

Stochastic Calculus for Finance
Brief Lecture Notes

Gautam Iyer



Gautam Iyer, 2020.

© 2020 by Gautam Iyer. This work is licensed under the *Creative Commons Attribution - Non Commercial - Share Alike 4.0 International License*. This means you may adapt and or redistribute this document **for non commercial purposes**, provided you give appropriate credit and re-distribute your work under the same licence. To view the full terms of this license, visit <http://creativecommons.org/licenses/by-nc-sa/4.0/> or send a letter to Creative Commons, PO Box 1866, Mountain View, CA 94042, USA.

A DRM free PDF of these notes will always be available free of charge at <http://www.math.cmu.edu/~gautam>. A self published print version at nominal cost may be made available for convenience. The L^AT_EX source is currently publicly hosted at GitLab: <https://gitlab.com/gi1242/cmu-mscf-944>.

These notes are provided as is, without any warranty and Carnegie Mellon University, the Department of Mathematical-Sciences, nor any of the authors are liable for any errors.

Preface

The purpose of these notes is to provide a rapid introduction to the Black-Scholes formula and the mathematical techniques used to derive it. Most mathematical concepts used are explained and motivated, but the complete rigorous proofs are beyond the scope of these notes. These notes were written in 2017 when I was teaching a seven week course in the Masters in Computational Finance program at Carnegie Mellon University.

The notes are somewhat minimal and mainly include material that was covered during the lectures itself. Only two sets of problems are included. These are problems that were used as a review for the midterm and final respectively. Supplementary problems and exams can be found on the course website: <http://www.math.cmu.edu/~gautam/sj/teaching/2018-19/944-scalc-finance1>.

For more comprehensive references and exercises, I recommend:

- (1) *Stochastic Calculus for Finance II* by Steven Shreve.
- (2) *The basics of Financial Mathematics* by Rich Bass.
- (3) *Introduction to Stochastic Calculus with Applications* by Fima C Klebaner.

Contents

Preface	iii
Chapter 1. Introduction	1
Chapter 2. Brownian motion, and an Introduction to Modern Probability	3
1. Scaling limit of random walks.	3
2. A brief review of probability	5
3. Independence of random variables	9
4. Conditional probability	11
5. Conditional expectation.	13
6. The Martingale Property	17
Chapter 3. Stochastic Integration	22
1. Motivation	22
2. The First Variation of Brownian motion	22
3. Quadratic Variation	23
4. Construction of the Itô integral	24
5. The Itô formula	27
6. A few examples using Itô's formula	31
7. Review Problems	33
8. The Black Scholes Merton equation.	34
9. Multi-dimensional Itô calculus.	40
Chapter 4. Risk Neutral Measures	48
1. The Girsanov Theorem.	48
2. Risk Neutral Pricing	51
3. The Black-Scholes formula	54
4. Review Problems	55

Introduction

The price of a stock is not a smooth function of time, and standard calculus tools can not be used to effectively model it. A commonly used technique is to model the price, S , as a *geometric Brownian motion*, given by the *stochastic differential equation (SDE)*

$$dS(t) = \alpha S(t) dt + \sigma S(t) dW(t),$$

where α and σ are parameters, and W is a Brownian motion. If $\sigma = 0$, this is simply the ordinary differential equation

$$dS(t) = \alpha S(t) dt \quad \text{or} \quad \partial_t S = \alpha S(t).$$

This is the price assuming it grows at a rate α . The σdW term models *noisy fluctuations* and the first goal of this course is to understand what this means. The mathematical tools required for this are Brownian motion, and Itô integrals, which we will develop and study.

An important point to note is that the above model can not be used to *predict* the price of S , because randomness is built into the model. Instead, we will use this model to *price securities*. Consider a *European call option* for a stock S with strike prices K and maturity T (i.e. this is the right to buy the asset S at price K at time T). Given the stock price $S(t)$ at some time $t \leq T$, what is a fair price for this option?

Seminal work of Black and Scholes computes the fair price of this option in terms of the time to maturity $T - t$, the stock price $S(t)$, the strike price K , the model parameters α, σ and the interest rate r . For notational convenience we suppress the explicit dependence on K, α, σ and let $c(t, x)$ represent the price of the option at time t given that the stock price is x . Clearly $c(T, x) = (x - K)^+$. For $t \leq T$, the Black-Scholes formula gives

$$c(t, x) = xN(d_+(T - t, x)) - Ke^{-r(T-t)}N(d_-(T - t, x))$$

where

$$d_{\pm}(\tau, x) \stackrel{\text{def}}{=} \frac{1}{\sigma\sqrt{\tau}} \left(\ln\left(\frac{x}{K}\right) + \left(r \pm \frac{\sigma^2}{2}\right)\tau \right).$$

Here r is the interest rate at which you can borrow or lend money, and N is the CDF of a standard normal random variable. (It might come as a surprise to you that the formula above is independent of α , the mean return rate of the stock.)

The second goal of this course is to understand the derivation of this formula. The main idea is to find a *replicating strategy*. If you've sold the above option, you hedge your bets by investing the money received in the underlying asset, and in an interest bearing account. Let $X(t)$ be the value of your portfolio at time t , of which you have $\Delta(t)$ invested in the stock, and $X(t) - \Delta(t)$ in the interest bearing account. If we are able to choose $X(0)$ and Δ in a way that would guarantee

$X(T) = (S(T) - K)^+$ almost surely, then $X(0)$ must be the fair price of this option. In order to find this strategy we will need to understand SDEs and the Itô formula, which we will develop subsequently.

The final goal of this course is to understand *risk neutral measures*, and use them to provide an elegant derivation of the Black-Scholes formula. If time permits, we will also study the *fundamental theorems of asset pricing*, which roughly state:

- (1) The *existence* of a risk neutral measure is equivalent to the market having no arbitrage (i.e. you can't make money without taking risk).
- (2) *Uniqueness* of a risk neutral measure is equivalent to the market having no arbitrage, and that all derivative securities can be hedged (i.e. for every derivative security we can find a replicating portfolio).

Brownian motion, and an Introduction to Modern Probability

1. Scaling limit of random walks.

Our first goal is to understand *Brownian motion*, which is used to model “noisy fluctuations” of stocks, and various other objects. This is named after the botanist Robert Brown, who observed that the microscopic movement of pollen grains appears random. Intuitively, Brownian motion can be thought of as a process that performs a random walk in continuous time.

We begin by describing Brownian motion as the scaling limit of discrete random walks. Let ξ_1, ξ_2, \dots , be a sequence of i.i.d. random variables which take on the values ± 1 with probability $1/2$. Define the time interpolated random walk $S(t)$ by setting $S(0) = 0$, and

$$(1.1) \quad S(t) = S(n) + (t - n)\xi_{n+1} \quad \text{when } t \in (n, n + 1].$$

Note $S(n) = \sum_1^n \xi_i$, and so at integer times S is simply a symmetric random walk with step size 1.

Our aim now is to rescale S so that it takes a random step at shorter and shorter time intervals, and then take the limit. In order to get a meaningful limit, we will have to compensate by also scaling the step size. Let $\varepsilon > 0$ and define

$$(1.2) \quad S_\varepsilon(t) = \alpha_\varepsilon S\left(\frac{t}{\varepsilon}\right),$$

where α_ε will be chosen below in a manner that ensures convergence of $S_\varepsilon(t)$ as $\varepsilon \rightarrow 0$. Note that S_ε now takes a random step of size α_ε after every ε time units.

To choose α_ε , we compute the variance of S_ε . Note first

$$\text{Var } S(t) = [t] + (t - [t])^2,$$

and¹ consequently

$$\text{Var } S_\varepsilon(t) = \alpha_\varepsilon^2 \left(\left[\frac{t}{\varepsilon} \right] + \left(\frac{t}{\varepsilon} - \left[\frac{t}{\varepsilon} \right] \right)^2 \right).$$

In order to get a “nice limit” of S_ε as $\varepsilon \rightarrow 0$, one would at least expect that $\text{Var } S_\varepsilon(t)$ converges as $\varepsilon \rightarrow 0$. From the above, we see that choosing

$$\alpha_\varepsilon = \sqrt{\varepsilon}$$

immediately implies

$$\lim_{\varepsilon \rightarrow 0} \text{Var } S_\varepsilon(t) = t.$$

¹Here $[x]$ denotes the greatest integer smaller than x . That is, $[x] = \max\{n \in \mathbb{Z} \mid n \leq x\}$.

THEOREM 1.1. *The processes $S_\varepsilon(t) \stackrel{\text{def}}{=} \sqrt{\varepsilon}S(t/\varepsilon)$ “converge” as $\varepsilon \rightarrow 0$. The limiting process, usually denoted by W , is called a (standard, one dimensional) Brownian motion.*

The proof of this theorem uses many tools from the modern theory of probability, and is beyond the scope of this course. The important thing to take away from this is that Brownian motion can be well approximated by a random walk that takes steps of variance ε on a time interval of size ε .

While the above construction provides good intuition as to what Brownian motion actually is, the scaling limit it is a somewhat unwieldy object to work with. We instead introduce an intrinsic characterization of Brownian motion, and we will shortly see that is both useful and mathematically convenient.

DEFINITION 1.2. A Brownian motion is a continuous process that has stationary independent increments.

Let us briefly explain the terms appearing in the above definition.

- (1) A *process* (aka stochastic process) is simply a collection of random variables $\{X(t) \mid 0 \leq t < \infty\}$. The index t usually represents time, and the process is often written as $\{X_t \mid 0 \leq t < \infty\}$ instead.
- (2) A *trajectory* (aka sample path) of a process is the outcome of one particular realization each of the random variables $X(t)$ viewed as function of time.
- (3) A process is called *continuous* if each of the trajectories are continuous. That is, for every $t > 0$ we have

$$(1.3) \quad \lim_{s \rightarrow t} X(s) = X(t).$$

- (4) An process is said to have *stationary increments* if for every $h \geq 0$, the distribution of $X(t+h) - X(t)$ does not depend on t .
- (5) A process is said to have *independent increments* if for every finite sequence of times $0 \leq t_0 < t_1 < \dots < t_N$, the random variables $X(t_0)$, $X(t_1) - X(t_0)$, $X(t_2) - X(t_1)$, \dots , $X(t_N) - X(t_{N-1})$ are all jointly independent.

For the process S in (1.1), note that for $n \in \mathbb{N}$, $S(n+1) - S(n) = X_{n+1}$ whose distribution *does not* depend on n as the variables $\{\xi_i\}$ were chosen to be independent and identically distributed. Similarly, $S(n+k) - S(n) = \sum_{n+1}^{n+k} \xi_i$ which has the same distribution as $\sum_1^k \xi_i$ and is independent of n .

However, if $t \in \mathbb{R}$ and is not necessarily an integer, $S(t+k) - S(t)$ will in general depend on t . So the process S (and also S_ε) do not have stationary (or independent) increments.

We claim, that the limiting process W does have stationary, independent, *normally distributed* increments. Suppose for some fixed $\varepsilon > 0$, both s and t are multiples of ε . In this case

$$S_\varepsilon(t) - S_\varepsilon(s) \sim \sqrt{\varepsilon} \sum_{i=1}^{\lfloor (t-s)/\varepsilon \rfloor} \xi_i \xrightarrow{\varepsilon \rightarrow 0} N(0, t-s),$$

by the central limit theorem. If s, t aren't multiples of ε as we will have in general, the first equality above is true up to a remainder which can easily be shown to vanish.

The above heuristic argument suggests that the limiting process W (from Theorem 1.1) satisfies $W(t) - W(s) \sim N(0, t-s)$. This certainly has independent

increments since $W(t+h) - W(t) \sim N(0, h)$ which is independent of t . Moreover, this also suggests that Brownian motion can be *equivalently* characterized as follows.

DEFINITION 1.3. A Brownian motion is a *continuous process* W such that:

- (1) W has independent increments, and
- (2) For $s < t$, $W(t) - W(s) \sim N(0, \sigma^2(t-s))$.

REMARK 1.4. A *standard* (one dimensional) Brownian motion is one for which $W(0) = 0$ and $\sigma = 1$.

2. A brief review of probability

In modern probability we usually start with a *probability triple* $(\Omega, \mathcal{G}, \mathbf{P})$.

- (1) Ω is a non-empty set called the *sample space*.
- (2) \mathcal{G} is a σ -algebra. This is a non-empty collection of events (subsets of Ω) of which the probability is known.
- (3) \mathbf{P} is a *probability measure*. For any event $A \in \mathcal{G}$, $\mathbf{P}(A)$ represents the probability of the event A occurring.

A subtle, but important, point in this framework is the in most case \mathcal{G} is usually not the collection of all subsets of Ω , but only a collection of some subsets of Ω . In fact, in most interesting examples, it is *impossible* to define the probability of arbitrary subsets of Ω consistently (i.e. in a manner that satisfies the required properties listed below), and we thus restrict ourselves to only talking about probabilities of elements of elements of the σ -algebra \mathcal{G} .

In order to be a probability space, the triple $(\Omega, \mathcal{G}, \mathbf{P})$ is required to satisfy certain properties. First the σ -algebra \mathcal{G} must satisfy the following:

- (1) It must be closed under compliments. That is, if $A \in \mathcal{G}$, then $A^c \in \mathcal{G}$.
- (2) It must be closed under *countable* unions. That is, if A_1, A_2, \dots are all elements of \mathcal{G} , then the union $\cup_1^\infty A_i$ is also an element of \mathcal{G} .

The precise mathematical definition of a σ -algebra is simply a non-empty collection of sets that satisfies the above two properties. Of course, if \mathcal{G} satisfies the above properties then one can quickly deduce the following:

- (3) The empty set \emptyset and the whole space Ω are elements of \mathcal{G} .
- (4) If A_1, A_2, \dots are all elements of \mathcal{G} , then the intersection $\cap_1^\infty A_i$ is also an element of \mathcal{G} .
- (5) If A, B are events in \mathcal{G} , then $A - B$ is also an event in \mathcal{G} .

The reason for requiring the above properties is that we expect \mathcal{G} is the collection of events of which the probability is known (or of which the probability can be deduced by performing repeated trials of some experiment). If you can deduce the probability of an event A , you should certainly be able to deduce the probability of A^c . Similarly, if you can deduce the probabilities of A, B , you should be able to deduce the probability of $A \cup B$ and $A \cap B$. The possibly surprising point above is that we don't require that \mathcal{G} be closed under finite unions, but we require it is closed under *countable* unions. The reason for this is that we would like our framework to allow us to perform repeated trials of an experiment and take limits.

Next, we turn our attention to the probability measure \mathbf{P} . We require that \mathbf{P} satisfies the following properties:

- (1) For each $A \in \mathcal{G}$, $\mathbf{P}(A) \in [0, 1]$. Moreover, $\mathbf{P}(\emptyset) = 0$, and $\mathbf{P}(\Omega) = 1$.

- (2) (*Countable additivity*) Given *pairwise disjoint* events $A_1, A_2, \dots \in \mathcal{G}$, we have

$$\mathbf{P}\left(\bigcup_{i=1}^{\infty} A_i\right) = \sum_{i=1}^{\infty} \mathbf{P}(A_i).$$

The above two properties are precisely the formal definition of a probability measure. Recall that in probability we require that the probability of mutually exclusive events add. The second property above is a generalization of this to countably many events.

Using the above properties, one can quickly verify that \mathbf{P} also satisfies the following properties:

- (1) $\mathbf{P}(A^c) = 1 - \mathbf{P}(A)$. More generally, if $A, B \in \mathcal{G}$ with $A \subseteq B$, then $\mathbf{P}(B - A) = \mathbf{P}(B) - \mathbf{P}(A)$.
- (2) If $A_1 \subseteq A_2 \subseteq A_3 \dots$ and each $A_i \in \mathcal{G}$ then $\mathbf{P}(\cup A_i) = \lim_{n \rightarrow \infty} \mathbf{P}(A_n)$.
- (3) If $A_1 \supseteq A_2 \supseteq A_3 \dots$ and each $A_i \in \mathcal{G}$ then $\mathbf{P}(\cap A_i) = \lim_{n \rightarrow \infty} \mathbf{P}(A_n)$.

We now describe *random variables* in the above context. In discrete probability, random variables are usually just real valued functions defined on the sample space. In our context, however, we need to be a bit more careful. If X is a random variable, then one should always be able to assign probabilities to questions such as “*Is X positive?*” or “*Does X belong to the interval (0, 1)?*”.

If X is simply a function from Ω to \mathbb{R} , then to compute the probability that X is positive, we should define $A = \{\omega \in \Omega \mid X(\omega) > 0\}$, and then compute $\mathbf{P}(A)$. This, however, is only possible if $A \in \mathcal{G}$; and since \mathcal{G} is *usually not* the entire power set of Ω , we should take care to ensure that all questions we might ask about the random variable X can be answered by only computing probabilities of events in \mathcal{G} , and not arbitrary subsets of Ω . For this reason, we define random variables as follows.

DEFINITION 2.1. A *random variable* is a \mathcal{G} -measurable function $X : \Omega \rightarrow \mathbb{R}$. That is, a random variable is a function $X : \Omega \rightarrow \mathbb{R}$ such that for every $\alpha \in \mathbb{R}$, the set $\{\omega \in \Omega \mid X(\omega) \leq \alpha\}$ is guaranteed to be an element of \mathcal{G} . (Such functions are also called \mathcal{G} -measurable, measurable with respect to \mathcal{G} , or simply measurable if the σ -algebra in question is clear from the context.)

REMARK 2.2. The argument ω is *always* suppressed when writing random variables. That is, the event $\{\omega \in \Omega \mid X(\omega) \leq \alpha\}$ is simply written as $\{X \leq \alpha\}$.

REMARK 2.3. Note for any random variable, $\{X > \alpha\} = \{X \leq \alpha\}^c$ which must also belong to \mathcal{G} since \mathcal{G} is closed under complements. One can check that for every $\alpha < \beta \in \mathbb{R}$ the events $\{X < \alpha\}$, $\{X \geq \alpha\}$, $\{X > \alpha\}$, $\{X \in (\alpha, \beta)\}$, $\{X \in [\alpha, \beta)\}$, $\{X \in (\alpha, \beta]\}$ and $\{X \in [\alpha, \beta]\}$ are all also elements of \mathcal{G} .

Thus to (for instance) compute the chance that X lies strictly between two real numbers α and β , we consider the event $\{X \in (\alpha, \beta)\}$. By Remark 2.3 this is guaranteed to be an element of \mathcal{G} , and thus we can compute the probability of it using \mathbf{P} . Hence, the quantity $\mathbf{P}(\{X \in (\alpha, \beta)\})$ is mathematically well defined, and represents the chance that the random variable X takes values in the interval (α, β) . For brevity, we almost always omit the outermost curly braces and write $\mathbf{P}(X \in (\alpha, \beta))$ for $\mathbf{P}(\{X \in (\alpha, \beta)\})$.

REMARK 2.4. One can check that if X, Y are random variables then so are $X \pm Y, XY, X/Y$ (when defined), $|X|, X \wedge Y$ and $X \vee Y$. In fact if $f: \mathbb{R} \rightarrow \mathbb{R}$ is any reasonably nice (more precisely, a Borel measurable) function, $f(X)$ is also a random variable.

EXAMPLE 2.5. Given $A \subseteq \Omega$, define *indicator function of A* by

$$\mathbf{1}_A(\omega) \stackrel{\text{def}}{=} \begin{cases} 1 & \omega \in A, \\ 0 & \omega \notin A. \end{cases}$$

One can check that $\mathbf{1}_A$ is a (\mathcal{G} -measurable) random variable if and only if $A \in \mathcal{G}$.

EXAMPLE 2.6. For $M \in \mathbb{N}$, $i \in \{1, \dots, M\}$, $a_i \in \mathbb{R}$ and $A_i \in \mathcal{G}$ be such that $A_i \cap A_j = \emptyset$ for $i \neq j$, and define

$$(2.1) \quad X \stackrel{\text{def}}{=} \sum_{i=1}^M a_i \mathbf{1}_{A_i}.$$

Then X is a (\mathcal{G} -measurable) random variable. (Such variables are called *simple random variables*.)

The next important concept concerning random variables is that of *expectation*, which we assume the reader is familiar with in the discrete setting. In the measure theoretic framework, the expectation of a random variable is the *Lebesgue integral*, and is denoted by²

$$\mathbf{E}X \stackrel{\text{def}}{=} \int_{\Omega} X d\mathbf{P}.$$

The precise construction of the Lebesgue integral, however, is too lengthy to be presented here, and we only present a brief summary.

If a random variable X only takes on finitely many values a_1, \dots, a_n , then the expectation of X is given by

$$(2.2) \quad \mathbf{E}X \stackrel{\text{def}}{=} \sum_{i=1}^n a_i \mathbf{P}(X = a_i).$$

This means that for any *simple random variable* of the form (2.1), the expectation is given by (2.2). For general random variables (i.e. random variables that are not simple), we can compute by expressing them as a limit of simple random variables. Namely, we can compute $\mathbf{E}X$ by

$$\mathbf{E}X = \lim_{n \rightarrow \infty} \mathbf{E} \left(\sum_{k=-n^2}^{n^2-1} \frac{k}{n} \mathbf{1}_{\{\frac{k}{n} \leq X < \frac{k+1}{n}\}} \right) = \lim_{n \rightarrow \infty} \sum_{k=-n^2}^{n^2-1} \frac{k}{n} \mathbf{P} \left(\frac{k}{n} \leq X < \frac{k+1}{n} \right),$$

for instance.

The above description, however, is only of theoretical importance and is not used to compute in practice. Here are a few computation rules and properties of expectations that will be useful later.

²If $A \in \mathcal{G}$ we define

$$\int_A Y d\mathbf{P} \stackrel{\text{def}}{=} \mathbf{E}(\mathbf{1}_A Y),$$

and when $A = \Omega$ we will often omit writing it.

- (1) (*Linearity*) If $\alpha \in \mathbb{R}$ and X, Y are random variables, then $\mathbf{E}(X + \alpha Y) = \mathbf{E}X + \alpha \mathbf{E}Y$.
- (2) (*Positivity*) If $X \geq 0$ almost surely,³ then $\mathbf{E}X \geq 0$. Moreover, if $X > 0$ almost surely, $\mathbf{E}X > 0$. Consequently, (using linearity) if $X \leq Y$ almost surely then $\mathbf{E}X \leq \mathbf{E}Y$.
- (3) (*Layer Cake formula*) If $X \geq 0$ almost surely, then

$$\mathbf{E}X = \int_0^\infty \mathbf{P}(X \geq t) dt.$$

More generally, if φ is an increasing differentiable function with $\varphi(0) = 0$ then

$$\mathbf{E}\varphi(X) = \int_0^\infty \varphi'(t) \mathbf{P}(X \geq t) dt.$$

- (4) (*Unconscious Statistician Formula*) If the probability density function of X is p , and f is any (Borel measurable) function, then

$$(2.3) \quad \mathbf{E}f(X) = \int_{-\infty}^\infty f(x)p(x) dx.$$

The proof of these properties goes beyond the scope of these notes. We do, however, make a few remarks. It turns out that the proof of positivity in this framework is immediate, however the proof of linearity is surprisingly not as straightforward as you would expect. While it is easy to verify linearity for simple random variables, for general random variables, the proof of linearity requires an approximation argument. The full proof of this involves either the *dominated* or *monotone* convergence theorem which guarantee $\lim \mathbf{E}X_n = \mathbf{E} \lim X_n$, under modest assumptions.

The layer cake formula can be proved by drawing a graph of X with Ω on the horizontal axis. Now $\mathbf{E}X$ should be the “area under the curve”, which is usually computed by slicing the region into vertical strips and adding up the area of each strip. If, instead, you compute the area by slicing the region into *horizontal* strips, you get exactly the layer cake formula!

Finally, unconscious statistician formula might already be familiar to you. In fact, the reason for this somewhat unusual name is that many people use this result “unconsciously” treating it as the definition, without realizing it is in fact a theorem that requires proof. To elaborate further, introductory (non-measure theoretic) probability courses usually stipulate that if a random variable X has density p_X , then

$$\mathbf{E}X = \int_{-\infty}^\infty xp_X(x) dx.$$

Thus if you set $Y = f(X)$ for some function f , we should have

$$\mathbf{E}Y = \int_{-\infty}^\infty yp_Y(y) dy.$$

If we could compute p_Y in terms of p_X and f , you could substitute it in the above formula, and obtain a formula for $\mathbf{E}Y$ in terms of p_X and f . Unfortunately, this isn’t easy to do. Namely, if f isn’t monotone, it isn’t easy to write down p_Y in terms of p_X . It turns out, however, that even though we can’t easily write down p_Y in terms of f and p_X , we can prove that $\mathbf{E}Y$ can be computed using (2.3).

³By $X \geq 0$ almost surely, we mean that $\mathbf{P}(X \geq 0) = 1$. More generally, we say an event occurs almost surely if the probability of it occurring is 1.

Since discussing these results and proofs further at this stage will lead us too far astray, we invite the curious to look them up in any standard measure theory book. The main point of this section was to introduce you to a framework which is capable of describing and studying the objects we will need for the remainder of the course.

We conclude this section by revisiting the notion of a *continuous process* defined in the previous section. Recall, our definition so far was that a process is simply a collection of random variables $\{X(t)\}_{t \geq 0}$, and a continuous process is a process whose trajectories are continuous. In our context, a process can now be thought of as a function

$$X: \Omega \times [0, \infty) \rightarrow \mathbb{R}.$$

For every fixed t , the function $\omega \mapsto X(\omega, t)$ is required to be a random variable (i.e. measurable with respect to \mathcal{G}). Since the ω is usually suppressed in probability, this random variable is simply denoted by $X(t)$.

The trajectory of X is now the slice of X for a fixed ω . Namely, for any fixed $\omega \in \Omega$, the function $t \mapsto X(\omega, t)$ is the trajectory of X . Saying a process has continuous trajectories means that for every $\omega \in \Omega$, the trajectory $t \mapsto X(\omega, t)$ is continuous as a function of t . Explicitly, this means for every $t \geq 0$ and $\omega \in \Omega$ we have

$$\lim_{s \rightarrow t} X(\omega, s) = X(\omega, t).$$

Following our convention of “never writing ω ”, this is exactly (1.3) as we had before.

3. Independence of random variables

Recall two events A, B are independent if $P(A | B) = P(A)$. This is of course immediately implies the multiplication law:

$$P(A \cap B) = P(A)P(B).$$

The notion of independence for random variables requires that *every* event that is observable from one is necessarily independent of *every* event that is observable from the other.

For example, suppose X and Y are two random variables. For any $a, b \in \mathbb{R}$, the event $\{X \in (a, b)\}$ can be observed using the random variable X . Similarly, any $c \in \mathbb{R}$, the event $\{Y > c\}$ can be observed using the random variable Y . If X and Y were independent, then the events $\{X \in (a, b)\}$ would necessarily be independent of the event $\{Y > c\}$. Of course, this is just an example and you can write down all sorts of other events (e.g. $X^2 - e^X < 15$, or $\sin(Y + 3) < .5$). No matter how you do it, if X and Y are independent, then any event observable from X alone must necessarily be independent of any event observable from Y alone.

Since the notion of “all events that can be observed from the random variable X ” will be useful later, we denote it by $\sigma(X)$.

DEFINITION 3.1. Let X be a random variable on $(\Omega, \mathcal{G}, \mathbf{P})$. We define *the σ -algebra generated by X* to be the σ algebra obtained by only using events that are observable using the random variable X .

One can mathematically prove that $\sigma(X)$ is generated by the events $\{X \leq \alpha\}$ for every $\alpha \in \mathbb{R}$. Namely, if a σ algebra contains the events $\{X \leq \alpha\}$ for every $\alpha \in \mathbb{R}$, then it must necessarily contain *all* events observable through the random variable X .

In particular, it will contain events of the form $\{X \in [\alpha, \beta)\}$, $e^{X+1} < \sin X$, or any other complicated formula that you can write down.

As mentioned above, the σ -algebra $\sigma(X)$ represents all the information one can obtain by observing X . To illustrate this, consider the following example: A card is drawn from a shuffled deck, and you win a dollar if it is red, and lose one if it is black. Now the likelihood of drawing any particular card is $1/52$. However, if you are blindfolded and only told the outcome of the game, you have no way to determine that each card is picked with probability $1/52$. The only thing you will be able to determine is that red cards are drawn as often as black ones.

This is captured by σ -algebra as follows. Let $\Omega = \{1, \dots, 52\}$ represent a deck of cards, $\mathcal{G} = \mathcal{P}(\Omega)$, and define $\mathbf{P}(A) = \text{card}(A)/52$. Let $R = \{1, \dots, 26\}$ represent the red cards, and $B = R^c$ represent the black cards. The outcome of the above game is now the random variable $X = \mathbf{1}_R - \mathbf{1}_B$, and you should check that $\sigma(X)$ is exactly $\{\emptyset, R, B, \Omega\}$.

With this, we can now revisit the notion of two random variables being independent.

DEFINITION 3.2. We say the random variables X_1, \dots, X_N are independent if for every $i \in \{1 \dots N\}$ and every $A_i \in \sigma(X_i)$ the events A_1, \dots, A_N are independent.

REMARK 3.3. Recall, a collection of events A_1, \dots, A_N is said to be independent if any sub collection $\{A_{i_1}, \dots, A_{i_k}\}$ satisfies the multiplication law

$$\mathbf{P}\left(\bigcap_{i=1}^k A_{i_k}\right) = \prod_{i=1}^k \mathbf{P}(A_i).$$

Note that this is a *stronger* condition than simply requiring

$$\mathbf{P}(A_1 \cap A_2 \cap \dots \cap A_N) = \mathbf{P}(A_1) \mathbf{P}(A_2) \dots \mathbf{P}(A_N).$$

In practice, one never tests independence of random variables using the above multiplication law.

PROPOSITION 3.4. Let X_1, \dots, X_N be N random variables. The following are equivalent:

- (1) The random variables X_1, \dots, X_N are independent.
- (2) For every $\alpha_1, \dots, \alpha_N \in \mathbb{R}$, we have

$$\mathbf{P}\left(\bigcap_{j=1}^N \{X_j \leq \alpha_j\}\right) = \prod_{j=1}^N \mathbf{P}(X_j \leq \alpha_j)$$

- (3) For every collection of bounded continuous functions f_1, \dots, f_N we have

$$\mathbf{E}\left[\prod_{j=1}^N f_j(X_j)\right] = \prod_{j=1}^N \mathbf{E}f_j(X_j).$$

- (4) For every $\xi_1, \dots, \xi_N \in \mathbb{R}$ we have

$$\mathbf{E} \exp\left(i \sum_{j=1}^N \xi_j X_j\right) = \prod_{j=1}^N \mathbf{E} \exp(i \xi_j X_j), \quad \text{where } i = \sqrt{-1}.$$

REMARK 3.5. It is instructive to explicitly check each of these implications when $N = 2$ and X_1, X_2 are simple random variables.

REMARK 3.6. The intuition behind the above result is as follows: Since the events $\{X_j \leq \alpha_j\}$ generate $\sigma(X_j)$, we expect the first two properties to be equivalent. Since $\mathbf{1}_{(-\infty, \alpha_j]}$ can be well approximated by continuous functions, we expect equivalence of the second and third properties. The last property is a bit more subtle: Since $\exp(a + b) = \exp(a) \exp(b)$, the third clearly implies the last property. The converse holds because of “completeness of the complex exponentials” or Fourier inversion, and again a through discussion of this will lead us too far astray.

REMARK 3.7. The third implication above implies that independent random variables are uncorrelated. Namely, if X, Y are independent random variables, then

$$(3.1) \quad \mathbf{E}(XY) = (\mathbf{E}X)(\mathbf{E}Y).$$

The converse, is of course false. Namely if (3.1) holds, there is no reason we should have

$$\mathbf{E}f(X)g(Y) = \mathbf{E}f(X)\mathbf{E}g(Y),$$

for every bounded continuous pair of functions f, g as required by the third part in Proposition 3.4. However, if (X, Y) is *jointly normal* and X, Y are uncorrelated, then the *normal correlation theorem* guarantees that X, Y are independent.

REMARK 3.8. If moment generating functions of the random variables are defined in an interval around 0, then one can test independence using real exponentials instead of the complex exponentials used in the last condition in Proposition 3.4. Explicitly, in this case X_1, \dots, X_N are independent if and only if for every t_1, \dots, t_N in some small interval containing 0 we have

$$\mathbf{E} \exp\left(\sum_{j=1}^N t_j X_j\right) = \prod_{j=1}^N \mathbf{E} \exp(t_j X_j).$$

EXAMPLE 3.9 (Covariance of Brownian motion). The independence of increments allows us to compute covariances of Brownian motion easily. Suppose W is a standard Brownian motion, and $s < t$. Then we know $W(s) \sim N(0, s)$, $W(t) - W(s) \sim N(0, t - s)$ and is independent of $W(s)$. Consequently $(W(s), W(t) - W(s))$ is jointly normal with mean 0 and covariance matrix $\begin{pmatrix} s & 0 \\ 0 & t-s \end{pmatrix}$. This implies that $(W(s), W(t))$ is a jointly normal random variable. Moreover we can compute the covariance by

$$\mathbf{E}W(s)W(t) = \mathbf{E}W(s)(W(t) - W(s)) + \mathbf{E}W(s)^2 = s.$$

In general if you don't assume $s < t$, the above immediately implies $\mathbf{E}W(s)W(t) = s \wedge t$.

4. Conditional probability

Our next goal is to understand *conditional probability*, and we do it directly here to help understanding. In the next section we will construct *conditional expectations* independently, and the reader may choose to skip this section.

Suppose you have an incomplete deck of cards which has 10 red cards, and 20 black cards. Suppose 5 of the red cards are *high cards* (i.e. ace, king, queen, jack or 10), and only 4 of the black cards are high. If a card is chosen at random, the *conditional probability* of it being high given that it is red is $1/2$, and the *conditional probability* of it being high given that it is black is $1/5$. Our aim is to encode both these facts into a single entity.

We do this as follows. Let R, B denote the set of all red and black cards respectively, and H denote the set of all high cards. A σ -algebra encompassing all the above information is exactly

$$\mathcal{G} \stackrel{\text{def}}{=} \{\emptyset, R, B, H, H^c, R \cap H, B \cap H, R \cap H^c, B \cap H^c, (R \cap H) \cup (B \cap H^c), (R \cap H^c) \cup (B \cap H), \Omega\}$$

and you can explicitly compute the probabilities of each of the above events. A σ -algebra encompassing only the color of cards is exactly

$$\mathcal{G} \stackrel{\text{def}}{=} \{\emptyset, R, B, \Omega\}.$$

Now we define the *conditional probability* of a card being high given the color to be the **random variable**

$$\mathbf{P}(H | \mathcal{C}) \stackrel{\text{def}}{=} \mathbf{P}(H | R)\mathbf{1}_R + \mathbf{P}(H | B)\mathbf{1}_B = \frac{1}{2}\mathbf{1}_R + \frac{1}{5}\mathbf{1}_B.$$

To emphasize:

- (1) What is given is the σ -algebra \mathcal{C} , and not just an event.
- (2) The conditional probability is now a \mathcal{C} -measurable random variable and not a number.

To see how this relates to $\mathbf{P}(H | R)$ and $\mathbf{P}(H | B)$ we observe

$$\int_R \mathbf{P}(H | \mathcal{C}) d\mathbf{P} \stackrel{\text{def}}{=} \mathbf{E}(\mathbf{1}_R \mathbf{P}(H | \mathcal{C})) = \mathbf{P}(H | R) \mathbf{P}(R).$$

The same calculation also works for B , and so we have

$$\mathbf{P}(H | R) = \frac{1}{\mathbf{P}(R)} \int_R \mathbf{P}(H | \mathcal{C}) d\mathbf{P} \quad \text{and} \quad \mathbf{P}(H | B) = \frac{1}{\mathbf{P}(B)} \int_B \mathbf{P}(H | \mathcal{C}) d\mathbf{P}.$$

Our aim is now to generalize this to a non-discrete scenario. The problem with the above identities is that if either R or B had probability 0, then the above would become meaningless. However, clearing out denominators yields

$$\int_R \mathbf{P}(H | \mathcal{C}) d\mathbf{P} = \mathbf{P}(H \cap R) \quad \text{and} \quad \int_B \mathbf{P}(H | \mathcal{C}) d\mathbf{P} = \mathbf{P}(H \cap B).$$

This suggests that the defining property of $\mathbf{P}(H | \mathcal{C})$ should be the identity

$$(4.1) \quad \int_C \mathbf{P}(H | \mathcal{C}) d\mathbf{P} = \mathbf{P}(H \cap C)$$

for every event $C \in \mathcal{C}$. Note $\mathcal{C} = \{\emptyset, R, B, \Omega\}$ and we have only checked (4.1) for $C = R$ and $C = B$. However, for $C = \emptyset$ and $C = \Omega$, (4.1) is immediate.

DEFINITION 4.1. Let $(\Omega, \mathcal{G}, \mathbf{P})$ be a probability space, and $\mathcal{F} \subseteq \mathcal{G}$ be a σ -algebra. Given $A \in \mathcal{G}$, we define the conditional probability of A given \mathcal{F} , denoted by $\mathbf{P}(A | \mathcal{F})$ to be an \mathcal{F} -measurable random variable that satisfies

$$(4.2) \quad \int_F \mathbf{P}(H | \mathcal{F}) d\mathbf{P} = \mathbf{P}(H \cap F) \quad \text{for every } F \in \mathcal{F}.$$

REMARK 4.2. Showing existence (and uniqueness) of the conditional probability isn't easy, and relies on the *Radon-Nikodym theorem*, which is beyond the scope of this course.

REMARK 4.3. It is crucial to require that $\mathbf{P}(H|\mathcal{F})$ is measurable with respect to \mathcal{F} . Without this requirement we could simply choose $\mathbf{P}(H|\mathcal{F}) = \mathbf{1}_H$ and (4.2) would be satisfied. However, note that if $H \in \mathcal{F}$, then the function $\mathbf{1}_F$ is \mathcal{F} -measurable, and in this case $\mathbf{P}(H|\mathcal{F}) = \mathbf{1}_F$.

REMARK 4.4. In general we can only expect (4.2) to hold for all events in \mathcal{F} , and it need not hold for events in \mathcal{G} ! Indeed, in the example above we see that

$$\int_H \mathbf{P}(H|\mathcal{C}) d\mathbf{P} = \frac{1}{2}\mathbf{P}(R \cap H) + \frac{1}{5}\mathbf{P}(B \cap H) = \frac{1}{2} \cdot \frac{5}{30} + \frac{1}{5} \cdot \frac{4}{30} = \frac{11}{100}$$

but

$$\mathbf{P}(H \cap H) = \mathbf{P}(H) = \frac{3}{10} \neq \frac{11}{100}.$$

REMARK 4.5. One situation where you can compute $\mathbf{P}(A|\mathcal{F})$ explicitly is when $\mathcal{F} = \sigma(\{F_i\})$ where $\{F_i\}$ is a pairwise disjoint collection of events whose union is all of Ω and $\mathbf{P}(F_i) > 0$ for all i . In this case

$$\mathbf{P}(A|\mathcal{F}) = \sum_i \frac{\mathbf{P}(A \cap F_i)}{\mathbf{P}(F_i)} \mathbf{1}_{F_i}.$$

5. Conditional expectation.

Conditional expectation arises when you have a random variable X , and want to *best approximate* it using only a (strict) subset of events. Precisely, suppose $\mathcal{F} \subseteq \mathcal{G}$ is a σ -sub-algebra of \mathcal{G} . That is, \mathcal{F} is a σ -algebra, and every event in \mathcal{F} is also an event in \mathcal{G} . Now to best approximate a (\mathcal{G} -measurable) random variable X using only events in \mathcal{F} , one would like to find an \mathcal{F} measurable random variable Z that minimizes

$$\mathbf{E}|X - Z|^2.$$

The minimizer is known as the *conditional expectation of X given \mathcal{F}* , and denoted by $\mathbf{E}(X|\mathcal{F})$. That is,

$$(5.1) \quad \mathbf{E}(X|\mathcal{F}) \stackrel{\text{def}}{=} \arg \min \{ \mathbf{E}|X - Z|^2 \mid Z \text{ is a } \mathcal{G}\text{-measurable random variable} \}.$$

While the above provides good intuition to the notion of conditional expectation, it is not as convenient to work with mathematically. For instance, the above requires $\mathbf{E}X^2 < \infty$, and we will often require conditional expectations of random variables that do not have this property.

To motivate the other definition of conditional expectation, we use the following example. Consider an incomplete deck of cards which has 10 red cards, of which 5 are high, and 20 black cards, of which 4 are high. Let X be the outcome of a game played through a dealer who pays you \$1 when a high card is drawn, and charges you \$1 otherwise. However, you are standing too far away from the dealer to tell whether the card drawn was high or not. You can only tell *the color*, and *whether or not you won*.

After playing this game often the only information you can deduce is that your expected return is 0 when a red card is drawn and $-3/5$ when a black card is drawn. That is, you approximate the game outcome X by the random variable

$$Y \stackrel{\text{def}}{=} 0\mathbf{1}_R - \frac{3}{5}\mathbf{1}_B,$$

where, as before R, B denote the set of all red and black cards respectively.

Note that the events you can deduce information about by playing this game (through the dealer) are exactly elements of the σ -algebra $\mathcal{C} = \{\emptyset, R, B, \Omega\}$. By construction, that your approximation Y is \mathcal{C} -measurable, and it is easy to verify that

$$(5.2) \quad Y = \arg \min \{ \mathbf{E}(X - Z)^2 \mid Z \text{ is a } \mathcal{C}\text{-measurable random variable} \}.$$

That is $Y = \mathbf{E}(X \mid \mathcal{C})$ according to the definition (5.1). In this case, we can also verify that Y has the same averages as X on all elements of \mathcal{C} . That is, for every $C \in \mathcal{C}$, we have⁴

$$(5.3) \quad \int_C Y d\mathbf{P} = \int_C X d\mathbf{P}.$$

It turns out that in general, one can show abstractly that any \mathcal{C} measurable random variable that satisfies (5.3), must in fact also be the minimizer in (5.2). We will thus use (5.3) to define conditional expectation.

DEFINITION 5.1. Let X be a \mathcal{G} -measurable random variable, and $\mathcal{F} \subseteq \mathcal{G}$ be a σ -sub-algebra. We define $\mathbf{E}(X \mid \mathcal{F})$, the *conditional expectation of X given \mathcal{F}* to be a *random variable* such that:

- (1) $\mathbf{E}(X \mid \mathcal{F})$ is \mathcal{F} -measurable.
- (2) For every $F \in \mathcal{F}$, we have the *partial averaging* identity:

$$(5.4) \quad \int_F \mathbf{E}(X \mid \mathcal{F}) d\mathbf{P} = \int_F X d\mathbf{P}.$$

REMARK 5.2. We can only expect (5.4) to hold for all events $F \in \mathcal{F}$. In general (5.4) *will not* hold for events $G \in \mathcal{G} - \mathcal{F}$.

REMARK 5.3. An equivalent way of phrasing (5.4) is to require

$$(5.5) \quad \mathbf{E}(XY) = \mathbf{E}(\mathbf{E}(X \mid \mathcal{F})Y)$$

for every \mathcal{F} measurable random variable Y . As before, we can only expect (5.5) to hold when Y is \mathcal{F} measurable. In general (5.5) *will not* hold when Y is not \mathcal{F} measurable.

REMARK 5.4. Choosing $F = \Omega$ we see $\mathbf{E}\mathbf{E}(X \mid \mathcal{F}) = \mathbf{E}X$.

REMARK 5.5. More concretely, suppose Y is another random variable and $\mathcal{F} = \sigma(Y)$. Then it turns out that one can find a special (non-random) function g such that $\mathbf{E}(X \mid \mathcal{F}) = g(Y)$. Moreover, the function g is characterized by the property that

$$\mathbf{E}(f(Y)X) = \mathbf{E}(f(Y)g(Y)).$$

for *every* bounded continuous function f .

REMARK 5.6. Under mild integrability assumptions one can show that conditional expectations exist. This requires the *Radon-Nikodym* theorem and goes beyond the scope of this course. If, however, $\mathcal{F} = \sigma(\{F_i\})$ where $\{F_i\}$ is a pairwise disjoint collection of events whose union is all of Ω and $\mathbf{P}(F_i) > 0$ for all i , then

$$\mathbf{E}(X \mid \mathcal{F}) = \sum_{i=1}^{\infty} \frac{\mathbf{1}_{F_i}}{\mathbf{P}(F_i)} \int_{F_i} X d\mathbf{P}.$$

⁴Recall $\int_C Y d\mathbf{P}$ is simply $\mathbf{E}(\mathbf{1}_C Y)$. That is $\int_C Y d\mathbf{P}$ is the expectation of the random variable which is Y on the event C , and 0 otherwise.

REMARK 5.7. Once existence is established it is easy to see that conditional expectations are unique. Namely, if Y is any \mathcal{F} -measurable random variable that satisfies

$$\int_F Y d\mathbf{P} = \int_F X d\mathbf{P} \quad \text{for every } F \in \mathcal{F},$$

then $Y = \mathbf{E}(X | F)$. Often, when computing the conditional expectation, we will “guess” what it is, and verify our guess by checking measurability and the above partial averaging identity.

PROPOSITION 5.8. *If X is \mathcal{F} -measurable, then $\mathbf{E}(X | \mathcal{F}) = X$. On the other hand, if X is independent⁵ of \mathcal{F} then $\mathbf{E}(X | \mathcal{F}) = \mathbf{E}X$.*

PROOF. If X is \mathcal{F} -measurable, then clearly the random variable X is both \mathcal{F} -measurable and satisfies the partial averaging identity. Thus by uniqueness, we must have $\mathbf{E}(X | \mathcal{F}) = X$.

Now consider the case when X is independent of \mathcal{F} . Suppose first $X = \sum a_i \mathbf{1}_{A_i}$ for finitely many sets $A_i \in \mathcal{G}$. Then for any $F \in \mathcal{F}$,

$$\int_F X d\mathbf{P} = \sum a_i \mathbf{P}(A_i \cap F) = \mathbf{P}(F) \sum a_i \mathbf{P}(A_i) = \mathbf{P}(F) \mathbf{E}X = \int_F \mathbf{E}X d\mathbf{P}.$$

Thus the constant random variable $\mathbf{E}X$ is clearly \mathcal{F} -measurable and satisfies the partial averaging identity. This forces $\mathbf{E}(X | \mathcal{F}) = \mathbf{E}X$. The general case when X is not simple follows by approximation. \square

The above fact has a generalization that is tremendously useful when computing conditional expectations. Intuitively, the general principle is to *average* quantities that are independent of \mathcal{F} , and *leave unchanged* quantities that are \mathcal{F} measurable. This is known as the independence lemma.

LEMMA 5.9 (Independence Lemma). *Suppose X, Y are two random variables such that X is independent of \mathcal{F} and Y is \mathcal{F} -measurable. Then if $f = f(x, y)$ is any function of two variables we have*

$$\mathbf{E}(f(X, Y) | \mathcal{F}) = g(Y),$$

where $g = g(y)$ is the function⁶ defined by

$$g(y) \stackrel{\text{def}}{=} \mathbf{E}f(X, y).$$

REMARK. If p_X is the probability density function of X , then the above says

$$\mathbf{E}(f(X, Y) | \mathcal{F}) = \int_{\mathbb{R}} f(x, Y) p_X(x) dx.$$

Indicating the ω dependence explicitly for clarity, the above says

$$\mathbf{E}(f(X, Y) | \mathcal{F})(\omega) = \int_{\mathbb{R}} f(x, Y(\omega)) p_X(x) dx.$$

⁵We say a random variable X is independent of σ -algebra \mathcal{F} if for every $A \in \sigma(X)$ and $B \in \mathcal{F}$ the events A and B are independent.

⁶To clarify, we are defining a *non-random* function $g = g(y)$ here when $y \in \mathbb{R}$ is any real number. Then, once we compute g , we substitute in $y = Y(= Y(\omega))$, where Y is the given random variable.

REMARK 5.10. Note we defined and motivated conditional expectations and conditional probabilities independently. They are however intrinsically related: Indeed, $\mathbf{E}(\mathbf{1}_A | \mathcal{F}) = \mathbf{P}(A | \mathcal{F})$, and this can be checked directly from the definition.

As we will see shortly, computing conditional expectations will be a very important part of pricing securities. Most of the time, all that is required to compute conditional expectations are the following properties.

PROPOSITION 5.11. *Conditional expectations satisfy the following properties.*

(1) (Linearity) *If X, Y are random variables, and $\alpha \in \mathbb{R}$ then*

$$\mathbf{E}(X + \alpha Y | \mathcal{F}) = \mathbf{E}(X | \mathcal{F}) + \alpha \mathbf{E}(Y | \mathcal{F}).$$

(2) (Positivity) *If $X \leq Y$, then $\mathbf{E}(X | \mathcal{F}) \leq \mathbf{E}(Y | \mathcal{F})$ (almost surely).*

(3) *If X is \mathcal{F} measurable and Y is an arbitrary (not necessarily \mathcal{F} -measurable) random variable then (almost surely)*

$$\mathbf{E}(XY | \mathcal{F}) = X \mathbf{E}(Y | \mathcal{F}).$$

(4) (Tower property) *If $\mathcal{E} \subseteq \mathcal{F} \subseteq \mathcal{G}$ are σ -algebras, then (almost surely)*

$$\mathbf{E}(X | \mathcal{E}) = \mathbf{E}\left(\mathbf{E}(X | \mathcal{F}) \Big| \mathcal{E}\right).$$

PROOF. The first property follows immediately from linearity. For the second property, set $Z = Y - X$ and observe

$$\int_{\mathbf{E}(Z|\mathcal{F})} \mathbf{E}(Z | \mathcal{F}) d\mathbf{P} = \int_{\mathbf{E}(Z|\mathcal{F})} Z d\mathbf{P} \geq 0,$$

which can only happen if $\mathbf{P}(\mathbf{E}(Z | \mathcal{F}) < 0) = 0$. The third property is easily checked for simple random variables, and follows in general by approximating. The tower property follows immediately from the definition. \square

As an illustration, of how the above properties come in handy, we show how they can be used to deduce (5.5).

PROPOSITION 5.12. *If $\mathcal{F} \subseteq \mathcal{G}$ is a σ -sub-algebra, X is a \mathcal{G} -measurable random variable, and Y is an \mathcal{F} -measurable random variable, then*

$$\mathbf{E}(XY) = \mathbf{E}(\mathbf{E}(X | \mathcal{F})Y)$$

PROOF. Using Remark 5.4 and then Proposition 5.11 part (3), we see

$$\mathbf{E}(XY) = \mathbf{E}(\mathbf{E}(XY | \mathcal{F})) = \mathbf{E}(Y \mathbf{E}(X | \mathcal{F})),$$

as desired. \square

Finally we conclude this section by showing that the conditional expectation of a random variable according to Definition 5.1 is precisely the minimizer as in (5.1).

PROPOSITION 5.13. *Let X be a square integrable \mathcal{G} -measurable random variable, and $\mathcal{F} \subseteq \mathcal{G}$ be a σ -sub-algebra of \mathcal{G} . Then amongst all \mathcal{F} -measurable random variables Z , the one that minimizes*

$$\mathbf{E}(X - Z)^2$$

is precisely $Z = \mathbf{E}(X | \mathcal{F})$.

PROOF. Since $\mathbf{E}(X | \mathcal{F})$ is known to be an \mathcal{F} -measurable random variable, we only need to show that for any (other) \mathcal{F} -measurable random variable Z we have

$$\mathbf{E}(X - Z)^2 \geq \mathbf{E}((X - \mathbf{E}(X | \mathcal{F}))^2).$$

To see this, note

$$\begin{aligned} \mathbf{E}(X - Z)^2 &= \mathbf{E}(X - \mathbf{E}(X | \mathcal{F}) + \mathbf{E}(X | \mathcal{F}) - Z)^2 \\ &= \mathbf{E}(X - \mathbf{E}(X | \mathcal{F}))^2 + \mathbf{E}((X | \mathcal{F}) - Z)^2 \\ &\quad + 2\mathbf{E}\left(\underbrace{(X - \mathbf{E}(X | \mathcal{F}))}_I \underbrace{(\mathbf{E}(X | \mathcal{F}) - Z)}_{II}\right) \end{aligned}$$

Since term II is \mathcal{F} measurable, we can use (5.5) to replace X with $\mathbf{E}(X | \mathcal{F})$ in term I . This yields

$$\begin{aligned} \mathbf{E}(X - Z)^2 &= \mathbf{E}(X - \mathbf{E}(X | \mathcal{F}))^2 + \mathbf{E}((X | \mathcal{F}) - Z)^2 \\ &\quad + 2\mathbf{E}\left((\mathbf{E}(X | \mathcal{F}) - \mathbf{E}(X | \mathcal{F}))(\mathbf{E}(X | \mathcal{F}) - Z)\right) \\ &= \mathbf{E}(X - \mathbf{E}(X | \mathcal{F}))^2 + \mathbf{E}((X | \mathcal{F}) - Z)^2 \geq \mathbf{E}(X - \mathbf{E}(X | \mathcal{F}))^2, \end{aligned}$$

as desired. \square

6. The Martingale Property

A martingale is “fair game”. Suppose you are playing a game and $M(t)$ is your cash stockpile at time t . As time progresses, you learn more and more information about the game. For instance, in blackjack getting a high card benefits the player more than the dealer, and a common card counting strategy is to have a “spotter” betting the minimum while counting the high cards. When the odds of getting a high card are favorable enough, the player will signal a “big player” who joins the table and makes large bets, as long as the high card count is favorable. Variants of this strategy have been shown to give the player up to a 2% edge over the house.

If a game is a martingale, then this extra information you have acquired *can not* help you going forward. That is, if you signal your “big player” at any point, you will not affect your expected return.

Mathematically this translates to saying that the *conditional expectation* of your stockpile at a later time given your present accumulated knowledge, is exactly the present value of your stockpile. Our aim in this section is to make this precise.

6.1. Adapted processes and filtrations. Let X be any stochastic process (for example Brownian motion). For any $t > 0$, we’ve seen before that $\sigma(X(t))$ represents the information you obtain by observing $X(t)$. Accumulating this over time gives us the *filtration*. To introduce this concept, we first need the notion of a σ algebra generated by a family of sets.

DEFINITION 6.1. Given a collection of sets A_α , where α belongs to some (possibly infinite) index set \mathcal{A} , we define $\sigma(\{A_\alpha\})$ to be the *smallest* σ -algebra that contains each of the sets A_α .

That is, if $\mathcal{G} = \sigma(\{A_\alpha\})$, then we must have each $A_\alpha \in \mathcal{G}$. Since \mathcal{G} is a σ -algebra, all sets you can obtain from these by taking complements, countable unions and

countable intersections intersections must also belong to \mathcal{G} .⁷ The fact that \mathcal{G} is the smallest σ -algebra containing each A_α also means that if \mathcal{G}' is any other σ -algebra that contains each A_α , then $\mathcal{G} \subseteq \mathcal{G}'$.

REMARK 6.2. The smallest σ -algebra under which X is a random variable (under which X is measurable) is exactly $\sigma(X)$. It turns out that $\sigma(X) = X^{-1}(\mathcal{B}) = \{X \in B \mid B \in \mathcal{B}\}$, where \mathcal{B} is the *Borel* σ -algebra on \mathbb{R} . Here \mathcal{B} is the *Borel* σ -algebra, defined to be the σ -algebra on \mathbb{R} generated by all open intervals.

DEFINITION 6.3. Given a stochastic process X , the *filtration generated by X* is the family of σ -algebras $\{\mathcal{F}_t^X \mid t \geq 0\}$ where

$$\mathcal{F}_t^X \stackrel{\text{def}}{=} \sigma\left(\bigcup_{s \leq t} \sigma(X_s)\right).$$

That is, \mathcal{F}_t^X is all events that can be observed using only the random variables X_s when $s \leq t$. Clearly each \mathcal{F}_t^X is a σ -algebra, and if $s \leq t$, $\mathcal{F}_s^X \subseteq \mathcal{F}_t^X$. A family of σ -algebras with this property is called a *filtration*.

DEFINITION 6.4. A *filtration* is a family of σ -algebras $\{\mathcal{F}_t \mid t \geq 0\}$ such that whenever $s \leq t$, we have $\mathcal{F}_s \subseteq \mathcal{F}_t$.

In our case, the filtration we work with will most often be the *Brownian filtration*, i.e. the filtration generated by Brownian motion. However, one can (and often needs to) consider more general filtrations. In this case the intuition we use is that the σ -algebra \mathcal{F}_t represents the information accumulated up to time t (i.e. all events whose probabilities can be deduced up to time t). When given a filtration, it is important that all stochastic processes we construct respect the flow of information, and do not look into the future. This is of course natural: trading / pricing strategies can not rely on the price at a later time, and gambling strategies do not know the outcome of the next hand. Mathematically this property is called *adapted*, and is defined as follows.

DEFINITION 6.5. A stochastic process X is said to be *adapted* to a filtration $\{\mathcal{F}_t \mid t \geq 0\}$ if for every t the random variable $X(t)$ is \mathcal{F}_t measurable (i.e. $\{X(t) \leq \alpha\} \in \mathcal{F}_t$ for every $\alpha \in \mathbb{R}$, $t \geq 0$).

Clearly a process X is adapted with respect to the filtration it generates $\{\mathcal{F}_t^X\}$.

6.2. Martingales. Recall, a martingale is a “fair game”. Using conditional expectations, we can now define this precisely.

DEFINITION 6.6. A stochastic process M is a martingale with respect to a filtration $\{\mathcal{F}_t\}$ if:

- (1) M is adapted to the filtration $\{\mathcal{F}_t\}$.
- (2) For any $s < t$ we have $\mathbf{E}(M(t) \mid \mathcal{F}_s) = M(s)$, almost surely.

⁷Usually \mathcal{G} contains *much more* than all countable unions, intersections and complements of the A_α 's. You might think you could keep including all sets you generate using countable unions and complements and arrive at all of \mathcal{G} . It turns out that to make this work, you will usually have to do this *uncountably* many times!

This won't be too important within the scope of these notes. However, if you read a rigorous treatment and find the authors using some fancy trick (using Dynkin systems or monotone classes) instead of a naive countable unions argument, then the above is the reason why.

REMARK 6.7. A *sub-martingale* is an adapted process M for which we have $\mathbf{E}(M(t) | \mathcal{F}_s) \geq M(s)$, and a *super-martingale* if $\mathbf{E}(M(t) | \mathcal{F}_s) \leq M(s)$. Thus $\mathbf{E}M(t)$ is an increasing function of time if M is a sub-martingale, constant in time if M is a martingale, and a decreasing function of time if M is a super-martingale.

REMARK 6.8. It is crucial to specify the filtration when talking about martingales, as it is certainly possible that a process is a martingale with respect to one filtration but not with respect to another. For our purposes the filtration will almost always be the *Brownian filtration* (i.e. the filtration generated by Brownian motion).

EXAMPLE 6.9. Let $\{\mathcal{F}_t\}$ be a filtration, $\mathcal{F}_\infty = \sigma(\cup_{t \geq 0} \mathcal{F}_t)$, and X be any \mathcal{F}_∞ -measurable random variable. The process $M(t) \stackrel{\text{def}}{=} \mathbf{E}(X_\infty | \mathcal{F}_t)$ is a martingale with respect to the filtration $\{\mathcal{F}_t\}$.

6.3. The martingale property of Brownian motion. In discrete time a random walk is a martingale, so it is natural to expect that in continuous time Brownian motion is a martingale as well.

THEOREM 6.10. *Let W be a Brownian motion, $\mathcal{F}_t = \mathcal{F}_t^W$ be the Brownian filtration. Brownian motion is a martingale with respect to this filtration.*

PROOF. By independence of increments, $W(t) - W(s)$ is certainly independent of $W(r)$ for any $r \leq s$. Since $\mathcal{F}_s = \sigma(\cup_{r \leq s} \sigma(W(r)))$ we expect that $W(t) - W(s)$ is independent of \mathcal{F}_s . Consequently

$$\mathbf{E}(W(t) | \mathcal{F}_s) = \mathbf{E}(W(t) - W(s) | \mathcal{F}_s) + \mathbf{E}(W(s) | \mathcal{F}_s) = 0 + W(s) = W(s). \quad \square$$

THEOREM 6.11. *Let W be a standard Brownian motion (i.e. a Brownian motion normalized so that $W(0) = 0$ and $\text{Var}(W(t)) = t$). For any $C_b^{1,2}$ function⁸ $f = f(t, x)$ the process*

$$M(t) \stackrel{\text{def}}{=} f(t, W(t)) - \int_0^t \left(\partial_t f(s, W(s)) + \frac{1}{2} \partial_x^2 f(s, W(s)) \right) ds$$

is a martingale (with respect to the Brownian filtration).

PROOF. This is an extremely useful fact about Brownian motion follows quickly from the Itô formula, which we will discuss later. However, at this stage, we can provide a simple, elegant and instructive proof as follows.

Adaptedness of M is easily checked. To compute $\mathbf{E}(M(t) | \mathcal{F}_r)$ we first observe

$$\mathbf{E}(f(t, W(t)) | \mathcal{F}_r) = \mathbf{E}(f(t, [W(t) - W(r)] + W(r)) | \mathcal{F}_r).$$

Since $W(t) - W(r) \sim N(0, t - r)$ and is independent of \mathcal{F}_r , the above conditional expectation can be computed by

$$\mathbf{E}(f(t, [W(t) - W(r)] + W(r)) | \mathcal{F}_r) = \int_{\mathbb{R}} f(t, y + W(r)) G(t - r, y) dy,$$

where

$$G(\tau, y) = \frac{1}{\sqrt{2\pi\tau}} \exp\left(\frac{-y^2}{2\tau}\right)$$

⁸Recall a function $f = f(t, x)$ is said to be $C^{1,2}$ if it is C^1 in t (i.e. differentiable with respect to t and $\partial_t f$ is continuous), and C^2 in x (i.e. twice differentiable with respect to x and $\partial_x f, \partial_x^2 f$ are both continuous). The space $C_b^{1,2}$ refers to all $C^{1,2}$ functions f for which and $f, \partial_t f, \partial_x f, \partial_x^2 f$ are all bounded functions.

is the density of $W(t) - W(r)$.

Similarly

$$\begin{aligned} \mathbf{E} \left(\int_0^t (\partial_t f(s, W(s)) + \frac{1}{2} \partial_x^2 f(s, W(s))) ds \mid \mathcal{F}_r \right) \\ = \int_0^r (\partial_t f(s, W(s)) + \frac{1}{2} \partial_x^2 f(s, W(s))) ds \\ + \int_r^t \int_{\mathbb{R}} (\partial_t f(s, y + W(r)) + \frac{1}{2} \partial_x^2 f(s, y + W(r))) G(s - r, y) dy ds \end{aligned}$$

Hence

$$\begin{aligned} \mathbf{E}(M(t) \mid \mathcal{F}_r) - M(r) &= \int_{\mathbb{R}} f(t, y + W(r)) G(t - r, y) dy \\ &\quad - \int_r^t \int_{\mathbb{R}} (\partial_t f(s, y + W(r)) + \frac{1}{2} \partial_x^2 f(s, y + W(r))) G(s - r, y) dy \\ &\quad - f(r, W(r)). \end{aligned}$$

We claim that the right hand side of the above vanishes. In fact, we claim the (deterministic) identity

$$\begin{aligned} f(r, x) &= \int_{\mathbb{R}} f(t, y + x) G(t - r, y) dy \\ &\quad - \int_r^t \int_{\mathbb{R}} (\partial_t f(s, y + x) + \frac{1}{2} \partial_x^2 f(s, y + x)) G(s - r, y) dy ds \end{aligned}$$

holds for any function f and $x \in \mathbb{R}$. For those readers who are familiar with PDEs, this is simply the Duhamel's principle for the heat equation. If you're unfamiliar with this, the above identity can be easily checked using the fact that $\partial_\tau G = \frac{1}{2} \partial_y^2 G$ and integrating the first integral by parts. We leave this calculation to the reader. \square

6.4. Stopping Times. For this section we assume that a filtration $\{\mathcal{F}_t\}$ is given to us, and fixed. When we refer to process being adapted (or martingales), we implicitly mean they are adapted (or martingales) with respect to this filtration.

Consider a game (played in continuous time) where you have the option to walk away at any time. Let τ be the *random* time you decide to stop playing and walk away. In order to respect the flow of information, you need to be able to decide whether you have stopped using only information up to the present. At time t , event $\{\tau \leq t\}$ is exactly when you have stopped and walked away. Thus, to respect the flow of information, we need to ensure $\{\tau \leq t\} \in \mathcal{F}_t$.

DEFINITION 6.12. A stopping time is a function $\tau: \Omega \rightarrow [0, \infty)$ such that for every $t \geq 0$ the event $\{\tau \leq t\} \in \mathcal{F}_t$.

A standard example of a stopping time is *hitting times*. Say you decide to liquidate your position once the value of your portfolio reaches a certain threshold. The time at which you liquidate is a hitting time, and under mild assumptions on the filtration, will always be a stopping time.

PROPOSITION 6.13. *Let X be an adapted continuous process, $\alpha \in \mathbb{R}$ and τ be the first time X hits α (i.e. $\tau = \inf\{t \geq 0 \mid X(t) = \alpha\}$). Then τ is a stopping time (if the filtration is right continuous).*

THEOREM 6.14 (Doob's optional sampling theorem). *If M is a martingale and τ is a bounded stopping time. Then the stopped process $M^\tau(t) \stackrel{\text{def}}{=} M(\tau \wedge t)$ is also a martingale. Consequently, $\mathbf{E}M(\tau) = \mathbf{E}M(\tau \wedge t) = \mathbf{E}M(0) = \mathbf{E}M(t)$ for all $t \geq 0$.*

REMARK 6.15. If instead of assuming τ is bounded, we assume M^τ is bounded the above result is still true.

The proof goes beyond the scope of these notes, and can be found in any standard reference. What this means is that if you're playing a fair game, then you can not hope to improve your odds by "quitting when you're ahead". Any rule by which you decide to stop, must be a stopping time and the above result guarantees that stopping a martingale still yields a martingale.

REMARK 6.16. Let W be a standard Brownian motion, τ be the first hitting time of W to 1. Then $\mathbf{E}W(\tau) = 1 \neq 0 = \mathbf{E}W(t)$. This is one situation where the optional sampling theorem doesn't apply (in fact, $\mathbf{E}\tau = \infty$, and W^τ is unbounded).

This example corresponds to the gambling strategy of walking away when you make your "million". The reason it's not a sure bet is because the time taken to achieve your winnings is finite almost surely, but very long (since $\mathbf{E}\tau = \infty$). In the mean time you might have incurred financial ruin and expended your entire fortune.

Suppose the price of a security you're invested in fluctuates like a martingale (say for instance Brownian motion). This is of course unrealistic, since Brownian motion can also become negative; but let's use this as a first example. You decide you're going to liquidate your position and walk away when either you're bankrupt, or you make your first million. What are your expected winnings? This can be computed using the optional sampling theorem.

PROBLEM 6.1. Let $a \geq 0$ and M be any continuous martingale with $M(0) = x \in (0, a)$. Let τ be the first time M hits either 0 or a . Compute $\mathbf{P}(M(\tau) = a)$ and your expected return $\mathbf{E}M(\tau)$.

Stochastic Integration

1. Motivation

Suppose $\Delta(t)$ is your position at time t on a security whose price is $S(t)$. If you only trade this security at times $0 = t_0 < t_1 < t_2 < \dots < t_n = T$, then the change in the value of your wealth up to time T is given by

$$X(t_n) - X(0) = \sum_{i=0}^{n-1} \Delta(t_i)(S(t_{i+1}) - S(t_i))$$

If you are trading this continuously in time, you'd expect that a "simple" limiting procedure should show that your wealth is given by the *Riemann-Stieltjes* integral:

$$X(T) - X(0) = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \Delta(t_i)(S(t_{i+1}) - S(t_i)) = \int_0^T \Delta(t) dS(t).$$

Here $P = \{0 = t_0 < \dots < t_n = T\}$ is a partition of $[0, T]$, and $\|P\| = \max\{t_{i+1} - t_i\}$.

This has been well studied by mathematicians, and it is well known that for the above limiting procedure to "work directly", you need S to have finite *first variation*. Recall, the *first variation* of a function is defined to be

$$V_{[0,T]}(S) \stackrel{\text{def}}{=} \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} |S(t_{i+1}) - S(t_i)|.$$

It turns out that almost *any* continuous martingale S *will not* have finite first variation. Thus to define integrals with respect to martingales, one has to do something 'clever'. It turns out that if X is adapted and S is an martingale, then the above limiting procedure works, and this was carried out by Itô (and independently by Doebelin).

2. The First Variation of Brownian motion

We begin by showing that the first variation of Brownian motion is infinite.

PROPOSITION 2.1. *If W is a standard Brownian motion, and $T > 0$ then*

$$\lim_{n \rightarrow \infty} \mathbf{E} \sum_{k=0}^{n-1} \left| W\left(\frac{k+1}{n}\right) - W\left(\frac{k}{n}\right) \right| = \infty.$$

REMARK 2.2. In fact

$$\lim_{n \rightarrow \infty} \sum_{k=0}^{n-1} \left| W\left(\frac{k+1}{n}\right) - W\left(\frac{k}{n}\right) \right| = \infty \quad \text{almost surely,}$$

but this won't be necessary for our purposes.

PROOF. Since $W((k+1)/n) - W(k/n) \sim N(0, 1/n)$ we know

$$\mathbf{E} \left| W\left(\frac{k+1}{n}\right) - W\left(\frac{k}{n}\right) \right| = \int_{\mathbb{R}} |x| G\left(\frac{1}{n}, x\right) dx = \frac{C}{\sqrt{n}},$$

where

$$C = \int_{\mathbb{R}} |y| e^{-y^2/2} \frac{dy}{\sqrt{2\pi}} = \mathbf{E}|N(0, 1)|.$$

Consequently

$$\sum_{k=0}^{n-1} \mathbf{E} \left| W\left(\frac{k+1}{n}\right) - W\left(\frac{k}{n}\right) \right| = \frac{Cn}{\sqrt{n}} \xrightarrow{n \rightarrow \infty} \infty. \quad \square$$

3. Quadratic Variation

It turns out that the *second* variation of any *square integrable* martingale is almost surely finite, and this is the key step in constructing the Itô integral.

DEFINITION 3.1. Let M be any process. We define the *quadratic variation* of M , denoted by $[M, M]$ by

$$[M, M](T) = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} (\Delta_i M)^2,$$

where $P = \{0 = t_1 < t_1 \cdots < t_n = T\}$ is a partition of $[0, T]$, and $\Delta_i M = M(t_{i+1}) - M(t_i)$.

PROPOSITION 3.2. *If W is a standard Brownian motion, then $[W, W](T) = T$ almost surely.*

PROOF. For simplicity, let's assume $t_i = Ti/n$. Note

$$\sum_{i=0}^{n-1} (\Delta_i W)^2 - T = \sum_{i=0}^{n-1} \left(W\left(\frac{(i+1)T}{n}\right) - W\left(\frac{iT}{n}\right) \right)^2 - T = \sum_{i=0}^{n-1} \xi_i,$$

where

$$\xi_i \stackrel{\text{def}}{=} (\Delta_i W)^2 - \frac{T}{n} = \left(W\left(\frac{(i+1)T}{n}\right) - W\left(\frac{iT}{n}\right) \right)^2 - \frac{T}{n}.$$

Note that ξ_i 's are i.i.d. with distribution $N(0, T/n)^2 - T/n$, and hence

$$\mathbf{E}\xi_i = 0 \quad \text{and} \quad \text{Var} \xi_i = \frac{T^2(\mathbf{E}N(0, 1)^4 - 1)}{n^2}.$$

Consequently

$$\text{Var} \left(\sum_{i=0}^{n-1} \xi_i \right) = \frac{T^2(\mathbf{E}N(0, 1)^4 - 1)}{n} \xrightarrow{n \rightarrow \infty} 0,$$

which shows

$$\sum_{i=0}^{n-1} \left(W\left(\frac{(i+1)T}{n}\right) - W\left(\frac{iT}{n}\right) \right)^2 - T = \sum_{i=0}^{n-1} \xi_i \xrightarrow{n \rightarrow \infty} 0. \quad \square$$

COROLLARY 3.3. *The process $M(t) \stackrel{\text{def}}{=} W(t)^2 - [W, W](t)$ is a martingale.*

PROOF. We see

$$\begin{aligned} \mathbf{E}(W(t)^2 - t | \mathcal{F}_s) &= \mathbf{E}((W(t) - W(s))^2 + 2W(s)(W(t) - W(s)) + W(s)^2 | \mathcal{F}_s) - t \\ &= W(s)^2 - s \end{aligned}$$

and hence $\mathbf{E}(M(t) | \mathcal{F}_s) = M(s)$. \square

The above wasn't a co-incidence. This property in fact characterizes the quadratic variation.

THEOREM 3.4. *Let M be a continuous martingale with respect to a filtration $\{\mathcal{F}_t\}$. Then $\mathbf{E}M(t)^2 < \infty$ if and only if $\mathbf{E}[M, M](t) < \infty$. In this case the process $M(t)^2 - [M, M](t)$ is also a martingale with respect to the same filtration, and hence $\mathbf{E}M(t)^2 - \mathbf{E}M(0)^2 = \mathbf{E}[M, M](t)$.*

The above is in fact a characterization of the quadratic variation of martingales.

THEOREM 3.5. *If $A(t)$ is any continuous, increasing, adapted process such that $A(0) = 0$ and $M(t)^2 - A(t)$ is a martingale, then $A = [M, M]$.*

The proof of these theorems are a bit technical and go beyond the scope of these notes. The results themselves, however, are extremely important and will be used subsequently.

REMARK 3.6. The intuition to keep in mind about the first variation and the quadratic variation is the following. Divide the interval $[0, T]$ into $T/\delta t$ intervals of size δt . If X has finite first variation, then on each subinterval $(k\delta t, (k+1)\delta t)$ the increment of X should be of order δt . Thus adding $T/\delta t$ terms of order δt will yield something finite.

On the other hand if X has finite quadratic variation, on each subinterval $(k\delta t, (k+1)\delta t)$ the increment of X should be of order $\sqrt{\delta t}$, so that adding $T/\delta t$ terms of the *square* of the increment yields something finite. Doing a quick check for Brownian motion (which has finite quadratic variation), we see

$$\mathbf{E}|W(t + \delta t) - W(t)| = \sqrt{\delta t} \mathbf{E}|N(0, 1)|,$$

which is in line with our intuition.

REMARK 3.7. If a continuous process has finite first variation, its quadratic variation will necessarily be 0. On the other hand, if a continuous process has finite (and non-zero) quadratic variation, its first variation will necessary be infinite.

4. Construction of the Itô integral

Let W be a standard Brownian motion, $\{\mathcal{F}_t\}$ be the Brownian filtration and D be an adapted process. We think of $D(t)$ to represent our position at time t on an asset whose spot price is $W(t)$.

LEMMA 4.1. *Let $P = \{0 = t_0 < t_1 < t_2 < \dots\}$ be an increasing sequence of times, and assume D is constant on $[t_i, t_{i+1})$ (i.e. the asset is only traded at times t_0, \dots, t_n). Let $I_P(T)$, defined by*

$$I_P(T) = \sum_{i=0}^{n-1} D(t_i) \Delta_i W + D(t_n)(W(T) - W(t_n)) \quad \text{if } T \in [t_n, t_{n+1}).$$

be your cumulative winnings up to time T . As before $\Delta_i W \stackrel{\text{def}}{=} W(t_{i+1}) - W(t_i)$. Then,

$$(4.1) \quad \mathbf{E}I_P(T)^2 = \mathbf{E} \left[\sum_{i=0}^n D(t_i)^2 (t_{i+1} - t_i) + D(t_n)^2 (T - t_n) \right] \quad \text{if } T \in [t_n, t_{n+1}).$$

Moreover, I_P is a martingale and

$$(4.2) \quad [I_P, I_P](T) = \sum_{i=0}^{n-1} D(t_i)^2 (t_{i+1} - t_i) + D(t_n)^2 (T - t_n) \quad \text{if } T \in [t_n, t_{n+1}).$$

This lemma, as we will shortly see, is the key to the construction of stochastic integrals (called Itô integrals).

PROOF. We first prove (4.1) with $T = t_n$ for simplicity. Note

$$(4.3) \quad \mathbf{E}I_P(t_n)^2 = \sum_{i=0}^{n-1} \mathbf{E}D(t_i)^2 (\Delta_i W)^2 + 2 \sum_{j=0}^{n-1} \sum_{i=0}^{j-1} \mathbf{E}D(t_i)D(t_j) (\Delta_i W)(\Delta_j W)$$

By the tower property

$$\begin{aligned} \mathbf{E}D(t_i)^2 (\Delta_i W)^2 &= \mathbf{E} \mathbf{E}(D(t_i)^2 (\Delta_i W)^2 \mid \mathcal{F}_{t_i}) \\ &= \mathbf{E}D(t_i)^2 \mathbf{E}((W(t_{i+1}) - W(t_i))^2 \mid \mathcal{F}_{t_i}) = \mathbf{E}D(t_i)^2 (t_{i+1} - t_i). \end{aligned}$$

Similarly we compute

$$\begin{aligned} \mathbf{E}D(t_i)D(t_j) (\Delta_i W)(\Delta_j W) &= \mathbf{E} \mathbf{E}(D(t_i)D(t_j) (\Delta_i W)(\Delta_j W) \mid \mathcal{F}_{t_j}) \\ &= \mathbf{E}D(t_i)D(t_j) (\Delta_i W) \mathbf{E}((W(t_{j+1}) - W(t_j)) \mid \mathcal{F}_{t_j}) = 0. \end{aligned}$$

Substituting these in (4.3) immediately yields (4.1) for $t_n = T$.

The proof that I_P is an martingale uses the same “tower property” idea, and is left to the reader to check. The proof of (4.2) is also similar in spirit, but has a few more details to check. The main idea is to let $A(t)$ be the right hand side of (4.2). Observe A is clearly a continuous, increasing, adapted process. Thus, if we show $M^2 - A$ is a martingale, then using Theorem 3.5 we will have $A = [M, M]$ as desired. The proof that $M^2 - A$ is an martingale uses the same “tower property” idea, but is a little more technical and is left to the reader. \square

Note that as $\|P\| \rightarrow 0$, the right hand side of (4.2) converges to the standard Riemann integral $\int_0^T D(t)^2 dt$. Itô realised he could use this to prove that I_P itself converges, and the limit is now called the Itô integral.

THEOREM 4.2. *If $\int_0^T D(t)^2 dt < \infty$ almost surely, then as $\|P\| \rightarrow 0$, the processes I_P converge to a continuous process I denoted by*

$$(4.4) \quad I(T) \stackrel{\text{def}}{=} \lim_{\|P\| \rightarrow 0} I_P(T) \stackrel{\text{def}}{=} \int_0^T D(t) dW(t).$$

This is known as the Itô integral of D with respect to W . If further

$$(4.5) \quad \mathbf{E} \int_0^T D(t)^2 dt < \infty,$$

then the process $I(T)$ is a martingale and the quadratic variation $[I, I]$ satisfies

$$[I, I](T) = \int_0^T D(t)^2 dt \quad \text{almost surely.}$$

REMARK 4.3. For the above to work, it is *crucial* that D is adapted, and is sampled at the left endpoint of the time intervals. That is, the terms in the sum are $D(t_i)(W(t_{i+1}) - W(t_i))$, and not $D(t_{i+1})(W(t_{i+1}) - W(t_i))$ or $\frac{1}{2}(D(t_i) + D(t_{i+1}))(W(t_{i+1}) - W(t_i))$, or something else.

Usually if the process is not adapted, there is no meaningful way to make sense of the limit. However, if you sample at different points, it still works out (usually) but what you get is *different* from the Itô integral (one example is the Stratonovich integral).

REMARK 4.4. The variable t used in (4.4) is a “dummy” integration variable. Namely one can write

$$\int_0^T D(t) dW(t) = \int_0^T D(s) dW(s) = \int_0^T D(r) dW(r),$$

or any other variable of your choice.

COROLLARY 4.5 (Itô Isometry). *If (4.5) holds then*

$$\mathbf{E}\left(\int_0^T D(t) dW(t)\right)^2 = \mathbf{E} \int_0^T D(t)^2 dt.$$

PROPOSITION 4.6 (Linearity). *If D_1 and D_2 are two adapted processes, and $\alpha \in \mathbb{R}$, then*

$$\int_0^T (D_1(t) + \alpha D_2(t)) dW(t) = \int_0^T D_1(t) dW(t) + \alpha \int_0^T D_2(t) dW(t).$$

REMARK 4.7. Positivity, however, is not preserved by Itô integrals. Namely if $D_1 \leq D_2$, there is no reason to expect $\int_0^T D_1(t) dW(t) \leq \int_0^T D_2(t) dW(t)$. Indeed choosing $D_1 = 0$ and $D_2 = 1$ we see that we can not possibly have $0 = \int_0^T D_1(t) dW(t)$ to be almost surely smaller than $W(T) = \int_0^T D_2(t) dW(t)$.

Recall, our starting point in these notes was modelling stock prices as *geometric Brownian motions*, given by the equation

$$dS(t) = \alpha S(t) dt + \sigma S(t) dW(t).$$

After constructing Itô integrals, we are now in a position to describe what this means. The above is simply shorthand for saying S is a process that satisfies

$$S(T) - S(0) = \int_0^T \alpha S(t) dt + \int_0^T \sigma S(t) dW(t).$$

The first integral on the right is a standard Riemann integral. The second integral, representing the noisy fluctuations, is the Itô integral we just constructed.

Note that the above is a little more complicated than the Itô integrals we will study first, since the process S (that we’re trying to define) also appears as an integrand on the right hand side. In general, such equations are called *Stochastic differential equations*, and are extremely useful in many contexts.

5. The Itô formula

Using the abstract “limit” definition of the Itô integral, it is hard to compute examples. For instance, what is

$$\int_0^T W(s) dW(s) ?$$

This, as we will shortly, can be computed easily using the Itô formula (also called the Itô-Doebelin formula).

Suppose b and σ are adapted processes. (In particular, they could but need not, be random). Consider a process X defined by

$$(5.1) \quad X(T) = X(0) + \int_0^T b(t) dt + \int_0^T \sigma(t) dW(t).$$

Note the first integral $\int_0^T b(t) dt$ is a regular Riemann integral that can be done directly. The second integral the Itô integral we constructed in the previous section.

DEFINITION 5.1. The process X is called an Itô process if $X(0)$ is deterministic (not random) and for all $T \geq 0$,

$$(5.2) \quad E \int_0^T \sigma(t)^2 dt < \infty \quad \text{and} \quad \int_0^T b(t) dt < \infty.$$

REMARK 5.2. The shorthand notation for (5.1) is to write

$$(5.1') \quad dX(t) = b(t) dt + \sigma(t) dW(t).$$

PROPOSITION 5.3. *The quadratic variation of X is*

$$(5.3) \quad [X, X](T) = \int_0^T \sigma(t)^2 dt \quad \text{almost surely.}$$

PROOF. Define B and M by

$$B(T) = \int_0^T b(t) dt \quad \text{and} \quad M(T) = \int_0^T \sigma(t) dW(t),$$

and let $P = \{0 = t_0 < t_1 < \dots < t_n = T\}$ be a partition of $[0, T]$, and $\|P\| = \max_i t_{i+1} - t_i$. Observe

$$\sum_{i=0}^{n-1} (\Delta_i X)^2 = \sum_{i=0}^{n-1} (\Delta_i M)^2 + \sum_{i=0}^{n-1} (\Delta_i B)^2 + 2 \sum_{i=0}^{n-1} (\Delta_i B)(\Delta_i M).$$

The first sum on the right converges (as $\|P\| \rightarrow 0$) to $[M, M](T)$, which we know is exactly $\int_0^T \sigma(t)^2 dt$. For the second sum, observe

$$(\Delta_i B)^2 = \left(\int_{t_i}^{t_{i+1}} b(s) ds \right)^2 \leq (\max |b|^2) (t_{i+1} - t_i)^2 \leq (\max |b|^2) \|P\| (t_{i+1} - t_i).$$

Hence

$$\left| \sum_{i=0}^{n-1} (\Delta_i B)^2 \right| \leq \|P\| (\max |b|^2) T \xrightarrow{\|P\| \rightarrow 0} 0.$$

For the third term, one uses the *Cauchy-Schwartz* inequality to observe

$$\left| \sum_{i=0}^{n-1} (\Delta_i B)(\Delta_i M) \right| \leq \left(\sum_{i=0}^{n-1} (\Delta_i B)^2 \right)^{1/2} \left(\sum_{i=0}^{n-1} (\Delta_i M)^2 \right)^{1/2} \xrightarrow{\|P\| \rightarrow 0} 0 \cdot [M, M](T) = 0.$$

□

REMARK 5.4. It's common to decompose $X = B + M$ where M is a martingale and B has finite first variation. Processes that can be decomposed in this form are called *semi-martingales*, and the decomposition is unique. The process M is called the martingale part of X , and B is called the *bounded variation* part of X .

PROPOSITION 5.5. *The semi-martingale decomposition of X is unique. That is, if $X = B_1 + M_1 = B_2 + M_2$ where B_1, B_2 are continuous adapted processes with finite first variation, and M_1, M_2 are continuous (square integrable) martingales, then $B_1 = B_2$ and $M_1 = M_2$.*

PROOF. Set $M = M_1 - M_2$ and $B = B_1 - B_2$, and note that $M = -B$. Consequently, M has finite first variation and hence 0 quadratic variation. This implies $\mathbf{E}M(t)^2 = \mathbf{E}[M, M](t) = 0$ and hence $M = 0$ identically, which in turn implies $B = 0$, $B_1 = B_2$ and $M_1 = M_2$. □

As an immediate consequence, we see that the sum of an Itô integral and Riemann integral is a martingale, if and only if the Riemann integral is identically 0.

PROPOSITION 5.6. *Suppose b, σ are two adapted processes satisfying (5.2), and let X be as in equation (5.1). Then X is a martingale if and only if b is identically 0.*

PROOF 1. If b is identically 0, then we already know that X is a martingale. To prove the converse, suppose X is a martingale. Then the process $B(T) = \int_0^T b(t) dt = X(T) - X(0) - \int_0^T \sigma(s) ds$ is the difference of two martingales, and so must itself be a martingale. Now the Itô decomposition of the process B can be expressed in two different ways: First $B = B + 0$, where B has bounded variation, and 0 is a martingale. Second $B = 0 + B$, where 0 has bounded variation, and B is a martingale. By Proposition 5.5, the bounded variation parts and martingale parts must be equal, showing $B = 0$ identically. □

PROOF 2. Here's an alternate, direct proof of Proposition 5.6 without relying on Proposition 5.5. Suppose X is a martingale. Then as above, the process $B(T) = \int_0^T b(t) dt$ must also be a martingale, and so we must have $\mathbf{E}(B(T+h) | \mathcal{F}_T) = B(T)$ for every $h > 0$. Thus

$$\begin{aligned} \int_0^T b(t) dt = B(T) &= \mathbf{E}(B(T+h) | \mathcal{F}_T) = \mathbf{E}\left(\int_0^{T+h} b(t) dt \mid F_T\right) \\ &= \int_0^T b(t) dt + \mathbf{E}\left(\int_T^{T+h} b(t) dt \mid F_T\right). \end{aligned}$$

This implies

$$\mathbf{E}\left(\int_T^{T+h} b(t) dt \mid F_T\right) = 0,$$

for every $h > 0$. Dividing both sides by h , and taking the limit as $h \rightarrow 0$ shows

$$\begin{aligned} 0 &= \lim_{h \rightarrow 0} \mathbf{E}\left(\frac{1}{h} \int_T^{T+h} b(t) dt \mid F_T\right) = \mathbf{E}\left(\lim_{h \rightarrow 0} \frac{1}{h} \int_T^{T+h} b(t) dt \mid F_T\right) \\ &= \mathbf{E}\left(b(T) \mid F_T\right) = b(T). \end{aligned}$$

Thus $b(T) = 0$ for every $T \geq 0$. This forces the process B to be identically 0, concluding the proof. □

Given an adapted process D , interpret X as the price of an asset, and D as our position on it. (We could either be long, or short on the asset so D could be positive or negative.)

DEFINITION 5.7. We define the integral of D with respect to X by

$$\int_0^T D(t) dX(t) \stackrel{\text{def}}{=} \int_0^T D(t)b(t) dt + \int_0^T D(t)\sigma(t) dW(t).$$

As before, $\int_0^T D dX$ represents the winnings or profit obtained using the trading strategy D .

REMARK 5.8. Note that the first integral on the right $\int_0^T D(t)b(t) dt$ is a regular Riemann integral, and the second one is an Itô integral. Recall that Itô integrals with respect to Brownian motion (i.e. integrals of the form $\int_0^t D(s) dW(s)$ are martingales). Integrals with respect to a general process X are only guaranteed to be martingales if X itself is a martingale (i.e. $b = 0$), or if the integrand is 0.

REMARK 5.9. If we define I_P by

$$I_P(T) = \sum_{i=0}^{n-1} D(t_i)(\Delta_i X) + D(t_n)(X(T) - X(t_n)) \quad \text{if } T \in [t_n, t_{n+1}),$$

then I_P converges to the integral $\int_0^T D(t) dX(t)$ defined above. This works in the same way as Theorem 4.2.

Suppose now $f(t, x)$ is some function. If X is differentiable as a function of t (which it most certainly is not), then the chain rule gives

$$\begin{aligned} f(T, X(T)) - f(0, X(0)) &= \int_0^T \partial_t (f(t, X(t))) dt \\ &= \int_0^T \partial_t f(t, X(t)) dt + \int_0^T \partial_x f(t, X(t)) \partial_t X(t) dt \\ &= \int_0^T \partial_t f(t, X(t)) dt + \int_0^T \partial_x f(t, X(t)) dX(t). \end{aligned}$$

Itô process are *almost never* differentiable as a function of time, and so the above has no chance of working. It turns out, however, that for Itô process you can make the above work by adding an *Itô correction* term. This is the celebrated Itô formula (more correctly the Itô-Doeblin¹ formula).

THEOREM 5.10 (Itô formula, aka Itô-Doeblin formula). *If $f = f(t, x)$ is $C^{1,2}$ function² then*

¹W. Doeblin was a French-German mathematician who was drafted for military service during the second world war. During the war he wrote down his mathematical work and sent it in a sealed envelope to the French Academy of Sciences, because he did not want it to “fall into the wrong hands”. When he was about to be captured by the Germans he burnt his mathematical notes and killed himself.

The sealed envelope was opened in 2000 which revealed that he had a treatment of stochastic Calculus that was essentially equivalent to Itô’s. In posthumous recognition, Itô’s formula is now referred to as the Itô-Doeblin formula by many authors.

²Recall a function $f = f(t, x)$ is said to be $C^{1,2}$ if it is C^1 in t (i.e. differentiable with respect to t and $\partial_t f$ is continuous), and C^2 in x (i.e. twice differentiable with respect to x and $\partial_x f, \partial_x^2 f$ are both continuous).

$$(5.4) \quad f(T, X(T)) - f(0, X(0)) = \int_0^T \partial_t f(t, X(t)) dt + \int_0^T \partial_x f(t, X(t)) dX(t) \\ + \frac{1}{2} \int_0^T \partial_x^2 f(t, X(t)) d[X, X](t).$$

REMARK 5.11. To clarify notation, $\partial_t f(t, X(t))$ means the following: differentiate $f(t, x)$ with respect to t (treating x as a constant), and then substitute $x = X(t)$. Similarly $\partial_x f(t, X(t))$ means differentiate $f(t, x)$ with respect to x , and then substitute $x = X(t)$. Finally $\partial_x^2 f(t, X(t))$ means take the second derivative of the function $f(t, x)$ with respect to x , and the substitute $x = X(t)$.

REMARK 5.12. In short hand differential form, this is written as

$$(5.4') \quad df(t, X(t)) = \partial_t f(t, X(t)) dt + \partial_x f(t, X(t)) dX(t) \\ + \frac{1}{2} \partial_x^2 f(t, X(t)) d[X, X](t).$$

The term $\frac{1}{2} \partial_x^2 f d[X, X](t)$ is an “extra” term, and is often referred to as the *Itô correction term*. The Itô formula is simply a version of the *chain rule* for stochastic processes.

REMARK 5.13. Substituting what we know about X from (5.1) and (5.3) we see that (5.4) becomes

$$f(T, X(T)) - f(0, X(0)) = \int_0^T (\partial_t f(t, X(t)) + \partial_x f(t, X(t)) b(t)) dt \\ + \int_0^T \partial_x f(t, X(t)) \sigma(t) dW(t) + \frac{1}{2} \int_0^T \partial_x^2 f(t, X(t)) \sigma(t)^2 dt.$$

The second integral on the right is an Itô integral (and hence a martingale). The other integrals are regular Riemann integrals which yield processes of finite variation.

While a complete rigorous proof of the Itô formula is technical, and beyond the scope of this course, we provide a quick heuristic argument that illustrates the main idea clearly.

INTUITION BEHIND THE ITÔ FORMULA. Suppose that the function f is only a function of x and doesn't depend on t , and X is a standard Brownian motion (i.e. $b = 0$ and $\sigma = 1$). In this case proving Itô's formula reduces to proving

$$f(W(T)) - f(W(0)) = \int_0^T f'(W(t)) dW(t) + \frac{1}{2} \int_0^T f''(W(t)) dt.$$

Let $P = \{0 = t_0 < t_1 < \dots < t_n = T\}$ be a partition of $[0, T]$. Taylor expanding f to second order gives

$$(5.5) \quad f(W(T)) - f(W(0)) = \sum_{i=0}^{n-1} f(W(t_{i+1})) - f(W(t_i)) \\ = \sum_{i=0}^{n-1} f'(W(t_i)) (\Delta_i W) + \frac{1}{2} \sum_{i=0}^{n-1} f''(W(t_i)) (\Delta_i W)^2 + \frac{1}{2} \sum_{i=0}^{n-1} o((\Delta_i W)^2),$$

where the last sum on the right is the remainder from the Taylor expansion.

Note the first sum on the right of (5.5) converges to the Itô integral

$$\int_0^T f'(W(t)) dW(t).$$

For the second sum on the right of (5.5), note

$$f''(W(t_i))(\Delta_i W)^2 = f''(W(t_i))(t_{i+1} - t_i) + f''(W(t_i))[(\Delta_i W)^2 - (t_{i+1} - t_i)].$$

After summing over i , first term on the right converges to the Riemann integral $\int_0^T f''(W(t)) dt$. The second term on the right is similar to what we had when computing the quadratic variation of W . The variance of $\xi_i \stackrel{\text{def}}{=} (\Delta_i W)^2 - (t_{i+1} - t_i)$ is of order $(t_{i+1} - t_i)^2$. Thus we expect that the second term above, when summed over i , converges to 0.

Finally each summand in the remainder term (the last term on the right of (5.5)) is smaller than $(\Delta_i W)^2$. (If, for instance, f is three times continuously differentiable in x , then each summand in the remainder term is of order $(\Delta_i W)^3$.) Consequently, when summed over i this should converge to 0 \square

6. A few examples using Itô's formula

Technically, as soon as you know Itô's formula you can "jump right in" and derive the Black-Scholes equation. However, because of the importance of Itô's formula, we work out a few simpler examples first.

EXAMPLE 6.1. Compute the quadratic variation of $W(t)^2$.

SOLUTION. Let $f(t, x) = x^2$. Then, by Itô's formula,

$$\begin{aligned} d(W(t)^2) &= df(t, W(t)) \\ &= \partial_t f(t, W(t)) dt + \partial_x f(t, W(t)) dW(t) + \frac{1}{2} \partial_x^2 f(t, W(t)) dt \\ &= 2W(t) dW(t) + dt. \end{aligned}$$

Or, in integral form,

$$W(T)^2 - W(0)^2 = W(T)^2 = 2 \int_0^T W(t) dW(t) + T.$$

Now the second term on the right has finite first variation, and won't affect our computations for quadratic variation. The first term is a martingale whose quadratic variation is $\int_0^T W(t)^2 dt$, and so

$$[W^2, W^2](T) = 4 \int_0^T W(t)^2 dt. \quad \square$$

REMARK 6.2. Note the above also tells you

$$2 \int_0^T W(t) dW(t) = W(T)^2 - T.$$

EXAMPLE 6.3. Let $M(t) = W(t)$ and $N(t) = W(t)^2 - t$. We know M and N are martingales. Is MN a martingale?

SOLUTION. Note $M(t)N(t) = W(t)^3 - tW(t)$. By Itô's formula,

$$d(MN) = -W(t) dt + (3W(t)^2 - t) dW(t) + 3W(t) dt.$$

Or in integral form

$$M(t)N(t) = \int_0^t 2W(s) ds + \int_0^t (3W(s)^2 - s) dW(s).$$

Now the second integral on the right is a martingale, but the first integral most certainly is not. So MN can not be a martingale. \square

REMARK 6.4. Note, above we changed the integration variable from t to s . This is perfectly legal – the variable with which you integrate with respect to is a dummy variable (just line regular Riemann integrals) and you can replace it what your favourite (unused!) symbol.

REMARK 6.5. It's worth pointing out that the Itô integral $\int_0^t \Delta(s) dW(s)$ is always a martingale (under the finiteness condition (4.5)). However, the Riemann integral $\int_0^t b(s) ds$ is only a martingale if $b = 0$ identically.

PROPOSITION 6.6. *If $f = f(t, x)$ is $C_b^{1,2}$ then the process*

$$M(t) \stackrel{\text{def}}{=} f(t, W(t)) - \int_0^t \left(\partial_t f(s, W(s)) + \frac{1}{2} \partial_x^2 f(s, W(s)) \right) ds$$

is a martingale.

REMARK 6.7. We'd seen this earlier, and the proof involved computing the conditional expectations directly and checking an algebraic identity involving the density of the normal distribution. With Itô's formula, the proof is “immediate”.

PROOF. By Itô's formula (in integral form)

$$\begin{aligned} & f(t, W(t)) - f(0, W(0)) \\ &= \int_0^t \partial_t f(s, W(s)) ds + \int_0^t \partial_x f(s, W(s)) dW(s) + \frac{1}{2} \int_0^t \partial_x^2 f(s, W(s)) ds \\ &= \int_0^t \left(\partial_t f(s, W(s)) + \frac{1}{2} \partial_x^2 f(s, W(s)) \right) ds + \int_0^t \partial_x f(s, W(s)) dW(s). \end{aligned}$$

Substituting this we see

$$M(t) = f(0, W(0)) + \int_0^t \partial_x f(s, W(s)) dW(s),$$

which is a martingale. \square

REMARK 6.8. Note we said $f \in C_b^{1,2}$ to “cover our bases”. Recall for Itô integrals to be martingales, we need the finiteness condition (4.5) to hold. This will certainly be the case if $\partial_x f$ is bounded, which is why we made this assumption. The result above is of course true under much more general assumptions.

EXAMPLE 6.9. Let $X(t) = t \sin(W(t))$. Is $X^2 - [X, X]$ a martingale?

SOLUTION. Let $f(t, x) = t \sin(x)$. Observe $X(t) = f(t, W(t))$, $\partial_t f = \sin x$, $\partial_x f = t \cos x$, and $\partial_x^2 f = -t \sin x$. Thus by Itô's formula,

$$\begin{aligned} dX(t) &= \partial_t f(t, W(t)) dt + \partial_x f(t, W(t)) dW(t) + \frac{1}{2} \partial_x^2 f(t, W(t)) d[W, W](t) \\ &= \sin(W(t)) dt + t \cos(W(t)) dW(t) - \frac{1}{2} t \sin(W(t)) dt, \end{aligned}$$

and so

$$d[X, X](t) = t^2 \cos^2(W(t)) dt.$$

Now let $g(x) = x^2$ and apply Itô's formula to compute $dg(X(t))$. This gives

$$dX(t)^2 = 2X(t) dX(t) + d[X, X](t)$$

and so

$$\begin{aligned} d(X(t)^2 - [X, X]) &= 2X(t) dX(t) \\ &= 2t \sin(t) \left(\sin(W(t)) - \frac{t \sin(W(t))}{2} \right) dt + 2t \sin(t) (t \cos(W(t))) dW(t). \end{aligned}$$

Since the dt term above isn't 0, $X(t)^2 - [X, X]$ can not be a martingale. \square

Recall we said earlier (Theorem 3.4) that for any martingale M , $M^2 - [M, M]$ is a martingale. In the above example X is not a martingale, and so there is no contradiction when we show that $X^2 - [X, X]$ is not a martingale. If M is a martingale, Itô's formula can be used to “prove”³ that $M^2 - [M, M]$ is a martingale.

PROPOSITION 6.10. *Let $M(t) = \int_0^t \sigma(s) dW(s)$. Then $M^2 - [M, M]$ is a martingale.*

PROOF. Let $N(t) = M(t)^2 - [M, M](t)$. Observe that by Itô's formula,

$$d(M(t)^2) = 2M(t) dM(t) + d[M, M](t).$$

Hence

$$dN = 2M(t) dM(t) + d[M, M](t) - d[M, M](t) = 2M(t) \sigma(t) dW(t).$$

Since there is no “ dt ” term and Itô integrals are martingales, N is a martingale. \square

7. Review Problems

PROBLEM 7.1. If $0 \leq r < s < t$, compute $\mathbf{E}(W(r)W(s)W(t))$.

PROBLEM 7.2. Define the processes X, Y, Z by

$$X(t) = \int_0^{W(t)} e^{-s^2} ds, \quad Y(t) = \exp\left(\int_0^t W(s) ds\right), \quad Z(t) = tX(t)^2$$

Decompose each of these processes as the sum of a martingale and a process of finite first variation. What is the quadratic variation of each of these processes?

³We used the fact that $M^2 - [M, M]$ is a martingale crucially in the construction of Itô integrals, and hence in proving Itô's formula. Thus proving $M^2 - [M, M]$ is a martingale using the Itô's formula is circular and not a valid proof. It is however instructive, and helps with building intuition, which is why it is presented here.

PROBLEM 7.3. Define the processes X, Y by

$$X(t) \stackrel{\text{def}}{=} \int_0^t W(s) ds, \quad Y(t) \stackrel{\text{def}}{=} \int_0^t W(s) dW(s).$$

Given $0 \leq s < t$, compute the conditional expectations $\mathbf{E}(X(t)|\mathcal{F}_s)$, and $\mathbf{E}(Y(t)|\mathcal{F}_s)$.

PROBLEM 7.4. Let $M(t) = \int_0^t W(s) dW(s)$. Find a function f such that

$$E(t) \stackrel{\text{def}}{=} \exp\left(M(t) - \int_0^t f(s, W(s)) ds\right)$$

is a martingale.

PROBLEM 7.5. Suppose $\sigma = \sigma(t)$ is a deterministic (i.e. non-random) process, and X is the Itô process defined by

$$X(t) = \int_0^t \sigma(u) dW(u).$$

- (a) Given $\lambda, s, t \in \mathbb{R}$ with $0 \leq s < t$ compute $\mathbf{E}(e^{\lambda(X(t)-X(s))} | \mathcal{F}_s)$.
- (b) If $r \leq s$ compute $\mathbf{E} \exp(\lambda X(r) + \mu(X(t) - X(s)))$.
- (c) What is the joint distribution of $(X(r), X(t) - X(s))$?
- (d) (*Lévy's criterion*) If $\sigma(u) = \pm 1$ for all u , then show that X is a standard Brownian motion.

PROBLEM 7.6. Define the process X, Y by

$$X = \int_0^t s dW(s), \quad Y = \int_0^t W(s) ds.$$

Find a formula for $\mathbf{E}X(t)^n$ and $\mathbf{E}Y(t)^n$ for any $n \in \mathbb{N}$.

PROBLEM 7.7. Let $M(t) = \int_0^t W(s) dW(s)$. For $s < t$, is $M(t) - M(s)$ independent of \mathcal{F}_s ? Justify.

PROBLEM 7.8. Determine whether the following identities are true or false, and justify your answer.

- (a) $e^{2t} \sin(2W(t)) = 2 \int_0^t e^{2s} \cos(2W(s)) dW(s)$.
- (b) $|W(t)| = \int_0^t \text{sign}(W(s)) dW(s)$. (Recall $\text{sign}(x) = 1$ if $x > 0$, $\text{sign}(x) = -1$ if $x < 0$ and $\text{sign}(x) = 0$ if $x = 0$.)

8. The Black Scholes Merton equation.

The price of an asset with a constant rate of return α is given by

$$dS(t) = \alpha S(t) dt.$$

To account for noisy fluctuations we model stock prices by adding the term $\sigma S(t) dW(t)$ to the above:

$$(8.1) \quad dS(t) = \alpha S(t) dt + \sigma S(t) dW(t).$$

The parameter α is called the *mean return rate* or the *percentage drift*, and the parameter σ is called the *volatility* or the *percentage volatility*.

DEFINITION 8.1. A stochastic process S satisfying (8.1) above is called a *geometric Brownian motion*.

The reason S is called a geometric Brownian motion is as follows. Set $Y = \ln S$ and observe

$$dY(t) = \frac{1}{S(t)} dS(t) - \frac{1}{2S(t)^2} d[S, S](t) = \left(\alpha - \frac{\sigma^2}{2} \right) dt + \sigma dW(t).$$

If $\alpha = \sigma^2/2$ then $Y = \ln S$ is simply a Brownian motion.

We remark, however, that our application of Itô's formula above is not completely justified. Indeed, the function $f(x) = \ln x$ is *not* differentiable at $x = 0$, and Itô's formula requires f to be at least C^2 . The reason the application of Itô's formula here is valid is because the process S *never* hits the point $x = 0$, and at all other points the function f is infinitely differentiable.

The above also gives us an explicit formula for S . Indeed,

$$\ln\left(\frac{S(t)}{S(0)}\right) = \left(\alpha - \frac{\sigma^2}{2}\right)t + \sigma W(t),$$

and so

$$(8.2) \quad S(t) = S(0) \exp\left(\left(\alpha - \frac{\sigma^2}{2}\right)t + \sigma W(t)\right).$$

Now consider a European call option with strike price K and maturity time T . This is a security that allows you the option (not obligation) to buy S at price K and time T . Clearly the price of this option at time T is $(S(T) - K)^+$. Our aim is to compute the *arbitrage free*⁴ price of such an option at time $t < T$.

Black and Scholes realised that the price of this option at time t should only depend on the asset price $S(t)$, and the current time t (or more precisely, the time to maturity $T - t$), and of course the model parameters α, σ . In particular, the option price does not depend on the price history of S .

THEOREM 8.2. *Suppose we have an arbitrage free financial market consisting of a money market account with constant return rate r , and a risky asset whose price is given by S . Consider a European call option with strike price K and maturity T .*

(1) *If $c = c(t, x)$ is a function such that at any time $t \leq T$, the arbitrage free price of this option is $c(t, S(t))$, then c satisfies*

$$(8.3) \quad \partial_t c + rx \partial_x c + \frac{\sigma^2 x^2}{2} \partial_x^2 c - rc = 0 \quad x > 0, t < T,$$

$$(8.4) \quad c(t, 0) = 0 \quad t \leq T,$$

$$(8.5) \quad c(T, x) = (x - K)^+ \quad x \geq 0.$$

(2) *Conversely, if c satisfies (8.3)–(8.5) then $c(t, S(t))$ is the arbitrage free price of this option at any time $t \leq T$.*

REMARK 8.3. Since α, σ and T are fixed, we suppress the explicit dependence of c on these quantities.

⁴In an arbitrage free market, we say p is the arbitrage free price of a non traded security if given the opportunity to trade the security at price p , the market is still arbitrage free. (Recall a financial market is said to be *arbitrage free* if there doesn't exist a self-financing portfolio X with $X(0) = 0$ such that at some $t > 0$ we have $X(t) \geq 0$ and $P(X(t) > 0) > 0$.)

REMARK 8.4. The above result assumes the following:

- (1) The market is *frictionless* (i.e. there are no transaction costs).
- (2) The asset is liquid and fractional quantities of it can be traded.
- (3) The borrowing and lending rate are both r .

REMARK 8.5. Even though the asset price $S(t)$ is random, the function c is a deterministic (non-random) function. The option price, however, is $c(t, S(t))$, which is certainly random.

REMARK 8.6. Equation (8.3)–(8.5) are the *Black Scholes Merton PDE*. This is a *partial differential equation*, which is a differential equation involving derivatives with respect to more than one variable. Equation (8.3) governs the evolution of c for $x \in (0, \infty)$ and $t < T$. Equation (8.5) specifies the terminal condition at $t = T$, and equation (8.4) specifies a boundary condition at $x = 0$.

To be completely correct, one also needs to add a boundary condition as $x \rightarrow \infty$ to the system (8.3)–(8.5). When x is very large, the call option is deep in the money, and will very likely end in the money. In this case the replicating portfolio should be long one share of the asset and short $e^{-r(T-t)}K$, the discounted strike price, in cash. This means that when x is very large, $c(x, t) \approx x - Ke^{-r(T-t)}$, and hence a boundary condition at $x = \infty$ can be obtained by supplementing (8.4) with

$$(8.4') \quad \lim_{x \rightarrow \infty} (c(t, x) - (x - Ke^{-r(T-t)})) = 0.$$

REMARK 8.7. The system (8.3)–(8.5) can be solved explicitly using standard calculus by substituting $y = \ln x$ and converting it into the *heat equation*, for which the solution is explicitly known. This gives the Black-Scholes-Merton formula

$$(8.6) \quad c(t, x) = xN(d_+(T-t, x)) - Ke^{-r(T-t)}N(d_-(T-t, x))$$

where

$$(8.7) \quad d_{\pm}(\tau, x) \stackrel{\text{def}}{=} \frac{1}{\sigma\sqrt{\tau}} \left(\ln\left(\frac{x}{K}\right) + \left(r \pm \frac{\sigma^2}{2}\right)\tau \right),$$

and

$$(8.8) \quad N(x) \stackrel{\text{def}}{=} \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-y^2/2} dy,$$

is the CDF of a standard normal variable.

Even if you're unfamiliar with the techniques involved in arriving at the solution above, you can certainly check that the function c given by (8.6)–(8.7) above satisfies (8.3)–(8.5). Indeed, this is a direct calculation that only involves patience and a careful application of the chain rule. We will, however, derive (8.6)–(8.7) later using risk neutral measures.

We will prove Theorem 8.2 by using a *replicating portfolio*. This is a portfolio (consisting of cash and the risky asset) that has exactly the same cash flow at maturity as the European call option that needs to be priced. Specifically, let $X(t)$ be the value of the replicating portfolio and $\Delta(t)$ be the number of shares of the asset held. The remaining $X(t) - S(t)\Delta(t)$ will be invested in the money market account with return rate r . (It is possible that $\Delta(t)S(t) > X(t)$, in which means we borrow money from the money market account to invest in stock.) For a replicating portfolio, the trading strategy Δ should be chosen in a manner that ensures that we have the same cash flow as the European call option. That is, we must have

$X(T) = (S(T) - K)^+ = c(T, S(T))$. Now the arbitrage free price is precisely the value of this portfolio.

REMARK 8.8. Through the course of the proof we will see that given the function c , the number of shares of S the replicating portfolio should hold is given by the *delta hedging rule*

$$(8.9) \quad \Delta(t) = \partial_x c(t, S(t)).$$

REMARK 8.9. Note that there is no α dependence in the system (8.3)–(8.5), and consequently the formula (8.6) does not depend on α . At first sight, this might appear surprising. (In fact, Black and Scholes had a hard time getting the original paper published because the community couldn't believe that the final formula is independent of α .) The fact that (8.6) is independent of α can be heuristically explained by the fact that the replicating portfolio also holds the same asset: thus a high mean return rate will help both an investor holding a call option and an investor holding the replicating portfolio. (Of course this isn't the entire story, as one has to actually write down the dependence and check that an investor holding the call option benefits exactly as much as an investor holding the replicating portfolio. This is done below.)

PROOF OF THEOREM 8.2 PART 1. If $c(t, S(t))$ is the arbitrage free price, then, by definition

$$(8.10) \quad c(t, S(t)) = X(t),$$

where $X(t)$ is the value of a replicating portfolio. Since our portfolio holds $\Delta(t)$ shares of S and $X(t) - \Delta(t)S(t)$ in a money market account, the evolution of the value of this portfolio is given by

$$\begin{aligned} dX(t) &= \Delta(t) dS(t) + r(X(t) - \Delta(t)S(t)) dt \\ &= (rX(t) + (\alpha - r)\Delta(t)S(t)) dt + \sigma\Delta(t)S(t) dW(t). \end{aligned}$$

Also, by Itô's formula we compute

$$\begin{aligned} dc(t, S(t)) &= \partial_t c(t, S(t)) dt + \partial_x c(t, S(t)) dS(t) + \frac{1}{2} \partial_x^2 c(t, S(t)) d[S, S](t), \\ &= \left(\partial_t c + \alpha S \partial_x c + \frac{1}{2} \sigma^2 S^2 \partial_x^2 c \right) dt + \partial_x c \sigma S dW(t) \end{aligned}$$

where we suppressed the $(t, S(t))$ argument in the last line above for convenience.

Equating $dc(t, S(t)) = dX(t)$ gives

$$\begin{aligned} (rX(t) + (\alpha - r)\Delta(t)S(t)) dt + \sigma\Delta(t)S(t) dW(t) \\ = \left(\partial_t c + \alpha S \partial_x c + \frac{1}{2} \sigma^2 S^2 \partial_x^2 c \right) dt + \partial_x c \sigma S dW(t). \end{aligned}$$

Using uniqueness of the semi-martingale decomposition (Proposition 5.5) we can equate the dW and the dt terms respectively. Equating the dW terms gives the delta hedging rule (8.9). Writing $S(t) = x$ for convenience, equating the dt terms and using (8.10) gives (8.3). Since the payout of the option is $(S(T) - K)^+$ at maturity, equation (8.5) is clearly satisfied.

Finally if $S(t_0) = 0$ at one particular time, then we must have $S(t) = 0$ at all times, otherwise we would have an arbitrage opportunity. (This can be checked directly from the formula (8.2) of course.) Consequently the arbitrage free price of

the option when $S = 0$ is 0, giving the *boundary condition* (8.4). Hence (8.3)–(8.5) are all satisfied, finishing the proof. \square

PROOF OF THEOREM 8.2 PART 2. For the converse, we suppose c satisfies the system (8.3)–(8.5). Choose $\Delta(t)$ by the delta hedging rule (8.9), and let X be a portfolio with initial value $X(0) = c(0, S(0))$ that holds $\Delta(t)$ shares of the asset at time t and the remaining $X(t) - \Delta(t)S(t)$ in cash. We claim that X is a replicating portfolio (i.e. $X(T) = (S(T) - K)^+$ almost surely) and $X(t) = c(t, S(t))$ for all $t \leq T$. Once this is established $c(t, S(t))$ is the arbitrage free price as desired.

To show X is a replicating portfolio, first claim that $X(t) = c(t, S(t))$ for all $t < T$. To see this, let $Y(t) = e^{-rt}X(t)$ be the discounted value of X . (That is, $Y(t)$ is the value of $X(t)$ converted to cash at time $t = 0$.) By Itô's formula, we compute

$$\begin{aligned} dY(t) &= -rY(t) dt + e^{-rt} dX(t) \\ &= e^{-rt}(\alpha - r)\Delta(t)S(t) dt + e^{-rt}\sigma\Delta(t)S(t) dW(t). \end{aligned}$$

Similarly, using Itô's formula, we compute

$$d(e^{-rt}c(t, S(t))) = e^{-rt}\left(-rc + \partial_t c + \alpha S \partial_x c + \frac{1}{2}\sigma^2 S^2 \partial_x^2 c\right) dt + e^{-rt}\partial_x c \sigma S dW(t).$$

Using (8.3) this gives

$$d(e^{-rt}c(t, S(t))) = e^{-rt}(\alpha - r)S \partial_x c dt + e^{-rt}\partial_x c \sigma S dW(t) = dY(t),$$

since $\Delta(t) = \partial_x c(t, S(t))$ by choice. This forces

$$\begin{aligned} e^{-rt}X(t) &= X(0) + \int_0^t dY(s) = X(0) + \int_0^t d(e^{-rs}c(s, S(s))) \\ &= X(0) + e^{-rt}c(t, S(t)) - c(0, S(0)) = e^{-rt}c(t, S(t)), \end{aligned}$$

since we chose $X(0) = c(0, S(0))$. This forces $X(t) = c(t, S(t))$ for all $t < T$, and by continuity also for $t = T$. Since $c(T, S(T)) = (S(T) - K)^+$ we have $X(T) = (S(T) - K)^+$ showing X is a replicating portfolio, concluding the proof. \square

REMARK 8.10. In order for the application of Itô's formula to be valid above, we need $c \in C^{1,2}$. This is certainly false at time T , since $c(T, x) = (x - K)^+$ which is not even differentiable, let alone twice continuously differentiable. However, if c satisfies the system (8.3)–(8.5), then it turns out that for every $t < T$ the function c will be infinitely differentiable with respect to x . This is why our proof first shows that $c(t, S(t)) = X(t)$ for $t < T$ and not directly that $c(t, S(t)) = X(t)$ for all $t \leq T$.

REMARK 8.11 (Put Call Parity). The same argument can be used to compute the arbitrage free price of European put options (i.e. the option to sell at the strike price, instead of buying). However, once the price of the price of a call option is computed, the *put call parity* can be used to compute the price of a put.

Explicitly let $p = p(t, x)$ be a function such that at any time $t \leq T$, $p(t, S(t))$ is the arbitrage free price of a European put option with strike price K . Consider a portfolio X that is long a call and short a put (i.e. buy one call, and sell one put). The value of this portfolio at time $t < T$ is

$$X(t) = c(t, S(t)) - p(t, S(t))$$

and at maturity we have⁵

$$X(T) = (S(T) - K)^+ - (K - S(T))^+ = S(T) - K.$$

This payoff can be replicated using a portfolio that holds one share of the asset and borrows Ke^{-rT} in cash (with return rate r) at time 0. Thus, in an arbitrage free market, we should have

$$c(t, S(t)) - p(t, S(t)) = X(t) = S(t) - Ke^{-r(T-t)}.$$

Writing x for $S(t)$ this gives the *put call parity* relation

$$c(t, x) - p(t, x) = x - Ke^{-r(T-t)}.$$

Using this the price of a put can be computed from the price of a call.

We now turn to understanding properties of c . The partial derivatives of c with respect to t and x measure the sensitivity of the option price to changes in the time to maturity and spot price of the asset respectively. These are called “the Greeks”:

(1) The *delta* is defined to be $\partial_x c$, and is given by

$$\partial_x c = N(d_+) + xN'(d_+)d'_+ - Ke^{-r\tau}N'(d_-)d'_-.$$

where $\tau = T - t$ is the time to maturity. Recall $d_{\pm} = d_{\pm}(\tau, x)$, and we suppressed the (τ, x) argument above for notational convenience. Using the formulae (8.6)–(8.8) one can verify

$$d'_+ = d'_- = \frac{1}{x\sigma\sqrt{\tau}} \quad \text{and} \quad xN'(d_+) = Ke^{-r\tau}N'(d_-),$$

and hence the delta is given by

$$\partial_x c = N(d_+).$$

Recall that the *delta hedging rule* (equation (8.9)) explicitly tells you that the replicating portfolio should hold precisely $\partial_x c(t, S(t))$ shares of the risky asset and the remainder in cash.

(2) The *gamma* is defined to be $\partial_x^2 c$, and is given by

$$\partial_x^2 c = N'(d_+)d'_+ = \frac{1}{x\sigma\sqrt{2\pi\tau}} \exp\left(\frac{-d_+^2}{2}\right).$$

(3) Finally the *theta* is defined to be $\partial_t c$, and simplifies to

$$\partial_t c = -rKe^{-r\tau}N(d_-) - \frac{\sigma x}{2\sqrt{\tau}}N'(d_+)$$

PROPOSITION 8.12. *The function $c(t, x)$ is convex and increasing as a function of x , and is decreasing as a function of t .*

PROOF. This follows immediately from the fact that $\partial_x c > 0$, $\partial_x^2 c > 0$ and $\partial_t c < 0$. \square

⁵A *forward contract* requires the holder to buy the asset at price K at maturity. The value of this contract at maturity is exactly $S(T) - K$, and so a portfolio that is long a call and short a put has exactly the same cash flow as a forward contract.

REMARK 8.13 (Hedging a short call). Suppose you sell a call option valued at $c(t, x)$, and want to create a replicating portfolio. The delta hedging rule calls for $x\partial_x c(t, x)$ of the portfolio to be invested in the asset, and the rest in the money market account. Consequently the value of your money market account is

$$c(t, x) - x\partial_x c = xN(d_+) - Ke^{-r\tau}N(d_-) - xN(d_+) = -Ke^{-r\tau}N(d_-) < 0.$$

Thus to properly hedge a short call you will have to borrow from the money market account and invest it in the asset. As $t \rightarrow T$ you will end up selling shares of the asset if $x < K$, and buying shares of it if $x > K$, so that at maturity you will hold the asset if $x > K$ and not hold it if $x < K$. To hedge a long call you do the opposite.

REMARK 8.14 (Delta neutral and Long Gamma). Suppose at some time t the price of a stock is x_0 . We short $\partial_x c(t, x_0)$ shares of this stock buy the call option valued at $c(t, x_0)$. We invest the balance $M = x_0\partial_x c(t, x_0) - c(t, x_0)$ in the money market account. Now if the stock price changes to x , and we do not change our position, then the value of our portfolio will be

$$\begin{aligned} c(t, x) - \partial_x c(t, x_0)x + M &= c(t, x) - x\partial_x c(t, x_0) + x_0\partial_x c(t, x_0) - c(t, x_0) \\ &= c(t, x) - (c(t, x_0) + (x - x_0)\partial_x c(t, x_0)). \end{aligned}$$

Note that the line $y = c(t, x_0) + (x - x_0)\partial_x c(t, x_0)$ is the equation for the tangent to the curve $y = c(t, x)$ at the point $(x_0, c(t, x_0))$. For this reason the above portfolio is called *delta neutral*.

Note that any convex function lies entirely above its tangent. Thus, under instantaneous changes of the stock price (both rises and falls), we will have

$$c(t, x) - \partial_x c(t, x_0)x + M > 0, \quad \text{both for } x > x_0 \text{ and } x < x_0.$$

For this reason the above portfolio is called *long gamma*.

Note, even though under instantaneous price changes the value of our portfolio always rises, this is *not* an arbitrage opportunity. The reason for this is that as time increases c decreases since $\partial_t c < 0$. The above instantaneous argument assumed c is constant in time, which it most certainly is not!

9. Multi-dimensional Itô calculus.

Finally we conclude this chapter by studying Itô calculus in higher dimensions. Let X, Y be Itô process. We typically expect X, Y will have finite and non-zero quadratic variation, and hence both the increments $X(t + \delta t) - X(t)$ and $Y(t + \delta t) - Y(t)$ should typically be of size $\sqrt{\delta}$. If we multiply these and sum over some finite interval $[0, T]$, then we would have roughly $T/\delta t$ terms each of size δt , and expect that this converges as $\delta t \rightarrow 0$. The limit is called the joint quadratic variation.

DEFINITION 9.1. Let X and Y be two Itô processes. We define the *joint quadratic variation* of X, Y , denoted by $[X, Y]$ by

$$[X, Y](T) = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} (X(t_{i+1}) - X(t_i))(Y(t_{i+1}) - Y(t_i)),$$

where $P = \{0 = t_1 < t_1 \cdots < t_n = T\}$ is a partition of $[0, T]$.

Using the identity

$$4ab = (a + b)^2 - (a - b)^2$$

we quickly see that

$$(9.1) \quad [X, Y] = \frac{1}{4}([X + Y, X + Y] - [X - Y, X - Y]).$$

Using this and the properties we already know about quadratic variation, we can quickly deduce the following.

PROPOSITION 9.2 (Product rule). *If X and Y are two Itô processes then*

$$(9.2) \quad d(XY) = X dY + Y dX + d[X, Y].$$

PROOF. By Itô's formula

$$\begin{aligned} d(X + Y)^2 &= 2(X + Y) d(X + Y) + d[X + Y, X + Y] \\ &= 2X dX + 2Y dY + 2X dY + 2Y dX + d[X + Y, X + Y]. \end{aligned}$$

Similarly

$$d(X - Y)^2 = 2X dX + 2Y dY - 2X dY - 2Y dX + d[X - Y, X - Y].$$

Since

$$4d(XY) = d(X + Y)^2 - d(X - Y)^2,$$

we obtain (9.2) as desired. \square

As with quadratic variation, processes of finite variation do not affect the joint quadratic variation.

PROPOSITION 9.3. *If X is an Itô process, and B is a continuous adapted process with finite variation, then $[X, B] = 0$.*

PROOF. Note $[X \pm B, X \pm B] = [X, X]$ and hence $[X, B] = 0$. \square

With this, we can state the higher dimensional Itô formula. Like the one dimensional Itô formula, this is a generalization of the chain rule and has an extra correction term that involves the joint quadratic variation.

THEOREM 9.4 (Itô-Doeblin formula). *Let X_1, \dots, X_n be n Itô processes and set $X = (X_1, \dots, X_n)$. Let $f: [0, \infty) \times \mathbb{R}^n$ be C^1 in the first variable, and C^2 in the remaining variables. Then*

$$\begin{aligned} f(T, X(T)) &= f(0, X(0)) + \int_0^T \partial_t f(t, X(t)) dt + \sum_{i=1}^N \int_0^T \partial_i f(t, X(t)) dX_i(t) \\ &\quad + \frac{1}{2} \sum_{i,j=1}^N \int_0^T \partial_i \partial_j f(t, X(t)) d[X_i, X_j](t), \end{aligned}$$

REMARK 9.5. Here we think of $f = f(t, x_1, \dots, x_n)$, often abbreviated as $f(t, x)$. The $\partial_i f$ appearing in the Itô formula above is the partial derivative of f with respect to x_i . As before, the $\partial_t f$ and $\partial_i f$ terms above are from the usual chain rule, and the last term is the extra Itô correction.

REMARK 9.6. In differential form Itô's formula says

$$d(f(t, X(t))) = \partial_t f(t, X(t)) dt + \sum_{i=1}^n \partial_i f(t, X(t)) dX_i(t) + \frac{1}{2} \sum_{i,j=1}^n \partial_i \partial_j f(t, X(t)) d[X_i, X_j](t).$$

For compactness, we will often omit the $(t, X(t))$ and write the above as

$$d(f(t, X(t))) = \partial_t f dt + \sum_{i=1}^n \partial_i f dX_i(t) + \frac{1}{2} \sum_{i,j=1}^n \partial_i \partial_j f d[X_i, X_j](t).$$

REMARK 9.7. We will most often use this in two dimensions. In this case, writing X and Y for the two processes, the Itô formula reduces to

$$d(f(t, X(t), Y(t))) = \partial_t f dt + \partial_x f dX(t) + \partial_y f dY(t) + \frac{1}{2} (\partial_x^2 f d[X, X](t) + 2\partial_x \partial_y f d[X, Y](t) + \partial_y^2 f d[Y, Y](t)).$$

INTUITION BEHIND THE ITÔ FORMULA. Let's assume we only have two Itô processes X, Y and $f = f(x, y)$ doesn't depend on t . Let $P = \{0 = t_0 < t_1 \cdots < t_m = T\}$ be a partition of the interval $[0, T]$ and write

$$f(X(T), Y(T)) - f(X(0), Y(0)) = \sum_{i=0}^{m-1} f(\xi_{i+1}) - f(\xi_i),$$

where we write $\xi_i = (X(t_i), Y(t_i))$ for compactness. Now by Taylor's theorem,

$$f(\xi_{i+1}) - f(\xi_i) = \partial_x f(\xi_i) \Delta_i X + \partial_y f(\xi_i) \Delta_i Y + \frac{1}{2} \left(\partial_x^2 f(\xi_i) (\Delta_i X)^2 + 2\partial_x \partial_y f(\xi_i) \Delta_i X \Delta_i Y + \partial_y^2 f(\xi_i) (\Delta_i Y)^2 \right) + \text{higher order terms.}$$

Here $\Delta_i X = X(t_{i+1}) - X(t_i)$ and $\Delta_i Y = Y(t_{i+1}) - Y(t_i)$. Summing over i , the first two terms converge to $\int_0^T \partial_x f(t) dX(t)$ and $\int_0^T \partial_y f(t) dY(t)$ respectively. The terms involving $(\Delta_i X)^2$ should to $\int_0^T \partial_x^2 f d[X, X]$ as we had with the one dimensional Itô formula. Similarly, the terms involving $(\Delta_i Y)^2$ should to $\int_0^T \partial_y^2 f d[Y, Y]$ as we had with the one dimensional Itô formula. For the cross term, we can use the identity (9.1) and quickly check that it converges to $\int_0^T \partial_x \partial_y f d[X, Y]$. The higher order terms are typically of size $(t_{i+1} - t_i)^{3/2}$ and will vanish as $\|P\| \rightarrow 0$. \square

The most common use of the multi-dimensional Itô formula is when the Itô processes are specified as a combination of Itô integrals with respect to different Brownian motions. Thus our next goal is to find an effective way to compute the joint quadratic variations in this case.

We've seen earlier (Theorems 3.4–3.5) that the quadratic variation of a martingale M is the unique increasing process that make $M^2 - [M, M]$ a martingale. A similar result holds for the joint quadratic variation.

PROPOSITION 9.8. *Suppose M, N are two continuous martingales with respect to a common filtration $\{\mathcal{F}_t\}$ such that $\mathbf{E}M(t)^2, \mathbf{E}N(t)^2 < \infty$.*

- (1) The process $MN - [M, N]$ is also a martingale with respect to the same filtration.
- (2) Moreover, if A is any continuous adapted process with finite first variation such that $A(0) = 0$ and $MN - A$ is a martingale with respect to $\{\mathcal{F}_t\}$, then $A = [M, N]$.

PROOF. The first part follows immediately from Theorem 3.4 and the fact that $4(MN - [M, N]) = (M + N)^2 - [M + N, M + N] - ((M - N)^2 - [M - N, M - N])$. The second part follows from the first part and uniqueness of the semi-martingale decomposition (Proposition 5.5). \square

PROPOSITION 9.9 (Bi-linearity). *If X, Y, Z are three Itô processes and $\alpha \in \mathbb{R}$ is a (non-random) constant, then*

$$(9.3) \quad [X, Y + \alpha Z] = [X, Y] + \alpha[X, Z].$$

PROOF. Let L, M and N be the martingale part in the Itô decomposition of X, Y and Z respectively. Clearly

$$L(M + \alpha N) - ([L, M] + \alpha[L, N]) = (LM - [L, M]) + \alpha(LN - [L, N]),$$

which is a martingale. Thus, since $[L, M] + \alpha[L, N]$ is also continuous adapted and increasing, by Proposition 9.8 we must have $[L, M + \alpha N] = [L, M] + \alpha[L, N]$. Since the joint quadratic variation of Itô processes can be computed in terms of their martingale parts alone, we obtain (9.3) as desired. \square

For integrals with respect to Itô processes, we can compute the joint quadratic variation explicitly.

PROPOSITION 9.10. *Let X_1, X_2 be two Itô processes, σ_1, σ_2 be two adapted processes and let I_j be the integral defined by $I_j(t) = \int_0^t \sigma_j(s) dX_j(s)$ for $j \in \{1, 2\}$. Then*

$$[I_1, I_2](t) = \int_0^t \sigma_1(s)\sigma_2(s) d[X_1, X_2](s).$$

PROOF. Let P be a partition and, as above, let $\Delta_i X = X(t_{i+1}) - X(t_i)$ denote the increment of a process X . Since

$$I_j(T) = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \sigma_j(t_i) \Delta_i X_j, \quad \text{and} \quad [X_1, X_2] = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \Delta_i X_1 \Delta_i X_2,$$

we expect that $\sigma_j(t_i) \Delta_i(X_j)$ is a good approximation for $\Delta_i I_j$, and $\Delta_i X_1 \Delta_i X_2$ is a good approximation for $\Delta_i [X_1, X_2]$. Consequently, we expect

$$\begin{aligned} [I_i, I_j](T) &= \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \Delta_i I_1 \Delta_i I_2 = \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \sigma_1(t_i) \Delta_i X_1 \sigma_2(t_i) \Delta_i X_2 \\ &= \lim_{\|P\| \rightarrow 0} \sum_{i=0}^{n-1} \sigma_1(t_i) \sigma_2(t_i) \Delta_i [X_1, X_2] = \int_0^T \sigma_1(t) \sigma_2(t) d[X_1, X_2](t), \end{aligned}$$

as desired. \square

PROPOSITION 9.11. *Let M, N be two continuous martingales with respect to a common filtration $\{\mathcal{F}_t\}$ such that $\mathbf{E}M(t)^2 < \infty$ and $\mathbf{E}N(t)^2 < \infty$. If M, N are independent, then $[M, N] = 0$.*

REMARK 9.12. If X and Y are independent, we know $\mathbf{E}XY = \mathbf{E}X\mathbf{E}Y$. However, we *need not* have $\mathbf{E}(XY | \mathcal{F}) = \mathbf{E}(X | \mathcal{F})\mathbf{E}(Y | \mathcal{F})$. So we can not prove the above result by simply saying

$$(9.4) \quad \mathbf{E}(M(t)N(t) | \mathcal{F}_s) = \mathbf{E}(M(t) | \mathcal{F}_s)\mathbf{E}(N(t) | \mathcal{F}_s) = M(s)N(s)$$

because M and N are independent. Thus MN is a martingale, and hence $[M, N] = 0$ by Proposition 9.8.

This reasoning is incorrect, even though the conclusion is correct. If you're not convinced, let me add that there exist martingales that are *not continuous* which are independent and have nonzero joint quadratic variation. The above argument, if correct, would certainly also work for martingales that are not continuous. The error in the argument is that the first equality in (9.4) need not hold even though M and N are independent.

PROOF. Let $P = \{0 = t_0 < t_1 < \dots < t_n = T\}$ be a partition of $[0, T]$, $\Delta_i M = M(t_{i+1}) - M(t_i)$ and $\Delta_i N = N(t_{i+1}) - N(t_i)$. Observe

$$(9.5) \quad \mathbf{E}\left(\sum_{i=0}^{n-1} \Delta_i M \Delta_i N\right)^2 = \mathbf{E}\sum_{i=0}^{n-1} (\Delta_i M)^2 (\Delta_i N)^2 + 2\mathbf{E}\sum_{j=0}^{n-1} \sum_{i=0}^{j-1} \Delta_i M \Delta_i N \Delta_j M \Delta_j N.$$

We claim the cross term vanishes because of independence of M and N . Indeed,

$$\begin{aligned} \mathbf{E}\Delta_i M \Delta_i N \Delta_j M \Delta_j N &= \mathbf{E}(\Delta_i M \Delta_j M) \mathbf{E}(\Delta_i N \Delta_j N) \\ &= \mathbf{E}(\Delta_i M \mathbf{E}(\Delta_j M | \mathcal{F}_{t_j})) \mathbf{E}(\Delta_i N \Delta_j N) = 0. \end{aligned}$$

Thus from (9.5)

$$\mathbf{E}\left(\sum_{i=0}^{n-1} \Delta_i M \Delta_i N\right)^2 = \mathbf{E}\sum_{i=0}^{n-1} (\Delta_i M)^2 (\Delta_i N)^2 \leq \mathbf{E}\left(\max_i (\Delta_i M)^2\right) \left(\sum_{i=0}^{n-1} (\Delta_i N)^2\right)$$

As $\|P\| \rightarrow 0$, $\max_i \Delta_i M \rightarrow 0$ because M is continuous, and $\sum_i (\Delta_i N)^2 \rightarrow [N, N](T)$. Thus we expect⁶

$$\mathbf{E}[M, N](T)^2 = \lim_{\|P\| \rightarrow 0} \mathbf{E}\left(\sum_{i=0}^{n-1} \Delta_i M \Delta_i N\right)^2 = 0,$$

finishing the proof. □

REMARK 9.13. The converse is false. If $[M, N] = 0$, it does *not* mean that M and N are independent. For example, if

$$M(t) = \int_0^t \mathbf{1}_{\{W(s) < 0\}} dW(s), \quad \text{and} \quad N(t) = \int_0^t \mathbf{1}_{\{W(s) \geq 0\}} dW(s),$$

⁶For this step we need to use $\lim_{\|P\| \rightarrow 0} \mathbf{E}(\dots) = \mathbf{E} \lim_{\|P\| \rightarrow 0} (\dots)$. To make this rigorous we need to apply the Lebesgue dominated convergence theorem. This is done by first assuming M and N are bounded, and then choosing a *localizing sequence* of stopping times, and a full discussion goes beyond the scope of these notes.

then clearly $[M, N] = 0$. However,

$$M(t) + N(t) = \int_0^t 1 dW(s) = W(t),$$

and with a little work one can show that M and N are not independent.

DEFINITION 9.14. We say $W = (W_1, W_2, \dots, W_d)$ is a standard d -dimensional Brownian motion if:

- (1) Each coordinate W_i is a standard (1-dimensional) Brownian motion.
- (2) If $i \neq j$, the processes W_i and W_j are independent.

When working with a multi-dimensional Brownian motion, we usually choose the filtration to be that generated by *all* the coordinates.

DEFINITION 9.15. Let W be a d -dimensional Brownian motion. We define the filtration $\{\mathcal{F}_t^W\}$ by

$$\mathcal{F}_t^W = \sigma\left(\bigcup_{\substack{s \leq t, \\ i \in \{1, \dots, d\}}} \sigma(W_i(s))\right)$$

With $\{\mathcal{F}_t^W\}$ defined above note that:

- (1) Each coordinate W_i is a martingale with respect to $\{\mathcal{F}_t^W\}$.
- (2) For every $s < t$, the increment of each coordinate $W_i(t) - W_i(s)$ is independent of $\{\mathcal{F}_s^W\}$.

REMARK 9.16. Since W_i is independent of W_j when $i \neq j$, we know $[W_i, W_j] = 0$ if $i \neq j$. When $i = j$, we know $d[W_i, W_j] = dt$. We often express this concisely as

$$d[W_i, W_j](t) = \mathbf{1}_{\{i=j\}} dt.$$

An extremely important fact about Brownian motion is that the converse of the above is also true.

THEOREM 9.17 (Lévy). *If $M = (M_1, M_2, \dots, M_d)$ is a continuous martingale such that $M(0) = 0$ and*

$$d[M_i, M_j](t) = \mathbf{1}_{\{i=j\}} dt,$$

then M is a d -dimensional Brownian motion.

PROOF. The main idea behind the proof is to compute the moment generating function (or characteristic function) of M , in the same way as in Problem 7.5. This can be used to show that $M(t) - M(s)$ is independent of \mathcal{F}_s and $M(t) \sim N(0, tI)$, where I is the $d \times d$ identity matrix. \square

EXAMPLE 9.18. If W is a 2-dimensional Brownian motion, then show that

$$B = \int_0^t \frac{W_1(s)}{|W(t)|} dW_1(s) + \int_0^t \frac{W_2(s)}{|W(t)|} dW_2(s),$$

is also a Brownian motion.

PROOF. Since B is the sum of two Itô integrals, it is clearly a continuous martingale. Thus to show that B is a Brownian motion, it suffices to show that $[B, B](t) = t$. For this, define

$$X(t) = \int_0^t \frac{W_1(s)}{|W(t)|} dW_1(s) \quad \text{and} \quad Y(t) = \int_0^t \frac{W_2(s)}{|W(t)|} dW_2(s),$$

and note

$$\begin{aligned} d[B, B](t) &= d[X + Y, X + Y](t) = d[X, X](t) + d[Y, Y](t) + 2d[X, Y](t) \\ &= \left(\frac{W_1(t)^2}{|W(t)|^2} + \frac{W_2(t)^2}{|W(t)|^2} \right) dt + 0 = dt. \end{aligned}$$

So by Lévy's criterion, B is a Brownian motion. \square

EXAMPLE 9.19. Let W be a 2-dimensional Brownian motion and define

$$X = \ln(|W|^2) = \ln(W_1^2 + W_2^2).$$

Compute dX . Is X a martingale?

SOLUTION. This is a bit tricky. First, if we set $f(x) = \ln|x|^2 = \ln(x_1^2 + x_2^2)$, then it is easy to check

$$\partial_i f = \frac{2x_i}{|x|^2} \quad \text{and} \quad \partial_1^2 f + \partial_2^2 f = 0.$$

Consequently,

$$dX(t) = \frac{2W_1(t)}{|W|^2} dW_1(t) + \frac{2W_2(t)}{|W|^2} dW_2(t).$$

With this one would be tempted to say that since there are no dt terms above, X is a martingale. This, however, is false! Martingales have constant expectation, but

$$\begin{aligned} \mathbf{E}X(t) &= \frac{1}{2\pi t} \iint_{\mathbb{R}^2} \ln(x_1^2 + x_2^2) \exp\left(-\frac{x_1^2 + x_2^2}{2t}\right) dx_1 dx_2 \\ &= \frac{1}{2\pi} \iint_{\mathbb{R}^2} \ln(t(y_1^2 + y_2^2)) \exp\left(-\frac{y_1^2 + y_2^2}{2}\right) dy_1 dy_2 \\ &= \ln t + \frac{1}{2\pi} \iint_{\mathbb{R}^2} \ln(y_1^2 + y_2^2) \exp\left(-\frac{y_1^2 + y_2^2}{2}\right) dy_1 dy_2 \xrightarrow{t \rightarrow \infty} \infty. \end{aligned}$$

Thus $\mathbf{E}X(t)$ is not constant in t , and so X can not be a martingale. \square

REMARK 9.20. We have repeatedly used the fact that Itô integrals are martingales. The example above obtains X as an Itô integral, but can not be a martingale. The reason this doesn't contradict Theorem 4.2 is that in order for Itô integral $\int_0^t \sigma(s) dW(s)$ to be defined, we only need the finiteness condition $\int_0^t \sigma(s)^2 ds < \infty$ almost surely. However, for an Itô integral to be a martingale, we need the stronger condition $\mathbf{E} \int_0^t \sigma(s)^2 ds < \infty$ (given in (4.5)) to hold. This is precisely what fails in the previous example. The process X above is an example of a *local martingale* that is not a martingale, and we will encounter a similar situation when we study exponential martingales and risk neutral measures.

EXAMPLE 9.21. Let $f = f(t, x_1, \dots, x_d) \in C^{1,2}$ and W be a d -dimensional Brownian motion. Then Itô's formula gives

$$d(f(t, W(t))) = \left(\partial_t f(t, W(t)) + \frac{1}{2} \Delta f(t, W(t)) \right) dt + \sum_{i=1}^d \partial_i f(t, W(t)) dW_i(t).$$

Here $\Delta f = \sum_{i=1}^d \partial_i^2 f$ is the Laplacian of f .

EXAMPLE 9.22. Consider a d -dimensional Brownian motion W , and n Itô processes X_1, \dots, X_n which we write (in differential form) as

$$dX_i(t) = b_i(t) dt + \sum_{k=1}^d \sigma_{i,k}(t) dW_k(t),$$

where each b_i and $\sigma_{i,j}$ are adapted processes. For brevity, we will often write b for the vector process (b_1, \dots, b_n) , σ for the *matrix* process $(\sigma_{i,j})$ and X for the n -dimensional Itô process (X_1, \dots, X_n) .

Now to compute $[X_i, X_j]$ we observe that $d[W_i, W_j] = dt$ if $i = j$ and 0 otherwise. Consequently,

$$d[X_i, X_j](t) = \sum_{k,l=1}^d \sigma_{i,k} \sigma_{j,l} \mathbf{1}_{\{k=l\}} dt = \sum_{k=1}^d \sigma_{i,k}(t) \sigma_{j,k}(t) dt.$$

Thus if f is any $C^{1,2}$ function, Itô formula gives

$$d(f(t, X(t))) = \left(\partial_t f + \sum_{i=1}^n b_i \partial_i f \right) dt + \sum_{i=1}^n \sigma_i \partial_i f dW_i(t) + \frac{1}{2} \sum_{i,j=1}^n a_{i,j} \partial_i \partial_j f dt$$

where

$$a_{i,j}(t) = \sum_{k=1}^N \sigma_{i,k}(t) \sigma_{j,k}(t).$$

In matrix notation, the matrix $a = \sigma \sigma^T$, where σ^T is the transpose of the matrix σ .

Risk Neutral Measures

Our aim in this section is to show how risk neutral measures can be used to price derivative securities. The key advantage is that under a risk neutral measure the discounted hedging portfolio becomes a martingale. Thus the price of any derivative security can be computed by conditioning the payoff at maturity. We will use this to provide an elegant derivation of the Black-Scholes formula, and discuss the fundamental theorems of asset pricing.

1. The Girsanov Theorem.

DEFINITION 1.1. Two probability measures \mathbf{P} and $\tilde{\mathbf{P}}$ are said to be equivalent if for every event A , $\mathbf{P}(A) = 0$ if and only if $\tilde{\mathbf{P}}(A) = 0$.

EXAMPLE 1.2. Let Z be a random variable such that $\mathbf{E}Z = 1$ and $Z > 0$. Define a new measure $\tilde{\mathbf{P}}$ by

$$(1.1) \quad \tilde{\mathbf{P}}(A) = \mathbf{E}Z1_A = \int_A Z d\mathbf{P}.$$

for every event A . Then \mathbf{P} and $\tilde{\mathbf{P}}$ are equivalent.

REMARK 1.3. The assumption $\mathbf{E}Z = 1$ above is required to guarantee $\tilde{\mathbf{P}}(\Omega) = 1$.

DEFINITION 1.4. When $\tilde{\mathbf{P}}$ is defined by (1.1), we say

$$d\tilde{\mathbf{P}} = Z d\mathbf{P} \quad \text{or} \quad Z = \frac{d\tilde{\mathbf{P}}}{d\mathbf{P}},$$

and Z is called the *density* of $\tilde{\mathbf{P}}$ with respect to \mathbf{P} .

THEOREM 1.5 (Radon-Nikodym). *Two measures \mathbf{P} and $\tilde{\mathbf{P}}$ are equivalent if and only if there exists a random variable Z such that $\mathbf{E}Z = 1$, $Z > 0$ and $\tilde{\mathbf{P}}$ is given by (1.1).*

The proof of this requires a fair amount of machinery from Lebesgue integration and goes beyond the scope of these notes. (This is exactly the result that is used to show that conditional expectations exist.) However, when it comes to risk neutral measures, it isn't essential since in most of our applications the density will be explicitly chosen.

Suppose now $T > 0$ is fixed, and Z is a martingale. Define a new measure $\tilde{\mathbf{P}} = \tilde{\mathbf{P}}_T$ by

$$d\tilde{\mathbf{P}} = d\tilde{\mathbf{P}}_T = Z(T)d\mathbf{P}.$$

We will denote expectations and conditional expectations with respect to the new measure by $\tilde{\mathbf{E}}$. That is, given a random variable X ,

$$\tilde{\mathbf{E}}X = \mathbf{E}Z(T)X = \int Z(T)X d\mathbf{P}.$$

Also, given a σ -algebra \mathcal{F} , $\tilde{\mathbf{E}}(X | \mathcal{F})$ is the unique \mathcal{F} -measurable random variable such that

$$(1.2) \quad \int_F \tilde{\mathbf{E}}(X | \mathcal{F}) d\tilde{\mathbf{P}} = \int_F X d\tilde{\mathbf{P}},$$

holds for all \mathcal{F} measurable events F . Of course, equation (1.2) is equivalent to requiring

$$(1.2') \quad \int_F Z(T) \tilde{\mathbf{E}}(X | \mathcal{F}) d\mathbf{P} = \int_F Z(T) X d\mathbf{P},$$

for all \mathcal{F} measurable events F .

The main goal of this section is to prove the Girsanov theorem.

THEOREM 1.6 (Cameron, Martin, Girsanov). *Let $b(t) = (b_1(t), b_2(t), \dots, b_d(t))$ be a d -dimensional adapted process, W be a d -dimensional Brownian motion, and define*

$$\tilde{W}(t) = W(t) + \int_0^t b(s) ds.$$

Let Z be the process defined by

$$Z(t) = \exp\left(-\int_0^t b(s) \cdot dW(s) - \frac{1}{2} \int_0^t |b(s)|^2 ds\right),$$

and define a new measure $\tilde{\mathbf{P}} = \tilde{\mathbf{P}}_T$ by $d\tilde{\mathbf{P}} = Z(T) d\mathbf{P}$. If Z is a martingale then \tilde{W} is a Brownian motion under the measure $\tilde{\mathbf{P}}$ up to time T .

REMARK 1.7. Above

$$b(s) \cdot dW(s) \stackrel{\text{def}}{=} \sum_{i=1}^d b_i(s) dW_i(s) \quad \text{and} \quad |b(s)|^2 = \sum_{i=1}^d b_i(s)^2.$$

REMARK 1.8. Note $Z(0) = 1$, and if Z is a martingale then $\mathbf{E}Z(T) = 1$ ensuring $\tilde{\mathbf{P}}$ is a probability measure. You might, however, be puzzled at need for the assumption that Z is a martingale. Indeed, let $M(t) = \int_0^t b(s) \cdot dW(s)$, and $f(t, x) = \exp(-x - \frac{1}{2} \int_0^t b(s)^2 ds)$. Then, by Itô's formula,

$$\begin{aligned} dZ(t) &= d(f(t, M(t))) = \partial_t f dt + \partial_x f dM(t) + \frac{1}{2} \partial_x^2 f d[M, M](t) \\ &= -\frac{1}{2} Z(t) |b(t)|^2 dt - Z(t) b(t) \cdot dW(t) + \frac{1}{2} Z(t) |b(t)|^2 dt, \end{aligned}$$

and hence

$$(1.3) \quad dZ(t) = -Z(t) b(t) \cdot dW(t).$$

Thus you might be tempted to say that Z is always a martingale, assuming it explicitly is unnecessary. However, we recall from Chapter 3, Theorem 4.2 in that Itô integrals are only guaranteed to be martingales under the square integrability condition

$$(1.4) \quad \mathbf{E} \int_0^T |Z(s) b(s)|^2 ds < \infty.$$

Without this finiteness condition, Itô integrals are only *local martingales*, whose expectation need not be constant, and so $\mathbf{E}Z(T) = 1$ is not guaranteed. Indeed,

there are many examples of processes b where the finiteness condition (1.4) does not hold and we have $\mathbf{E}Z(T) < 1$ for some $T > 0$.

REMARK 1.9. In general the process Z above is always a super-martingale, and hence $\mathbf{E}Z(T) \leq 1$. Two conditions that guarantee Z is a martingale are the Novikov and Kazamaki conditions: If either

$$\mathbf{E} \exp\left(\frac{1}{2} \int_0^t |b(s)|^2 ds\right) < \infty \quad \text{or} \quad \mathbf{E} \exp\left(\frac{1}{2} \int_0^t b(s) \cdot dW(s)\right) < \infty,$$

then Z is a martingale and hence $\mathbf{E}Z(T) = 1$ for all $T > 0$. Unfortunately, in many practical situations these conditions do not apply, and you have to show Z is a martingale by hand.

REMARK 1.10. The components b_1, \dots, b_d of the process b are not required to be independent. Yet, under the new measure, the process \tilde{W} is a Brownian motion and hence has *independent* components.

The main idea behind the proof of the Girsanov theorem is the following: Clearly $[\tilde{W}_i, \tilde{W}_j] = [W_i, W_j] = \mathbf{1}_{\{i=j\}}t$. Thus if we can show that \tilde{W} is a martingale *with respect to the new measure $\tilde{\mathbf{P}}$* , then Lévy's criterion will guarantee \tilde{W} is a Brownian motion. We now develop the tools required to check when processes are martingales under the new measure.

LEMMA 1.11. *Let $0 \leq s \leq t \leq T$. If X is a \mathcal{F}_t -measurable random variable then*

$$(1.5) \quad \tilde{\mathbf{E}}(X | \mathcal{F}_s) = \frac{1}{Z(s)} \mathbf{E}(Z(t)X | \mathcal{F}_s)$$

PROOF. Let $A \in \mathcal{F}_s$ and observe that

$$\begin{aligned} \int_A \tilde{\mathbf{E}}(X | \mathcal{F}_s) d\tilde{\mathbf{P}} &= \int_A Z(T) \tilde{\mathbf{E}}(X | \mathcal{F}_s) d\tilde{\mathbf{P}} \\ &= \int_A \mathbf{E}(Z(T) \tilde{\mathbf{E}}(X | \mathcal{F}_s) | \mathcal{F}_s) d\mathbf{P} = \int_A Z(s) \tilde{\mathbf{E}}(X | \mathcal{F}_s) d\mathbf{P}. \end{aligned}$$

Also,

$$\begin{aligned} \int_A \tilde{\mathbf{E}}(X | \mathcal{F}_s) d\tilde{\mathbf{P}} &= \int_A X d\tilde{\mathbf{P}} = \int_A X Z(T) d\mathbf{P} = \int_A \mathbf{E}(X Z(T) | \mathcal{F}_t) d\mathbf{P} \\ &= \int_A Z(t) X d\mathbf{P} = \int_A \mathbf{E}(Z(t) X | \mathcal{F}_s) d\mathbf{P} \end{aligned}$$

Thus

$$\int_A Z(s) \tilde{\mathbf{E}}(X | \mathcal{F}_s) d\tilde{\mathbf{P}} = \int_A \mathbf{E}(Z(t) X | \mathcal{F}_s) d\mathbf{P},$$

for every \mathcal{F}_s measurable event A . Since the integrands are both \mathcal{F}_s measurable this forces them to be equal, giving (1.5) as desired. \square

LEMMA 1.12. *An adapted process M is a martingale under $\tilde{\mathbf{P}}$ if and only if MZ is a martingale under \mathbf{P} .*

PROOF. Suppose first MZ is a martingale with respect to \mathbf{P} . Then

$$\tilde{\mathbf{E}}(M(t) | \mathcal{F}_s) = \frac{1}{Z(s)} \mathbf{E}(Z(t)M(t) | \mathcal{F}_s) = \frac{1}{Z(s)} Z(s)M(s) = M(s),$$

showing M is a martingale with respect to $\tilde{\mathbf{P}}$.

Conversely, suppose M is a martingale with respect to $\tilde{\mathbf{P}}$. Then

$$\mathbf{E}(M(t)Z(t) \mid \mathcal{F}_s) = Z(s)\tilde{\mathbf{E}}(M(t) \mid \mathcal{F}_s) = Z(s)M(s),$$

and hence ZM is a martingale with respect to \mathbf{P} . \square

PROOF OF THEOREM 1.6. Clearly \tilde{W} is continuous and

$$d[\tilde{W}_i, \tilde{W}_j](t) = d[W_i, W_j](t) = \mathbf{1}_{\{i=j\}} dt.$$

Thus if we show that each \tilde{W}_i is a martingale (under $\tilde{\mathbf{P}}$), then by Lévy's criterion, \tilde{W} will be a Brownian motion under $\tilde{\mathbf{P}}$.

We now show that each \tilde{W}_i is a martingale under $\tilde{\mathbf{P}}$. By Lemma 1.12, \tilde{W}_i is a martingale under $\tilde{\mathbf{P}}$ if and only if $Z\tilde{W}_i$ is a martingale under \mathbf{P} . To show $Z\tilde{W}_i$ is a martingale under \mathbf{P} , we use the product rule and (1.3) to compute

$$\begin{aligned} d(Z\tilde{W}_i) &= Z d\tilde{W}_i + \tilde{W}_i dZ + d[Z, \tilde{W}_i] \\ &= Z dW_i + Zb_i dt - \tilde{W}_i Zb \cdot dW - b_i Z dt = Z dW_i - \tilde{W}_i Zb \cdot dW. \end{aligned}$$

Thus $Z\tilde{W}_i$ is a martingale¹ under \mathbf{P} , and by Lemma 1.12, \tilde{W}_i is a martingale under $\tilde{\mathbf{P}}$. This finishes the proof. \square

2. Risk Neutral Pricing

Consider a stock whose price is modelled by a generalized geometric Brownian motion

$$dS(t) = \alpha(t)S(t) dt + \sigma(t)S(t) dW(t),$$

where $\alpha(t)$, $\sigma(t)$ are the (time dependent) mean return rate and volatility respectively. Here α and σ are no longer constant, but allowed to be *adapted processes*. We will, however, assume $\sigma(t) > 0$.

Suppose an investor has access to a money market account with variable interest rate $R(t)$. Again, the interest rate R need not be constant, and is allowed to be any adapted process. Define the *discount process* D by

$$D(t) = \exp\left(-\int_0^t R(s) ds\right),$$

and observe

$$dD(t) = -D(t)R(t) dt.$$

Since the price of one share of the money market account at time t is $1/D(t)$ times the price of one share at time 0, it is natural to consider the *discounted stock price* DS .

DEFINITION 2.1. A *risk neutral measure* is a measure $\tilde{\mathbf{P}}$ that is equivalent to \mathbf{P} under which the discounted stock price process $D(t)S(t)$ is a martingale.

REMARK 2.2. It turns out that the *existence* of a risk neutral measure is equivalent to there being no arbitrage opportunity in the market. Moreover, *uniqueness* of a risk neutral measure is equivalent to both the absence of an arbitrage opportunity, and that every derivative security can be hedged. These are the *fundamental theorems of asset pricing*.

¹Technically, we have to check the square integrability condition to ensure that $Z\tilde{W}_i$ is a martingale, and not a *local martingale*. This, however, follows quickly from the Cauchy-Schwartz inequality and our assumption.

Using the Girsanov theorem, we can compute the risk neutral measure explicitly. Observe

$$\begin{aligned} d(D(t)S(t)) &= -RDS dt + D dS = (\alpha - R)DS dt + DS\sigma dW(t) \\ &= \sigma(t)D(t)S(t)\left(\theta(t) dt + dW(t)\right) \end{aligned}$$

where

$$\theta(t) \stackrel{\text{def}}{=} \frac{\alpha(t) - R(t)}{\sigma(t)}$$

is known as the *market price of risk*.

Define a new process \tilde{W} by

$$d\tilde{W}(t) = \theta(t) dt + dW(t),$$

and observe

$$(2.1) \quad d(D(t)S(t)) dt = \sigma(t)D(t)S(t) d\tilde{W}(t).$$

PROPOSITION 2.3. *If Z is the process defined by*

$$Z(t) = \exp\left(-\int_0^t \theta(s) dW(s) - \frac{1}{2} \int_0^t \theta(s)^2 ds\right),$$

then the measure $\tilde{\mathbf{P}} = \tilde{\mathbf{P}}_T$ defined by $d\tilde{\mathbf{P}} = Z(T) d\mathbf{P}$ is a risk neutral measure.

PROOF. By the Girsanov theorem 1.6 we know \tilde{W} is a Brownian motion under $\tilde{\mathbf{P}}$. Thus using (2.1) we immediately see that the discounted stock price is a martingale. \square

Our next aim is to develop *risk neutral pricing formula*.

THEOREM 2.4 (Risk Neutral Pricing formula). *Let $V(T)$ be a \mathcal{F}_T -measurable random variable that represents the payoff of a derivative security, and let $\tilde{\mathbf{P}} = \tilde{\mathbf{P}}_T$ be the risk neutral measure above. The arbitrage free price at time t of a derivative security with payoff $V(T)$ and maturity T is given by*

$$(2.2) \quad V(t) = \tilde{\mathbf{E}}\left(\exp\left(-\int_t^T R(s) ds\right)V(T) \mid \mathcal{F}_t\right).$$

REMARK 2.5. It is important to note that the price $V(t)$ above is the actual arbitrage free price of the security, and there is no alternate “risk neutral world” which you need to teleport to in order to apply this formula. The risk neutral measure is simply a tool that is used in the above formula, which gives the arbitrage free price under the *standard* measure.

As we will see shortly, the reason for this formula is that under the risk neutral measure, the discounted replicating portfolio becomes a martingale. To understand why this happens we note

$$(2.3) \quad dS(t) = \alpha(t)S(t) dt + \sigma(t)S(t) dW(t) = R(t)S(t) dt + \sigma(t)S(t) d\tilde{W}.$$

Under the standard measure \mathbf{P} this isn’t much use, since \tilde{W} isn’t a martingale. However, under the risk neutral measure, the process \tilde{W} is a Brownian motion and hence certainly a martingale. Moreover, S becomes a geometric Brownian motion under $\tilde{\mathbf{P}}$ with mean return rate of S exactly the same as that of the money market account. The fact that S and the money market account have exactly the same

mean return rate (under $\tilde{\mathbf{P}}$) is precisely what makes the replicating portfolio (or any self-financing portfolio for that matter) a martingale (under $\tilde{\mathbf{P}}$).

LEMMA 2.6. *Let Δ be any adapted process, and $X(t)$ be the wealth of an investor with that holds $\Delta(t)$ shares of the stock and the rest of his wealth in the money market account. If there is no external cash flow (i.e. the portfolio is self financing), then the discounted portfolio $D(t)X(t)$ is a martingale under $\tilde{\mathbf{P}}$.*

PROOF. We know

$$dX(t) = \Delta(t) dS(t) + R(t)(X(t) - \Delta(t)S(t)) dt.$$

Using (2.3) this becomes

$$\begin{aligned} dX(t) &= \Delta RS dt + \Delta \sigma S d\tilde{W} + RX dt - R\Delta S dt \\ &= RX dt + \Delta \sigma S d\tilde{W}. \end{aligned}$$

Thus, by the product rule,

$$\begin{aligned} d(DX) &= D dX + X dD + d[D, X] = -RDX dt + DRX dt + D\Delta \sigma S d\tilde{W} \\ &= D\Delta \sigma S d\tilde{W}. \end{aligned}$$

Since \tilde{W} is a martingale under $\tilde{\mathbf{P}}$, DX must be a martingale under $\tilde{\mathbf{P}}$. \square

PROOF OF THEOREM 2.4. Suppose $X(t)$ is the wealth of a replicating portfolio at time t . Then by definition we know $V(t) = X(t)$, and by the previous lemma we know DX is a martingale under $\tilde{\mathbf{P}}$. Thus

$$V(t) = X(t) = \frac{1}{D(t)} D(t)X(t) = \frac{1}{D(t)} \tilde{\mathbf{E}}(D(T)X(T) \mid \mathcal{F}_t) = \tilde{\mathbf{E}}\left(\frac{D(T)V(T)}{D(t)} \mid \mathcal{F}_t\right),$$

which is precisely (2.2). \square

REMARK 2.7. Our proof assumes that a security with payoff $V(T)$ has a replicating portfolio. This is true in general because of the *martingale representation theorem*, which guarantees any martingale (with respect to the Brownian filtration) can be expressed as an Itô integral with respect to Brownian motion. Recall, we already know that Itô integrals are martingales. The martingale representation theorem is a partial converse.

Now clearly the process Y defined by

$$Y(t) = \tilde{\mathbf{E}}\left(\exp\left(-\int_0^T R(s) ds\right) V(T) \mid \mathcal{F}_t\right),$$

is a martingale. Thus, the *martingale representation theorem* can be used to express this as an Itô integral (with respect to \tilde{W}). With a little algebraic manipulation one can show that $D(t)^{-1}Y(t)$ is the wealth of a self financing portfolio. Since the terminal wealth is clearly $V(T)$, this must be a replicating portfolio.

REMARK 2.8. If $V(T) = f(S(T))$ for some function f and R is not random, then the *Markov Property* guarantees $V(t) = c(t, S(t))$ for some function c . Equating $c(t, S(t)) = X(t)$, the wealth of a replicating portfolio and using Itô's formula, we immediately obtain the *Delta hedging rule*

$$(2.4) \quad \Delta(t) = \partial_x c(t, S(t)).$$

If, however, that if V is not of the form $f(S(T))$ for some function f , then the option price will in general depend on the *entire history* of the stock price, and not only the spot price $S(t)$. In this case we will not (in general) have the delta hedging rule (2.4).

3. The Black-Scholes formula

Recall our first derivation of the Black-Scholes formula only obtained a PDE. The Black-Scholes formula is the solution to this PDE, which we simply wrote down without motivation. The risk neutral pricing formula can be used to derive the Black-Scholes formula quickly, and independently of our previous derivation. We carry out this calculation in this section.

Suppose σ and R are deterministic constants, and for notational consistency, set $r = R$. The risk neutral pricing formula says that the price of a European call is

$$c(t, S(t)) = \tilde{\mathbf{E}}\left(e^{-r(T-t)}(S(T) - K)^+ \mid \mathcal{F}_t\right),$$

where K is the strike price. Since

$$S(t) = S(0) \exp\left(\left(r - \frac{\sigma^2}{2}\right)t + \sigma \tilde{W}(t)\right),$$

and \tilde{W} is a Brownian motion under $\tilde{\mathbf{P}}$, we see

$$\begin{aligned} c(t, S(t)) &= e^{-r\tau} \tilde{\mathbf{E}}\left(\left[S(0) \exp\left(\left(r - \frac{\sigma^2}{2}\right)T + \sigma \tilde{W}(T)\right) - K\right]^+ \mid \mathcal{F}_t\right) \\ &= e^{-r\tau} \tilde{\mathbf{E}}\left(\left[S(t) \exp\left(\left(r - \frac{\sigma^2}{2}\right)\tau + \sigma(\tilde{W}(T) - \tilde{W}(t))\right) - K\right]^+ \mid \mathcal{F}_t\right) \\ &= \frac{e^{-r\tau}}{\sqrt{2\pi}} \int_{\mathbb{R}} \left[S(t) \exp\left(\left(r - \frac{\sigma^2}{2}\right)\tau + \sigma\sqrt{\tau}y\right) - K\right]^+ e^{-y^2/2} dy, \end{aligned}$$

by the independence lemma. Here $\tau = T - t$.

Now set $S(t) = x$,

$$d_{\pm}(\tau, x) \stackrel{\text{def}}{=} \frac{1}{\sigma\sqrt{\tau}} \left(\ln\left(\frac{x}{K}\right) + \left(r \pm \frac{\sigma^2}{2}\right)\tau \right),$$

and

$$N(x) = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^x e^{-y^2/2} dy = \frac{1}{\sqrt{2\pi}} \int_{-x}^{\infty} e^{-y^2/2} dy.$$

Observe

$$\begin{aligned} c(t, x) &= \frac{1}{\sqrt{2\pi}} \int_{-d_-}^{\infty} x \exp\left(\frac{-\sigma^2\tau}{2} + \sigma\sqrt{\tau}y - \frac{y^2}{2}\right) dy - e^{-r\tau} KN(d_-) \\ &= \frac{1}{\sqrt{2\pi}} \int_{-d_-}^{\infty} x \exp\left(\frac{-(y - \sigma\sqrt{\tau})^2}{2}\right) dy - e^{-r\tau} KN(d_-) \\ &= xN(d_+) - e^{-r\tau} KN(d_-), \end{aligned}$$

which is precisely the Black-Scholes formula.

4. Review Problems

PROBLEM 4.1. Let f be a deterministic function, and define

$$X(t) \stackrel{\text{def}}{=} \int_0^t f(s)W(s) ds.$$

Find the distribution of X .

PROBLEM 4.2. Suppose σ, τ, ρ are three deterministic functions and M and N are two martingales with respect to a common filtration $\{\mathcal{F}_t\}$ such that $M(0) = N(0) = 0$, and

$$d[M, M](t) = \sigma(t) dt, \quad d[N, N](t) = \tau(t) dt, \quad \text{and} \quad d[M, N](t) = \rho(t) dt.$$

- (a) Compute the joint moment generating function $\mathbf{E} \exp(\lambda M(t) + \mu N(t))$.
 (b) (*Lévy's criterion*) If $\sigma = \tau = 1$ and $\rho = 0$, show that (M, N) is a two dimensional Brownian motion.

PROBLEM 4.3. Consider a financial market consisting of a risky asset and a money market account. Suppose the return rate on the money market account is r , and the price of the risky asset, denoted by S , is a geometric Brownian motion with mean return rate α and volatility σ . Here r, α and σ are all deterministic constants. Compute the arbitrage free price of derivative security that pays

$$V(T) = \frac{1}{T} \int_0^T S(t) dt$$

at maturity T . Also compute the trading strategy in the replicating portfolio.

PROBLEM 4.4. Let $X \sim N(0, 1)$, and $a, \alpha, \beta \in \mathbb{R}$. Define a new measure $\tilde{\mathbf{P}}$ by

$$d\tilde{\mathbf{P}} = \exp(\alpha X + \beta) d\mathbf{P}.$$

Find α, β such that $X + a \sim N(0, 1)$ under $\tilde{\mathbf{P}}$.

PROBLEM 4.5. Let $x_0, \mu, \theta, \sigma \in \mathbb{R}$, and suppose X is an Itô process that satisfies

$$dX(t) = \theta(\mu - X(t)) dt + \sigma dW(t),$$

with $X(0) = x_0$.

- (a) Find functions $f = f(t)$ and $g = g(s, t)$ such that

$$X(t) = f(t) + \int_0^t g(s, t) dW(s).$$

The functions f, g may depend on the parameters x_0, θ, μ and σ , but should not depend on X .

- (b) Compute $\mathbf{E}X(t)$ and $\text{cov}(X(s), X(t))$ explicitly.

PROBLEM 4.6. Let M be a martingale, and φ be a convex function. Must the process $\varphi(M)$ be a martingale, sub-martingale, or a super-martingale? If yes, explain why. If no, find a counter example.

PROBLEM 4.7. Let $\theta \in \mathbb{R}$ and define

$$Z(t) = \exp\left(\theta W(t) - \frac{\theta^2 t}{2}\right).$$

Given $0 \leq s < t$, and a function f , find a function such that

$$\mathbf{E}(f(Z(t)) \mid \mathcal{F}_s) = g(Z(s)).$$

Your formula for the function g can involve f , s , t and integrals, but not the process Z or expectations.

PROBLEM 4.8. Let W be a Brownian motion, and define

$$B(t) = \int_0^t \text{sign}(W(s)) dW(s).$$

- (a) Show that B is a Brownian motion.
 (b) Is there an adapted process σ such that

$$W(t) = \int_0^t \sigma(s) dB(s)?$$

If yes, find it. If no, explain why.

- (c) Compute the joint quadratic variation $[B, W]$.
 (d) Are B and W uncorrelated? Are they independent? Justify.

PROBLEM 4.9. Let W be a Brownian motion. Does there exist an equivalent measure $\tilde{\mathbf{P}}$ under which the process $tW(t)$ is a Brownian motion? Prove it.

PROBLEM 4.10. Suppose M is a continuous process such that both M and M^2 are martingales. Must M be constant in time? Prove it, or find a counter example.