle_t, \quad (29)$$

the convergence being in probability.

Proof. (i) $(M + N)_t^2 - \langle M + N, M + N \rangle_t$ is a continuous local martingale and by adding and subtracting terms it is equal to

$$\underbrace{M_t^2 - \langle M, M \rangle_t}_{\text{l.mat}} + \underbrace{N_t^2 - \langle N, N \rangle_t}_{\text{l.mat}} + \underbrace{2(M_t N_t - \frac{1}{2}(\langle M + N, M + N \rangle_t - \langle M, M \rangle_t - \langle N, N \rangle_t))}_{\text{hence also a l.mat}}$$

Uniqueness follows from Theorem 6.13.

(iv) Note that

$$(M_t + N_t - M_s - N_s)^2 - (M_t - M_s)^2 - (N_t - N_s)^2 = 2(M_t - M_s)(N_t - N_s).$$

The asserted convergence then follows from Theorem 6.14

(ii) Both properties follow from (iv). Symmetry is obvious from the definition in (27).

(iii) Follows from (iv). \square

Definition 6.21 Two continuous local martingales M, N , are said to be orthogonal if $\langle M, N \rangle = 0$.

For example, if B and B' are independent Brownian motions, then $\langle B, B' \rangle = 0$.

Remark 6.22 It follows that if M and N are two martingales bounded in L^2 and with $M_0 N_0 = 0$ a.s., then $(M_t N_t - \langle M, N \rangle_t, t \geq 0)$ is a uniformly integrable martingale. In particular, for any stopping time τ ,

$$\mathbb{E}[M_\tau N_\tau] = \mathbb{E}[\langle M, N \rangle_\tau]. \quad (30)$$

Take $M, N \in \mathbf{H}^2$, which we recall had the norm $\|M\|_{\mathbf{H}^2}^2 = \mathbb{E}[\langle M, M \rangle_\infty] = \mathbb{E}[M_\infty^2]$. Then we see that this norm is consistent with the inner product on $\mathbf{H}^2 \times \mathbf{H}^2$ given by $\mathbb{E}[\langle M, N \rangle_\infty] = \mathbb{E}[M_\infty N_\infty]$ and, by the usual Cauchy-Schwarz inequality,

$$\mathbb{E}[|\langle M, N \rangle_\infty|] = \mathbb{E}[|M_\infty N_\infty|] \leq \sqrt{\mathbb{E}[\langle M \rangle_\infty] \mathbb{E}[\langle N \rangle_\infty]}.$$

Actually, it is easy to obtain an almost sure version of this inequality, using that

$$|\sum (M_{t_{i+1}} - M_{t_i})(N_{t_{i+1}} - N_{t_i})| \leq \sqrt{\sum (M_{t_{i+1}} - M_{t_i})^2} \sqrt{\sum (N_{t_{i+1}} - N_{t_i})^2}$$

and taking limits to deduce that

$$|\langle M, N \rangle_t| \leq \sqrt{\langle M \rangle_t} \sqrt{\langle N \rangle_t}.$$

It's often convenient to have a more general version of this inequality.

Theorem 6.23 (Kunita-Watanabe inequality) *Let M, N be continuous local martingales and K, H two measurable processes. Then for all $0 \leq t \leq \infty$,*

$$\int_0^t |H_s| |K_s| |d\langle M, N \rangle_s| \leq \left(\int_0^t H_s^2 d\langle M \rangle_s \right)^{1/2} \left(\int_0^t K_s^2 d\langle N \rangle_s \right)^{1/2} \text{ a.s.} \quad (31)$$

We omit the proof which approximates H, K by simple functions and then essentially uses the Cauchy-Schwarz inequality for sums noted above.

6.5 Continuous semimartingales

Definition 6.24 *A stochastic process $X = (X_t : t \geq 0)$ is called a continuous semimartingale if it can be written as*

$$X_t = X_0 + M_t + A_t, \quad t \geq 0 \quad (32)$$

where M is a continuous local martingale, A is a continuous process of finite variation, and $M_0 = A_0 = 0$ a.s..

The decomposition is unique (up to modification on a set of measure zero). It should be remembered that there is a filtration (\mathcal{F}_t) and a probability measure \mathbb{P} implicit in our definition.

Proposition 6.25 *A continuous semimartingale is of finite quadratic variation and in the notation above $\langle X, X \rangle = \langle M, M \rangle$.*

Proof. Fix $t \geq 0$ and consider a sequence of partitions of $[0, t]$, $\pi_m = \{0 = t_0 < t_1 < \dots < t_{n_m} = t\}$ with $\text{mesh}(\pi_m) \rightarrow 0$ as $m \rightarrow \infty$. Then

$$\sum_{i=1}^{n_m} (X_{t_i} - X_{t_{i-1}})^2 = \underbrace{\sum_{i=1}^{n_m} (M_{t_i} - M_{t_{i-1}})^2}_{(i)} + \underbrace{\sum_{i=1}^{n_m} (A_{t_i} - A_{t_{i-1}})^2}_{(ii)} + 2 \underbrace{\sum_{i=1}^{n_m} (M_{t_i} - M_{t_{i-1}})(A_{t_i} - A_{t_{i-1}})}_{(iii)}.$$

It follows from the properties of M and A that, as $m \rightarrow \infty$,

$$\begin{aligned} (i) &\rightarrow \langle M, M \rangle_t, \\ (ii) &\leq \sup_{1 \leq i \leq n_m} |A_{t_i} - A_{t_{i-1}}| \cdot V_t(A) \rightarrow 0 \quad \text{a.s.}, \\ (iii) &\leq \sup_{1 \leq i \leq n_m} |M_{t_i} - M_{t_{i-1}}| \cdot V_t(A) \rightarrow 0 \quad \text{a.s.} \end{aligned}$$

□

If X, Y are two continuous semimartingales, we can define their co-variation $\langle X, Y \rangle$ via the polarisation formula that we used for martingales. If $X_t = X_0 + M_t + A_t$ and $Y_t = Y_0 + N_t + A'_t$, then $\langle X, Y \rangle_t = \langle M, N \rangle_t$.

7 Stochastic Integration

At the beginning of the course we argued that whereas classically differential equations take the form

$$dX(t) = a(t, X(t))dt,$$

in many settings, the dynamics of the physical quantity in which we are interested may also have a random component and so perhaps takes the form

$$dX_t = a(t, X_t)dt + b(t, X_t)dB_t.$$

We actually understand equations like this in the integral form:

$$X_t - X_0 = \int_0^t a(s, X_s)ds + \int_0^t b(s, X_s)dB_s.$$

If a is nice enough, then the first term has a classical interpretation. It is the second term, or rather a generalisation of it, that we want to make sense of now.

The first approach will be to mimic what we usually do for construction of the Lebesgue integral, namely work out how to integrate simple functions and then extend to general functions through passage to the limit. We'll then provide a very slick, but not at all intuitive, approach that nonetheless gives us some 'quick wins' in proving properties of the integral.

7.1 Stochastic integral w.r.t. L^2 -bounded martingales

Remark on Notation: We are going to use the notation $\varphi \bullet M$ for the (Itô) stochastic integral of φ with respect to M . This is not universally accepted notation; many authors would write $\int_0^t \varphi_s dM_s$ for $(\varphi \bullet M)_t$. Moreover, for emphasis, when the integrator is stochastic, we have used ' \bullet ' in place of the ' \cdot ' that we used for the Stieltjes integral.

We're going to develop a theory of integration w.r.t. martingales in H^2 . Recall that H_0^2 is the space of continuous martingales M , zero at zero, which are bounded in L^2 . It is a Hilbert space with the inner product $\langle M, N \rangle_{H^2} = \mathbb{E}[M_\infty N_\infty]$ and induced norm

$$\|M\|_{H^2} = \sqrt{\mathbb{E}[M_\infty^2]} = \sqrt{\mathbb{E}[\langle M \rangle_\infty]}.$$

(In a very real sense we are identifying H^2 with L^2 .)

Define \mathcal{E} to be the space of simple bounded process of the form

$$\varphi_t = \sum_{i=0}^m \varphi^{(i)} \mathbf{1}_{(t_i, t_{i+1}]}(t), \quad t \geq 0, \quad (33)$$

for some $m \in \mathbb{N}$, $0 \leq t_0 < t_1 < \dots < t_{m+1}$ and where $\varphi^{(i)}$ are bounded \mathcal{F}_{t_i} -measurable random variables. Define the stochastic integral $\varphi \bullet M$ of φ in (33) with respect to $M \in H^2$ via

$$(\varphi \bullet M)_t := \sum_{i=0}^m \varphi^{(i)} (M_{t \wedge t_{i+1}} - M_{t \wedge t_i}), \quad t \geq 0. \quad (34)$$

If we write $M_t^i := \varphi^{(i)}(M_{t \wedge t_{i+1}} - M_{t \wedge t_i})$ then clearly $M^i \in \mathbf{H}^2$ and so $\varphi \bullet M$ is a martingale. Moreover, since for $i \neq j$ the intervals $(t_i, t_{i+1}]$ and $(t_j, t_{j+1}]$ are disjoint, $M_t^i M_t^j$ is a martingale and hence $\langle M^i, M^j \rangle_t = 0$. Using the bilinearity of the bracket process then yields

$$\langle \varphi \bullet M \rangle_t = \sum_{i=0}^m \langle M^i \rangle_t = \sum_{i=0}^m \left(\varphi^{(i)} \right)^2 (\langle M \rangle_{t_{i+1} \wedge t} - \langle M \rangle_{t_i \wedge t}) = \int_0^t \varphi_s^2 d\langle M \rangle_s, \quad t \geq 0. \quad (35)$$

We already used the notation that if K is progressively measurable and A is of finite variation, then

$$(K \cdot A)_t = \int_0^t K_s(\omega) dA_s(\omega), \quad t \geq 0.$$

In that notation

$$\langle \varphi \bullet M \rangle = \varphi^2 \cdot \langle M \rangle.$$

More generally, for $N \in \mathbf{H}^2$,

$$\begin{aligned} \langle \varphi \bullet M, N \rangle_t &= \sum_{i=0}^m \langle M^i, N \rangle_t = \sum_{i=0}^m \varphi^{(i)} (\langle M, N \rangle_{t_{i+1} \wedge t} - \langle M, N \rangle_{t_i \wedge t}) \\ &= \int_0^t \varphi_s d\langle M, N \rangle_s = (\varphi \cdot \langle M, N \rangle)_t. \end{aligned} \quad (36)$$

Proposition 7.1 *Let $M \in \mathbf{H}^2$. The mapping $\varphi \mapsto \varphi \bullet M$ is a linear map from \mathcal{E} to \mathbf{H}_0^2 . Moreover,*

$$\|\varphi \bullet M\|_{\mathbf{H}^2}^2 = \mathbb{E} \left[\int_0^\infty \varphi_t^2 d\langle M \rangle_t \right]. \quad (37)$$

The proof is easy - we just need to show linearity. But given $\varphi, \psi \in \mathcal{E}$, we use a refinement of the partitions on which they are constant to write them as simple functions with respect to the same partition and the result is trivial.

Remark 7.2 *If we were considering left continuous martingales, then it would be important that the processes in \mathcal{E} are left continuous.*

We are expecting an L^2 -theory - we have already found an expression for the ' L^2 -norm' of $\varphi \bullet M$. Let us define the appropriate spaces more carefully.

Definition 7.3 *Given $M \in \mathbf{H}^2$ we denote by $L^2(M)$ the space of progressively measurable processes K such that*

$$\|K\|_{L^2(M)}^2 := \mathbb{E} \left[\int_0^\infty K_t^2 d\langle M \rangle_t \right] < +\infty. \quad (38)$$

$L^2(M)$ is a Hilbert space, with inner product

$$H, K \mapsto \mathbb{E} \left[\int_0^\infty H_t K_t d\langle M \rangle_t \right] = \mathbb{E} [(HK \cdot \langle M \rangle)_\infty].$$

We have $\mathcal{E} \subseteq L^2(M)$ and (35) tells us that the map $\mathcal{E} \rightarrow H_0^2$ given by $\varphi \mapsto \varphi \bullet M$ is a linear isometry. If we can show that the elementary functions are dense in $L^2(M)$, this observation will allow us to define integrals of functions from $L^2(M)$ with respect to M via a limiting procedure.

Proposition 7.4 *Let $M \in H^2$. Then \mathcal{E} is a dense vector subspace of $L^2(M)$.*

Proof. It is enough to show that if $K \in L^2(M)$ is orthogonal to φ for all $\varphi \in \mathcal{E}$, then $K = 0$ (as an element of $L^2(M)$). So suppose that $\langle K, \varphi \rangle_{L^2(M)} = 0$ for all $\varphi \in \mathcal{E}$. Let $X = K \cdot \langle M \rangle$, i.e. $X_t = \int_0^t K_u d\langle M \rangle_u$. This is well defined and by Cauchy-Schwarz

$$\mathbb{E}[|X_t|] \leq \mathbb{E} \left[\int_0^t |K_u| d\langle M \rangle_u \right] \leq \sqrt{\mathbb{E} \left[\int_0^t K_u^2 d\langle M \rangle_u \right]} \sqrt{\mathbb{E} \langle M \rangle_t} < +\infty$$

since $M \in H^2$ and $K \in L^2(M)$ (we took one of the functions to be identically one in Cauchy-Schwarz).

Taking $\varphi = \xi \mathbf{1}_{(s,t]} \in \mathcal{E}$, with $0 \leq s < t$ and ξ a bounded \mathcal{F}_s -measurable r.v., we have

$$0 = \langle K, \varphi \rangle_{L^2(M)} = \mathbb{E} \left[\xi \int_s^t K_u d\langle M \rangle_u \right] = \mathbb{E} [\xi (X_t - X_s)].$$

Since this holds for any \mathcal{F}_s -measurable bounded ξ , we conclude that $\mathbb{E}[(X_t - X_s) | \mathcal{F}_s] = 0$. In other words, X is a martingale. But X is also continuous and of finite variation and hence $X \equiv 0$ a.s. Thus $K = 0$ $d\langle M \rangle$ - a.e. a.s. and hence $K = 0$ in $L^2(M)$. \square

We now know that any $K \in L^2(M)$ is a limit of simple processes $\varphi^n \rightarrow K$. For each φ^n we can define the stochastic integral $\varphi^n \bullet M$. The isometry property then shows that $\varphi^n \bullet M$ converge in H^2 to some element that we denote $K \bullet M$ and which does not depend on the choice of approximating sequence φ^n .

Theorem 7.5 *Let $M \in H^2$. The mapping $\varphi \mapsto \varphi \bullet M$ from \mathcal{E} to H_0^2 defined in (34) has a unique extension to a linear isometry from $L^2(M)$ to H_0^2 which we denote $K \mapsto K \bullet M$.*

Remark 7.6 *For $K \in L^2(M)$, the martingale $K \bullet M$ is called the Itô stochastic integral of K with respect to M and is often written as $(K \bullet M)_t = \int_0^t K_u dM_u$. The isometry property may be then written as*

$$\|K \bullet M\|_{H^2}^2 = \mathbb{E} \left[\left(\int_0^\infty K_t dM_t \right)^2 \right] = \mathbb{E} \left[\int_0^\infty K_t^2 d\langle M \rangle_t \right] = \|K\|_{L^2(M)}^2. \quad (39)$$

Notice that if B is standard Brownian motion and we calculate $(B \bullet B)_t$, then

$$(B \bullet B)_t = \lim_{\delta(\pi) \rightarrow 0} \sum_{j=0}^{N(\pi)-1} B_{t_j} (B_{t_{j+1}} - B_{t_j}). \quad (40)$$

We also know already that the quadratic variation is

$$t = \lim_{\delta(\pi) \rightarrow 0} \sum_{j=0}^{N(\pi)-1} (B_{t_{j+1}} - B_{t_j})^2 = B_t^2 - B_0^2 - 2 \sum_{j=0}^{N(\pi)-1} B_{t_j} (B_{t_{j+1}} - B_{t_j}),$$

and so rearranging we find

$$\int_0^t B_s dB_s = \frac{1}{2} (B_t^2 - B_0^2 - t) = \frac{1}{2} (B_t^2 - t).$$

This is *not* what one would have predicted from classical integration theory (the extra term here comes from the quadratic variation).

Even more strangely, it *matters* that in (40) we took the *left* endpoint of the interval for evaluating the integrand. On the problem sheet, you are asked to evaluate

$$\lim_{\delta(\pi) \rightarrow 0} \sum B_{t_{j+1}} (B_{t_{j+1}} - B_{t_j}), \quad \text{and} \quad \lim_{\delta(\pi) \rightarrow 0} \sum \frac{B_{t_j} + B_{t_{j+1}}}{2} (B_{t_{j+1}} - B_{t_j}).$$

Each gives a different answer.

We can more generally define

$$\int_0^T f(B_s) \circ dB_s = \lim_{\delta(\pi) \rightarrow 0} \sum \left(\frac{f(B_{t_j}) + f(B_{t_{j+1}})}{2} \right) (B_{t_{j+1}} - B_{t_j}).$$

This so-called *Stratonovich integral* has the advantage that from the point of view of calculations, the rules of Newtonian calculus hold true. From a modelling perspective however, it can be the wrong choice. For example, suppose that we are modelling the change in a population size over time and we use $[t_i, t_{i+1})$ to represent the $(i+1)$ st generation. The change over (t_i, t_{i+1}) will be driven by the number of *adults*, so the population size at the *beginning* of the interval.

7.2 Intrinsic characterisation of stochastic integrals using the quadratic co-variation

We can also characterise the Itô integral in a slightly different way.

Theorem 7.7 *Let $M \in \mathbf{H}^2$. For any $K \in \mathbf{L}^2(M)$ there exists a unique element in \mathbf{H}_0^2 , denoted $K \bullet M$, such that*

$$\langle K \bullet M, N \rangle = K \cdot \langle M, N \rangle, \quad \forall N \in \mathbf{H}^2. \quad (41)$$

Furthermore, $\|K \bullet M\|_{\mathbf{H}^2} = \|K\|_{\mathbf{L}^2(M)}$ and the map

$$\mathbf{L}^2(M) \ni K \longrightarrow K \bullet M \in \mathbf{H}_0^2$$

is a linear isometry.

Proof. We first check uniqueness. Suppose that there are two such elements, X and X' . Then

$$\langle X, N \rangle - \langle X', N \rangle = \langle X - X', N \rangle \equiv 0, \quad \forall N \in \mathbf{H}^2.$$

Taking $N = X - X'$ we conclude, by Corollary 6.18, that $X = X'$.

Now let us verify (41) for the Itô integral.

Fix $N \in \mathbf{H}^2$. First note that for $K \in L^2(M)$ the Kunita-Watanabe inequality shows that

$$\mathbb{E} \left[\int_0^\infty |K_s| |d\langle M, N \rangle_s| \right] \leq \|K\|_{L^2(M)} \|N\|_{\mathbf{H}^2} < \infty$$

and thus the variable

$$\int_0^\infty K_s d\langle M, N \rangle_s = \left(K \cdot \langle M, N \rangle \right)_\infty$$

is well defined and in L^1 .

If K is an elementary process, in the notation of (34) and (35),

$$\langle K \bullet M, N \rangle = \sum_{i=0}^m \langle M^i, N \rangle$$

and

$$\langle M^i, N \rangle_t = \varphi^{(i)} \left(\langle M, N \rangle_{t_{i+1} \wedge t} - \langle M, N \rangle_{t_i \wedge t} \right),$$

so

$$\langle K \bullet M, N \rangle_t = \sum \varphi^{(i)} \left(\langle M, N \rangle_{t_{i+1} \wedge t} - \langle M, N \rangle_{t_i \wedge t} \right) = \int_0^t K_s d\langle M, N \rangle_s.$$

Now observe that the mapping $X \mapsto \langle X, N \rangle_\infty$ is continuous from \mathbf{H}^2 into L^1 . Indeed, by Kunita-Watanabe

$$\mathbb{E} \left[|\langle X, N \rangle| \right] \leq \mathbb{E} \left[\langle X, X \rangle_\infty \right]^{1/2} \mathbb{E} \left[\langle N, N \rangle_\infty \right]^{1/2} = \|N\|_{\mathbf{H}^2} \|X\|_{\mathbf{H}^2}.$$

So if K^n is a sequence in \mathcal{E} such that $K^n \rightarrow K$ in $L^2(M)$,

$$\langle K \bullet M, N \rangle_\infty = \lim_{n \rightarrow \infty} \langle K^n \bullet M, N \rangle_\infty = \lim_{n \rightarrow \infty} \left(K^n \cdot \langle M, N \rangle \right)_\infty = \left(K \cdot \langle M, N \rangle \right)_\infty,$$

where the convergence is in L^1 and the last equality is again a consequence of Kunita-Watanabe by writing

$$\mathbb{E} \left[\left| \int_0^\infty (K_s^n - K_s) d\langle M, N \rangle_s \right| \right] \leq \mathbb{E} \left[\langle N, N \rangle_\infty \right]^{1/2} \|K^n - K\|_{L^2(M)}.$$

We have thus obtained

$$\langle K \bullet M, N \rangle_\infty = \left(K \cdot \langle M, N \rangle \right)_\infty,$$

but replacing N by the stopped martingale N^t in this identity also gives

$$\langle K \bullet M, N \rangle_t = \left(K \cdot \langle M, N \rangle \right)_t$$

which completes the proof of (41). \square

We could write the relationship (41) as

$$\left\langle \int_0^\cdot K_s dM_s, N \right\rangle_t = \int_0^t K_s d\langle M, N \rangle_s;$$

that is, the stochastic integral ‘commutes’ with the bracket. One important consequence is that if $M \in \mathbb{H}^2$ and $K \in L^2(M)$, then applying (41) twice gives

$$\langle K \bullet M, K \bullet M \rangle = K \cdot \left(K \cdot \langle M, M \rangle \right) = K^2 \cdot \langle M, M \rangle.$$

In other words, the bracket process of $\int K_s dM_s$ is $\int K_s^2 d\langle M, M \rangle_s$. More generally, for N another martingale and $H \in L^2(N)$,

$$\left\langle \int_0^\cdot H_s dN_s, \int_0^\cdot K_s dM_s \right\rangle_t = \int_0^t H_s K_s d\langle M, N \rangle_s.$$

Proposition 7.8 (Associativity of stochastic integration) *Let $H \in L^2(M)$. If K is progressive, then $KH \in L^2(M)$ if and only if $K \in L^2(H \bullet M)$. In that case,*

$$(KH) \bullet M = K \bullet (H \bullet M).$$

(This is the analogue of what we already know for finite variation processes, where $K \cdot (H \cdot A) = (KH) \cdot A$.)

Proof.

$$\mathbb{E} \left[\int_0^\infty K_s^2 H_s^2 d\langle M, M \rangle_s \right] = \mathbb{E} \left[\int_0^\infty K_s^2 d\langle H \bullet M, H \bullet M \rangle_s \right],$$

which gives the first assertion.

For the second, for $N \in \mathbb{H}^2$ we write

$$\begin{aligned} \langle (KH) \bullet M, N \rangle &= KH \cdot \langle M, N \rangle = K \cdot (H \cdot \langle M, N \rangle) \\ &= K \cdot \langle H \bullet M, N \rangle = \langle K \bullet (H \bullet M), N \rangle, \end{aligned}$$

and by uniqueness in (41) this implies

$$(KH) \bullet M = K \bullet (H \bullet M).$$

\square

Recall that if $M \in \mathbb{H}^2$ and τ is a stopping time, then $M^\tau = (M_{t \wedge \tau}, t \geq 0)$ denotes the stopped process, which is itself a martingale and clearly $M^\tau \in \mathbb{H}^2$. For any $N \in \mathbb{H}^2$ we have

$$\langle M^\tau, N \rangle = \langle M, N \rangle^\tau = \mathbf{1}_{[0, \tau]} \cdot \langle M, N \rangle = \langle \mathbf{1}_{[0, \tau]} \bullet M, N \rangle,$$

so by uniqueness in Theorem 7.7, $\mathbf{1}_{[0, \tau]} \bullet M = M^\tau$.

In fact a much more general property holds true.

Proposition 7.9 (Stopped stochastic integrals) *Let $M \in \mathbb{H}^2$, $K \in \mathbb{L}^2(M)$ and τ a stopping time. Then*

$$(K \bullet M)^\tau = K \bullet M^\tau = K \mathbf{1}_{[0, \tau]} \bullet M.$$

Proof. We already argued above that the result holds for $K \equiv 1$.

Associativity says

$$K \bullet M^\tau = K \bullet (\mathbf{1}_{[0, \tau]} \bullet M) = K \mathbf{1}_{[0, \tau]} \bullet M.$$

Applying the same result to the martingale $K \bullet M$ we obtain

$$(K \bullet M)^\tau = \mathbf{1}_{[0, \tau]} \bullet (K \bullet M) = \mathbf{1}_{[0, \tau]} K \bullet M,$$

which gives the desired equalities. \square

7.3 Extensions: stochastic integration with respect to continuous semi-martingales

Definition 7.10 *For a continuous local martingale M , denote by $\mathbb{L}_{loc}^2(M)$ the space of progressively measurable processes K such that*

$$\forall t \geq 0 \quad \int_0^t K_s^2 d\langle M \rangle_s < +\infty \text{ a.s.}$$

Theorem 7.11 *Let M be a continuous local martingale. For any $K \in \mathbb{L}_{loc}^2(M)$ there exists a unique continuous local martingale, zero in zero, denoted $K \bullet M$ and called the Itô integral of K with respect to M , such that for any continuous local martingale N*

$$\langle K \bullet M, N \rangle = K \cdot \langle M, N \rangle. \quad (42)$$

If $M \in \mathbb{H}^2$ and $K \in \mathbb{L}^2(M)$ then this definition coincides with the previous one.

Proof. We only sketch the proof. Not surprisingly, we use a stopping argument.

For every $n \geq 1$, set

$$\tau_n = \inf \left\{ t \geq 0 : \int_0^t (1 + K_s^2) d\langle M \rangle_s \geq n \right\},$$

so that τ_n is a sequence of stopping times that increases to infinity. Since $\langle M^{\tau_n} \rangle_\infty = \langle M \rangle_{\tau_n} \leq n$, the stopped martingale M^{τ_n} is in \mathbb{H}^2 . Also

$$\int_0^\infty K_s^2 d\langle M^{\tau_n}, M^{\tau_n} \rangle_s = \int_0^{\tau_n} K_s^2 d\langle M, M \rangle_s \leq n,$$

so that $K \in \mathbb{L}^2(M^{\tau_n})$ and the definition of $K \bullet M^{\tau_n}$ makes sense. If $m > n$,

$$K \bullet M^{\tau_n} = (K \bullet M^{\tau_m})^{\tau_n}$$

so there is a unique process, that we denote $K \bullet M$ such that

$$(K \bullet M)^{\tau_n} = K \bullet M^{\tau_n}$$

and $(K \bullet M)_t = \lim_{n \rightarrow \infty} (K \bullet M^{\tau_n})_t$ and so, since $(K \bullet M^{\tau_n})$ is a martingale, the process $K \bullet M$ is a continuous local martingale with reducing sequence τ_n .

If N is a continuous local martingale (and without loss of generality $N_0 = 0$), we consider a reducing sequence

$$\tilde{\tau}_n = \inf\{t \geq 0 : |N_t| \geq n\} \quad \text{and set} \quad \rho_n := \tau_n \wedge \tilde{\tau}_n.$$

Then $N^{\rho_n} \in H_0^2$ and hence

$$\begin{aligned} \langle K \bullet M, N \rangle^{\rho_n} &= \langle (K \bullet M)^{\rho_n}, N^{\rho_n} \rangle \stackrel{\tau_n \geq \rho_n}{=} \langle (K \bullet M^{\tau_n})^{\rho_n}, N^{\rho_n} \rangle \stackrel{(28)}{=} \langle K \bullet M^{\tau_n}, N^{\rho_n} \rangle \\ &\stackrel{\text{Thm 7.7}}{=} K \cdot \langle M^{\tau_n}, N^{\rho_n} \rangle \stackrel{(28)}{=} K \cdot \langle M, N \rangle^{\rho_n} = (K \cdot \langle M, N \rangle)^{\rho_n}, \end{aligned}$$

so that $\langle K \bullet M, N \rangle = K \cdot \langle M, N \rangle$ as required. Uniqueness of $K \bullet M$ follows as in Theorem 7.7. \square

Naturally, we're going to define an integral with respect to a continuous semimartingale $X = X_0 + M + A$ as a sum of integrals w.r.t. M and w.r.t. A .

Definition 7.12 *We say that a progressively measurable process K is locally bounded if*

$$\text{a.s. } \forall t \geq 0 \quad \sup_{u \leq t} |K_u| < +\infty.$$

In particular, any adapted process with continuous sample paths is a locally bounded progressively measurable process.

If K is progressively measurable and locally bounded, then for any finite variation process A , almost surely,

$$\forall t \geq 0, \quad \int_0^t |K_s| |dA_s| < \infty$$

and, similarly, $K \in L_{\text{loc}}^2(M)$ for every continuous local martingale M .

Definition 7.13 *Let $X = X_0 + M + A$ be a continuous semimartingale and K a locally bounded process. The Itô stochastic integral of K with respect to X is the continuous semimartingale $K \bullet X$ defined by*

$$K \bullet X := K \bullet M + K \cdot A$$

often written

$$(K \bullet X)_t = \int_0^t K_s dX_s = \int_0^t K_s dM_s + \int_0^t K_s dA_s.$$

This integral inherits all the nice properties of the Stieltjes integral and the Itô integral that we have already derived (linearity, associativity, stopping etc.).

And of course, it is still the case for an elementary function $\varphi \in \mathcal{E}$ that

$$(\varphi \bullet X)_t = \sum_{i=1}^m \varphi^{(i)}(X_{t_{i+1} \wedge t} - X_{t_i \wedge t}).$$

We should also like to know how our integral behaves under limits.

Proposition 7.14 (Stochastic Dominated Convergence Theorem) *Let X be a continuous semimartingale and K^n a sequence of locally bounded processes with $K_t^n \rightarrow 0$ as $n \rightarrow \infty$ a.s. for all t . Further suppose that $|K_t^n| \leq K_t$ for all n where K is a locally bounded process. Then $K^n \bullet X$ converges to zero in probability and, more precisely,*

$$\forall t \geq 0 \quad \sup_{s \leq t} \left| \int_0^s K_u^n dX_u \right| \longrightarrow 0 \text{ in probability as } n \rightarrow \infty.$$

Proof. We can treat the finite variation part, $X_0 + A$, and the local martingale part, M , separately. For the first, note that

$$\begin{aligned} \left| \int_0^t K_u^n dA_u \right| &= \left| \int_0^t K_u^n dA_u^+ - \int_0^t K_u^n dA_u^- \right| \\ &\leq \int_0^t |K_u^n| dA_u^+ + \int_0^t |K_u^n| dA_u^- = \int_0^t |K_u^n| |dA_u|. \end{aligned}$$

The a.s. pointwise convergence of K^n to 0, together with the bound $|K^n| \leq K$, allow us to apply the (usual) Dominated Convergence Theorem to conclude that, for any $t > 0$, $\int_0^t |K_u^n| |dA_u|$ converges to 0 a.s. (in fact, as $\int_0^t |K_u^n| |dA_u|$ is non-decreasing in t , the convergence is uniform on any compact interval).

For the continuous local martingale part M , let (τ_m) be a reducing sequence such that $M^{\tau_m} \in H_0^2$ and $K \in L^2(M^{\tau_m})$. Then, by the Itô isometry,

$$\|K^n \bullet M^{\tau_m}\|_{H^2}^2 = \mathbb{E} \left[\left(\int_0^{\tau_m} K_t^n dM_t \right)^2 \right] = \mathbb{E} \left[\int_0^{\tau_m} (K_t^n)^2 \mathbf{1}_{[0, \tau_m]}(t) d\langle M \rangle_t \right] = \|K^n\|_{L^2(M^{\tau_m})}^2.$$

The right hand side tends to zero by the usual Dominated Convergence Theorem. For a fixed $t \geq 0$, and any given $\varepsilon > 0$, we may take m large enough that $\mathbb{P}[\tau_m \leq t] \leq \varepsilon/2$. We then have

$$\begin{aligned} \mathbb{P} \left[\sup_{s \leq t} |(K^n \bullet M)_s| > \varepsilon \right] &\leq \mathbb{P} \left[\sup_{s \leq t \wedge \tau_m} |(K^n \bullet M)_s| > \varepsilon \right] + \varepsilon/2 \\ &\leq \frac{1}{\varepsilon^2} \|K^n \bullet M^{\tau_m}\|_{H^2}^2 + \varepsilon/2 \leq \varepsilon, \end{aligned}$$

for n large enough. □

From this we can also confirm that even in their most general form our stochastic integrals can be thought of as limits of integrals of simple functions.

Proposition 7.15 *Let X be a continuous semimartingale and K a left-continuous locally bounded process. If π^n is a sequence of partitions of $[0, t]$ with mesh converging to zero then*

$$\sum_{t_i \in \pi^n} K_{t_i} (X_{t_{i+1}} - X_{t_i}) \longrightarrow \int_0^t K_s dX_s \text{ in probability as } n \rightarrow \infty.$$

7.4 Itô's formula and its applications

We already saw that the stochastic integral of Brownian motion with respect to itself did not behave as we would expect from Newtonian calculus. So what are the analogues of integration by parts and the chain rule for stochastic integrals?

Proposition 7.16 (Integration by parts) *If X and Y are two continuous semimartingales then*

$$\begin{aligned} X_t Y_t &= X_0 Y_0 + \int_0^t X_s dY_s + \int_0^t Y_s dX_s + \langle X, Y \rangle_t, \quad t \geq 0 \quad \text{a.s.} \\ &= X_0 Y_0 + (X \bullet Y)_t + (Y \bullet X)_t + \langle X, Y \rangle_t. \end{aligned} \quad (43)$$

Proof. Fix t and let π^n be a sequence of partitions of $[0, t]$ with mesh converging to zero. Note that

$$X_t Y_t - X_s Y_s = X_s (Y_t - Y_s) + Y_s (X_t - X_s) + (X_t - X_s)(Y_t - Y_s)$$

so for any n

$$\begin{aligned} X_t Y_t - X_0 Y_0 &= \sum_{t_i \in \pi^n} \left(X_{t_i} (Y_{t_{i+1}} - Y_{t_i}) + Y_{t_i} (X_{t_{i+1}} - X_{t_i}) + (X_{t_{i+1}} - X_{t_i})(Y_{t_{i+1}} - Y_{t_i}) \right) \\ &\longrightarrow (X \bullet Y)_t + (Y \bullet X)_t + \langle X, Y \rangle_t \quad \text{as } n \rightarrow \infty. \end{aligned}$$

□

Theorem 7.17 (Itô's formula) *Let X^1, \dots, X^d be continuous semimartingales and $F : \mathbb{R}^d \rightarrow \mathbb{R}$ a C^2 function. Then $(F(X_t^1, \dots, X_t^d) : t \geq 0)$ is a continuous semimartingale and a.s. for all $t \geq 0$*

$$\begin{aligned} F(X_t^1, \dots, X_t^d) &= F(X_0^1, \dots, X_0^d) + \sum_{i=1}^d \int_0^t \frac{\partial F}{\partial x^i} (X_s^1, \dots, X_s^d) dX_s^i \\ &\quad + \frac{1}{2} \sum_{1 \leq i, j \leq d} \int_0^t \frac{\partial^2 F}{\partial x^i \partial x^j} (X_s^1, \dots, X_s^d) d\langle X^i, X^j \rangle_s. \end{aligned} \quad (44)$$

In particular, for $d = 1$, we have

$$F(X_t) = F(X_0) + \int_0^t F'(X_s) dX_s + \frac{1}{2} \int_0^t F''(X_s) d\langle X \rangle_s.$$

Proof. Let $X^i = X_0^i + M^i + A^i$ be the semimartingale decomposition of X^i and denote by V^i the total variation process of A^i . Let

$$\tau_r^i = \inf\{t \geq 0 : |X_t^i| + V_t^i + \langle M^i \rangle_t > r\},$$

and $\tau_r = \min\{\tau_r^i, i = 1, \dots, d\}$. Then $(\tau_r)_{r \geq 0}$ is a family of stopping times with $\tau_r \uparrow \infty$. It is sufficient to prove (44) up to time τ_r . We will prove that the result holds for polynomials and then the full result follows by approximating C^2 functions by polynomials.

First note that it is obvious that the set of functions for which the formula holds is a vector space containing the functions $F \equiv 1$ and $F(x_1, \dots, x_d) = x_i$ for $i \leq d$.

We now check that if (44) holds for two functions F and G , then it holds for the product FG . Integration by parts yields

$$F_t G_t - F_0 G_0 = \int_0^t F_s dG_s + \int_0^t G_s dF_s + \langle F, G \rangle_t. \quad (45)$$

By associativity of stochastic integration, and because (44) holds for G ,

$$\int_0^t F_s dG_s = \sum_{i=1}^d \int_0^t F(X_s) \frac{\partial G_s}{\partial x^i} dX_s^i + \frac{1}{2} \sum_{1 \leq i, j \leq d} \int_0^t F(X_s) \frac{\partial^2 G_s}{\partial x^i \partial x^j} d\langle X^i, X^j \rangle_s,$$

with a similar expression for $\int_0^t G_s dF_s$. Using the fact that (44) holds for F and G , we also have

$$\langle F, G \rangle_t = \sum_{i=1}^d \sum_{j=1}^d \int_0^t \frac{\partial F_s}{\partial x^i} \frac{\partial G_s}{\partial x^j} d\langle X^i, X^j \rangle_s.$$

Substituting these into (45), we obtain Itô's formula for FG .

To pass to limits of sequences of polynomials, use the stochastic Dominated Convergence Theorem (and the fact that everything is nicely bounded up to time τ_r). \square

As a first application of this, suppose that M is a continuous local martingale and A is a process of finite variation. Then $\langle M, A \rangle \equiv 0$ and applying Itô's formula with $X^1 = M$ and $X^2 = A$ yields

$$\begin{aligned} F(M_t, A_t) &= F(M_0, A_0) + \int_0^t \frac{\partial F}{\partial m}(M_s, A_s) dM_s + \int_0^t \frac{\partial F}{\partial a}(M_s, A_s) dA_s \\ &\quad + \frac{1}{2} \int_0^t \frac{\partial^2 F}{\partial m^2}(M_s, A_s) d\langle M \rangle_s. \end{aligned}$$

Note that this gives us the semimartingale decomposition of $F(M_t, A_t)$ and we can, for example, read off the conditions on F under which we recover a local martingale. In particular, taking $F(x, y) = \exp(\lambda x - \frac{\lambda^2}{2} y)$ with $X^1 = M$ and $X^2 = \langle M, M \rangle$, we obtain:

Proposition 7.18 *Let M be a continuous local martingale and $\lambda \in \mathbb{R}$. Then*

$$\mathcal{E}^\lambda(M)_t := \exp\left(\lambda M_t - \frac{\lambda^2}{2}\langle M \rangle_t\right), \quad t \geq 0, \quad (46)$$

is a continuous local martingale. In fact the same holds true for any $\lambda \in \mathbb{C}$ with the real and imaginary parts being local martingales.

Proof. Computing the partial derivatives and simplifying gives:

$$\mathcal{E}^\lambda(M)_t = \mathcal{E}^\lambda(M)_0 + \int_0^t \frac{\partial}{\partial x} F^\lambda(M_s, \langle M \rangle_s) dM_s.$$

□

Note that we have $\frac{\partial}{\partial x} F(x, y) = \lambda F(x, y)$ so that we could have written this as

$$\mathcal{E}^\lambda(M)_t = \mathcal{E}^\lambda(M)_0 + \lambda \int_0^t \mathcal{E}^\lambda(M)_s dM_s$$

or in ‘differential form’ as

$$d\mathcal{E}^\lambda(M)_t = \lambda \mathcal{E}^\lambda(M)_t dM_t$$

which shows $\mathcal{E}^\lambda(M)$ solves the stochastic exponential differential equation driven by M : $dY_t = \lambda Y_t dM_t$.

Here is a beautiful application of exponential martingales:

Theorem 7.19 (Lévy’s characterisation of Brownian motion) *Let M be a continuous local martingale starting at zero. Then M is a standard Brownian motion if and only if $\langle M \rangle_t = t$ a.s. for all $t \geq 0$.*

Proof. We know that the quadratic variation of a Brownian motion B is given by $\langle B \rangle_t = t$.

Suppose M is a continuous local martingale starting in zero with $\langle M \rangle_t = t$ a.s. for all $t \geq 0$. Then, by Proposition 7.18,

$$\exp\left(i\xi M_t + \frac{\xi^2}{2}t\right), \quad t \geq 0$$

is a local martingale for any $\xi \in \mathbb{R}$ and, since it is bounded, it is a martingale. Let $0 \leq s < t$. We have

$$\mathbb{E}\left[\exp\left(i\xi M_t + \frac{\xi^2}{2}t\right) \middle| \mathcal{F}_s\right] = \exp\left(i\xi M_s + \frac{\xi^2}{2}s\right)$$

which we can rewrite as

$$\mathbb{E}\left[e^{i\xi(M_t - M_s)} \middle| \mathcal{F}_s\right] = e^{-\frac{\xi^2}{2}(t-s)}. \quad (47)$$

In other words, $M_t - M_s$ is centred Gaussian with variance $t - s$.

It follows also from (47) that for $A \in \mathcal{F}_s$,

$$\mathbb{E} \left[\mathbf{1}_A e^{i\xi(M_t - M_s)} \right] = \mathbb{P}[A] \mathbb{E} \left[e^{i\xi(M_t - M_s)} \right],$$

so fixing $A \in \mathcal{F}_s$ with $\mathbb{P}[A] > 0$ and writing $\mathbb{P}_A = \mathbb{P}[\cdot \cap A] / \mathbb{P}[A]$ (which is a probability measure on \mathcal{F}_s) for the conditional probability given A , we have that $M_t - M_s$ has the same distribution under \mathbb{P} as under \mathbb{P}_A and so $M_t - M_s$ is independent of \mathcal{F}_s and we have that M is an (\mathcal{F}_t) -Brownian motion. \square

So the quadratic variation is capturing all the information about M . This is surprising - recall that it is a special property of Gaussians that they are characterised by their means and the variance-covariance matrix, but in general we need to know much more. It turns out that what we just saw for Brownian motion has a powerful consequence for all continuous local martingales - they are characterised by their quadratic variation and, in fact, they are all time changes of Brownian motion.

Theorem 7.20 (Dambis, Dubins and Schwarz) *Let M be an $((\mathcal{F}_t), \mathbb{P})$ -continuous local martingale with $M_0 = 0$ and $\langle M \rangle_\infty = \infty$ a.s. Let $\tau_s := \inf\{t \geq 0 : \langle M \rangle_t > s\}$. Then the process B defined by $B_s := M_{\tau_s}$, is an $((\mathcal{F}_{\tau_s}), \mathbb{P})$ -Brownian motion and $M_t = B_{\langle M \rangle_t}$, $\forall t \geq 0$ a.s.*

Proof. Note that τ_s is the first hitting time of an open set (s, ∞) for an adapted process $\langle M \rangle$ with continuous sample paths, and hence τ_s is a stopping time (recall that (\mathcal{F}_t) is right-continuous). Further, $\langle M \rangle_\infty = \infty$ a.s. implies that $\tau_s < \infty$ a.s. The process $(\tau_s : s \geq 0)$ is non-decreasing and right-continuous (in fact $s \rightarrow \tau_s$ is the right-continuous inverse of $t \rightarrow \langle M \rangle_t$). Let $\mathcal{G}_s := \mathcal{F}_{\tau_s}$. Note that it satisfies the usual conditions. The process B is right continuous by continuity of M and right-continuity of τ . We have

$$\lim_{u \uparrow s} B_u = \lim_{u \uparrow s} M_{\tau_u} = M_{\tau_{s-}}.$$

But $[\tau_{s-}, \tau_s]$ is either a point or an interval of constancy of $\langle M \rangle$. The latter are known (exercise) to coincide a.s. with the intervals of constancy of M and hence $M_{\tau_{s-}} = M_{\tau_s} = B_s$ so that B has a.s. continuous paths. To conclude that B is a (\mathcal{G}_s) -Brownian motion, by Lévy's theorem, it remains to show that (B_s) and $(B_s^2 - s)$ are (\mathcal{G}_s) -local martingales.

Note that M^{τ_n} and $(M^{\tau_n})^2 - \langle M \rangle^{\tau_n}$ are uniformly integrable martingales. Taking $0 \leq u < s < n$ and applying the Optional Stopping Theorem we obtain

$$\mathbb{E}[B_s | \mathcal{G}_u] = \mathbb{E}[M_{\tau_s}^{\tau_n} | \mathcal{F}_{\tau_u}] = M_{\tau_u}^{\tau_n} = M_{\tau_u} = B_u$$

and

$$\mathbb{E}[B_s^2 - s | \mathcal{G}_u] = \mathbb{E}[(M_{\tau_s}^{\tau_n})^2 - \langle M \rangle_{\tau_s}^{\tau_n} | \mathcal{F}_{\tau_u}] = (M_{\tau_u}^{\tau_n})^2 - \langle M \rangle_{\tau_u}^{\tau_n} = (M_{\tau_u})^2 - \langle M \rangle_{\tau_u} = B_u^2 - u,$$

where we used continuity of $\langle M \rangle$ to write $\langle M \rangle_{\tau_u} = u$. It follows that B is indeed a (\mathcal{G}_s) -Brownian motion.

Finally, $B_{\langle M \rangle_t} = M_{\tau_{\langle M \rangle_t}} = M_t$, again since the intervals of constancy of M and of $\langle M \rangle$ coincide a.s. so that $s \rightarrow \tau_s$ is constant on $[t, \tau_{\langle M \rangle_t}]$. \square

A The Monotone Class Lemma/ Dynkin's $\pi - \lambda$ Theorem

There are multiple names used for this result (often with slightly different formulations).

Let E be an arbitrary set and let $\mathcal{P}(E)$ be the set of all subsets of E .

Definition A.1 A subset \mathcal{M} of $\mathcal{P}(E)$ is called a monotone class, or a Dynkin system/ λ -system, if

- i. $E \in \mathcal{M}$;
- ii. if $A, B \in \mathcal{M}$ and $A \subset B$, then $B \setminus A \in \mathcal{M}$;
- iii. if $(A_n)_{n \geq 0}$ is an increasing sequence of subsets of E such that $A_n \in \mathcal{M}$, then $\bigcup_{n \geq 0} A_n \in \mathcal{M}$.

The monotone class generated by an arbitrary subset \mathcal{C} of $\mathcal{P}(E)$ is

$$\mathcal{M}(\mathcal{C}) = \bigcap_{\mathcal{M} \text{ monotone class, } \mathcal{C} \subset \mathcal{M}} \mathcal{M}.$$

Equivalently, \mathcal{M} is a monotone class if

- i. $E \in \mathcal{M}$;
- ii. if $A, B \in \mathcal{M}$ and $A \subset B$, then $B \setminus A \in \mathcal{M}$;
- iii. if $(A_n)_{n \geq 0}$ is a sequence of subsets of E such that $A_i \cap A_j = \emptyset$ for $i \neq j$, then $\bigcup_{n \geq 0} A_n \in \mathcal{M}$.

Definition A.2 A collection \mathcal{I} of subsets of E such that $\emptyset \in \mathcal{I}$ and for all $A, B \in \mathcal{I}$, $A \cap B \in \mathcal{I}$ is called a π -system.

You may have seen the result expressed as:

Theorem A.3 (Dynkin's $\pi - \lambda$ Theorem) If P is a π -system and D is a λ -system such that $P \subseteq D$, then $\sigma(P) \subseteq \sigma(D)$.

Le Gall's (equivalent) formulation is:

Lemma A.4 (Monotone class lemma) If $\mathcal{C} \subset \mathcal{P}(E)$ is stable under finite intersections, then $\mathcal{M} = \sigma(\mathcal{C})$.

In other words, a Dynkin system which is also a π -system is a σ -algebra.

Here are some useful consequences:

- i. Let \mathcal{A} be a σ -field of E and let μ, ν be two probability measures on (E, \mathcal{A}) . Assume that there exists $\mathcal{C} \subset \mathcal{A}$ which is stable under finite intersections and such that $\sigma(\mathcal{C}) = \mathcal{A}$ and $\mu(A) = \nu(A)$ for every $A \in \mathcal{C}$, then $\mu = \nu$. (Use that $\mathcal{G} = \{A \in \mathcal{A} : \mu(A) = \nu(A)\}$ is a monotone class.)

- ii. Let $(X_i)_{i \in I}$ be an arbitrary collection of random variables, and let \mathcal{G} be a σ -field on some probability space. In order to show that the σ -fields $\sigma(X_i : i \in I)$ and \mathcal{G} are independent, it is enough to verify that $(X_{i_1}, \dots, X_{i_p})$ is independent of \mathcal{G} for any choice of the finite set $\{i_1, \dots, i_p\} \subset I$. (Observe that the class of all events that depend on a finite number of the X_i is stable under finite intersections and generates $\sigma(X_i, i \in I)$.)
- iii. Let $(X_i)_{i \in I}$ be an arbitrary collection of random variables and let Z be a bounded real variable. Let $i_0 \in I$. In order to verify that $\mathbb{E}[Z|X_i, i \in I] = \mathbb{E}[Z|X_{i_0}]$, it is enough to show that $\mathbb{E}[Z|X_{i_0}, X_{i_1}, \dots, X_{i_p}] = \mathbb{E}[Z|X_{i_0}]$ for any choice of the finite collection $\{i_1, \dots, i_p\} \subset I$. (Observe that the class of all events A such that $\mathbb{E}[\mathbf{1}_A Z] = \mathbb{E}[\mathbf{1}_A \mathbb{E}[Z|X_{i_0}]]$ is a monotone class.)

B Review of some basic probability

B.1 Convergence of random variables.

- a) $X_n \rightarrow X$ a.s. iff $P\{\omega : \lim_{n \rightarrow \infty} X_n(\omega) = X(\omega)\} = 1$.
- b) $X_n \rightarrow X$ in probability iff $\forall \varepsilon > 0, \lim_{n \rightarrow \infty} P\{|X_n - X| > \varepsilon\} = 0$.
- c) X_n converges to X in distribution (denoted $X_n \Rightarrow X$) iff $\lim_{n \rightarrow \infty} P\{X_n \leq x\} = P\{X \leq x\} \equiv F_X(x)$ for all x at which F_X is continuous.

Theorem B.1 *a) implies b) implies c).*

Proof. ($b \Rightarrow c$) Let $\varepsilon > 0$. Then

$$\begin{aligned} P\{X_n \leq x\} - P\{X \leq x + \varepsilon\} &= P\{X_n \leq x, X > x + \varepsilon\} - P\{X \leq x + \varepsilon, X_n > x\} \\ &\leq P\{|X_n - X| > \varepsilon\} \end{aligned}$$

and hence $\limsup P\{X_n \leq x\} \leq P\{X \leq x + \varepsilon\}$. Similarly, $\liminf P\{X_n \leq x\} \geq P\{X \leq x - \varepsilon\}$. Since ε is arbitrary, the implication follows. \square

B.2 Convergence in probability.

- a) If $X_n \rightarrow X$ in probability and $Y_n \rightarrow Y$ in probability then $aX_n + bY_n \rightarrow aX + bY$ in probability.
- b) If $\mathcal{Q} : \mathbb{R} \rightarrow \mathbb{R}$ is continuous and $X_n \rightarrow X$ in probability then $\mathcal{Q}(X_n) \rightarrow \mathcal{Q}(X)$ in probability.
- c) If $X_n \rightarrow X$ in probability and $X_n - Y_n \rightarrow 0$ in probability, then $Y_n \rightarrow X$ in probability.

Remark B.2 *(b) and (c) hold with convergence in probability replaced by convergence in distribution; however (a) is not in general true for convergence in distribution.*

B.3 Uniform Integrability

If X is an integrable random variable (that is $\mathbb{E}[|X|] < \infty$) and Λ_n is a sequence of sets with $\mathbb{P}[\Lambda_n] \rightarrow 0$, then $\mathbb{E}[|X\mathbf{1}_{\Lambda_n}|] \rightarrow 0$ as $n \rightarrow \infty$. (This is a consequence of the DCT since $|X|$ dominates $|X\mathbf{1}_{\Lambda_n}|$ and $|X\mathbf{1}_{\Lambda_n}| \rightarrow 0$ a.s.) Uniform integrability demands that this type of property holds *uniformly* for random variables from some class.

Definition B.3 (Uniform Integrability) A class \mathcal{C} of random variables is called uniformly integrable if given $\varepsilon > 0$ there exists $K \in (0, \infty)$ such that

$$\mathbb{E}[|X\mathbf{1}_{\{|X|>K\}}|] < \varepsilon \quad \text{for all } X \in \mathcal{C}.$$

There are two reasons why this definition is important:

- i. Uniform integrability is necessary and sufficient for passing to the limit under an expectation,
- ii. it is often easy to verify in the context of martingale theory.

Property 1 should be sufficient to guarantee that uniform integrability is interesting, but in fact uniform integrability is not often used in analysis where it is usually simpler to use the MCT or the DCT. It is only taken seriously in probability and that is because of 2.

Proposition B.4 Suppose that $\{X_\alpha, \alpha \in I\}$ is a uniformly integrable family of random variables on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Then

i.

$$\sup_{\alpha} \mathbb{E}[|X_\alpha|] < \infty,$$

ii.

$$\mathbb{P}[|X_\alpha| > N] \rightarrow 0 \quad \text{as } N \rightarrow \infty, \text{ uniformly in } \alpha.$$

iii.

$$\mathbb{E}[|X_\alpha|\mathbf{1}_\Lambda] \rightarrow 0 \quad \text{as } \mathbb{P}[\Lambda] \rightarrow 0, \text{ uniformly in } \alpha.$$

Conversely, either 1 and 3 or 2 and 3 implies uniform integrability.

B.4 The Convergence Theorems

Theorem B.5 (Bounded Convergence Theorem) Suppose that $X_n \Rightarrow X$ and that there exists a constant b such that $P(|X_n| \leq b) = 1$. Then $E[X_n] \rightarrow E[X]$.

Proof. Let $\{x_i\}$ be a partition of \mathbb{R} such that F_X is continuous at each x_i . Then

$$\sum_i x_i P\{x_i < X_n \leq x_{i+1}\} \leq E[X_n] \leq \sum_i x_{i+1} P\{x_i < X_n \leq x_{i+1}\}$$

and taking limits we have

$$\begin{aligned} \sum_i x_i P\{x_i < X \leq x_{i+1}\} &\leq \underline{\lim}_{n \rightarrow \infty} E[X_n] \\ &\leq \overline{\lim}_{n \rightarrow \infty} E[X_n] \leq \sum_i x_{i+1} P\{x_i < X \leq x_{i+1}\} \end{aligned}$$

As $\max |x_{i+1} - x_i| \rightarrow 0$, the left and right sides converge to $E[X]$, as required. \square

Lemma B.6 *Let $X \geq 0$ a.s. Then $\lim_{M \rightarrow \infty} E[X \wedge M] = E[X]$.*

Proof. Check the result first for X having a discrete distribution and then extend to general X by approximation. \square

Theorem B.7 (*Monotone Convergence Theorem.*) *Suppose $0 \leq X_n \leq X$ and $X_n \rightarrow X$ in probability. Then $\lim_{n \rightarrow \infty} E[X_n] = E[X]$.*

Proof. For $M > 0$

$$E[X] \geq E[X_n] \geq E[X_n \wedge M] \rightarrow E[X \wedge M]$$

where the convergence on the right follows from the bounded convergence theorem. It follows that

$$E[X \wedge M] \leq \liminf_{n \rightarrow \infty} E[X_n] \leq \limsup_{n \rightarrow \infty} E[X_n] \leq E[X]$$

and the result follows by Lemma B.6. \square

Lemma B.8 (*Fatou's lemma.*) *If $X_n \geq 0$ and $X_n \Rightarrow X$, then $\liminf E[X_n] \geq E[X]$.*

Proof. Since $E[X_n] \geq E[X_n \wedge M]$ we have

$$\liminf E[X_n] \geq \liminf E[X_n \wedge M] = E[X \wedge M].$$

By the Monotone Convergence Theorem $E[X \wedge M] \rightarrow E[X]$ and the result follows. \square

Theorem B.9 (*Dominated Convergence Theorem*) *Assume $X_n \Rightarrow X$, $Y_n \Rightarrow Y$, $|X_n| \leq Y_n$, and $E[Y_n] \rightarrow E[Y] < \infty$. Then $E[X_n] \rightarrow E[X]$.*

Proof. For simplicity, assume in addition that $X_n + Y_n \Rightarrow X + Y$ and $Y_n - X_n \Rightarrow Y - X$ (otherwise consider subsequences along which $(X_n, Y_n) \Rightarrow (X, Y)$). Then by Fatou's lemma $\liminf E[X_n + Y_n] \geq E[X + Y]$ and $\liminf E[Y_n - X_n] \geq E[Y - X]$. From these observations $\liminf E[X_n] + \lim E[Y_n] \geq E[X] + E[Y]$, and hence $\liminf E[X_n] \geq E[X]$. Similarly $\liminf E[-X_n] \geq E[-X]$ and $\limsup E[X_n] \leq E[X]$ \square

Lemma B.10 (*Markov's inequality*)

$$P\{|X| > a\} \leq E[|X|]/a, \quad a \geq 0.$$

Proof. Note that $|X| \geq aI_{\{|X| > a\}}$. Taking expectations proves the desired inequality. \square

B.5 Information and independence.

Information obtained by observations of the outcome of a random experiment is represented by a sub- σ -algebra \mathcal{D} of the collection of events \mathcal{F} . If $D \in \mathcal{D}$, then the observer “knows” whether or not the outcome is in D .

An S -valued random variable Y is *independent* of a σ -algebra \mathcal{D} if

$$P(\{Y \in B\} \cap D) = P\{Y \in B\}P(D), \forall B \in \mathcal{B}(S), D \in \mathcal{D}.$$

Two σ -algebras $\mathcal{D}_1, \mathcal{D}_2$ are independent if

$$P(D_1 \cap D_2) = P(D_1)P(D_2), \quad \forall D_1 \in \mathcal{D}_1, D_2 \in \mathcal{D}_2.$$

Random variables X and Y are independent if $\sigma(X)$ and $\sigma(Y)$ are independent, that is, if

$$P(\{X \in B_1\} \cap \{Y \in B_2\}) = P\{X \in B_1\}P\{Y \in B_2\}.$$

B.6 Conditional expectation.

Interpretation of conditional expectation in L_2 .

Problem: Approximate $X \in L_2$ using information represented by \mathcal{D} such that the mean square error is minimized, i.e., find the \mathcal{D} -measurable random variable Y that minimizes $E[(X - Y)^2]$.

Solution: Suppose Y is a minimizer. For any $\varepsilon \neq 0$ and any \mathcal{D} -measurable random variable $Z \in L_2$

$$E[|X - Y|^2] \leq E[|X - Y - \varepsilon Z|^2] = E[|X - Y|^2] - 2\varepsilon E[Z(X - Y)] + \varepsilon^2 E[Z^2].$$

Hence $2\varepsilon E[Z(X - Y)] \leq \varepsilon^2 E[Z^2]$. Since ε is arbitrary, $E[Z(X - Y)] = 0$ and hence

$$E[ZX] = E[ZY] \tag{48}$$

for every \mathcal{D} -measurable Z with $E[Z^2] < \infty$. □

With (48) in mind, for an integrable random variable X , the *conditional expectation* of X , denoted $E[X|\mathcal{D}]$, is the unique (up to changes on events of probability zero) random variable Y satisfying

A) Y is \mathcal{D} -measurable.

B) $\int_D X dP = \int_D Y dP$ for all $D \in \mathcal{D}$.

Note that Condition B is a special case of (48) with $Z = I_D$ (where I_D denotes the indicator function for the event D) and that Condition B implies that (48) holds for all bounded \mathcal{D} -measurable random variables. Existence of conditional expectations is a consequence of the Radon-Nikodym theorem.

C Convergence of backwards discrete time supermartingales

For a backwards supermartingale, Y_n is \mathcal{H}_n -measurable and for $m < n \leq 0$, $\mathbb{E}[Y_n | \mathcal{H}_m] \leq Y_m$. Notice that the σ -field $(\mathcal{H}_n)_{n \in -\mathbb{N}}$ gets ‘smaller and smaller’ as $n \rightarrow -\infty$.

If $(Y_n)_{n \in -\mathbb{N}}$ is a backwards supermartingale and if the sequence $(Y_n)_{n \in -\mathbb{N}}$ is bounded in L^1 , then $(Y_n)_{n \in -\mathbb{N}}$ converges almost surely and in L^1 as $n \rightarrow -\infty$.

D A very short primer in functional analysis

We start with a brief recall of basic notions of functional analysis leading to Hilbert spaces and identification of their dual.

E Normed vector spaces

We start with basics. A vector space V over \mathbb{R} is a set endowed with two binary operations: addition and multiplication by a scalar, which satisfy the natural axioms. We focus the discussion on real scalars as this is relevant for us but most of what follows applies to spaces over complex numbers (or more general fields).

Definition E.1 A norm $\|\cdot\|$ on a vector space V is a mapping from V to $[0, \infty)$ such that

- (i) for any $a \in \mathbb{R}$, $v \in V$, $\|av\| = |a|\|v\|$ (absolute homogeneity);
- (ii) for any $x, y \in V$, $\|x + y\| \leq \|x\| + \|y\|$ (triangle inequality);
- (iii) $\|v\| = 0$ if and only if v is the zero vector in V (separates points).

Note that a norm induces a metric on V through $d(x, y) = \|x - y\|$ and hence a topology on V . A norm is then a continuous function from V to \mathbb{R} . The space of continuous linear functions plays a special role:

Definition E.2 Given a normed vector space over reals $(V, \|\cdot\|_V)$, its dual V' is the space of all continuous linear maps (functionals) from V to \mathbb{R} . V' itself is a vector space over \mathbb{R} equipped with a norm

$$\|\phi\|_{V'} := \sup_{v \in V, \|v\|_V \leq 1} |\phi(v)|.$$

The classical examples of spaces to consider are spaces of sequences or of functions. Let $(\mathcal{S}, \mathbb{F}, \mu)$ be a measurable space endowed with a σ -finite measure. Then for a real valued measurable function f on \mathcal{S} we can consider

$$\|f\|_p := \left(\int_{\mathcal{S}} |f(x)|^p \mu(dx) \right)^{1/p}$$

and let $\mathcal{L}^p(\mathcal{S}, \mathbb{F}, \mu)$ be the space of such functions for which $\|f\|_p < \infty$. Observe that $\|\cdot\|_p$ is not yet a norm on \mathcal{L}^p – indeed it fails to satisfy (iii) in Definition E.1 since if $f = 0$ μ -a.e. but is not zero, e.g. $f = \mathbf{1}_A$ for a measurable $A \in \mathbb{F}$ with $\mu(A) = 0$, then still $\|f\|_p = 0$. We then say that $\|\cdot\|_p$ is a semi-norm on \mathcal{L}^p .

To remedy this, we consider the space $L^p(\mathcal{S}, \mathbb{F}, \mu)$ which is the quotient of \mathcal{L}^p with respect to the equivalence relation $f \sim g$ iff $f = g$ μ -a.e. Put differently, L^p is the space of equivalence classes of functions equal μ -a.e. and which are integrable with p^{th} power. Then $(L^p(\mathcal{S}, \mathbb{F}, \mu), \|\cdot\|_p)$ is a normed vector space for $p \geq 1$. The triangle inequality for $\|\cdot\|_p$ is simply the *Minkowski inequality*.

A more geometric notion of measuring the relation between vectors is given by an inner product.

Definition E.3 Given a vector space V over \mathbb{R} , a mapping $\langle \cdot, \cdot \rangle : V \times V \rightarrow \mathbb{R}$ is called an inner product if

- (i) it is bilinear and symmetric: $\langle ax + bz, y \rangle = a \langle x, y \rangle + b \langle z, y \rangle$ and $\langle x, y \rangle = \langle y, x \rangle$ for $a, b \in \mathbb{R}, x, y, z \in V$;
- (ii) for any $x \in V, \langle x, x \rangle \geq 0$;
- (iii) $\langle x, x \rangle = 0$ if and only if x is the zero vector in V .

This notion is very familiar on $V = \mathbb{R}^n$ where an inner product is given by

$$\langle x, y \rangle = x^T y = \sum_{i=1}^n x_i y_i.$$

An inner product satisfies the Cauchy-Schwartz inequality

Proposition E.4 An inner product on a vector space satisfies

$$|\langle x, y \rangle| \leq \sqrt{\langle x, x \rangle} \sqrt{\langle y, y \rangle}, \quad x, y \in V. \quad (49)$$

Proof. Let $\|x\| := \sqrt{\langle x, x \rangle}, x \in V$. Fix $x, y \in V$ and define a quadratic function $Q : \mathbb{R} \rightarrow \mathbb{R}$ by

$$Q(r) = \|x + ry\|^2 = \|y\|^2 r^2 + 2 \langle x, y \rangle r + \|x\|^2, \quad r \in \mathbb{R}$$

which clearly is non-negative and hence its discriminant has to be non-positive i.e.

$$4 \langle x, y \rangle^2 - 4 \|x\|^2 \|y\|^2 \leq 0 \text{ that is } |\langle x, y \rangle| \leq \|x\| \|y\|,$$

as required. We note also that equality holds if and only if the vectors x, y are linearly dependent i.e. $x = ry$ for some $r \in \mathbb{R}$. \square

The above implies that $\|x\| := \sqrt{\langle x, x \rangle}$ is a norm on V . We say that the norm is induced by an inner product. Among spaces L^p defined above only L^2 has norm which is induced by an inner product, namely by

$$\langle f, g \rangle = \int_{\mathcal{S}} f(x)g(x)\mu(dx). \quad (50)$$

F Banach spaces

We first define Cauchy sequences which embody the idea of a converging sequence when we do not know the limiting element.

Definition F.1 A sequence (x_n) of elements in a normed vector space $(X, \|\cdot\|)$ is called a Cauchy sequence if for any $\varepsilon > 0$ there exists $N \geq 1$ such that for all $n, m \geq N$ we have $\|x_n - x_m\| \leq \varepsilon$.

Definition F.2 A normed vector space $(X, \|\cdot\|)$ is complete if every Cauchy sequence converges to an element $x \in X$. It is then called a Banach space. Further, if the norm induces by an inner product then $(X, \|\cdot\|)$ is called a Hilbert space.

Naturally, the Euclidean space \mathbb{R}^d is a Banach space (and in fact a Hilbert space) with the norm $\|x\| = \sqrt{\sum_{i=1}^d x_i^2}$. This implies (reasoning for $d = 1$) that

Proposition F.3 If $(X, \|\cdot\|_X)$ is a normed vector space over \mathbb{R} then its dual $(X', \|\cdot\|_{X'})$ in Definition E.2 is a Banach space.

In many cases it is interesting to build linear functionals satisfying certain additional properties. This is often done using the Hahn-Banach theorem. It states in particular that a continuous linear functional defined on a linear subspace Y of X can be continuously extended to the whole of X without increasing its norm. A version of this is also known as the separating hyperplane theorem since it allows to separate two convex sets (one open) using an affine hyperplane.

An important step in studying continuous linear functionals on X is achieved by describing the structure of X' . We have

Proposition F.4 Let $(\mathcal{S}, \mathbb{F}, \mu)$ be a measurable space with a σ -finite measure. Then for any $p \geq 1$, $L^p(\mathcal{S}, \mathbb{F}, \mu)$ is a Banach space and for $p > 1$ its dual is equivalent to (isometric to) the space $L^q(\mathcal{S}, \mathbb{F}, \mu)$, where $1/p + 1/q = 1$.

In particular we see that L^2 is its own dual. This means that any continuous linear functional on L^2 can be identified with an element in L^2 . This property remains true for any Hilbert space:

Proposition F.5 Let $(X, \|\cdot\|)$ be a Hilbert space with the norm induced by an inner product, $\|x\| = \sqrt{\langle x, x \rangle}$. If $\phi : X \rightarrow \mathbb{R}$ is a continuous linear map then there exists an element $x_\phi \in X$ such that

$$\phi(y) = \langle x, y \rangle, \quad \forall y \in X.$$

In particular, if $\phi : L^2(\mathcal{S}, \mathbb{F}, \mu) \rightarrow \mathbb{R}$ is a continuous linear map then there exists an element $f_\phi \in L^2(\mathcal{S}, \mathbb{F}, \mu)$ such that

$$\phi(g) = \int_{\mathcal{S}} g(x) f_\phi(x) \mu(dx), \quad \forall g \in L^2(\mathcal{S}, \mathbb{F}, \mu).$$

Note that the inner product $\langle x, y \rangle$, or the integral $\int_{\mathcal{J}} g(x) f_{\phi}(x) \mu(dx)$, in the above statement is well defined by (49)–(50).

On a (separable) Hilbert space, we can also state an infinite dimensional analogue of the Pythagorean theorem. Recall that \mathbb{R}^n is a Hilbert space with inner product $\langle x, y \rangle = \sum_{i=1}^n x_i y_i$. This uses the canonical basis in \mathbb{R}^n but if we take any orthonormal basis in \mathbb{R}^n , say $(\varepsilon_1, \dots, \varepsilon_n)$, then

$$x = \sum_{i=1}^n \langle x, \varepsilon_i \rangle \varepsilon_i \quad \text{and hence} \quad \|x\|^2 = \sum_{i=1}^n \langle x, \varepsilon_i \rangle^2,$$

which is the Pythagorean theorem. The same reasoning gives $\langle x, y \rangle = \sum_{i=1}^n \langle x, \varepsilon_i \rangle \langle y, \varepsilon_i \rangle$. The infinite dimensional version is known as the

Proposition F.6 (Parseval's identity) *Let $(X, \|\cdot\|)$ be a separable Hilbert space with the norm induced by an inner product $(x, y) \rightarrow \langle x, y \rangle$ and let $(\varepsilon_n : n \geq 1)$ be an orthonormal basis of X . Then for any $x, y \in X$*

$$\langle x, y \rangle = \sum_{n \geq 1} \langle x, \varepsilon_n \rangle \langle y, \varepsilon_n \rangle, \quad \text{and in particular} \quad \|x\|^2 = \sum_{n \geq 1} \langle x, \varepsilon_n \rangle^2.$$

Finally we state one more result, which is crucial for the construction of the stochastic integral.

Proposition F.7 *Suppose $(X, \|\cdot\|_X)$ and $(Y, \|\cdot\|_Y)$ are two Banach spaces, $\mathcal{E} \subset X$ is a dense vector subspace in X and $I : \mathcal{E} \rightarrow Y$ is a linear isometry, i.e. a linear map which preserves the norm, $\|I(x)\|_Y = \|x\|_X$ for all $x \in \mathcal{E}$. Then I may be extended in an unique way to a linear isometry from X to Y .*

Proof. Take $x \in X$ and $x_n \rightarrow x$ with $x_n \in \mathcal{E}$. Then, by the isometry property, $(I(x_n))$ is Cauchy in Y since (x_n) is Cauchy in X . It follows that it converges to some element which we denote $I(x)$. Further, if we have two sequences in \mathcal{E} , (x_n) and (y_n) , both converging to $x \in X$ and giving rise to potentially two elements $I(x)$ and $I(x)'$ then we can build a third sequence $z_{2n} = x_n, z_{2n+1} = y_n$ which also converges to x and we see that $I(z_n)$ has to converge and the limit has to agree with both $I(x)$ and $I(x)'$, so that $I(x) = I(x)'$ is unique. It follows that we defined $I(x) \in Y$ uniquely for all $x \in X$. Further,

$$\|I(x)\|_Y = \lim_n \|I(x_n)\|_Y = \lim_n \|x_n\|_X = \|x\|_X$$

so that I is norm-preserving. Finally, if $x, y \in X$ then we can write them as limits of sequences of elements in \mathcal{E} , say (x_n) and (y_n) respectively. For any $a, b \in \mathbb{R}$, $ax_n + by_n \in \mathcal{E}$ since \mathcal{E} is a vector space, and then by the above and linearity of I on \mathcal{E} we have

$$I(ax + by) = \lim_n I(ax_n + by_n) = \lim_n (aI(x_n) + bI(y_n)) = aI(x) + bI(y),$$

so that I is linear on X as required. \square