Financial Mathematics

Spring 2003

Richard F. Bass
Department of Mathematics
University of Connecticut

0. Introduction.

In this course we will study mathematical finance. Mathematical finance is **not** about predicting the price of a stock. What it is about is figuring out the price of options and derivatives.

The most familiar type of option is the option to buy a stock at a given price at a given time. For example, suppose Microsoft is currently selling today at $40 per share. A European call option is something I can buy that gives me the right to buy a share of Microsoft at some future date. To make up an example, suppose I have an option that allows me to buy a share of Microsoft for $50 in three months time, but does not compel me to do so. If Microsoft happens to be selling at $45 in three months time, the option is worthless. I would be silly to buy a share for $50 when I could call my broker and buy it for $45. So I would choose not to exercise the option. On the other hand, if Microsoft is selling for $60 three months from now, the option would be quite valuable. I could exercise the option and buy a share for $50. I could then turn around and sell the share on the open market for $60 and make a profit of $10 per share. Therefore this stock option I possess has some value. There is some chance it is worthless and some chance that it will lead me to a profit. The basic question is: how much is the option worth today?

The huge impetus in financial derivatives was the seminal paper of Black and Scholes in 1973. Although many researchers had studied this question, Black and Scholes gave a definitive answer, and a great deal of research has been done since. These are not just academic questions; today the market in financial derivatives is **larger** than the market in stock securities. In other words, more money is invested in options on stocks than in stocks themselves.

Options have been around for a long time. The earliest ones were used by manufacturers and food producers to hedge their risk. A farmer might agree to sell a bushel of wheat at a fixed price six months from now rather than take a chance on the vagaries of market prices. Similarly a steel refinery might want to lock in the price of iron ore at a fixed price.

The sections of these notes can be grouped into 5 categories. The first is elementary probability. Although someone who has had a course in undergraduate probability will be familiar with some of this, we will talk about a number of topics that are not usually covered in such a course: σ-fields, conditional expectations, martingales. The second category is the binomial asset pricing model. This is just about the simplest model of a stock that one can imagine, and this will provide a case where we can see most of the major ideas of mathematical finance, but in a very simple setting. Then we will turn to advanced probability, that is, ideas such as Brownian motion, stochastic integrals, stochastic differential equations, Girsanov tranformation. Although to do this rigorously requires measure theory, we can still learn enough to understand and work with these concepts. We then return to finance and work with the continuous model. We will derive the Black-Scholes formula, see the Fundamental Theorem of Asset Pricing, work with equivalent martingale measures, and the like. The fifth main category is term structure models, which means models of interest rate behavior.

1. Review of elementary probability.

Let's begin by recalling some of the definitions and basic concepts of elementary probability. We will only work with discrete models at first.

We start with an arbitrary set, called the probability space, which we will denote by Ω, the Greek letter "capital omega." We are given a class $\mathcal{F}$ of subsets of Ω. These are called events. We require $\mathcal{F}$ to be a σ-field. This means that

(1) $\emptyset \in \mathcal{F}$,
(2) $\Omega \in \mathcal{F}$,
(3) $A \in \mathcal{F}$ implies $A^c \in \mathcal{F}$, and
(4) $A_1, A_2, \ldots \in \mathcal{F}$ implies both $\cup_{i=1}^{\infty} A_i \in \mathcal{F}$ and $\cap_{i=1}^{\infty} A_i \in \mathcal{F}$.

Here $A^c = \{\omega \in \Omega : \omega \notin A\}$ denotes the complement of A. $\emptyset$ denotes the empty set, that is, the set with no elements. We will use without special comment the usual notations of $\cup$ (union), $\cap$ (intersection), $\subset$ (contained in), $\in$ (is an element of).

Typically, in an elementary probability course, $\mathcal{F}$ will consist of all subsets of Ω, but we will later need to distinguish between various σ-fields. Here is an example. Suppose one tosses a coin two times and lets Ω denote all possible outcomes. So $\Omega = \{HH, HT, TH, TT\}$. A typical σ-field $\mathcal{F}$ would be the one consisting of all subsets (of which there are 16). In this case it is trivial to show that $\mathcal{F}$ is a σ-field, since every subset is in $\mathcal{F}$. But if we let $\mathcal{G} = \{\emptyset, \Omega, \{HH, HT\}, \{TH, TT\}\}$, then $\mathcal{G}$ is also a σ-field. One has to check the definition, but to illustrate, the event $\{HH, HT\}$ is in $\mathcal{G}$, so we require the complement of that set to be in $\mathcal{G}$ as well. But the complement is $\{TH, TT\}$ and that event is indeed in $\mathcal{G}$.

One point of view which we will explore much more fully later on is that the σ-field tells you what events you know. In this example, $\mathcal{F}$ is the σ-field where you "know" everything, while $\mathcal{G}$ is the σ-field where you "know" only the result of the first toss but not the second. We won't try to be precise here, but to try to add to the intuition, suppose one knows whether an event in $\mathcal{F}$ has happened or not for a particular outcome. We would then know which of the events $\{HH\}, \{HT\}, \{TH\}$, or $\{TT\}$ has happened and so would know what the two tosses of the coin showed. On the other hand, if we know which events in $\mathcal{G}$ happened, we would only know whether the event $\{HH, HT\}$ happened, which means we would know that the first toss was a heads, or we would know whether the event $\{TH, TT\}$ happened, in which case we would know that the first toss was a tails. But there is no way to tell what happened on the second toss from knowing which events in $\mathcal{G}$ happened. Much more on this later.

The third basic ingredient is a function $\mathbb{P}$ on $\mathcal{F}$ satisfying

(1) if $A \in \mathcal{F}$, then $0 \le \mathbb{P}(A) \le 1$,
(2) $\mathbb{P}(\Omega) = 1$, and
(3) if $A_1, A_2, \ldots \in \mathcal{F}$ are pairwise disjoint, then $\mathbb{P}(\cup_{i=1}^\infty A_i) = \sum_{i=1}^\infty \mathbb{P}(A_i)$.

$\mathbb{P}$ is called a probability or probability measure.

There are a number of conclusions one can draw from this definition. As one example, if $A \subset B$, then $\mathbb{P}(A) \le \mathbb{P}(B)$ and $\mathbb{P}(A^c) = 1 - \mathbb{P}(A)$.

Someone who has had real analysis will realize that a σ-field is the same thing as a σ-algebra and a probability is a measure of total mass one.

A random variable (abbreviated r.v.) is a function X from Ω to $\mathbb{R}$, the reals. To be more precise, to be a r.v. X must also be measurable, which means that $\{\omega : X(\omega) \ge a\} \in \mathcal{F}$ for all reals a.

The notion of measurability has a simple definition but is a bit subtle. If we take the point of view that we know all the events in $\mathcal{G}$, then if Y is $\mathcal{G}$-measurable, then we know Y.

Here is an example. In the example where we tossed a coin two times, let X be the number of heads in the two tosses. Then X is $\mathcal{F}$ measurable but not $\mathcal{G}$ measurable. To see this, let us consider $A_a = \{\omega \in \Omega : X(\omega) \ge a\}$. This event will equal

$$\begin{cases} \Omega & \text{if } a \le 0; \\ \{HH, HT, TH\} & \text{if } 0 < a \le 1; \\ \{HH\} & \text{if } 1 < a \le 2; \\ \emptyset & \text{if } 2 < a. \end{cases}$$

For example, if $a = \frac{3}{2}$, then the event where the number of heads is $\frac{3}{2}$ or greater is the event where we had two heads, namely, $\{HH\}$. Now observe that for each a the event A_a

is in $\mathcal{F}$ because $\mathcal{F}$ contains all subsets of Ω. Therefore X is measurable with respect to $\mathcal{F}$. However it is not true that A_a is in $\mathcal{G}$ for every value of a – take $a = \frac{3}{2}$ as just one example. So X is not measurable with respect to the σ-field $\mathcal{G}$.

A discrete r.v. is one where $\mathbb{P}(\omega : X(\omega) = a) = 0$ for all but countably many a's. In defining sets one usually omits the ω; thus $(X = x)$ is the same as $\{\omega : X(\omega) = x\}$.

In the discrete case, to check measurability with respect to a σ-field $\mathcal{F}$, it is enough that $(X = a) \in \mathcal{F}$ for all reals a. The reason for this is if $x_1, x_2, \dots$ are the values of x for which $\mathbb{P}(X = x) \neq 0$, then we can write $(X \geq a) = \cup_{x_i \geq a}(X = x_i)$ and we have a countable union. So if $(X = x_i) \in \mathcal{F}$, then $(X \geq a) \in \mathcal{F}$.

Given a discrete r.v. X, the expectation or mean is defined by

$$\mathbb{E}\, X = \sum_x x\mathbb{P}(X = x)$$

provided the sum converges. If X only takes finitely many values, then this is a finite sum and of course it will converge. This is the situation that we will consider for quite some time. However, if X can take an infinite number of values (but countable), convergence needs to be checked.

There is an alternate definition of expectation which is equivalent in the discrete setting. Set

$$\mathbb{E}\, X = \sum_{\omega \in \Omega} X(\omega)\mathbb{P}(\{\omega\}).$$

To see that this is the same, we have

$$\begin{aligned}
\sum_x x\mathbb{P}(X = x) &= \sum_x x \sum_{\{\omega\in\Omega : X(\omega) = x\}} \mathbb{P}(\{\omega\}) \\
&= \sum_x \sum_{\{\omega\in\Omega : X(\omega) = x\}} X(\omega)\mathbb{P}(\{\omega\}) \\
&= \sum_{\omega\in\Omega} X(\omega)\mathbb{P}(\{\omega\}).
\end{aligned}$$

The advantage of the second definition is that some properties of expectation, such as $\mathbb{E}\,(X + Y) = \mathbb{E}\, X + \mathbb{E}\, Y$, are immediate, while with the first definition they require quite a bit of proof.

Two events A and B are independent if $\mathbb{P}(A \cap B) = \mathbb{P}(A)\mathbb{P}(B)$. Two random variables X and Y are independent if $\mathbb{P}(X \in A, Y \in B) = \mathbb{P}(X \in A)\mathbb{P}(X \in B)$ for all A and B that are subsets of the reals. The comma in the expression on the left hand side means "and." The extension of this definition to the case of more than two events or random variables is obvious.

Two σ-fields $\mathcal{F}$ and $\mathcal{G}$ are independent if A and B are independent whenever $A \in \mathcal{F}$ and $B \in \mathcal{G}$. A r.v. X and a σ-field $\mathcal{G}$ are independent if $\mathbb{P}((X \in A) \cap B) = \mathbb{P}(X \in A)\mathbb{P}(B)$ whenever A is a subset of the reals and $B \in \mathcal{G}$.

As an example, suppose we toss a coin two times and we define the σ-fields $\mathcal{G}_1 = \{\emptyset, \Omega, \{HH, HT\}, \{TH, TT\}\}$ and $\mathcal{G}_2 = \{\emptyset, \Omega, \{HH, TH\}, \{HT, TT\}\}$. Then $\mathcal{G}_1$ and $\mathcal{G}_2$ are independent if $\mathbb{P}(HH) = \mathbb{P}(HT) = \mathbb{P}(TH) = \mathbb{P}(TT) = \frac{1}{4}$. (Here we are writing $\mathbb{P}(HH)$ when a more accurate way would be to write $\mathbb{P}(\{HH\})$.) An easy way to understand this is that if we look at an event in $\mathcal{G}_1$ that is not $\emptyset$ or Ω, then that is the event that the first toss is a heads or it is the event that the first toss is a tails. Similarly, a set other than $\emptyset$ or Ω in $\mathcal{G}_2$ will be the event that the second toss is a heads or that the second toss is a tails.

If two r.v.s X and Y are independent, we have the multiplication theorem, which says that $\mathbb{E}\,(XY) = (\mathbb{E}\,X)(\mathbb{E}\,Y)$ provided all the expectations are finite.

Suppose $X_1, \ldots, X_n$ are n independent r.v.s, such that for each one $\mathbb{P}(X_i = 1) = p$, $\mathbb{P}(X_i = 0) = 1 - p$, where $p \in [0, 1]$. The random variable $S_n = \sum_{i=1}^n X_i$ is called a binomial r.v., and represents, for example, the number of successes in n trials, where the probability of a success is p. An important result in probability is that

$$\mathbb{P}(S_n = k) = \frac{n!}{k!(n-k)!} p^k (1-p)^{n-k}.$$

We close this section with a definition of conditional probability. The probability of A given B, written $\mathbb{P}(A \mid B)$ is defined by

$$\frac{\mathbb{P}(A \cap B)}{\mathbb{P}(B)},$$

provided $\mathbb{P}(B) \neq 0$. The conditional expectation of X given B is defined to be

$$\frac{\mathbb{E}\,[X; B]}{\mathbb{P}(B)},$$

provided $\mathbb{P}(B) \neq 0$. The notation $\mathbb{E}\,[X; B]$ means $\mathbb{E}\,[X 1_B]$, where $1_B(\omega)$ is 1 if $\omega \in B$ and 0 otherwise. Another way of writing $\mathbb{E}\,[X; B]$ is

$$\mathbb{E}\,[X; B] = \sum_{\omega \in B} X(\omega)\mathbb{P}(\{\omega\}).$$

2. Conditional expectation.

Suppose we have 200 men and 100 women, 70 of the men are smokers, and 50 of the women are smokers. If a person is chosen at random, then the conditional probability that the person is a smoker given that it is a man is 70 divided by 200, or 35%, while the conditional probability the person is a smoker given that it is a women is 50 divided by 100, or 50%. We will want to be able to encompass both facts in a single entity.

The way to do that is to make conditional probability a random variable rather than a number. To reiterate, we will make conditional probabilities random. Let M, W be man, woman, respectively, and S, S^c smoker and nonsmoker, respectively. We have

$$\mathbb{P}(S \mid M) = .35, \qquad \mathbb{P}(S \mid W) = .50.$$

We introduce the random variable

$$(.35)1_M + (.50)1_W$$

and use that for our conditional probability. So on the set M its value is .35 and on the set W its value is .50.

We need to give this random variable a name, so what we do is let $\mathcal{G}$ be the σ-field consisting of $\{\emptyset, \Omega, M, W\}$ and denote this random variable $\mathbb{P}(S \mid \mathcal{G})$. Thus we are going to talk about the conditional probability of an event given a σ-field.

What is the precise definition? Suppose there exist finitely (or countably) many sets $B_1, B_2, \ldots$, all having positive probability, such that they are pairwise disjoint, Ω is equal to their union, and $\mathcal{G}$ is the σ-field one obtains by taking all finite or countable unions of the B_i. Then the conditional probability of A given $\mathcal{G}$ is

$$\mathbb{P}(A \mid \mathcal{G}) = \sum_i \frac{\mathbb{P}(A \cap B_i)}{\mathbb{P}(B_i)} 1_{B_i}(\omega).$$

In short, on the set B_i the conditional probability is equal to $\mathbb{P}(A \mid B_i)$.

Not every σ-field can be so represented, so this definition will need to be extended when we get to continuous models. σ-fields that can be represented in this way are called finitely (or countably) generated and are said to be generated by the sets $B_1, B_2, \ldots$.

Let's look at another example. Suppose Ω consists of the possible results when we toss a coin three times: HHH, HHT, etc. Let $\mathcal{F}_3$ denote all subsets of Ω. Let $\mathcal{F}_1$ consist of the sets $\emptyset, \Omega, \{HHH, HHT, HTH, HTT\}$, and $\{THH, THT, TTH, TTT\}$. So $\mathcal{F}_1$ consists of those events that can be determined by knowing the result of the first toss. We want to let $\mathcal{F}_2$ denote those events that can be determined by knowing the first two tosses. This will include the sets $\emptyset, \Omega, \{HHH, HHT\}, \{HTH, HTT\}, \{THH, THT\}, \{TTH, TTT\}$. This is not enough to make $\mathcal{F}_2$ a σ-field, so we add to $\mathcal{F}_2$ all sets that can be obtained by taking unions of these sets.

Suppose we tossed the coin independently and suppose that it was fair. Let us calculate $\mathbb{P}(A \mid \mathcal{F}_1), \mathbb{P}(A \mid \mathcal{F}_2)$, and $\mathbb{P}(A \mid \mathcal{F}_3)$ when A is the event $\{HHH\}$. First the conditional probability given $\mathcal{F}_1$. Let $C_1 = \{HHH, HHT, HTH, HTT\}$ and $C_2 = \{THH, THT, TTH, TTT\}$. On the set C_1 the conditional probability is $\mathbb{P}(A \cap C_1)/\mathbb{P}(C_1) = \mathbb{P}(HHH)/\mathbb{P}(C_1) = \frac{1}{8}/\frac{1}{2} = \frac{1}{4}$. On the set C_2 the conditional probability is $\mathbb{P}(A \cap C_2)/\mathbb{P}(C_2) = \mathbb{P}(\emptyset)/\mathbb{P}(C_2) = 0$. Therefore $\mathbb{P}(A \mid \mathcal{F}_1) = (.25)1_{C_1}$. This is plausible – the probability of getting three heads given the first toss is $\frac{1}{4}$ if the first toss is a heads and 0 otherwise.

Next let us calculate $\mathbb{P}(A \mid \mathcal{F}_2)$. Let $D_1 = \{HHH, HHT\}, D_2 = \{HTH, HTT\}, D_3 = \{THH, THT\}, D_4 = \{TTH, TTT\}$. So $\mathcal{F}_2$ is the σ-field consisting of all possible unions of some of the D_i's. $\mathbb{P}(A \mid D_1) = \mathbb{P}(HHH)/\mathbb{P}(D_1) = \frac{1}{8}/\frac{1}{4} = \frac{1}{2}$. Also, as above, $\mathbb{P}(A \mid D_i) = 0$ for $i = 2, 3, 4$. So $\mathbb{P}(A \mid \mathcal{F}_2) = (.50)1_{D_1}$. This is again plausible – the probability of getting three heads given the first two tosses is $\frac{1}{2}$ if the first two tosses were heads and 0 otherwise.

What about conditional expectation? Given a random variable X, we define

$$\mathbb{E}[X \mid \mathcal{G}] = \sum_i \frac{\mathbb{E}[X; B_i]}{\mathbb{P}(B_i)} 1_{B_i}.$$

This is the obvious definition, and it agrees with what we had before because $\mathbb{E}[1_A \mid \mathcal{G}] = \mathbb{P}(A \mid \mathcal{G})$.

We now turn to some properties of conditional expectation. Some of the following propositions may seem a bit technical. In fact, they are! However, these properties are crucial to what follows and there is no choice but to master them.

Proposition 2.1. *$\mathbb{E}[X \mid \mathcal{G}]$ is $\mathcal{G}$ measurable, that is, if $Y = \mathbb{E}[X \mid \mathcal{G}]$, then $(Y > a)$ is a set in $\mathcal{G}$ for each real a.*

Proof. By the definition,

$$Y = \mathbb{E}[X \mid \mathcal{G}] = \sum_i \frac{\mathbb{E}[X; B_i]}{\mathbb{P}(B_i)} 1_{B_i} = \sum_i b_i 1_{B_i}$$

if we set $b_i = \mathbb{E}[X; B_i]/\mathbb{P}(B_i)$. The set $(Y \geq a)$ is a union of some of the B_i, namely, those B_i for which $b_i \geq a$. But the union of any collection of the B_i is in $\mathcal{G}$. $\square$

An example might help. Suppose

$$Y = 2 \cdot 1_{B_1} + 3 \cdot 1_{B_2} + 6 \cdot 1_{B_3} + 4 \cdot 1_{B_4}$$

and $a = 3.5$. Then $(Y \geq a) = B_3 \cup B_4$, which is in $\mathcal{G}$.

Proposition 2.2. *If $C \in \mathcal{G}$ and $Y = \mathbb{E}[X \mid \mathcal{G}]$, then $\mathbb{E}[Y; C] = \mathbb{E}[X; C]$.*

Proof. Since $Y = \sum \frac{\mathbb{E}[X;B_i]}{\mathbb{P}(B_i)} 1_{B_i}$ and the B_i are disjoint, then

$$\mathbb{E}[Y; B_j] = \frac{\mathbb{E}[X; B_j]}{\mathbb{P}(B_j)} \mathbb{E}\, 1_{B_j} = \mathbb{E}[X; B_j].$$

Now if $C = B_{j_1} \cup \cdots \cup B_{j_n} \cup \cdots$, summing the above over the j_k gives $\mathbb{E}[Y; C] = \mathbb{E}[X; C]$. $\square$

Let us look at the above example for this proposition, and let us do the case where $C = B_2$. Note $1_{B_2} 1_{B_2} = 1_{B_2}$ because the product is $1 \cdot 1 = 1$ if ω is in B_2 and 0 otherwise. On the other hand, it is not possible for an ω to be in more than one of the B_i, so $1_{B_2} 1_{B_i} = 0$ if $i \neq 2$. Mutiplying Y in the above example by 1_{B_2}, we see that

$$\begin{aligned} \mathbb{E}[Y; C] &= \mathbb{E}[Y; B_2] = \mathbb{E}[Y 1_{B_2}] = \mathbb{E}[3 \cdot 1_{B_2}] \\ &= 3\mathbb{E}[1_{B_2}] = 3\mathbb{P}(B_2). \end{aligned}$$

However the number 3 is not just any number; it is $\mathbb{E}[X; B_2]/\mathbb{P}(B_2)$. So

$$3\mathbb{P}(B_2) = \frac{\mathbb{E}[X; B_2]}{\mathbb{P}(B_2)} \mathbb{P}(B_2) = \mathbb{E}[X; B_2] = \mathbb{E}[X; C],$$

just as we wanted. If $C = B_1 \cup B_4$, for example, we then write

$$\begin{aligned} \mathbb{E}[X; C] &= \mathbb{E}[X 1_C] = \mathbb{E}[X(1_{B_2} + 1_{B_4})] \\ &= \mathbb{E}[X 1_{B_2}] + \mathbb{E}[X 1_{B_4}] = \mathbb{E}[X; B_2] + \mathbb{E}[X; B_4]. \end{aligned}$$

By the first part, this equals $\mathbb{E}[Y; B_2] + \mathbb{E}[Y; B_4]$, and we undo the above string of equalities but with Y instead of X to see that this is $\mathbb{E}[Y; C]$.

If a r.v. Y is $\mathcal{G}$ measurable, then for any a we have $(Y = a) \in \mathcal{G}$ which means that $(Y = a)$ is the union of one of more of the B_i. Since the B_i are disjoint, it follows that Y must be constant on each B_i.

Again let us look at an example. Suppose Z takes only the values $1, 3, 4, 7$. Let $D_1 = (Z = 1), D_2 = (Z = 3), D_3 = (Z = 4), D_4 = (Z = 7)$. Note that we can write

$$Z = 1 \cdot 1_{D_1} + 3 \cdot 1_{D_2} + 4 \cdot 1_{D_3} + 7 \cdot 1_{D_4}.$$

To see this, if $\omega \in D_2$, for example, the right hand side will be $0 + 3 \cdot 1 + 0 + 0$, which agrees with $Z(\omega)$. Now if Z is $\mathcal{G}$ measurable, then $(Z \geq a) \in \mathcal{G}$ for each a. Take $a = 7$, and we see $D_4 \in \mathcal{G}$. Take $a = 4$ and we see $D_3 \cup D_4 \in \mathcal{G}$. Taking $a = 3$ shows $D_2 \cup D_3 \cup D_4 \in \mathcal{G}$. Now $D_3 = (D_3 \cup D_4) \cap D_4^c$, so since $\mathcal{G}$ is a σ-field, $D_3 \in \mathcal{G}$. Similarly $D_2, D_1 \in \mathcal{G}$. Because sets in $\mathcal{G}$ are unions of the B_i's, we must have Z constant on the B_i's. For example, if it so happened that $D_1 = B_1, D_2 = B_2 \cup B_4, D_3 = B_3 \cup B_6 \cup B_7$, and $D_4 = B_5$, then

$$Z = 1 \cdot 1_{B_1} + 3 \cdot 1_{B_2} + 4 \cdot 1_{B_3} + 3 \cdot 1_{B_4} + 7 \cdot 1_{B_5} + +4 \cdot 1_{B_6} + 4 \cdot 1_{B_7}.$$

We still restrict ourselves to the discrete case. In this context, the properties given in Propositions 2.1 and 2.2 uniquely determine $\mathbb{E}[X \mid \mathcal{G}]$.

Proposition 2.3. *Suppose Z is $\mathcal{G}$ measurable and $\mathbb{E}[Z;C] = \mathbb{E}[X;C]$ whenever $C \in \mathcal{G}$. Then $Z = \mathbb{E}[X \mid \mathcal{G}]$.*

Proof. Since Z is $\mathcal{G}$ measurable, then Z must be constant on each B_i. Let the value of Z on B_i be z_i. So $Z = \sum_i z_i 1_{B_i}$. Then

$$z_i \mathbb{P}(B_i) = \mathbb{E}[Z;B_i] = \mathbb{E}[X;B_i],$$

or $z_i = \mathbb{E}[X;B_i]/\mathbb{P}(B_i)$ as required. □

The following propositions contain the main facts about this new definition of conditional expectation that we will need.

Proposition 2.4. *(1) If $X_1 \geq X_2$, then $\mathbb{E}[X_1 \mid \mathcal{G}] \geq \mathbb{E}[X_2 \mid \mathcal{G}]$.*
(2) $\mathbb{E}[aX_1 + bX_2 \mid \mathcal{G}] = a\mathbb{E}[X_1 \mid \mathcal{G}] + b\mathbb{E}[X_2 \mid \mathcal{G}]$.
(3) If X is $\mathcal{G}$ measurable, then $\mathbb{E}[X \mid \mathcal{G}] = X$.
(4) $\mathbb{E}[\mathbb{E}[X \mid \mathcal{G}]] = \mathbb{E}X$.
(5) If X is independent of $\mathcal{G}$, then $\mathbb{E}[X \mid \mathcal{G}] = \mathbb{E}X$.

If we were to think of $\mathbb{E}[X \mid \mathcal{G}]$ as the best prediction of X given $\mathcal{G}$, we can give an interpretation of (1)-(5). (1) says that if X_1 is larger than X_2, then the predicted value of X_1 should be larger than the predicted value of X_2. (2) says that the predicted value of $X_1 + X_2$ should be the sum of the predicted values. (3) says that if we know $\mathcal{G}$ and X is $\mathcal{G}$ measurable, then we know X and our best prediction of X is X itself. (4) says that the average of the predicted value of X should be the predicted value of X. (5) says that if knowing $\mathcal{G}$ gives us no additional information on X, then the best prediction for the value of X is just $\mathbb{E}X$.

Proof. (1) and (2) are immediate from the definition. To prove (3), note that if $Z = X$, then Z is $\mathcal{G}$ measurable and $\mathbb{E}[X;C] = \mathbb{E}[Z;C]$ for any $C \in \mathcal{G}$; this is trivial. By Proposition 2.3 it follows that $Z = \mathbb{E}[X \mid \mathcal{G}]$; this proves (3). To prove (4), if we let $C = \Omega$ and $Y = \mathbb{E}[X \mid \mathcal{G}]$, then $\mathbb{E}Y = \mathbb{E}[Y;C] = \mathbb{E}[X;C] = \mathbb{E}X$.

Last is (5). Let $Z = \mathbb{E}X$. Z is constant, so clearly $\mathcal{G}$ measurable. By the independence, if $C \in \mathcal{G}$, then $\mathbb{E}[X;C] = \mathbb{E}[X 1_C] = (\mathbb{E}X)(\mathbb{E}1_C) = (\mathbb{E}X)(\mathbb{P}(C))$. But $\mathbb{E}[Z;C] = (\mathbb{E}X)(\mathbb{P}(C))$ since Z is constant. By Proposition 2.3 we see $Z = \mathbb{E}[X \mid \mathcal{G}]$. □

Propostion 2.5. *If Z is $\mathcal{G}$ measurable, then $\mathbb{E}[XZ \mid \mathcal{G}] = Z\mathbb{E}[X \mid \mathcal{G}]$.*

Proof. Note that $Z\mathbb{E}[X \mid \mathcal{G}]$ is $\mathcal{G}$ measurable, so by Proposition 2.3 we need to show its expectation over sets C in $\mathcal{G}$ is the same as that of XZ. As in the proof of Proposition

2.2, it suffices to consider only the case when C is one of the B_i. Now Z is $\mathcal{G}$ measurable, hence it is constant on B_i; let its value be z_i. Then

$$\mathbb{E}\,[Z\mathbb{E}\,[X \mid \mathcal{G}]; B_i] = \mathbb{E}\,[z_i\mathbb{E}\,[X \mid \mathcal{G}]; B_i] = z_i\mathbb{E}\,[\mathbb{E}\,[X \mid \mathcal{G}]; B_i] = z_i\mathbb{E}\,[X; B_i] = \mathbb{E}\,[XZ; B_i]$$

as desired. □

This proposition says that as far as conditional expectations with respect to a σ-field $\mathcal{G}$ go, $\mathcal{G}$-measurable random variables act like constants: they can be taken inside or outside the conditional expectation at will.

Proposition 2.6. *If $\mathcal{H} \subset \mathcal{G} \subset \mathcal{F}$, then*

$$\mathbb{E}\,[\mathbb{E}\,[X \mid \mathcal{H}] \mid \mathcal{G}] = \mathbb{E}\,[X \mid \mathcal{H}] = \mathbb{E}\,[\mathbb{E}\,[X \mid \mathcal{G}] \mid \mathcal{H}].$$

Proof. $\mathbb{E}\,[X \mid \mathcal{H}]$ is $\mathcal{H}$ measurable, hence $\mathcal{G}$ measurable, since $\mathcal{H} \subset \mathcal{G}$. The left hand equality now follows by Proposition 2.4(3). To get the right hand equality, let W be the right hand expression. It is $\mathcal{H}$ measurable, and if $C \in \mathcal{H} \subset \mathcal{G}$, then

$$\mathbb{E}\,[W; C] = \mathbb{E}\,[\mathbb{E}\,[X \mid \mathcal{G}]; C] = \mathbb{E}\,[X; C]$$

as required. □

In words, if we are predicting a prediction of X given limited information, this is the same as a single prediction given the least amount of information.

If Y is a discrete random variables, that is, it takes only countably many values $y_1, y_2, \ldots$, we let $B_i = (Y = y_i)$. These will be disjoint sets whose union is Ω. If $\sigma(Y)$ is the collection of all unions of the B_i, then $\sigma(Y)$ is a σ-field, and is called the σ-field generated by Y. It is easy to see that this is the smallest σ-field with respect to which Y is measurable. We write $\mathbb{E}\,[X \mid Y]$ for $\mathbb{E}\,[X \mid \sigma(Y)]$.

3. Martingales.

Suppose we have a sequence of σ-fields $\mathcal{F}_1 \subset \mathcal{F}_2 \subset \mathcal{F}_3 \cdots$. An example would be repeatedly tossing a coin and letting $\mathcal{F}_k$ be the sets that can be determined by the first k tosses. Another example is to let $\mathcal{F}_k$ be the events that are determined by the values of a stock at times 1 through k. A third example is to let $X_1, X_2, \ldots$ be a sequence of random variables and let $\mathcal{F}_k$ be the σ-field generated by $X_1, \ldots, X_k$, the smallest σ-field with respect to which $X_1, \ldots, X_k$ are measurable.

A r.v. X is integrable if $\mathbb{E}\,|X| < \infty$. Given an increasing sequence of σ-fields $\mathcal{F}_n$, a sequence of r.v.'s X_n is adapted if X_n is $\mathcal{F}_n$ measurable for each n.

A martingale M_n is a sequence of random variables such that M_n is integrable for all n, M_n is adapted to $\mathcal{F}_n$, and

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = M_n \tag{3.1}$$

for each n. Martingales will be ubiquitous in financial math.

The word "martingale" comes from the piece of a horse's bridle that runs from the horse's head to its chest. It keeps the horse from raising its head too high. It turns out that martingales in probability cannot get too large.

Here is an example. Let $X_1, X_2, \ldots$ be a sequence of independent r.v.'s with mean 0 that are independent. Set $\mathcal{F}_n = \sigma(X_1, \ldots, X_n)$, the σ-field generated by $X_1, \ldots, X_n$. Let $M_n = \sum_{i=1}^n X_i$. Then

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = X_1 + \cdots + X_n + \mathbb{E}\,[X_{n+1} \mid \mathcal{F}_n] = M_n + \mathbb{E}\,X_{n+1} = M_n,$$

where we used the independence.

Another example: suppose in the above that the X_k all have variance 1, and let $M_n = S_n^2 - n$, where $S_n = \sum_{i=1}^n X_i$. We compute

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = \mathbb{E}\,[S_n^2 + 2X_{n+1}S_n + X_{n+1}^2 \mid \mathcal{F}_n] - (n+1).$$

We have $\mathbb{E}\,[S_n^2 \mid \mathcal{F}_n] = S_n^2$ since S_n is $\mathcal{F}_n$ measurable. $\mathbb{E}\,[2X_{n+1}S_n \mid \mathcal{F}_n] = 2S_n\mathbb{E}\,[X_{n+1} \mid \mathcal{F}_n] = 2S_n\mathbb{E}\,X_{n+1} = 0$. And $\mathbb{E}\,[X_{n+1}^2 \mid \mathcal{F}_n] = \mathbb{E}\,X_{n+1}^2 = 1$. Substituting, we obtain $\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = M_n$, or M_n is a martingale.

A third example: Suppose you start with a dollar and you are tossing a fair coin independently. If it turns up heads you double your fortune, tails you go broke. This is "double or nothing." Let M_n be your fortune at time n. To formalize this, let $X_1, X_2, \ldots$ be independent r.v.'s that are equal to 2 with probability $\frac{1}{2}$ and 0 with probability $\frac{1}{2}$. Then $M_n = X_1 \cdots X_n$. To compute the conditional expectation, note $\mathbb{E}\,X_{n+1} = 1$. Then

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = M_n\mathbb{E}\,[X_{n+1} \mid \mathcal{F}_n] = M_n\mathbb{E}\,X_{n+1} = M_n,$$

using the independence.

A final example for now: let $\mathcal{F}_1, \mathcal{F}_2, \ldots$ be given and let X be a fixed r.v. Let $M_n = \mathbb{E}\,[X \mid \mathcal{F}_n]$. We have

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] = \mathbb{E}\,[\mathbb{E}\,[X \mid \mathcal{F}_{n+1}] \mid \mathcal{F}_n] = \mathbb{E}\,[X \mid \mathcal{F}_n] = M_n.$$

4. Properties of martingales.

When it comes to discussing American options, we will need the concept of stopping times. A mapping τ from Ω into the nonnegative integers is a stopping time if $(\tau = k) \in \mathcal{F}_k$ for each k.

An example is $\tau = \min\{k : S_k \geq A\}$. This is a stopping time because $(\tau = k) = (S_1, \ldots, S_{k-1} < A, S_k \geq A) \in \mathcal{F}_k$. We can think of a stopping time as the first time something happens. $\sigma = \max\{k : S_k \geq A\}$, the last time, is not a stopping time.

Here is an intuitive description of a stopping time. If I tell you to drive to the city limits and then drive until you come to the second stop light after that, you know when you get there that you have arrived; you don't need to have been there before or to look ahead. But if I tell you to drive until you come to the second stop light **before** the city limits, either you must have been there before or else you have to go past where you are supposed to stop, continue on to the city limits, and then turn around and come back two stop lights. You don't know when you first get to the second stop light before the city limits that you get to stop there. The first set of instructions forms a stopping time, the second set does not.

If M_n is an adapted sequence of integrable r.v.'s with

$$\mathbb{E}\,[M_{n+1} \mid \mathcal{F}_n] \geq M_n$$

for each n, then M_n is a submartingale. (Is the "$\geq$" is replaced by "$=$", of course, that is what we called a martingale; if the "$\geq$" is replaced by "$\leq$", we call M_n a supermartingale.)

Our first result is Jensen's inequality.

Proposition 4.1. *If g is convex, then*

$$g(\mathbb{E}\,[X \mid \mathcal{G}]) \leq \mathbb{E}\,[g(X) \mid \mathcal{G}]$$

provided all the expectations exist.

For ordinary expectations rather than conditional expectations, this is still true. That is, if g is convex and the expectations exist, then

$$g(\mathbb{E}\,X) \leq \mathbb{E}\,[g(X)].$$

We already know some special cases of this: when $g(x) = |x|$, this says $|\mathbb{E}\,X| \leq \mathbb{E}\,|X|$; when $g(x) = x^2$, this says $(\mathbb{E}\,X)^2 \leq \mathbb{E}\,X^2$, which we know because $\mathbb{E}\,X^2 - (\mathbb{E}\,X)^2 = \mathbb{E}\,(X - \mathbb{E}\,X)^2 \geq 0$.

Proof. If g is convex, then the graph of g lies above all the tangent lines. Even if g does not have a derivative at x_0, there is a line passing through x_0 which lies beneath the graph of g. So for each x_0 there exists $c(x_0)$ such that

$$g(x) \geq g(x_0) + c(x_0)(x - x_0).$$

Apply this with $x = X(\omega)$ and $x_0 = \mathbb{E}[X \mid \mathcal{G}](\omega)$. We then have

$$g(X) \geq g(\mathbb{E}[X \mid \mathcal{G}]) + c(\mathbb{E}[X \mid \mathcal{G}])(X - \mathbb{E}[X \mid \mathcal{G}]).$$

One can check that c can be chosen so that $c(\mathbb{E}[X \mid \mathcal{G}])$ is $\mathcal{G}$ measurable.

Now take the conditional expectation with respect to $\mathcal{G}$. The first term on the right is $\mathcal{G}$ measurable, so remains the same. The second term on the right is equal to

$$c(\mathbb{E}[X \mid \mathcal{G}])\mathbb{E}[X - \mathbb{E}[X \mid \mathcal{G}] \mid \mathcal{G}] = 0.$$

□

One reason we want Jensen's inequality is to show that a convex function applied to a martingale yields a submartingale.

Proposition 4.2. *If M_n is a martingale and g is convex, then $g(M_n)$ is a submartingale, provided all the expectations exist.*

Proof. By Jensen's inequality,

$$\mathbb{E}[g(M_{n+1}) \mid \mathcal{F}_n] \geq g(\mathbb{E}[M_{n+1} \mid \mathcal{F}_n]) = g(M_n).$$

□

If M_n is a martingale, then $\mathbb{E}\, M_n = \mathbb{E}[\mathbb{E}[M_{n+1} \mid \mathcal{F}_n]] = \mathbb{E}\, M_{n+1}$. So $\mathbb{E}\, M_0 = \mathbb{E}\, M_1 = \cdots = \mathbb{E}\, M_n$. Doob's optional stopping theorem says the same thing holds when fixed times n are replaced by stopping times.

Theorem 4.3. *Suppose K is a positive integer, N is a stopping time such that $N \leq K$ a.s., and M_n is a martingale. Then*

$$\mathbb{E}\, M_N = \mathbb{E}\, M_K.$$

Here, to evaluate M_N, one first finds $N(\omega)$ and then evaluates $M_\cdot(\omega)$ for that value of N.

Proof. We have

$$\mathbb{E}\, M_N = \sum_{k=0}^{K} \mathbb{E}[M_N; N = k].$$

If we show that the k-th summand is $\mathbb{E}[M_n; N = k]$, then the sum will be

$$\sum_{k=0}^{K} \mathbb{E}[M_n; N = k] = \mathbb{E}\, M_n$$

as desired. We have

$$\mathbb{E}\,[M_N; N = k] = \mathbb{E}\,[M_k; N = k]$$

by the definition of M_N. Now $(N = k)$ is $\mathcal{F}_k$ measurable, so by Proposition 2.2 and the fact that $M_k = \mathbb{E}\,[M_{k+1} \mid \mathcal{F}_k]$,

$$\mathbb{E}\,[M_k; N = k] = \mathbb{E}\,[M_{k+1}; N = k].$$

We have $(N = k) \in \mathcal{F}_k \subset \mathcal{F}_{k+1}$. Since $M_{k+1} = \mathbb{E}\,[M_{k+2} \mid \mathcal{F}_{k+1}]$, Proposition 2.2 tells us that

$$\mathbb{E}\,[M_{k+1}; N = k] = \mathbb{E}\,[M_{k+2}; N = k].$$

We continue, using $(N = k) \in \mathcal{F}_k \subset \mathcal{F}_{k+1} \subset \mathcal{F}_{k+2}$, and we obtain

$$\mathbb{E}\,[M_N; N = k] = \mathbb{E}\,[M_k; N = k] = \mathbb{E}\,[M_{k+1}; N = k] = \cdots = \mathbb{E}\,[M_n; N = k].$$

$\square$

If we change the equalities in the above to inequalities, the same result holds for submartingales.

As a corollary we have two of Doob's inequalities:

Theorem 4.4. *(a) If M_n is a nonnegative submartingale,*

$$\mathbb{P}(\max_{k \le n} M_k \ge \lambda) \le \frac{1}{\lambda} \mathbb{E}\, M_n.$$

(b) $$\mathbb{E}\,(\max_{k \le n} M_k^2) \le 4 \mathbb{E}\, M_n^2.$$

Proof. Set $M_{n+1} = M_n$. It is easy to see that the sequence $M_1, M_2, \ldots, M_{n+1}$ is also a martingale. Let $N = \min\{k : M_k \ge \lambda\} \wedge (n + 1)$, the first time that M_k is greater than or equal to λ, where $a \wedge b = \min(a, b)$. Then

$$\mathbb{P}(\max_{k \le n} M_k \ge \lambda) = \mathbb{P}(N \le n)$$

and if $N \le n$, then $M_N \ge \lambda$. Now

$$\begin{aligned}\mathbb{P}(\max_{k \le n} M_k \ge \lambda) &= \mathbb{E}\,[1_{(N \le n)}] \le \mathbb{E}\left[\frac{M_N}{\lambda}; N \le n\right] \\ &= \frac{1}{\lambda}\mathbb{E}\,[M_{N \wedge n}; N \le n] \le \frac{1}{\lambda}\mathbb{E}\, M_{N \wedge n}.\end{aligned} \tag{4.1}$$

Finally, since M_n is a submartingale, $\mathbb{E}\, M_{N \wedge n} \le \mathbb{E}\, M_n$.

We now look at (b). Let us write M^* for $\max_{k \le n} M_k$. We have

$$\mathbb{E}\,[M_{N\wedge n}; N \le n] = \sum_{k=0}^{\infty} \mathbb{E}\,[M_{k\wedge n}; N = k].$$

Arguing as in the proof of Theorem 4.3,

$$\mathbb{E}\,[M_{k\wedge n}; N = k] \le \mathbb{E}\,[M_n; N = k],$$

and so

$$\mathbb{E}\,[M_{N\wedge n}; N \le n] \le \sum_{k=0}^{\infty} \mathbb{E}\,[M_n; N = k] = \mathbb{E}\,[M_n; N \le n].$$

The last expression is at most $\mathbb{E}\,[M_n; M^* \ge \lambda]$. If we multiply (4.1) by 2λ and integrate over λ from 0 to ∞, we obtain

$$\begin{aligned}
\int_0^\infty 2\lambda \mathbb{P}(M^* \ge \lambda) d\lambda &\le 2\int_0^\infty \mathbb{E}\,[M_n : M^* \ge \lambda]\\
&= 2\mathbb{E}\int_0^\infty M_n 1_{(M^*\ge\lambda)} d\lambda\\
&= 2\mathbb{E}\Big[M_n \int_0^{M^*} d\lambda\Big]\\
&= 2\mathbb{E}\,[M_n M^*].
\end{aligned}$$

Using Cauchy-Schwarz, this is bounded by

$$2(\mathbb{E}\, M_n^2)^{1/2}(\mathbb{E}\,(M^*)^2)^{1/2}.$$

On the other hand,

$$\begin{aligned}
\int_0^\infty 2\lambda \mathbb{P}(M^* \ge \lambda) d\lambda &= \mathbb{E}\int_0^\infty 2\lambda 1_{(M^*\ge\lambda)} d\lambda\\
&= \mathbb{E}\int_0^{M^*} 2\lambda\, d\lambda = \mathbb{E}\,(M^*)^2.
\end{aligned}$$

We therefore have

$$\mathbb{E}\,(M^*)^2 \le 2(\mathbb{E}\, M_n^2)^{1/2}(\mathbb{E}\,(M^*)^2)^{1/2}.$$

Suppose $\mathbb{E}\,(M^*)^2 < \infty$. We divide both sides by $(\mathbb{E}\,(M^*)^2)^{1/2}$ and square both sides. (When $\mathbb{E}\,(M^*)^2$ is infinite, there is a way to circumvent the division by infinity.) $\square$

The last result we want is that bounded martingales converge. (The hypothesis of boundedness can be weakened.)

Theorem 4.5. *Suppose M_n is a martingale bounded in absolute value by K. That is, $|M_n| \le K$ for all n. Then $\lim_{n\to\infty} M_n$ exists a.s.*

Proof. Since M_n is bounded, it can't tend to $+\infty$ or $-\infty$. The only possibility is that it might oscillate. Let $a < b$ be two rationals. What might go wrong is that M_n might be larger than b infinitely often and less than a infinitely often. If we show the probability of this is 0, then taking the union over all pairs of rationals (a, b) shows that almost surely M_n cannot oscillate, and hence must converge.

Fix a, b and let $S_1 = \min\{k : M_k \le a\}$, $T_1 = \min\{k > S_1 : M_k \ge b\}$, $S_2 = \min\{k > T_1 : M_k \le a\}$, and so on. Let $U_n = \max\{k : T_k \le n\}$. U_n is called the number of upcrossings up to time n.

We can write

$$2K \ge M_n - M_0 = \sum_{k=1}^{n}(M_{S_{k+1}\wedge n} - M_{T_k\wedge n}) + \sum_{k=1}^{\infty}(M_{T_k\wedge n} - M_{S_k\wedge n}) + (M_{S_1\wedge n} - M_0).$$

Now take expectations. The expectation of the first sum on the right and the last term are zero by optional stopping. The middle term is larger than $(b-a)U_n$, so we conclude

$$(b-a)\mathbb{E}\, U_n \le 2K.$$

Let $n \to \infty$ to see that $\mathbb{E}\max_n U_n < \infty$, which implies $\max_n U_n < \infty$ a.s., which is what we needed. $\square$

5. The one step binomial asset pricing model.

Let us begin by giving the simplest possible model of a stock and see how a European call option should be valued in this context.

Suppose we have a single stock whose price is S_0. Let d and u be two numbers with $0 < d < 1 < u$. Here "d" is a mnemonic for "down" and "u" for "up." After one time unit the stock price will be either uS_0 with probability P or else dS_0 with probability Q, where $P + Q = 1$. Instead of purchasing shares in the stock, you can also put your money in the bank where one will earn interest at rate r. Alternatives to the bank are money market funds or bonds; the key point is that these are considered to be risk-free.

A European call option in this context is the option to buy one share of the stock at time 1 at price K. K is called the strike price. Let S_1 be the price of the stock at time 1. If S_1 is less than K, then the option is worthless at time 1. If S_1 is greater than K, you can use the option at time 1 to buy the stock at price K, immediately turn around and sell the stock for price S_1 and make a profit of $S_1 - K$. So the value of the option at time 1 is

$$V_1 = (S_1 - K)^+,$$

where x^+ is $\max(x, 0)$. The principal question to be answered is: what is the value V_0 of the option at time 0? In other words, how much should one pay for a European call option with strike price K?

It is possible to buy a negative number of shares of a stock. This is equivalent to selling shares of a stock you don't have and is called selling short. If you sell one share of stock short, then at time 1 you must buy one share at whatever the market price is at that time and turn it over to the person that you sold the stock short to. Similarly you can buy a negative number of options, that is, sell an option.

You can also deposit a negative amount of money in the bank, which is the same as borrowing. We assume that you can borrow at the same interest rate r, not exactly a totally realistic assumption. One way to make it seem more realistic is to assume you have a large amount of money on deposit, and when you borrow, you simply withdraw money from that account.

We are looking at the simplest possible model, so we are going to allow only one time step: one makes an investment, and looks at it again one day later.

Let's suppose the price of a European call option is V_0 and see what conditions one can put on V_0. Suppose you start out with V_0 dollars. One thing you could do is buy one option. The other thing you could do is use the money to buy Δ_0 shares of stock. If $V_0 > \Delta_0 S_0$, there will be some money left over and you put that in the bank. If $V_0 < \Delta_0 S_0$, you do not have enough money to buy the stock, and you make up the shortfall by borrowing money from the bank. In either case, at this point you have $V_0 - \Delta_0 S_0$ in the bank and Δ_0 shares of stock.

If the stock goes up, at time 1 you will have

$$\Delta_0 u S_0 + (1 + r)(V_0 - \Delta_0 S_0),$$

and if it goes down,

$$\Delta_0 d S_0 + (1 + r)(V_0 - \Delta_0 S_0).$$

We have not said what Δ_0 should be. Let us do that now. Let $V_1^u = (uS_0 - K)^+$ and $V_1^d = (dS_0 - K)^+$. Note these are deterministic quantities, i.e., not random. Let

$$\Delta_0 = \frac{V_1^u - V_1^d}{uS_0 - dS_0},$$

and we will also need

$$W_0 = \frac{1}{1 + r}\Big[\frac{1 + r - d}{u - d} V_1^u + \frac{u - (1 + r)}{u - d} V_1^d\Big].$$

After some algebra, we see that if the stock goes up and you had bought stock instead of the option you would now have

$$V_1^u + (1 + r)(V_0 - W_0),$$

while if the stock went down, you would now have

$$V_1^d + (1+r)(V_0 - W_0).$$

Let's check the first of these, the second being similar. We need to show

$$\Delta_0 u S_0 + (1+r)(V_0 - \Delta_0 S_0) = V_1^u + (1+r)(V_0 - W_0). \tag{5.1}$$

The left hand side is equal to

$$\Delta_0 S_0(u - (1+r)) + (1+r)V_0 = \frac{V_1^u - V_1^d}{u-d}(u-(1+r)) + (1+r)V_0. \tag{5.2}$$

The right hand side of (5.1) is equal to

$$V_1^u - \Big[\frac{1+r-d}{u-d}V_1^u + \frac{u-(1+r)}{u-d}V_1^d\Big] + (1+r)V_0. \tag{5.3}$$

Now check that the coefficients of V_0, of V_1^u, and of V_1^d agree in (5.2) and (5.3).

Suppose that $V_0 > W_0$. What you want to do is come along with no money, sell one option for V_0 dollars, use the money to buy Δ_0 shares, and put the rest in the bank (or borrow if necessary). If the buyer of your option wants to exercise the option, you give him one share of stock and sell the rest. If he doesn't want to exercise the option, you sell your shares of stock and pocket the money. Remember it is possible to have a negative number of shares. You will have cleared $(1+r)(V_0 - W_0)$, whether the stock went up or down, with no risk.

If $V_0 < W_0$, you just do the opposite: sell Δ_0 shares of stock short, buy one option, and deposit or make up the shortfall from the bank. This time, you clear $(1+r)(W_0 - V_0)$, whether the stock goes up or down.

Now most people believe that you can't make a profit on the stock market without taking a risk. The name for this is "no free lunch," or "arbitrage opportunities do not exist." The only way to avoid this is if $V_0 = W_0$. In other words, we have shown that the only reasonable price for the European call option is W_0.

The "no arbitrage" condition is not just a reflection of the belief that one cannot get something for nothing. It also represents the belief that the market is freely competitive. The way it works is this: suppose $W_0 = \$3$. Suppose you could sell options at a price $V_0 = \$5$; this is larger than W_0 and you would earn $V_0 - W_0 = \$2$ per option without risk. Then someone else would observe this and decide to sell the same option at a price less than V_0 but larger than W_0, say \$4. This person would still make a profit, and customers would go to him and ignore you because they would be getting a better deal. But then a third person would decide to sell the option for less than your competition but more than W_0, say at \$3.50. This would continue as long as any one would try to sell an option above price W_0.

We will examine this problem of pricing options in more complicated contexts, and while doing so, it will become apparent where the formulas for Δ_0 and W_0 came from. At this point, we want to make a few observations.

Remark 5.1. First of all, if $1+r > u$, one would never buy stock, since one can always do better by putting money in the bank. So we may suppose $1+r < u$. We always have $1+r \geq 1 > d$. If we set

$$\overline{p} = \frac{1+r-d}{u-d}, \qquad \overline{q} = \frac{u-(1+r)}{u-d},$$

then $\overline{p}, \overline{q} \geq 0$ and $\overline{p}+\overline{q} = 1$. Thus $\overline{p}$ and $\overline{q}$ act like probabilities, but they have nothing to do with P and Q. Note also that the price $V_0 = W_0$ does not depend on P or Q. It does depend on $\overline{p}$ and $\overline{q}$, which seems to suggest that there is an underlying probability which controls the option price and is not the one that governs the stock price.

Remark 5.2. There is nothing special about European call options in our argument above. One could let V_1^u and V_d^1 be any two values of any option, which are paid out if the stock goes up or down, respectively. The above analysis shows we can exactly duplicate the result of buying any option V by instead buying some shares of stock. If in some model one can do this for any option, the market is called *complete* in this model.

Remark 5.3. If we let $\overline{\mathbb{P}}$ be the probability so that $S_1 = uS_0$ with probability $\overline{p}$ and $S_1 = dS_0$ with probability $\overline{q}$ and we let $\overline{E}$ be the corresponding expectation, then some algebra shows that

$$V_0 = \frac{1}{1+r}\overline{\mathbb{E}}V_1.$$

This will be generalized later.

Remark 5.4. If one buys one share of stock at time 0, then one expects at time 1 to have $(Pu+Qd)S_0$. One then divides by $1+r$ to get the value of the stock in today's dollars. Suppose instead of P and Q being the probabilities of going up and down, they were in fact $\overline{p}$ and $\overline{q}$. One would then expect to have $(\overline{p}u+\overline{q}d)S_0$ and then divide by $1+r$. Substituting the values for $\overline{p}$ and $\overline{q}$, this reduces to S_0. In other words, if $\overline{p}$ and $\overline{q}$ were the correct probabilities, one would expect to have the same amount of money one started with. When we get to the binomial asset pricing model with more than one step, we will see that the generalization of this fact is that the stock price at time n is a martingale, still with the assumption that $\overline{p}$ and $\overline{q}$ are the correct probabilities. This is a special case of the *fundamental theorem of finance*: there always exists some probability, not necessarily the one you observe, under which the stock price is a martingale.

Remark 5.5. Our model allows after one time step the possibility of the stock going up or going down, but only these two options. What if instead there are 3 (or more) possibilities. Suppose for example, that the stock goes up a factor u with probability P, down a factor d with probability Q, and remains constant with probability R, where $P+Q+R=1$. The corresponding price of a European call option would be $(uS_0-K)^+$, $(dS_0-K)^+$, or $(S_0-K)^+$. If one could replicate this outcome by buying and selling shares of the stock, then the "no arbitrage" rule would give the exact value of the call option in this model. But, except in very special circumstances, one cannot do this, and the theory falls apart. One has three equations one wants to satisfy, in terms of V_1^u, V_1^d, and V_1^c. (The "c" is a mnemonic for "constant.") There are however only two variables, Δ_0 and V_0 at your disposal, and most of the time three equations in two unknowns cannot be solved.

6. The multi-step binomial asset pricing model.

In this section we will obtain a formula for the pricing of options when there are n time steps, but each time the stock can only go up by a factor u or down by a factor d. The "Black-Scholes" formula we will obtain is already a nontrivial result that is useful.

We assume the following.

(1) Unlimited short selling of stock
(2) Unlimited borrowing
(3) No transaction costs
(4) Our buying and selling is on a small enough scale that it does not affect the market.

We need to set up the probability model. Ω will be all sequences of length n of H's and T's. S_0 will be a fixed number and we define $S_k(\omega)=u^jd^{k-j}S_0$ if the first k elements of a given $\omega\in\Omega$ has j occurrences of H and $k-j$ occurrences of T. (What we are doing is saying that if the j-th element of the sequence making up ω is an H, then the stock price goes up by a factor u; if T, then down by a factor d.) $\mathcal{F}_k$ will be the σ-field generated by $S_0,\ldots,S_k$.

Let

$$\overline{p}=\frac{(1+r)-d}{u-d},\qquad \overline{q}=\frac{u-(1+r)}{u-d}$$

and define $\overline{\mathbb{P}}(\omega)=\overline{p}^j\overline{q}^{n-j}$ if ω has j appearances of H and $n-j$ appearances of T. It is not hard to see that under $\overline{\mathbb{P}}$ the random variables S_{k+1}/S_k are independent and equal to u with probability $\overline{p}$ and d with probability $\overline{q}$. To see this, let $Y_k=S_k/S_{k-1}$. Then $\mathbb{P}(Y_1=y_1,\ldots,Y_n=y_n)=\overline{p}^j\overline{q}^{n-j}$, where j is the number of the y_k that are equal to u. On the other hand, this is equal to $\mathbb{P}(Y_1=y_1)\cdots\mathbb{P}(Y_n=y_n)$. Let $\overline{\mathbb{E}}$ denote the expectation corresponding to $\overline{\mathbb{P}}$.

The $\overline{\mathbb{P}}$ we construct may not be the true probabilities of going up or down. That doesn't matter - it will turn out that using the principle of "no arbitrage," it is $\overline{\mathbb{P}}$ that governs the price.

Our first result is the fundamental theorem of finance in the current context.

Proposition 6.1. *Under $\overline{\mathbb{P}}$ the discounted stock price $(1+r)^{-k}S_k$ is a martingale.*

Proof. Since the random variable S_{k+1}/S_k is independent of $\mathcal{F}_k$, we have

$$\overline{\mathbb{E}}\,[(1+r)^{-(k+1)}S_{k+1} \mid \mathcal{F}_k] = (1+r)^{-k}S_k(1+r)^{-1}\overline{\mathbb{E}}\,[S_{k+1}/S_k \mid \mathcal{F}_k].$$

Using the independence the conditional expectation on the right is equal to

$$\overline{\mathbb{E}}\,[S_{k+1}/S_k] = \overline{p}u + \overline{q}d = 1 + r.$$

Substituting yields the proposition. □

Let Δ_k be the number of shares held between times k and $k+1$. We require Δ_k to be $\mathcal{F}_k$ measurable. $\Delta_0, \Delta_1, \ldots$ is called the portfolio process. Let W_0 be the amount of money you start with and let W_k be the amount of money you have at time k. W_k is the wealth process. Then

$$W_{k+1} = \Delta_k S_{k+1} + (1+r)[W_k - \Delta_k S_k].$$

Note that in the case where $r = 0$ we have

$$W_{k+1} - W_k = \Delta_k(S_{k+1} - S_k),$$

or

$$W_{k+1} = \sum_{i=0}^{k} \Delta_i(S_{i+1} - S_i).$$

This is a discrete version of a stochastic integral. Since

$$\overline{\mathbb{E}}\,[W_{k+1} - W_k \mid \mathcal{F}_k] = \Delta_k \overline{\mathbb{E}}\,[S_{k+1} - S_k \mid \mathcal{F}_k] = 0,$$

it follows that W_k is a martingale. More generally

Proposition 6.2. *Under $\overline{\mathbb{P}}$ the discounted wealth process $(1+r)^{-k}W_k$ is a martingale.*

Proof. We have

$$(1+r)^{-(k+1)}W_{k+1} = (1+r)^{-k}W_k + \Delta_k[(1+r)^{-(k+1)}S_{k+1} - (1+r)^{-k}S_k],$$

and so

$$\begin{aligned}\overline{\mathbb{E}}\,[\Delta_k[(1+r)^{-(k+1)}S_{k+1} - (1+r)^{-k}S_k \mid \mathcal{F}_k]\\ = \Delta_k\overline{\mathbb{E}}\,[(1+r)^{-(k+1)}S_{k+1} - (1+r)^{-k}S_k \mid \mathcal{F}_k] = 0.\end{aligned}$$

The result follows. □

Our next result is that the binomial model is complete. It is easy to lose the idea in the algebra, so first let us try to see why the theorem is true.

For simplicity suppose $r = 0$. Let $V_k = \mathbb{E}\,[V \mid \mathcal{F}_k]$; we saw that V_k is a martingale. We want to construct a portfolio process so that $W_n = V$. We will do it inductively by arranging matters so that $W_k = V_k$ for all k. Recall that W_k is also a martingale.

Suppose we have $W_k = V_k$ at time k and we want to find Δ_k so that $W_{k+1} = V_{k+1}$. At the $(k+1)$-st step there are only two possible changes for the price of the stock and so since V_{k+1} is $\mathcal{F}_{k+1}$ measurable, only two possible values for V_{k+1}. We need to choose Δ_k so that $W_{k+1} = V_{k+1}$ for each of these two possibilities. We only have one parameter, Δ_k, to play with to match up two numbers, which may seem like an overconstrained system of equations. But both V and W are martingales, which is why the system can be solved.

Now let us turn to the details.

Theorem 6.3. *The binomial asset pricing model is complete.*

Proof. Let

$$V_k = (1+r)^k\overline{\mathbb{E}}\,[(1+r)^{-n}V \mid \mathcal{F}_k]$$

so that $(1+r)^{-k}V_k$ is a martingale. If $\omega = (t_1, \ldots, t_n)$, where each t_i is an H or T, let

$$\Delta_k(\omega) = \frac{V_{k+1}(t_1, \ldots, t_k, H, t_{k+2}, \ldots, t_n) - V_{k+1}(t_1, \ldots, t_k, T, t_{k+2}, \ldots, t_n)}{S_{k+1}(t_1, \ldots, t_k, H, t_{k+2}, \ldots, t_n) - S_{k+1}(t_1, \ldots, t_k, T, t_{k+2}, \ldots, t_n)}.$$

Set $W_0 = V_0$, and we will show by induction that the wealth process at time $k+1$ equals V_{k+1}.

The first thing to show is that Δ_k is $\mathcal{F}_k$ measurable. Neither S_{k+1} nor V_{k+1} depends on $t_{k+2}, \ldots, t_n$. So Δ_k depends only on the variables $t_1, \ldots, t_k$, hence is $\mathcal{F}_k$ measurable.

Now $t_{k+2}, \ldots, t_n$ play no role in the rest of the proof, and $t_1, \ldots, t_k$ will be fixed, so we drop the t's from the notation.

We know $(1+r)^{-k}V_k$ is a martingale under $\overline{\mathbb{P}}$ so that

$$\begin{aligned}V_k &= \overline{\mathbb{E}}\,[(1+r)^{-1}V_{k+1} \mid \mathcal{F}_k]\\ &= \frac{1}{1+r}[\overline{p}V_{k+1}(H) + \overline{q}V_{k+1}(T)].\end{aligned}$$

We now suppose $W_k = V_k$ and want to show $W_{k+1}(H) = V_{k+1}(H)$ and $W_{k+1}(T) = V_{k+1}(T)$. Then using induction we have $W_n = V_n = V$ as required. We show the first equality, the second being similar.

$$\begin{aligned}
W_{k+1}(H) &= \Delta_k S_{k+1}(H) + (1+r)[W_k - \Delta_k S_k] \\
&= \Delta_k [uS_k - (1+r)S_k] + (1+r)V_k \\
&= \frac{V_{k+1}(H) - V_{k+1}(T)}{(u-d)S_k} S_k[u-(1+r)] + \overline{p}V_{k+1}(H) + \overline{q}V_{k+1}(T) \\
&= V_{k+1}(H).
\end{aligned}$$

We are done. □

Finally, we obtain the Black-Scholes formula in this context. Let V be any option that is $\mathcal{F}_n$-measurable. The one we have in mind is the European call, for which $V = (S_n - K)^+$, but the argument is the same for any option whatsoever.

Theorem 6.4. *The value of the option V at time 0 is $V_0 = (1+r)^{-n}\overline{\mathbb{E}}\, V$.*

Proof. We can construct a portfolio process Δ_k so that if we start with $W_0 = (1+r)^{-n}\overline{\mathbb{E}}\, V$, then the wealth at time n will equal V, no matter what the market does in between. If we could buy or sell the option V at a price other than W_0, we could obtain a riskless profit. That is, if the option V could be sold at a price c_0 larger than W_0, we would sell the option for c_0 dollars, use W_0 to buy and sell stock according to the portfolio process Δ_k, have a net worth of $V + (1+r)^n(c_0 - W_0)$ at time n, meet our obligation to the buyer of the option by using V dollars, and have a net profit, at no risk, of $(1+r)^n(c_0 - W_0)$. If c_0 were less than W_0, we would do the same except buy an option, hold $-\Delta_k$ shares at time k, and again make a riskless profit. By the "no arbitrage" rule, that can't happen, so the price of the option V must be W_0. □

Remark 6.5. Note that the proof of Theorem 6.4 tells you precisely what hedging strategy (i.e., what portfolio process to use).

In the binomial asset pricing model, there is no difficulty computing the price of a European call. We have

$$\overline{\mathbb{E}}\,(S_n - K)^+ = \sum_x (x-K)^+ \overline{\mathbb{P}}(S_n = x)$$

and

$$\mathbb{P}(S_n = x) = \binom{n}{k} \overline{p}^k \overline{q}^{n-k}$$

if $x = u^k d^{n-k} S_0$. Therefore the price of the European call is

$$\sum_{k=0}^{n} (u^k d^{n-k} S_0 - K)^+ \binom{n}{k} \overline{p}^k \overline{q}^{n-k}.$$

The formula in Theorem 6.4 holds for exotic options as well. Suppose

$$V = \max_{i=1,\ldots,n} S_i - \min_{j=1,\ldots,n} S_j.$$

In other words, you sell the stock for the maximum value it takes during the first n time steps and you buy at the minimum value the stock takes; you are allowed to wait until time n and look back to see what the maximum and minimum were. You can even do this if the maximum comes before the minimum. This V is still $\mathcal{F}_n$ measurable, so the theory applies. Naturally, such a "buy low, sell high" option is very desirable, and the price of such a V will be quite high. It is interesting that even without using options, you can duplicate the operation of buying low and selling high by holding an appropriate number of shares Δ_k at time k, where you do not look into the future to determine Δ_k.

7. American options.

An American option is one where you can exercise the option any time before some fixed time T. For example, on a European call, one can only use it to buy a share of stock at the expiration time T, while for an American call, at any time before time T, one can decide to pay K dollars and obtain a share of stock.

Let us give an informal argument on how to price an American call, giving a more rigorous argument in a moment. One can always wait until time T to exercise an American call, so the value must be at least as great as that of a European call. On the other hand, suppose you decide to exercise early. You pay K dollars, receive one share of stock, and your wealth is $S_t - K$. You hold onto the stock, and at time T you have one share of stock worth S_T, and for which you paid K dollars. So your wealth is $S_T - K \le (S_T - K)^+$. In fact, we have strict inequality, because you lost the interest on your K dollars that you would have received if you had waited to exercise until time T. Therefore an American call is worth no more than a European call, and hence its value must be the same as that of a European call.

This argument does not work for puts, because selling stock gives you some money on which you will receive interest, so it may be advantageous to exercise early. (A put is the option to sell a stock at a price K at time T.)

Here is the more rigorous argument. Let $g(x)$ be convex with $g(0) = 0$. Certainly $g(x) = (x - K)^+$ is such a function. We have

$$g(\lambda x) = g(\lambda x + (1-\lambda) \cdot 0) \le \lambda g(x) + (1-\lambda) g(0) = \lambda g(x).$$

By Jensen's inequality,

$$\begin{aligned}\overline{\mathbb{E}}\,[(1+r)^{-(k+1)}g(S_{k+1}) \mid \mathcal{F}_k] &= (1+r)^{-k}\overline{\mathbb{E}}\Big[\frac{1}{1+r}g(S_{k+1}) \mid \mathcal{F}_k\Big]\\ &\ge (1+r)^{-k}\overline{\mathbb{E}}\Big[g\Big(\frac{1}{1+r}S_{k+1}\Big) \mid \mathcal{F}_k\Big]\\ &\ge (1+r)^{-k}g\Big(\overline{\mathbb{E}}\Big[\frac{1}{1+r}S_{k+1} \mid \mathcal{F}_k\Big]\Big)\\ &= (1+r)^{-k}g(S_k).\end{aligned}$$

So $(1+r)^{-k}g(S_k)$ is a submartingale. By optional stopping,

$$\overline{\mathbb{E}}\,[(1+r)^{-\tau}g(S_\tau)] \le \overline{\mathbb{E}}\,[(1+r)^{-n}g(S_n)],$$

so $\tau \equiv n$ always does best.

8. Continuous random variables.

We are now going to start working towards continuous times and stocks that can take any positive number as a value, so we need to prepare by extending some of our definitions.

Given any random variable X, we can approximate it by r.v's X_n that are discrete. We let

$$X_n = \sum_{i=-n2^n}^{n2^n} \frac{i}{2^n} 1_{(i/2^n \le X < (i+1)/2^n)}.$$

In words, if $X(\omega)$ lies between $-n$ and n, we let $X_n(\omega)$ be the closest value $i/2^n$ that is less than or equal to $X(\omega)$. For ω where $|X(\omega)| > n$ we set $X_n(\omega) = 0$. Clearly the X_n are discrete, and approximate X. In fact, on the set where $|X| \le n$, we have that $|X(\omega) - X_n(\omega)| \le 2^{-n}$.

For reasonable X we are going to define $\mathbb{E}\,X = \lim \mathbb{E}\,X_n$. There are some things one wants to prove, but all this has been worked out in measure theory and the theory of the Lebesgue integral. Let us confine ourselves here to showing this definition is the same as the usual one when X has a density.

Recall X has a density f_X if

$$\mathbb{P}(X \in [a,b]) = \int_a^b f_X(x)dx$$

for all a and b. In this case

$$\mathbb{E}\,X = \int_{-\infty}^{\infty} x f_X(x)dx.$$

With our definition of X_n we have

$$\mathbb{P}(X_n = i/2^n) = \mathbb{P}(X \in [i/2^n, (i+1)/2^n)) = \int_{i/2^n}^{(i+1)/2^n} f_X(x)dx.$$

Then

$$\mathbb{E}\, X_n = \sum_i \frac{i}{2^n}\mathbb{P}(X_n = i/2^n) = \sum_i \int_{i/2^n}^{(i+1)/2^n} \frac{i}{2^n} f_X(x)dx.$$

Since x differs from $i/2^n$ by at most $1/2^n$ when $x \in [i/2^n, (i+1)/2^n)$, this will tend to $\int x f_X(x)dx$, unless the contribution to the integral for $|x| \geq n$ does not go to 0 as $n \to \infty$. As long as $\int |x| f_X(x)dx < \infty$, one can show that this contribution does indeed go to 0.

We also need an extension of the definition of conditional probability. A r.v. is $\mathcal{G}$ measurable if $(X > a) \in \mathcal{G}$ for every a. How do we define $\mathbb{E}\,[Z \mid \mathcal{G}]$ when $\mathcal{G}$ is not generated by a countable collection of disjoint sets?

Again, there is a completely worked out theory that holds in all cases. Let us give a definition that is equivalent that works except for a very few cases. Let us suppose that for each n the σ-field $\mathcal{G}_n$ is finitely generated. This means that $\mathcal{G}_n$ is generated by finitely many disjoint sets $B_{n1}, \ldots, B_{nm_n}$. So for each n, the number of B_{ni} is finite but arbitrary, the B_{ni} are disjoint, and their union is Ω. Suppose also that $\mathcal{G}_1 \subset \mathcal{G}_2 \subset \cdots$. Now $\cup_n \mathcal{G}_n$ will not in general be a σ-field, but suppose $\mathcal{G}$ is the smallest σ-field that contains all the $\mathcal{G}_n$. Finally, define $\mathbb{P}(A \mid \mathcal{G}) = \lim \mathbb{P}(A \mid \mathcal{G}_n)$.

This is a fairly general set-up. For example, let Ω be the real line and let $\mathcal{G}_n$ be generated by the sets $(-\infty, n), [n, \infty)$ and $[i/2^n, (i+1)/2^n)$. Then $\mathcal{G}$ will contain every interval that is closed on the left and open on the right, hence $\mathcal{G}$ must be the σ-field that one works with when one talks about Lebesgue measure on the line.

The question that one might ask is: how does one know the limit exists? Since the $\mathcal{G}_n$ increase, we know that $M_n = \mathbb{P}(A \mid \mathcal{G}_n)$ is a martingale with respect to the $\mathcal{G}_n$. It is certainly bounded above by 1 and bounded below by 0, so by the martingale convergence theorem, it must have a limit as $n \to \infty$.

Once one has a definition of conditional probability, one defines conditional expectation by what one expects. If X is discrete, one can write X as $\sum_j a_j 1_{A_j}$ and then one defines

$$\mathbb{E}\,[X \mid \mathcal{G}] = \sum_j a_j \mathbb{P}(A_j \mid \mathcal{G}_n).$$

If the X is not discrete, one approximates as above. One has to worry about convergence, but everything does go through.

With this extended definition of conditional expectation, do all the properties of Section 2 hold? The answer is yes, and the proofs are by taking limits of the discrete approximations.

We will be talking about stochastic processes. Previously we discussed sequences $S_1, S_2, \ldots$ of r.v.'s. Now we want to talk about processes Y_t for $t \geq 0$. We typically let $\mathcal{F}_t$ be the smallest σ-field with respect to which Y_s is measurable for all $s \leq t$. As you might imagine, there are a few technicalities one has to worry about. We will try to avoid thinking about them as much as possible.

A continuous time martingale (or submartingale) is what one expects: M_t is integrable, adapted to $\mathcal{F}_t$, and if $s < t$, then $\mathbb{E}[M_t \mid \mathcal{F}_s] = M_s$. The analogues of Doob's theorems go through. The way to prove these is to observe that $M_{k/2^n}$ is a discrete time martingale, and then to take limits as $n \to \infty$.

9. Brownian motion.

Let S_n be a simple symmetric random walk. This means that $Y_k = S_k - S_{k-1}$ equals $+1$ with probability $\frac{1}{2}$, equals -1 with probability $\frac{1}{2}$, and is independent of Y_j for $j < k$. We notice that $\mathbb{E}\, S_n = 0$ while $\mathbb{E}\, S_n^2 = \sum_{i=1}^n \mathbb{E}\, Y_i^2 + \sum_{i \neq j} \mathbb{E}\, Y_i Y_j = n$ using the fact that $\mathbb{E}[Y_i Y_j] = (\mathbb{E}\, Y_i)(\mathbb{E}\, Y_j) = 0$.

Define $X_t^n = S_{nt}/\sqrt{n}$ if nt is an integer and by linear interpolation for other t. If nt is an integer, $\mathbb{E}\, X_t^n = 0$ and $\mathbb{E}\,(X_t^n)^2 = t$. It turns out X_t^n does not converge for any ω.

However there is another kind of convergence, called weak convergence, that takes place. There exists a process Z_t such that for each k, each $t_1 < t_2 < \cdots < t_k$, and each $a_1 < b_1, a_2 < b_2, \ldots, a_k < b_k$, we have

(1) The paths of Z_t are continuous as a function of t.

(2) $\mathbb{P}(X_{t_1}^n \in [a_1, b_1], \ldots, X_{t_k}^n \in [a_k, b_k]) \to \mathbb{P}(Z_{t_1} \in [a_1, b_1], \ldots, Z_{t_k} \in [a_k, b_k])$.

The limit Z_t is called a Brownian motion starting at 0. It has the following properties.

(1) $\mathbb{E}\, Z_t = 0$.

(2) $\mathbb{E}\, Z_t^2 = t$.

(3) $Z_t - Z_s$ is independent of $\mathcal{F}_s = \sigma(Z_r, r \leq s)$.

(4) $Z_t - Z_s$ has the distribution of a normal random variable with mean 0 and variance $t - s$. This means

$$\mathbb{P}(Z_t - Z_s \in [a, b]) = \int_a^b \frac{1}{\sqrt{2\pi(t-s)}} e^{-y^2/2(t-s)} dy.$$

(This result follows from the central limit theorem.)

(5) The map $t \to Z_t(\omega)$ is continuous for each ω.

10. Markov properties of Brownian motion.

Fix r and let $W_t = Z_{t+r} - Z_r$. Clearly the map $t \to W_t$ is continuous since the same is true for Z. Since $W_t - W_s = Z_{t+r} - Z_{s+r}$, then the distribution of $W_t - W_s$ is normal with mean zero and variance $(t + r) - (s + r)$. One can also check the other parts of the definition and we see that W_t is also a Brownian motion. This is a version of the Markov property. We will prove the following stronger result, which is a version of the strong Markov property.

A stopping time in the continuous framework is a r.v. T taking values in $[0, \infty)$ such that $(T > t) \in \mathcal{F}_t$ for all t. To make a satisfactory theory, one needs that the $\mathcal{F}_t$ be what is called right continuous: $\mathcal{F}_t = \cap_{\varepsilon>0}\mathcal{F}_{t+\varepsilon}$, but this is fairly technical and we will ignore it.

If T is a stopping time, $\mathcal{F}_T$ is the collection of events A such that $A \cap (T > t) \in \mathcal{F}_t$ for all t.

Proposition 10.1. *If X_t is a Brownian motion and T is a bounded stopping time, then $X_{T+t} - X_T$ is a mean 0 variance t random variable and is independent of $\mathcal{F}_T$.*

Proof. Let T_n be defined by $T_n(\omega) = (k+1)/2^n$ if $T(\omega) \in [k/2^n, (k+1)/2^n)$. It is easy to check that T_n is a stopping time. Let f be continuous and $A \in \mathcal{F}_T$. Then $A \in \mathcal{F}_{T_n}$ as well. We have

$$\begin{aligned}
\mathbb{E}\,[f(X_{T_n+t} - X_{T_n}); A] &= \sum \mathbb{E}\,[f(X_{\frac{k}{2^n}+t} - X_{\frac{k}{2^n}}); A \cap T_n = k/2^n] \\
&= \sum \mathbb{E}\,[f(X_{\frac{k}{2^n}+t} - X_{\frac{k}{2^n}})]\mathbb{P}(A \cap T_n = k/2^n) \\
&= \mathbb{E}\,f(X_t)\mathbb{P}(A).
\end{aligned}$$

Let $n \to \infty$, so

$$\mathbb{E}\,[f(X_{T+t} - X_T); A] = \mathbb{E}\,f(X_t)\mathbb{P}(A).$$

Taking limits this equation holds for all bounded f.

If we take $A = \Omega$ and $f = 1_B$, we see that $X_{T+t} - X_T$ has the same distribution as X_t, which is that of a mean 0 variance t normal random variable. If we let $A \in \mathcal{F}_T$ be arbitrary and $f = 1_B$, we see that

$$\mathbb{P}(X_{T+t} - X_T \in B, A) = \mathbb{P}(X_t \in B)\mathbb{P}(A) = \mathbb{P}(X_{T+t} - X_T \in B)\mathbb{P}(A),$$

which implies that $X_{T+t} - X_T$ is independent of $\mathcal{F}_T$. $\square$

This proposition says: if you want to predict X_{T+t}, you could do it knowing all of $\mathcal{F}_T$ or just knowing X_T. Since $X_{T+t} - X_T$ is independent of $\mathcal{F}_T$, the extra information given in $\mathcal{F}_T$ does you no good at all.

We need a way of expressing the Markov and strong Markov properties that will generalize to other processes.

Let W_t be a Brownian motion. Consider the process $W_t^x = x + W_t$, Brownian motion started at x. Define Ω' to be set of continuous functions on $[0,\infty)$, let $X_t(\omega) = \omega(t)$, and let the σ-field be the one generated by the X_t. Define $\mathbb{P}^x$ on $(\Omega', \mathcal{F}')$ by

$$\mathbb{P}^x(X_{t_1} \in A_1, \ldots, X_{t_n} \in A_n) = \mathbb{P}(W_{t_1}^x \in A_1, \ldots, W_{t_n}^x \in A_n).$$

What we have done is gone from one probability space Ω with many processes W_t^x to one process X_t with many probability measures $\mathbb{P}^x$.

Proposition 10.2. *If $s < t$ and f is bounded or nonnegative, then*

$$\mathbb{E}^{\,x}[f(X_t) \mid \mathcal{F}_s] = \mathbb{E}^{\,X_s}[f(X_{t-s})], \qquad \text{a.s.}$$

The right hand side is to be interpreted as follows. Define $\varphi(x) = \mathbb{E}^{\,x} f(X_{t-s})$. Then $\mathbb{E}^{\,X_s} f(X_{t-s})$ means $\varphi(X_s(\omega))$. One often writes $P_t f(x)$ for $\mathbb{E}^{\,x} f(X_t)$.

Before proving this, recall from undergraduate analysis that every bounded function is the limit of linear combinations of functions e^{iux}, $u \in \mathbb{R}$. This follows from using the inversion formula for Fourier transforms. There are various slightly different formulas for the Fourier transform. We use $\widehat{f}(u) = \int e^{iux} f(x)\,dx$. If f is smooth enough and has compact support, then one can recover f by the formula $f(x) = \frac{1}{2\pi}\int e^{-iux}\widehat{f}(u)\,du$. We can first approximate this improper integral by $\frac{1}{2\pi}\int_{-N}^{N} e^{-iux}\widehat{f}(u)\,du$ by taking N large enough, and then approximate the approximation by using Riemann sums. Thus we can approximate $f(x)$ by a linear combination of terms of the form $e^{iu_j x}$. Also, bounded functions can be approximated by smooth functions with compact support.

Proof. Let $f(x) = e^{iux}$. Then

$$\begin{aligned}\mathbb{E}^{\,x}[e^{iuX_t} \mid \mathcal{F}_s] &= e^{iuX_s}\mathbb{E}\,[e^{iu(X_t - X_s)} \mid \mathcal{F}_s] \\ &= e^{iuX_s}e^{-u^2(t-s)/2}.\end{aligned}$$

On the other hand,

$$\varphi(y) = \mathbb{E}^{\,y}[f(X_{t-s})] = \mathbb{E}\,[e^{iu(W_{t-s}+y)}] = e^{iuy}e^{-u^2(t-s)/2}.$$

So $\varphi(X_s) = \mathbb{E}^{\,x}[e^{iuX_t} \mid \mathcal{F}_s]$. Using linearity and taking limits, we have the lemma for all f. □

This formula generalizes: If $s < t < u$, then

$$\mathbb{E}^{\,x}[f(X_t)g(X_u) \mid \mathcal{F}_s] = \mathbb{E}^{\,X_s}[f(X_{t-s})g(X_{u-s})],$$

and so on for functions of X at more times.

Using Proposition 10.1, the statement and proof of Proposition 10.2 can be extended to stopping times.

Proposition 10.3. *If T is a bounded stopping time, then*
$$\mathbb{E}^x[f(X_{T+t}) \mid \mathcal{F}_T] = \mathbb{E}^{X_T}[f(X_t)].$$

11. Stochastic integrals.

If one wants to consider the (deterministic) integral $\int_0^t f(s)\,dg(s)$, where f and g are continuous and g is differentiable, we can define it analogously to the usual Riemann integral as the limit of Riemann sums $\sum_{i=1}^n f(s_i)[g(s_i)-g(s_{i-1})]$, where $s_1 < s_2 < \cdots < s_n$ is a partition of $[0,t]$. This is known as the Riemann-Stieltjes integral. One can show (using the mean value theorem, for example) that
$$\int_0^t f(s)\,dg(s) = \int_0^t f(s)g'(s)\,ds.$$
If we were to take $f(s) = 1_{[0,a]}(s)$, one would expect the following:
$$\int_0^t 1_{[0,a]}(s)\,dg(s) = \int_0^t 1_{[0,a]}(s)g'(s)\,ds = \int_0^a g'(s)\,ds = g(a) - g(0).$$
Note that although we use the fact that g is differentiable in the intermediate stages, the first and last terms make sense for any g.

We now want to replace g by a Brownian path and f by a random integrand. The expression $\int f(s)\,dW(s)$ does not make sense as a Riemann-Stieltjes integral because it is a fact that $W(s)$ is not differentiable as a function of t. We need to define the expression by some other means. We will show that it can be defined as the limit in L^2 of Riemann sums. The resulting integral is called a stochastic integral.

Let us consider a very special case first. Suppose f is continuous and deterministic (i.e., does not depend on ω). Suppose we take a Riemann sum approximation $\sum f(\frac{i}{2^n})[W(\frac{i+1}{2^n}) - W(\frac{i}{2^n})]$. If we take the difference of two successive approximations we have terms like
$$\sum_{i\ odd} [f(i/2^{n+1}) - f((i+1)/2^{n+1})][W((i+1)/2^{n+1}) - W(i/2^{n+1})].$$
This has mean zero. By the independence, the second moment is
$$\sum [f(i/2^{n+1}) - f((i+1)/2^{n+1})]^2(1/2^{n+1}).$$
This will be small if f is continuous. So by taking a limit in L^2 we obtain a nontrivial limit.

We now turn to the general case. Let W_t be a Brownian motion. We will only consider integrands H_s such that H_s is $\mathcal{F}_s$ measurable for each s. We will construct $\int_0^t H_s\,dW_s$ for all H with $\mathbb{E}\int_0^t H_s^2\,ds < \infty$.

Before we proceed we will need to define the quadratic variation of a continuous martingale. We will use the following theorem without proof because in our applications we can construct the desired increasing process directly.

Theorem 11.1. *Suppose M_t is a continuous martingale such that $\mathbb{E}\, M_t^2 < \infty$ for all t. There exists one and only one increasing process A_t that is adapted to $\mathcal{F}_t$, has continuous paths, and $A_0 = 0$ such that $M_t^2 - A_t$ is a martingale.*

The simplest example of such a martingale is Brownian motion. If W_t is a Brownian motion, let $L_t = W_t^2 - t$. Note

$$\begin{aligned}\mathbb{E}\,[L_t \mid \mathcal{F}_s] &= \mathbb{E}\,[((W_t - W_s) + W_s)^2 \mid \mathcal{F}_s] - t \\ &= \mathbb{E}\,[(W_t - W_s)^2 \mid \mathcal{F}_s] - 2\mathbb{E}\,[W_s(W_t - W_s) \mid \mathcal{F}_s] + \mathbb{E}\,[W_s^2 \mid \mathcal{F}_s] - t.\end{aligned}$$

The first term, using independence, is $\mathbb{E}\,[(W_t - W_s)^2] = t - s$, the second term is $W_s\mathbb{E}\,[W_t - W_s \mid \mathcal{F}_s] = W_s\mathbb{E}\,[W_t - W_s] = 0$, and so we have

$$\mathbb{E}\,[L_t \mid \mathcal{F}_s] = (t - s) + W_s^2 - t = L_s,$$

or L_t is a martingale. So in the case of Brownian motion, the increasing process $A_t = t$.

We use the notation $\langle M \rangle_t$ for the increasing process given in Theorem 11.1. We will see that in the case of stochastic integrals, where $N_t = \int_0^t H_s dW_s$, it turns out that $\langle N \rangle_t = \int_0^t H_s^2 ds$.

If K is bounded and $\mathcal{F}_a$ measurable, let $N_t = K(W_{t\wedge b} - W_{t\wedge a})$. Part of the statement of the next proposition is that $\langle N \rangle_t$ exists.

Lemma 11.2. *N_t is a continuous martingale, $\mathbb{E}\, N_\infty^2 = \mathbb{E}\,[K^2(b - a)]$ and*

$$\langle N \rangle_t = \int_0^t K^2 1_{[a,b]}(s)\, ds.$$

Proof. The continuity is clear. Let us look at $\mathbb{E}\,[N_t \mid \mathcal{F}_s]$. In the case $a < s < t < b$, this is equal to

$$\mathbb{E}\,[K(W_t - W_a) \mid \mathcal{F}_s] = K\mathbb{E}\,[(W_t - W_a) \mid \mathcal{F}_s] = K(W_s - W_a) = N_s.$$

In the case $s < a < t < b$, $\mathbb{E}\,[N_t \mid \mathcal{F}_s]$ is equal to

$$\mathbb{E}\,[K(W_t - W_a) \mid \mathcal{F}_s] = \mathbb{E}\,[K\mathbb{E}\,[W_t - W_a \mid \mathcal{F}_a] \mid \mathcal{F}_s] = 0 = N_s.$$

The other possibilities for where s and t are can be done similarly.

Recall $W_t^2 - t$ is a martingale. For $\mathbb{E}\, N_\infty^2$, we have

$$\begin{aligned}\mathbb{E}\, N_\infty^2 &= \mathbb{E}\,[K^2\mathbb{E}\,[(W_b - W_a)^2 \mid \mathcal{F}_a]] = \mathbb{E}\,[K^2\mathbb{E}\,[W_b^2 - W_a^2 \mid \mathcal{F}_a]] \\ &= \mathbb{E}\,[K^2\mathbb{E}\,[b - a \mid \mathcal{F}_a]] = \mathbb{E}\,[K^2(b - a)].\end{aligned}$$

For $\langle N \rangle_t$, we need to show

$$\begin{aligned}\mathbb{E}\,[K^2(W_{t\wedge b} - W_{t\wedge a})^2 - K^2(t\wedge b - t\wedge a) \mid \mathcal{F}_s]& \\ &= K^2(W_{s\wedge b} - W_{s\wedge a})^2 - K^2(s\wedge b - s\wedge a).\end{aligned}$$

We do this by checking all the cases. □

H_s is said to be simple if it can be written in the form $\sum_{j=1}^J H_j 1_{[a_j,b_j]}(s)$, where H_j is $\mathcal{F}_{s_j}$ measurable and bounded. Define

$$N_t = \int_0^t H_s\, dW_s = \sum_{j=1}^J H_j(W_{b_j\wedge t} - W_{a_j\wedge t}).$$

Proposition 11.3. *N_t is a continuous martingale, $\mathbb{E}\, N_\infty^2 = \mathbb{E}\int_0^\infty H_s^2\, ds$, and $\langle N \rangle_t = \int_0^t H_s^2\, ds$.*

Proof. We may rewrite H so that the intervals $[a_j, b_j]$ satisfy $a_1 \le b_1 \le a_2 \le b_2 \le \cdots \le b_j$. It is then clear that N_t is a martingale.

We have

$$\mathbb{E}\, N_\infty^2 = \mathbb{E}\Big[\sum H_j^2(W_{b_j} - W_{a_j})^2\Big] + 2\mathbb{E}\Big[\sum_{i<j} H_iH_j(W_{b_i} - W_{a_i})(W_{b_j} - W_{a_j})\Big].$$

The cross terms vanish, because when we condition on $\mathcal{F}_{a_j}$, we have

$$\mathbb{E}\,[H_iH_j(W_{b_i} - W_{a_i})\mathbb{E}\,[(W_{b_j} - W_{a_j}) \mid \mathcal{F}_{a_j}] = 0.$$

For the diagonal terms

$$\begin{aligned}\mathbb{E}\,[H_j^2(W_{b_j} - W_{a_j})^2] &= \mathbb{E}\,[H_j^2\mathbb{E}\,[(W_{b_j} - W_{a_j})^2 \mid \mathcal{F}_{a_j}]] \\ &= \mathbb{E}\,[H_j^2\mathbb{E}\,[W_{b_j}^2 - W_{a_j}^2 \mid \mathcal{F}_{a_j}]] \\ &= \mathbb{E}\,[H_j^2\mathbb{E}\,[b_j - a_j \mid \mathcal{F}_{a_j}]] \\ &= \mathbb{E}\,[H_j^2([b_j - a_j)].\end{aligned}$$

So $\mathbb{E}\, N_\infty^2 = \mathbb{E}\int_0^\infty H_s^2 ds$. □

Now suppose H_s is adapted and $\mathbb{E}\int_0^\infty H_s^2\, ds < \infty$. Let us define some norms; these are just for technical purposes and can be ignored if wished. Let

$$\|H_s\|_H = \Big(\mathbb{E}\int_0^\infty H_s^2 ds\Big)^{1/2}.$$

This is the L^2 norm on $[0,\infty)\times\Omega$ with respect to the measure $\nu_H(A) = \mathbb{E}\int_0^\infty 1_A ds$. Define another norm by

$$\|Y\|_S = (\mathbb{E}\,[\sup_t |Y_t|^2])^{1/2}.$$

Using some results from measure theory, we can choose H_s^n simple such that $\mathbb{E}\int_0^\infty (H_s^n - H_s)^2\,ds \to 0$. (This is the fact that simple functions are dense in L^2.) By Doob's inequality we have

$$\begin{aligned}\mathbb{E}\Big[\sup_t\Big(\int_0^t (H_s^n - H_s^m)\,dW_s\Big)^2\Big] &\le 4\mathbb{E}\Big(\int_0^\infty (H_s^n - H_s^m)\,dW_s\Big)^2\\ &= 4\mathbb{E}\int_0^\infty (H_s^n - H_s^m)^2\,ds \to 0.\end{aligned}$$

One can show that the norm $\|Y\|_S$ is complete, so there exists a process N_t such that $\sup_t[\int_0^t H_s^n\,dW_s - N_t] \to 0$ in L^2.

If H_s^n and $H_s^{n\prime}$ are two sequences converging to H, then $\mathbb{E}\,(\int_0^t (H_s^n - H_s^{n\prime})\,dW_s)^2 = \mathbb{E}\int_0^t (H_s^n - H_s^{n\prime})^2\,ds \to 0$, or the limit is independent of which sequence H^n we choose. It is easy to see, because of the L^2 convergence, that N_t is a martingale, $\mathbb{E}\,N_t^2 = \mathbb{E}\int_0^t H_s^2\,ds$, and $\langle N\rangle_t = \int_0^t H_s^2\,ds$. Because $\sup_t[\int_0^t H_s^n\,dW_s - N_t] \to 0$ in L^2, one can show there exists a subsequence such that the convergence takes place almost surely. So with probability one, N_t has continuous paths. We write $N_t = \int_0^t H_s\,dW_s$ and call N_t the stochastic integral of H with respect to W.

We discuss some extensions of the definition. First of all, if we replace W_t by a continuous martingale M_t and H_s is adapted with $\mathbb{E}\int_0^t H_s^2 d\langle M\rangle_s < \infty$, we can duplicate everything we just did with ds replaced by $d\langle M\rangle_s$ and get a stochastic integral. In particular, if $d\langle M\rangle_s = K_s^2 ds$, we replace ds by $K_s^2 ds$.

There are some other extensions of the definition that are not hard. If $\int_0^\infty H_s^2\langle M\rangle_s < \infty$ but without the expectation being finite, we can define the stochastic integral by looking at $M_{t\wedge T_N}$ for suitable stopping times T_N and then letting $T_N \to \infty$.

A process A_t is of bounded variation if the paths of A_t have bounded variation. This means that one can write $A_t = A_t^+ - A_t^-$, where A_t^+ and A_t^- have paths that are increasing. $|A|_t$ is then defined to be $A_t^+ + A_t^-$. A semimartingale is the sum of a martingale and a process of bounded variation. If $\int_0^\infty H_s^2\,d\langle M\rangle_s + \int_0^\infty |H_s|\,|dA_s| < \infty$ and $X_t = M_t + A_t$, we define

$$\int_0^t H_s\,dX_s = \int_0^t H_s\,dM_s + \int_0^t H_s\,dA_s,$$

where the first integral on the right is a stochastic integral and the second is a Riemann-Stieltjes or Lebesgue-Stieltjes integral. For a semimartingale, we define $\langle X\rangle_t = \langle M_t\rangle$.

Given two semimartingales X and Y we define $\langle X, Y\rangle_t$ by what is known as polarization:

$$\langle X, Y\rangle_t = \tfrac{1}{2}[\langle X+Y\rangle_t - \langle X\rangle_t - \langle Y\rangle_t].$$

As an example, if $X_t = \int_0^t H_s dW_s$ and $Y_t = \int_0^t K_s dW_s$, then $(X+Y)_t = \int_0^t (H_s + K_s) dW_s$, so

$$\langle X+Y \rangle_t = \int_0^t (H_s+K_s)^2 ds = \int_0^t H_s^2 ds + \int_0^t 2H_s K_s ds + \int_0^t K_s^2 ds.$$

Since $\langle X \rangle_t = \int_0^t H_s^2 ds$ with a similar formula for $\langle Y \rangle_t$, we conclude

$$\langle X, Y \rangle_t = \int_0^t H_s K_s ds.$$

What does a stochastic integral mean? If one thinks of the derivative of Z_t as being a white noise, then $\int_0^t H_s dZ_s$ is like a filter that increases or decreases the volume by a factor H_s.

For us, an interpretation is that Z_t represents a stock price. Then $\int_0^t H_s dZ_s$ represents our profit (or loss) if we hold H_s shares at time s. This can be seen most easily if $H_s = G1_{[a,b]}$. So we buy $G(\omega)$ shares at time a and sell them at time b. The stochastic integral represents our profit or loss.

Since we are in continuous time, we are allowed to buy and sell continuously and instantaneously. What we are not allowed to do is look into the future to make our decisions, which is where the H_s adapted condition comes in.

12. Ito's formula.

Suppose W_t is a Brownian motion and $f : \mathbb{R} \to \mathbb{R}$ is a C^2 function, that is, f and its first two derivatives are continuous. Ito's formula, which is sometime known as the change of variables formula, says that

$$f(W_t) - f(W_0) = \int_0^t f'(W_s) ds + \tfrac{1}{2} \int_0^t f''(W_s) ds.$$

Compare this with the fundamental theorem of calculus:

$$f(t) - f(0) = \int_0^t f'(s) ds.$$

In Ito's formula we have a second order term to carry along.

The idea behind the proof is quite simple. By Taylor's theorem.

$$\begin{aligned} f(W_t) - f(W_0) &= \sum_{i=0}^{n-1} [f(W_{(i+1)t/n}) - f(W_{it/n})] \\ &\approx \sum_{i=1}^{n-1} f'(W_{it/n})(W_{(i+1)t/n} - W_{it/n}) \\ &\qquad + \tfrac{1}{2} \sum_{i=0}^{n-1} f''(W_{it/n})(W_{(i+1)t/n} - W_{it/n})^2. \end{aligned}$$

The first sum on the right is approximately the stochastic integral and the second is approximately the quadratic variation.

For a more general semimartingale $X_t = M_t + A_t$, Ito's formula reads

Theorem 12.1. *If $f \in C^2$, then*

$$f(X_t) - f(X_0) = \int_0^t f'(X_s) dXs + \tfrac{1}{2} \int_0^t f''(X_s) d\langle M \rangle_s.$$

Let us look at an example. Let W_t be Brownian motion, $X_t = \sigma W_t - \sigma^2 t/2$, and $f(x) = e^x$. Then $\langle X \rangle_t = \langle \sigma W \rangle_t = \sigma^2 t$ and

$$\begin{aligned} e^{\sigma W_t - \sigma^2 t/2} &= 1 + \int_0^t e^{X_s} \sigma dW_s - \tfrac{1}{2} \int_0^t e^{X_s} \tfrac{1}{2} \sigma^2 ds \\ &\quad + \tfrac{1}{2} \int_0^t e^{X_s} \sigma^2 ds \\ &= 1 + \int_0^t e^{X_s} \sigma dW_s. \end{aligned} \tag{12.1}$$

For a semimartingale $X_t = M_t + A_t$ we set $\langle X \rangle_t = \langle M \rangle_t$. Given two semimartingales X, Y, we define

$$\langle X, Y \rangle_t = \tfrac{1}{2} [\langle X + Y \rangle_t - \langle X \rangle_t - \langle Y \rangle_t].$$

Proposition 12.2.

$$X_t Y_t = X_0 Y_0 + \int_0^t X_s dY_s + \int_0^t Y_s dX_s + \langle X, Y \rangle_t.$$

Proof. Applying Ito's formula with $f(x) = x^2$ to $X_t + Y_t$, we obtain

$$(X_t + Y_t)^2 = (X_0 + Y_0)^2 + 2 \int_0^t (X_s + Y_s)(dX_s + dY_s) + \langle X + Y \rangle_t.$$

Applying Ito's formula with $f(x) = x^2$ to X and to Y, then

$$X_t^2 = X_0^2 + 2 \int_0^t X_s dX_s + \langle X \rangle_t$$

and

$$Y_t^2 = Y_0^2 + 2 \int_0^t Y_s dY_s + \langle Y \rangle_t.$$

Then some algebra and the fact that

$$X_t Y_t = \tfrac{1}{2}[(X_t + Y_t)^2 - X_t^2 - Y_t^2]$$

yields the formula. $\square$

There is a multidimensional version of Ito's formula: if $X_t = (X_t^1, \ldots, X_t^d)$ is a vector, each component of which is a semimartingale, and $f \in C^2$, then

$$f(X_t) - f(X_0) = \sum_{i=1}^{d} \int_0^t \frac{\partial f}{\partial x_i}(X_s) dX_s^i) + \tfrac{1}{2} \sum_{i,j=1}^{d} \int_0^t \frac{\partial^2 f}{\partial x_i^2}(X_s) d\langle X^i, X^j \rangle_s.$$

The following application of Ito's formula, known as Lévy's theorem, is important.

Theorem 12.3. *Suppose M_t is a continuous martingale with $\langle M \rangle_t = t$. Then M_t is a Brownian motion.*

Before proving this, recall from undergraduate probability that the moment generating function of a r.v. X is defined by $m_X(a) = \mathbb{E}\, e^{aX}$ and that if two random variables have the same moment generating function, they have the same law. This is also true if we replace a by iu. In this case we have $\varphi_X(u) = \mathbb{E}\, e^{iuX}$ and φ_X is called the characteristic function of X. The reason for looking at the characteristic function is that φ_X always exists, whereas $m_X(a)$ might be infinite. The one special case we will need is that if X is a normal r.v. with mean 0 and variance t, then $\varphi_X(u) = e^{-u^2 t/2}$. This follows from the formula for $m_X(a)$ with a replaced by iu (this can be justified rigorously).

Proof. Apply Ito's formula with $f(x) = e^{iux}$. Then

$$e^{iuM_t} = 1 + \int_0^t iu e^{iuM_s} dM_s + \tfrac{1}{2} \int_0^t (-u^2) e^{iuM_s} d\langle M \rangle_s.$$

Taking expectations and using $\langle M \rangle_s = s$ and the fact that a stochastic integral is a martingale, hence has 0 expectation, we have

$$\mathbb{E}\, e^{iuM_t} = 1 - \frac{u^2}{2} \int_0^t e^{iuM_s} ds.$$

Let $J(t) = \mathbb{E}\, e^{iuM_t}$. The equation can be rewritten

$$J(t) = 1 - \frac{u^2}{2} \int_0^t J(s) ds.$$

So $J'(t) = -\frac{1}{2} u^2 J(t)$ with $J(0) = 1$. The solution to this elementary ODE is $J(t) = e^{-u^2 t/2}$, which shows that M_t is a mean 0 variance t normal r.v.

If $A \in \mathcal{F}_s$ and we do the same argument with M_t replaced by $M_{s+t} - M_s$, we have

$$e^{iu(M_{s+t}-M_s)} = 1 + \int_0^t iue^{iu(M_{s+r}-M_s)} dM_r + \tfrac{1}{2}\int_0^t (-u^2) e^{iu(M_{s+r}-M_s)} d\langle M\rangle_r.$$

Multiply this by 1_A and take expectations. Since a stochastic integral is a martingale, the stochastic integral term again has expectation 0. If we let $K(t) = \mathbb{E}\,[e^{iu(M_{t+s}-M_t)}; A]$, we now arrive at $K'(t) = -\frac{1}{2}u^2 K(t)$ with $K(0) = \mathbb{P}(A)$, so

$$K(t) = \mathbb{P}(A) e^{-u^2 t/2}.$$

Therefore

$$\mathbb{E}\left[e^{iu(M_{t+s}-M_s)}; A\right] = \mathbb{E}\, e^{iu(M_{t+s}-M_s)}\mathbb{P}(A). \tag{12.2}$$

If f is a nice function and $\widehat{f}$ is its Fourier transform, replace u in the above by $-u$, multiply by $\widehat{f}(u)$, and integrate over u. (To do the integral, we approximate the integral by a Riemann sum and then take limits.) We then have

$$\mathbb{E}\,[f(M_{s+t} - M_s); A] = \mathbb{E}\,[f((M_{s+t} - M_s)]\mathbb{P}(A).$$

By taking limits we have this for $f = 1_B$, so

$$\mathbb{P}(M_{s+t} - M_s \in B, A) = \mathbb{P}(M_{s+t} - M_s \in B)\mathbb{P}(A).$$

This implies that $M_{s+t} - M_s$ is independent of $\mathcal{F}_s$.

Note $\mathrm{Var}\,(M_t - M_s) = t - s$; take $A = \Omega$ in (12.2). □

13. The Girsanov theorem.

Suppose $\mathbb{P}$ is a probability measure and

$$dX_t = dW_t + \mu(X_t)dt.$$

Let

$$M_t = \exp\Big(- \int_0^t \mu(X_s)dW_s - \int_0^t \mu(X_s)^2 ds/2\Big).$$

Then as we have seen before, by Ito's formula, M_t is a martingale.

Now let us define a new probability by setting

$$\mathbb{Q}(A) = \mathbb{E}\,[M_t; A] \tag{13.1}$$

if $A \in \mathcal{F}_t$. We had better be sure this $\mathbb{Q}$ is well defined. If $A \in \mathcal{F}_s \subset \mathcal{F}_t$, then $\mathbb{E}\,[M_t; A] = \mathbb{E}\,[M_s; A]$ because M_t is a martingale.

What the Girsanov theorem says is

Theorem 13.1. *Under $\mathbb{Q}$, X_t is a Brownian motion.*

There is a more general version.

Theorem 13.2. *If X_t is a martingale under $\mathbb{P}$, then under $\mathbb{Q}$ the process $X_t - D_t$ is a martingale where*

$$D_t = \int_0^t \frac{1}{M_s} d\langle X, M\rangle_s.$$

$\langle X\rangle_t$ is the same under both $\mathbb{P}$ and $\mathbb{Q}$.

Let us see how Theorem 13.1 can be used. Let S_t be the stock price, and suppose

$$dS_t = \sigma S_t dW_t + \mu S_t dt.$$

Define

$$M_t = e^{(-\mu/\sigma)(W_t) - (\mu^2/2\sigma^2)t}.$$

Then from (12.1) M_t is a martingale and

$$M_t = 1 + \int_0^t \Big(-\frac{\mu}{\sigma}\Big) M_s dW_s.$$

Let $X_t = W_t$. Then

$$\langle X, M\rangle_t = \int_0^t \Big(-\frac{\mu}{\sigma}\Big) M_s ds = -\int_0^t M_s \frac{\mu}{\sigma} ds.$$

Therefore

$$\int_0^t \frac{1}{M_s} d\langle X, M\rangle_s = -\int_0^t \frac{\mu}{\sigma} ds = -(\mu/\sigma)t.$$

Define $\mathbb{Q}$ by (13.1). By Theorem 13.2, under $\mathbb{Q}$ the process $\widetilde{W}_t = W_t + (\mu/\sigma)t$ is a martingale. Hence

$$dS_t = \sigma S_t (dW_t + (\mu/\sigma)dt) = \sigma S_t d\widetilde{W}_t,$$

or

$$S_t = S_0 + \int_0^t \sigma S_s d\widetilde{W}_s$$

is a martingale. So we have found a probability under which the asset price is a martingale. This means that $\mathbb{Q}$ is the risk-neutral probability, which we have been calling $\overline{\mathbb{P}}$.

Let us give another example of the use of the Girsanov theorem. Suppose $X_t = W_t + \mu t$. We want to compute the probability that X_t exceeds the level a by time t_0.

We first need the probability that a Brownian motion crosses a level a by time t_0. Any path that crosses a but is at level $x < a$ at time t_0 has a corresponding path determined by reflecting across level a at the first time the Brownian motion hits a; the reflected path will end up at $a+(a-x)=2a-x$. This is known as the reflection principle, and can be written, informally, by

$$\mathbb{P}(\sup_{s\le t_0} W_s \ge a, W_{t_0}=x)=\mathbb{P}(W_{t_0}=2a-x).$$

Now let W_t be a Brownian motion under $\mathbb{P}$. Let $d\mathbb{Q}/d\mathbb{P}=M_t=e^{\mu W_t-\mu^2 t/2}$. Let $Y_t=W_t-\mu t$. Theorem 13.1 says that under $\mathbb{Q}$, Y_t is a Brownian motion. We have $W_t=Y_t+\mu t$.

Let $A=(\sup_{s\le t_0} W_s\ge a)$. We want to calculate

$$\mathbb{P}(\sup_{s\le t_0}(W_s+\mu s)\ge a).$$

W_t is a Brownian motion under $\mathbb{P}$ while Y_t is a Brownian motion under $\mathbb{Q}$. So this probability is equal to

$$\mathbb{Q}(\sup_{s\le t_0}(Y_s+\mu s)\ge a).$$

This in turn is equal to

$$\mathbb{Q}(\sup_{s\le t_0} W_s\ge a)=\mathbb{Q}(A).$$

Now we use the expression for M_t:

$$\begin{aligned}\mathbb{Q}(A)&=\mathbb{E}_{\mathbb{P}}[e^{\mu W_{t_0}-\mu^2 t_0/2};A]\\&=\int_{-\infty}^{\infty} e^{\mu x-\mu^2 t_0/2}\mathbb{P}(\sup_{s\le t_0} W_s\ge a, W_{t_0}=x)dx\\&=e^{-\mu^2 t_0/2}\Big[\int_{-\infty}^{a}\frac{1}{\sqrt{2\pi t_0}}e^{\mu x}e^{-(2a-x)^2/2t_0}\,dx+\int_a^\infty \frac{1}{\sqrt{2\pi t_0}}e^{\mu x}e^{-x^2/2t_0}\,dx.\Big]\end{aligned}$$

Now for the proofs of Theorems 13.1 and 13.2.

Proof of Theorem 13.2. Assume without loss of generality that $X_0=0$. Then if $A\in\mathcal{F}_s$,

$$\begin{aligned}\mathbb{E}_{\mathbb{Q}}[X_t;A]&=\mathbb{E}_{\mathbb{P}}[M_tX_t;A]\\&=\mathbb{E}_P\Big[\int_0^t M_r\,dX_r;A\Big]+\mathbb{E}_{\mathbb{P}}\Big[\int_0^t X_r\,dM_r;A\Big]+\mathbb{E}_{\mathbb{P}}[\langle X,M\rangle_t;A]\\&=\mathbb{E}_P\Big[\int_0^s M_r\,dX_r;A\Big]+\mathbb{E}_{\mathbb{P}}\Big[\int_0^s X_r\,dM_r;A\Big]+\mathbb{E}_{\mathbb{P}}[\langle X,M\rangle_t;A]\\&=\mathbb{E}_{\mathbb{Q}}[X_s;A]+\mathbb{E}\,[\langle X,M\rangle_t-\langle X,M\rangle_s;A].\end{aligned}$$

Here we used the fact that stochastic integrals with respect to the martingales X and M are again martingales.

On the other hand,

$$\begin{aligned}\mathbb{E}_P[\langle X,M\rangle_t - \langle X,M\rangle_s; A] &= \mathbb{E}_{\mathbb{P}}\Big[\int_s^t d\langle X,M\rangle_r; A\Big]\\ &= \mathbb{E}_{\mathbb{P}}\Big[\int_s^t M_r\, dD_r; A\Big]\\ &= \mathbb{E}_{\mathbb{P}}\Big[\int_s^t \mathbb{E}_{\mathbb{P}}[M_t \mid \mathcal{F}_r]\, dD_t; A\Big]\\ &= \mathbb{E}_{\mathbb{P}}\Big[\int_s^t M_t\, dD_r; A\Big]\\ &= \mathbb{E}_{\mathbb{P}}[(D_t - D_s)M_t; A]\\ &= \mathbb{E}_{\mathbb{Q}}[D_t - D_s; A].\end{aligned}$$

The quadratic variation proof is similar. □

Proof of Theorem 13.1. From our formula for M we have $dM_t = -M_t\mu(X_t)dW_t$, so $d\langle X,M\rangle_t = -M_t\mu(X_t)dt$. Hence by Theorem 13.2 we see that under $\mathbb{Q}$, X_t is a continuous martingale with $\langle X\rangle_t = t$. By Lévy's theorem, this means that X is a Brownian motion under $\mathbb{Q}$. □

To help understand what is going on, let us give another proof of Theorem 13.1 along the lines of the proof of Theorem 13.2.

Proof of Theorem 13.1, second version. Using Ito's formula with $f(x) = e^x$,

$$M_t = 1 - \int_0^t \mu(X_r)M_r dW_r.$$

So

$$\langle W,M\rangle_t = -\int_0^t \mu(X_r)M_r\, dr.$$

Since $\mathbb{Q}(A) = \mathbb{E}_{\mathbb{P}}[M_t; A]$, it is not hard to see that

$$\mathbb{E}_{\mathbb{Q}}[W_t; A] = \mathbb{E}_{\mathbb{P}}[M_t W_t; A].$$

By Ito's product formula this is

$$\mathbb{E}_{\mathbb{P}}\Big[\int_0^t M_r dW_r; A\Big] + \mathbb{E}_{\mathbb{P}}\Big[\int_0^t W_r dM_r; A\Big] + \mathbb{E}_{\mathbb{P}}\Big[\langle W,M\rangle_t; A\Big].$$

Since $\int_0^t M_r dW_r$ and $\int_0^t W_r dM_r$ are stochastic integrals with respect to martingales, they are themselves martingales. Thus the above is equal to

$$\mathbb{E}_{\mathbb{P}}\Big[\int_0^s M_r dW_r; A\Big] + \mathbb{E}_{\mathbb{P}}\Big[\int_0^s W_r dM_r; A\Big] + \mathbb{E}_{\mathbb{P}}\Big[\langle W,M\rangle_t; A\Big].$$

Using the product formula again, this is

$$\mathbb{E}_{\mathbb{P}}[M_s W_s; A] + \mathbb{E}_{\mathbb{P}}[\langle W, M\rangle_t - \langle W, M\rangle_s; A] = \mathbb{E}_{\mathbb{Q}}[W_s; A] + \mathbb{E}_{\mathbb{P}}[\langle W, M\rangle_t - \langle W, M\rangle_s; A].$$

The last term on the right is equal to

$$\begin{aligned}\mathbb{E}_{\mathbb{P}}\Big[\int_s^t d\langle W, M\rangle_r; A\Big] &= \mathbb{E}_{\mathbb{P}}\Big[-\int_s^t M_r \mu(X_r) dr; A\Big] = \mathbb{E}_{\mathbb{P}}\Big[-\int_s^t \mathbb{E}\,[M_t \mid \mathcal{F}_r]\mu(X_r) dr; A\Big]\\ &= \mathbb{E}_{\mathbb{P}}\Big[-\int_s^t M_t \mu(X_r) dr; A\Big] = \mathbb{E}_{\mathbb{Q}}\Big[-\int_s^t \mu(X_r) dr; A\Big]\\ &= -\mathbb{E}_{\mathbb{Q}}\Big[\int_0^t \mu(X_r)\, dr; A\Big] + \mathbb{E}_{\mathbb{Q}}\Big[\int_0^s \mu(X_r)\, dr; A\Big].\end{aligned}$$

Therefore

$$\mathbb{E}_{\mathbb{Q}}\Big[W_t + \int_0^t \mu(X_r) dr; A\Big] = \mathbb{E}_{\mathbb{Q}}\Big[W_s + \int_0^s \mu(X_r) dr; A\Big],$$

which shows X_t is a martingale with respect to $\mathbb{Q}$.

Similarly, $X_t^2 - t$ is a martingale with respect to $\mathbb{Q}$. By Lévy's theorem, X_t is a Brownian motion. $\square$

14. Stochastic differential equations.

Let W_t be a Brownian motion. We are interested in the existence and uniqueness for stochastic differential equations (SDEs) of the form

$$dX_t = \sigma(X_t)\, dW_t + b(X_t)\, dt, \qquad X_0 = 0.$$

This means X_t satisfies

$$X_t = x_0 + \int_0^t \sigma(X_s)\, dW_s + \int_0^t b(X_s)\, ds. \tag{14.1}$$

Here W_t is a Brownian motion, and (14.1) holds for almost every ω.

We have to make some assumptions on σ and b. We assume they are Lipschitz, which means:

$$|\sigma(x) - \sigma(y)| \le c|x-y|, \qquad |b(x) - b(y)| \le c|x-y|$$

for some constant c. We also suppose that σ and b grow at most linearly, which means:

$$|\sigma(x)| \le c(1+|x|), \qquad |b(x)| \le c(1+|x|).$$

Theorem 14.1. *There exists one and only one solution to (14.1).*

The idea of the proof is Picard iteration, which is how existence and uniqueness for ordinary differential equations is proved. Let us illustrate the uniqueness part, and for simplicity, assume b is identically 0.

Proof of uniqueness. If X and Y are two solutions,

$$X_t - Y_t = \int_0^t [\sigma(X_s) - \sigma(Y_s)] dW_s.$$

So

$$\mathbb{E}\,|X_t - Y_t|^2 = \mathbb{E}\int_0^t |\sigma(X_s) - \sigma(Y_s)|^2 ds \le c\int_0^t \mathbb{E}\,|X_s - Y_s|^2 ds,$$

using the Lipschitz hypothesis on σ. If we let $g(t) = \mathbb{E}\,|X_t - Y_t|^2$, we have

$$g(t) \le c\int_0^t g(s)\,ds.$$

Then

$$g(t) \le c\int_0^t \left[c\int_0^s g(r)\,dr\right] ds.$$

g is easily seen to be bounded on finite intervals, and iteration implies

$$g(t) \le At^n/n!$$

for each n, which implies g must be 0. □

The above theorem also works in higher dimensions. We want to solve

$$dX_t^i = \sum_{j=1}^d \sigma_{ij}(X_s) dW_s^j + b_i(X_s) ds, \qquad i = 1, \ldots, d.$$

If all of the σ_{ij} and b_i are Lipschitz and grow at most linearly, we have uniqueness for the solution.

Suppose one wants to solve

$$dZ_t = aZ_t\,dW_t + bZ_t\,dt.$$

Note that this equation is linear in Z_t, and it turns out that linear equations are almost the only ones that have an explicit solution. In this case we can write down the explicit solution and then verify that it satisfies the SDE. The uniqueness result shows that we have in fact found the solution.

Let

$$Z_t = Z_0 e^{aW_t - a^2 t/2 + bt}.$$

We will verify that this is correct by using Ito's formula. Let $X_t = aW_t - a^2t/2 + bt$. Then X_t is a semimartingale with martingale part aW_t and $\langle X\rangle_t = a^2 t$. $Z_t = e^{X_t}$. By Ito's formula with $f(x) = e^x$,

$$\begin{aligned} Z_t &= Z_0 + \int_0^t e^{X_s} dX_s + \tfrac{1}{2}\int_0^t e^{X_s} a^2 ds \\ &= Z_0 + \int_0^t aZ_s dW_s - \int_0^t \frac{a^2}{2} Z_s ds + \int_0^t b\, ds \\ &\qquad + \tfrac{1}{2}\int_0^t a^2 Z_s ds \\ &= \int_0^t aZ_s dW_s + \int_0^t bZ_s ds. \end{aligned}$$

This is the integrated form of the equation we wanted to solve.

If we let X_t^x denote the solution to

$$X_t^x = x + \int_0^t \sigma(X_s^x) dW_s + \int_0^t b(X_s^x) ds,$$

so that X_t^x is the solution of the SDE started at x, we can define new probabilities by

$$\mathbb{P}^x(X_{t_1} \in A_1, \ldots, X_{t_n} \in A_n) = \mathbb{P}(X_{t_1}^x \in A_1, \ldots, X_{t_n}^x \in A_n).$$

This is similar to what we did in defining $\mathbb{P}^x$ for Brownian motion, but here we do not have translation invariance. One can show that when there is uniqueness, the family $(\mathbb{P}^x, X_t)$ satisfies the strong Markov property.

15. Continuous time financial models.

The most common model by far in finance is that the security price is based on a Brownian motion. One does not want to say the price is some multiple of Brownian motion for two reasons. First, of all, a Brownian motion can become negative, which doesn't make sense for stock prices. Second, if one invests \$1,000 in a stock selling for \$1 and it goes up to \$2, one has the same profit as if one invests \$1,000 in a stock selling for \$100 and it goes up to \$200. It is the proportional increase one wants.

Therefore one sets $\Delta S_t/S_t$ to be the quantity related to a Brownian motion. Different stocks have different volatilities σ (consider a high-tech stock versus a pharmaceutical).

In addition, one expects a mean rate of return μ on ones investment that is positive (otherwise, why not just put the money in the bank?). In fact, one expects the mean rate of return to be higher than the risk-free interest rate r because one expects something in return for undertaking risk.

So the model that is used is to let the stock price be modeled by the SDE

$$dS_t/S_t = \sigma dW_t + \mu dt,$$

or what looks better,

$$dS_t = \sigma S_t dW_t + \mu S_t dt. \tag{15.1}$$

Fortunately this SDE is one of those that can be solved explicitly.

Proposition 15.1. *The solution to (15.1) is given by*

$$S_t = S_0 e^{\sigma W_t + (\mu - (\sigma^2/2)t)}. \tag{15.2}$$

Proof. Using Theorem 14.1 there will only be one solution, so we need to verify that S_t as given in (15.2) satisfies (15.1). We already did this, but it is important enough that we will do it again. Let us first assume $S_0 = 1$. Let $X_t = \sigma W_t + (\mu - (\sigma^2/2)t$, let $f(x) = e^x$, and apply Ito's formula. We obtain

$$\begin{aligned} S_t = e^{X_t} &= e^{X_0} + \int_0^t e^{X_s} dX_s + \tfrac{1}{2}\int_0^t e^{X_s} d\langle X\rangle_s \\ &= 1 + \int_0^t S_s \sigma dW_s + \int_0^t S_s(\mu - \tfrac{1}{2}\sigma^2) ds \\ &\qquad + \tfrac{1}{2}\int_0^t S_s \sigma^2 ds \\ &= 1 + \int_0^t S_s \sigma dW_s + \int_0^t S_s \mu ds, \end{aligned}$$

which is (15.1). If $S_0 \neq 0$, just multiply both sides by S_0. □

If one purchases Δ_0 shares (possibly a negative number) at time t_0, then changes the investment to Δ_1 shares at time t_1, etc., then ones wealth at time t will be

$$X_{t_0} + \Delta_0(S_{t_1} - S_{t_0}) + \Delta_1(S_{t_2} - S_{t_1}) + \cdots + \Delta_i(S_{t_{i+1}} - S_{t_i}).$$

But this is the same as

$$X_{t_0} + \int_0^t \Delta(s) dS_s,$$

where we have $t \geq t_{i+1}$ and $\Delta(s) = \Delta_i$ if $t_i \leq s < t_{i+1}$. In other words, our wealth is given by a stochastic integral with respect to the stock price. The requirement that the integrand of a stochastic integral be adapted is very natural: we cannot base the number of shares we own at time s on information that will not be available until the future.

The continuous time model of finance is then that the security price is given by (15.1) (often called geometric Brownian motion), that there are no transaction costs, but one can trade as many shares as one wants and vary the amount held in a continuous fashion. This clearly is not the way the market actually works, for example, stock prices are discrete, but this model has proved to be a very good one.

16. Martingale representation theorem.

In this section we want to show that every random variable that is $\mathcal{F}_t$ measurable can be written as a stochastic integral of Brownian motion. In the next section we use this to show that under the model of geometric Brownian motion the market is complete. This means that no matter what option one comes up with, one can exactly replicate the result (no matter what the market does) by buying and selling shares of stock.

In mathematical terms, we let $\mathcal{F}_t$ be the σ-field generated by $W_s, s \leq t$. From (15.2) we see that $\mathcal{F}_t$ is also the same as the σ-field generated by $S_s, s \leq t$, so it doesn't matter which one we work with. We want to show that if V is $\mathcal{F}_t$ measurable, then there exists H_s adapted such that

$$V = V_0 + \int H_s dW_s, \tag{16.1}$$

where V_0 is a constant.

We first need the following.

Proposition 16.1. *Suppose*

$$V_t^n = V_0^n + \int_0^t H_s^n dW_s$$

and

$$\mathbb{E}\,|V_t^n - V_t|^2 \to 0$$

for each t with the H^n adapted. Then there exists H_s adapted so that

$$V_t = V_0 + \int_0^t H_s dW_s.$$

What this proposition says is that if we can duplicate a sequence of options V_n and $V_n \to V$, then we can duplicate V.

Proof. By our assumptions,

$$\mathbb{E}\,|(V_t^n - V_0^n) - (V_t^m - V_0^m)|^2 \to 0$$

as $n, m \to \infty$. So

$$\mathbb{E}\Big|\int_0^t (H_s^n - H_s^m) dW_s\Big|^2 \to 0.$$

From our formulas for stochastic integrals, this means

$$\mathbb{E}\int_0^t |H_s^n - H_s^m|^2 ds \to 0.$$

This says that H_s^n is a Cauchy sequence in the space L^2 (with respect to the norm $\|\cdot\|_H$ given in Section 11). We will assume that you know that L^2 is complete or are willing to believe it, so there exists H_s such that

$$\mathbb{E}\int_0^t |H_s^n - H_s|^2 ds \to 0.$$

In particular $H_s^n \to H_s$, and this implies H_s is adapted.

Let $U_t = \int_0^t H_s dW_s$. Then as above,

$$\mathbb{E}\,|(V_t^n - V_0^n) - U_t|^2 = \mathbb{E}\int_0^t (H_s^n - H_s)^2 ds \to 0.$$

Therefore $U_t = V_t - V_0$, and U has the desired form. □

Next we show our result for a particular collection of options.

Proposition 16.2. *If g is bounded,*

$$g(W_t) = c + \int_0^t H_s dW_s$$

for an integrand H_s that is adapted and some constant c.

Proof. By Ito's formula with $X_s = -iuW_s + u^2 s/2$ and $f(x) = e^x$,

$$\begin{aligned} e^{X_t} &= 1 + \int_0^t e^{X_s}(-iu) dW_s + \int_0^t e^{X_s}(u^2/2) ds \\ &\quad + \tfrac{1}{2}\int_0^t e^{X_s}(-iu)^2 ds \\ &= 1 - iu\int_0^t e^{X_s} dW_s. \end{aligned}$$

If we multiply both sides by $e^{-u^2t/2}$, which is a constant and hence adapted, we obtain

$$e^{-iuW_t} = c_u + \int_0^t H_s^u dW_s \tag{16.2}$$

for an appropriate constant c_u and integrand H^u.

If f is a smooth function (e.g., C^∞ with compact support), then its Fourier transform $\widehat{f}$ will also be very nice. So if we multiply (16.2) by $\widehat{f}(u)$ and integrate over u from $-\infty$ to ∞, we obtain

$$f(W_t) = c + \int_0^t H_s dW_s$$

for some constant c and some adapted integrand H. (We implicitly used Proposition 16.1, because we approximate our integral by Riemann sums, and then take a limit.) Now using Proposition 16.1 we take limits and obtain the proposition. $\square$

An almost identical proof shows that if f is bounded, then

$$f(W_t - W_s) = c + \int_s^t H_r dW_r$$

for some c and H_r.

Theorem 16.3. *If V is $\mathcal{F}_t$ measurable and $\mathbb{E}\,V^2 < \infty$, then there exists a constant c and an adapted integrand H_s such that*

$$V = c + \int_0^t H_s dW_s.$$

Proof. We will show that all V of the form

$$V = f_1(W_{t_1} - W_{t_0}) f_2(W_{t_2} - W_{t_1}) \cdots f_n(W_{t_n} - W_{t_{n-1}})$$

can be so represented. If we take linear combinations of these, then the linear combinations can also be so represented. Since every V that is $\mathcal{F}_t$ measurable can be written as a limit of such linear combinations, Proposition 16.1 then implies the result.

The argument is by induction; let us do the case $n = 2$ for clarity. So we suppose

$$V = f(W_t) g(W_u - W_t).$$

From Proposition 16.2 we now have that

$$f(W_t) = c + \int_0^t H_s dW_s, \qquad g(W_u - W_t) = d + \int_t^u K_s dW_s.$$

Set $\overline{H}_r = H_r$ if $0 \le s < t$ and 0 otherwise. Set $\overline{K}_r = K_r$ if $s \le r < t$ and 0 otherwise. Let $X_s = c + \int_0^s \overline{H}_r dW_r$ and $Y_s = d + \int_0^s \overline{K}_r dW_r$. Then

$$\langle X, Y \rangle_s = \int_0^s \overline{H}_r \overline{K}_r dr = 0.$$

Then by the Ito product formula,

$$\begin{aligned} X_s Y_s &= X_0 Y_0 + \int_0^s X_r dY_r + \int_0^s Y_r dX_r \\ &\quad + \langle X, Y \rangle_s \\ &= cd + \int_0^s [X_r \overline{K}_r + Y_r \overline{H}_r] dW_r. \end{aligned}$$

If we now take $s = u$, that is exactly what we wanted. Note that $X_r \overline{K}_r + Y_r \overline{H}_r$ is 0 if $r > u$; this is needed to do the general induction step. □

17. Completeness.

Now let S_t be a geometric Brownian motion. As we mentioned in the last section, if $S_t = S_0 \exp(\sigma W_t + (\mu - \sigma^2/2)t)$, then given S_t we can determine W_t and vice versa, so the σ fields generated by S_t and W_t are the same. Recall S_t satisfies

$$dS_t = \sigma S_t dW_t + \mu S_t dt.$$

Define a new probability $\overline{\mathbb{P}}$ by

$$\frac{d\overline{\mathbb{P}}}{d\mathbb{P}} = M_t = \exp(aW_t - a^2 t/2).$$

By the Girsanov theorem,

$$\widetilde{W}_t = W_t - at$$

is a Brownian motion under $\overline{\mathbb{P}}$. So

$$dS_t = \sigma S_t d\widetilde{W}_t + \sigma S_t a dt + \mu S_t dt.$$

If we choose $a = -\mu/\sigma$, we then have

$$dS_t = \sigma S_t d\widetilde{W}_t. \tag{17.1}$$

Since $\widetilde{W}_t$ is a Brownian motion under $\overline{\mathbb{P}}$, then S_t must be a martingale, since it is a stochastic integral of a Brownian motion. We can rewrite (17.1) as

$$d\widetilde{W}_t = \sigma^{-1} S_t^{-1} dS_t. \tag{17.2}$$

Given a $\mathcal{F}_t$ measurable variable V, we know by Theorem 16.3 that there exists adapted H_s such that

$$V = c + \int_0^t H_s d\widetilde{W}_s.$$

But then using (17.2) we have

$$V = c + \int_0^t H_s \sigma^{-1} S_s^{-1} dS_s.$$

We have therefore proved

Theorem 17.1. *If S_t is a geometric Brownian motion and V is $\mathcal{F}_t$ measurable, then there exist a constant c and an adapted process K_s such that*

$$V = c + \int_0^t K_s dS_s.$$

Moreover, there is a probability $\overline{\mathbb{P}}$ under which S_t is a martingale.

The probability $\overline{\mathbb{P}}$ is called the risk-neutral measure. Under $\overline{\mathbb{P}}$ the stock price is a martingale.

18. Black-Scholes formula, I.

We can now derive the formula for the price of any option. If V is $\mathcal{F}_t$ measurable, we have by Theorem 17.1 that

$$V = c + \int_0^t K_s dS_s, \tag{18.1}$$

and under $\overline{\mathbb{P}}$, the process S_s is a martingale.

Theorem 18.1. *The price of V must be $\overline{\mathbb{E}} V$.*

Proof. This is the "no arbitrage" principle again. Suppose the price of the option V at time 0 is W. Starting with 0 dollars, we can sell the option V for W dollars, and use the W dollars to buy and trade shares of the stock. In fact, if we use c of those dollars, and invest according to the strategy of holding K_s shares at time s, then at time t we will have

$$e^{rt}(W_0 - c) + V$$

dollars. At time t the buyer of our option exercises it and we use V dollars to meet that obligation. That leaves us a profit of $e^{rt}(W_0 - c)$ if $W_0 > c$, without any risk. Therefore W_0 must be less than or equal to c. If $W_0 < c$, we just reverse things: we buy the option instead of sell it, and hold $-K_s$ shares of stock at time s. By the same argument, since we can't get a riskless profit, we must have $W_0 \geq c$, or $W_0 = c$.

Finally, under $\overline{\mathbb{P}}$ the process S_t is a martingale. So taking expectations in (18.1), we obtain

$$\overline{\mathbb{E}}V = c.$$

$\square$

Note that there is a slight difference with the approach we used in the binomial asset model. There we showed that $(1+r)^{-k}S_k$ was a martingale under $\overline{\mathbb{P}}$. We could do the analogue to that here, but it is slightly simpler to make S_t a martingale under $\overline{\mathbb{P}}$ and to incorporate the interest rate into the definition of V. So, for example, in the case of pricing a European call, we let

$$V = e^{-rt}(S_t - K)^+.$$

We can think of this as saying the value of V at time t is $(S_t - K)^+$ in terms of the value of the dollar at time t. In terms of present day dollars the value of V is $e^{-rt}(S_t - K)^+$.

The formula in the statement of Theorem 18.1. is amenable to calculation. Suppose we have the standard European option, where $V = e^{-rt}(S_t - K)^+$. Recall that under $\overline{\mathbb{P}}$ the stock price satisfies

$$dS_t = \sigma S_t d\widetilde{W}_t,$$

where $\widetilde{W}_t$ is a Brownian motion under $\overline{\mathbb{P}}$. So then

$$S_t = S_0 e^{\sigma \widetilde{W}_t - \sigma^2 t/2}.$$

Hence

$$\begin{aligned}\overline{\mathbb{E}}V &= \overline{\mathbb{E}}[e^{-rt}(S_t - K)^+] \\ &= \mathbb{E}\left[e^{-rt}[S_0 e^{\sigma \widetilde{W}_t - (\sigma^2/2)t} - K]^+\right].\end{aligned} \tag{18.2}$$

We know the density of $\widetilde{W}_t$ is just $(2\pi t)^{-1}e^{-y^2/2}$, so we can do the calculations and end up with the famous Black-Scholes formula:

$$W_0 = x\Phi(g(x,t)) - Ke^{-rt}\Phi(h(x,t)),$$

where $\Phi(z) = \frac{1}{\sqrt{2\pi}}\int_{-\infty}^{z} e^{-y^2/2}\,dy$, $x = S_0$,

$$g(x,t) = \frac{\log(x/K) + (r + \sigma^2/2)t}{\sigma\sqrt{t}},$$

$$h(x,t) = g(x,t) - \sigma\sqrt{t}.$$

It is of considerable interest that the final formula depends on σ but is completely independent of μ. The reason for that can be explained as follows. Under $\overline{\mathbb{P}}$ the process S_t satisfies $dS_t = \sigma S_t d\widetilde{W}_t$, where $\widetilde{W}_t$ is a Brownian motion. Therefore, similarly to formulas we have already done,

$$S_t = S_0 e^{\sigma \widetilde{W}_t - \sigma^2/2},$$

and there is no μ present here. (We used the Girsanov formula to get rid of the μ.) The price of the option V is

$$\overline{\mathbb{E}} e^{-rt}[S_t - K]^+, \tag{18.3}$$

which is independent of μ since S_t is. Therefore we can use this instead of (18.2), i.e., we can assume μ is zero, and the calculations become much simpler.

19. Black-Scholes formula, II.

Here is a second approach to the Black-Scholes formula. This approach works for European calls and several other options, but does not work in the generality that the first approach does. On the other hand, it allows one to computes what the equivalent strategy of buying or selling stock should be to duplicate the outcome of the given option.

Let V_t be the value of the portfolio and assume $V_t = f(S_t, T-t)$ for all t, where f is some function that is sufficiently smooth. We also want $V_T = (S_T - K)^+$.

Recall Ito's formula. The multivariate version is

$$f(X_t) = f(X_0) + \int_0^t \sum_{i=1}^d f_{x_i}(X_s)\, dX_s^i + \frac{1}{2} \int_0^t \sum_{i,j=1}^d f_{x_i x_j}(X_s)\, d\langle X^i, X^j \rangle_s.$$

Here $X_t = (X_t^1, \ldots, X_t^d)$ and f_{x_i} denotes the partial derivative of f in the x_i direction, and similarly for the second partial derivatives.

We apply this with $d = 2$ and $X_t = (S_t, T-t)$. From the SDE that S_t solves, $\langle X^1 \rangle_t = \sigma^2 S_t^2\, dt$, $\langle X^2 \rangle_t = 0$ (since $T-t$ is of bounded variation and hence has no martingale part), and $\langle X^1, X^2 \rangle_t = 0$. Also, $dX_t^2 = -dt$. Then

$$\begin{aligned} V_t - V_0 &= f(S_t, T-t) - f(S_0, T-t) \\ &= \int_0^t f_x(S_u, T-u)\, dS_u - \int_0^t f_s(S_u, T-u)\, du \\ &\quad + \tfrac{1}{2} \int_0^t \sigma^2 S_u^2 f_{xx}(S_u, T-u)\, du. \end{aligned} \tag{19.1}$$

On the other hand,

$$V_t - V_0 = \int_0^t a_u\, dS_u + \int_0^t b_u\, d\beta_u. \tag{19.2}$$

Since the value of the portfolio is

$$V_t = a_t S_t + b_t B_t,$$

we must have

$$b_t = (V_t - a_t S_t)/\beta_t. \tag{19.3}$$

Also, recall

$$\beta_t = \beta_0 e^{rt}. \tag{19.4}$$

We must therefore have

$$a_t = f_x(S_t, T-t) \tag{19.5}$$

and

$$r[f(S_t, T-t) - S_t f_x(S_t, T-t)] = -f_s(S_t, T-t) + \frac{1}{2}\sigma^2 S_t^2 f_{xx}(S_t, T-t) \tag{19.6}$$

for all t and all S_t. (19.6) leads to the parabolic PDE

$$f_s = \tfrac{1}{2}\sigma^2 x^2 f_{xx} + rxf_x - rf, \qquad (x,s) \in (0,\infty) \times [0,T), \tag{19.7}$$

and

$$f(x,0) = (x-K)^+. \tag{19.8}$$

Solving this equation for f, $f(x,T)$ is what V_0 should be, i.e., the cost of setting up the equivalent portfolio. Equation (19.5) shows what the trading strategy should be. In the next section we show how to solve this PDE.

20. Solving PDE.

Without going into the theory of PDE, let us look at how to solve some simple PDE using probability. Let us consider

$$f_t = af_{xx} + bf_x, \qquad f(x,0) = g(x). \tag{20.1}$$

Here f is a function of x and t, a and b are given functions of x and g is also given. The above equation is known as the Cauchy problem.

Proposition 20.1. *Let $A = \sqrt{2a}$ and let X_t be the solution to*

$$dX_t = A(X_t)dW_t + b(X_t)dt. \tag{20.2}$$

The solution to the above equation is given by

$$f(x,t) = \mathbb{E}^x g(X_t).$$

Proof. Fix t_0 and let $M_t = f(X_t, t_0 - t)$. We first show M_t is a martingale. By Ito's formula,

$$\begin{aligned} dM_s &= f_x(X_s, t_0 - s)dX_s - f_t(X_s, t_0 - s)ds + \tfrac{1}{2} f_{xx}(X_s, t_0 - s)A^2(X_s)ds \qquad (20.3) \\ &= f_x(X_s, t_0 - s)A(X_s)dW_s + f_x(X_s, t_0 - s)b(X_s)ds + \tfrac{1}{2} f_{xx}(X_s, t_0 - s)A^2(X_s)ds \\ &\qquad - f_t(X_s, t_0 - s)ds. \end{aligned}$$

Since f solves (20.1), then

$$dM_t = f_x(X_s, t_0 - s)A(X_s)dW_s,$$

which is a stochastic integral of a Brownian motion, hence a martingale.

Now $\mathbb{E}^x M_0 = f(x, t_0)$ and $\mathbb{E}^x M_{t_0} = \mathbb{E}^x f(X_{t_0}, 0) = \mathbb{E}^x g(X_{t_0})$. Since martingales have constant expectation,

$$f(x, t_0) = \mathbb{E}^x g(X_{t_0}).$$

Since t_0 is arbitrary, the proposition is proved. □

To solve the equation

$$f_t = af_{xx} + bf_x + cf, \qquad f(x, 0) = g(x),$$

similar methods show that the solution is given by

$$f(x, t) = \mathbb{E}^x \Big[g(X_t) e^{\int_0^t c(X_s)ds} \Big],$$

where X_t is the solution to (20.2). (This is known as the Feynman-Kac formula.) To see this, if we let

$$N_t = M_t e^{\int_0^t c(X_s)ds},$$

where $M_t = f(X_t, t_0 - t)$, then the Ito product formula yields

$$dN_t = M_t e^{\int_0^t c(X_s)ds} c(X_t)dt + e^{\int_0^t c(X_s)ds} dM_t.$$

Using (20.3) and the fact that $af_{xx} + bf_x + cf = 0$, we see that that the dt term is 0 and N_t is a martingale. Using $\mathbb{E}\, N_0 = \mathbb{E}\, N_{t_0}$ leads to the desired representation of the solution.

Let us look at an example:

$$f_t = \tfrac{1}{2}\sigma^2 x^2 f_{xx} + rxf_x - rf,$$

which is the PDE that arises in Black-Scholes. Here $a = \frac{1}{2}\sigma^2 x^2$ so that $A = \sigma x$, $b = rx$, and $c = -r$. The SDE to be solved, then, is

$$dX_t = \sigma X_t dW_t + rX_t dt, \qquad X_0 = x.$$

We know the solution to this is

$$X_t = xe^{\sigma W_t - \sigma^2 t/2 + rt}.$$

Hence

$$f(x,t) = \mathbb{E}^{\,x}\Big[g(X_t)e^{-rt}\Big] = \mathbb{E}\Big[e^{-rt} g\Big(xe^{\sigma W_t - \sigma^2 t/2 + rt}\Big)\Big].$$

Since we know the density of W_t, we can get an explicit expression (as an integral) for $f(x,t)$.

21. The fundamental theorem of finance.

In Section 17, we showed there was a probability measure under which S_t was a martingale. This is true very generally. Let S_t be the price of a security. We will suppose S_t is a continuous semimartingale, and can be written $S_t = M_t + A_t$.

The NFLVR condition ("no free lunch with vanishing risk") is that there do not exist a fixed time T, $\varepsilon, b > 0$, and H_n (that are adapted and satisfy the appropriate integrability conditions) such that

$$\int_0^T H_n(s)\, dS_s > -\frac{1}{n}, \qquad \text{a.s.}$$

for all t and

$$\mathbb{P}\Big(\int_0^T H_n(s)\, dS_s > b\Big) > \varepsilon.$$

Here T, b, ε do not depend on n. The condition says that one can with positive probability ε make a profit of b and with a loss no larger than $1/n$.

$\mathbb{Q}$ is an equivalent martingale measure if $\mathbb{Q}$ is a probability measure, $\mathbb{Q}$ is equivalent to $\mathbb{P}$, and S_t is a martingale under $\mathbb{Q}$.

Theorem 21.1. *If S_t is a continuous semimartingale and the NFLVR conditions holds, then there exists an equivalent martingale measure $\mathbb{Q}$.*

The proof is rather technical and involves some heavy-duty measure theory, so we will only point examine a part of it. Suppose that we happened to have $S_t = W_t + f(t)$, where $f(t)$ is a deterministic function. To obtain the equivalent martingale measure, we would want to let

$$M_t = e^{-\int_0^t f'(s) dW_s - \int_0^t (f'(s))^2 ds}.$$

In order for M_t to make sense, we need f to be differentiable. A result from measure theory says that if f is not differentiable, then we can find a subset A of $[0, \infty)$ such that $\int_0^t 1_A(s)ds = 0$ but the amount of increase of f over the set A is positive. (This is not a very precise statement.) Then if we hold $H_s = 1_A(s)$ shares at time s, our net profit is

$$\int_0^t H_s dS_s = \int_0^t 1_A(s)dW_s + \int_0^t 1_A(s)\, df(s).$$

The second term would be positive since this is the amount of increase of f over the set A. The first term is 0, since $\mathbb{E}\,(\int_0^t 1_A(s)dW_s)^2 = \int_0^t 1_A(s)^2 ds = 0$. So our net profit is nonrandom and positive, or in other words, we have made a net gain without risk. This contradicts "no arbitrage."

Sometime Theorem 21.1 is called the first fundamental theorem of asset pricing. The second fundamental theorem is the following.

Theorem 21.2. *The equivalent martingale measure is unique if and only if the market is complete.*

We will not prove this.

22. American puts.

The proper valuation of American puts is one of the important unsolved problems in mathematical finance. Recall that a European put pays out $(K - S_T)^+$ at time T, while an American put allows one to exercise early. If one exercises an American put at time $t < T$, one receives $(K - S_t)^+$. Then during the period $[t, T]$ one receives interest, and the amount one has is $(K - S_t)^+ e^{r(T-t)}$. In today's dollars that is the equivalent of $(K - S_t)^+ e^{-rt}$. One wants to find a rule, known as the exercise policy, for when to exercise the put, and then one wants to see what the value is for that policy. Since one cannot look into the future, one is in fact looking for a stopping time τ that maximizes

$$\overline{\mathbb{E}}\, e^{-r\tau}(K - S_\tau)^+.$$

There is no good theoretical solution to finding the stopping time τ, although good approximations exist. We will, however, discuss just a bit of the theory of optimal stopping, which reworks the problem into another form.

Let G_t denote the amount you will receive at time t. For American puts, we set

$$G_t = e^{-rt}(K - S_t)^+.$$

Our problem is to maximize $\overline{\mathbb{E}}\, G_\tau$ over all stopping times τ.

We first need

Proposition 22.1. *If S and T are bounded stopping times with $S \le T$ and M is a martingale, then*

$$\mathbb{E}[M_T \mid \mathcal{F}_S] = M_S.$$

Proof. Let $A \in \mathcal{F}_S$. Define U by

$$U(\omega) = \begin{cases} S(\omega) & \text{if } \omega \in A, \\ T(\omega) & \text{if } \omega \notin A. \end{cases}$$

It is easy to see that U is a stopping time, so by Doob's optional stopping theorem,

$$\mathbb{E}\, M_0 = \mathbb{E}\, M_U = \mathbb{E}[M_S; A] + \mathbb{E}[M_T; A^c].$$

Also,

$$\mathbb{E}\, M_0 = \mathbb{E}\, M_T = \mathbb{E}[M_T; A] + \mathbb{E}[M_T; A^c].$$

Taking the difference, $\mathbb{E}[M_T; A] = \mathbb{E}[M_s; A]$, which is what we needed to show. $\square$

Given two supermartingales X_t and Y_t, it is routine to check that $X_t \wedge Y_t$ is also a supermartingale. Also, if X_t^n are supermartingales with $X_t^n \downarrow X_t$, one can check that X_t is again a supermartingale. With these facts, one can show that given a process such as G_t, there is a least supermartingale larger than G_t.

So we define W_t to be a supermartingale (with respect to $\overline{\mathbb{P}}$, of course) such that $W_t \ge G_t$ a.s for each t and if Y_t is another supermartingale with $Y_t \ge G_t$ for all t, then $W_t \le Y_t$ for all t. We set $\overline{\tau} = \inf\{t : W_t = G_t\}$. We will show that $\overline{\tau}$ is the solution to the problem of finding the optimal stopping time. Of course, computing W_t and $\overline{\tau}$ is another problem entirely.

Let

$$\mathcal{T}_t = \{\tau : \tau \text{ is a stopping time}, t \le \tau \le T\}.$$

Let

$$V_t = \sup_{\tau \in \mathcal{T}_t} \overline{\mathbb{E}}[G_\tau \mid \mathcal{F}_t].$$

Proposition 22.2. *V_t is a supermartingale and $V_t \ge G_t$ for all t.*

Proof. The fixed time t is a stopping time in $\mathcal{T}_t$, so $V_t \ge \overline{\mathbb{E}}[G_t \mid \mathcal{F}_t] = G_t$, or $V_t \ge G_t$. so we only need to show that V_t is a supermartingale.

Suppose $s < t$. Let π be the stopping time in $\mathcal{T}_t$ for which $V_t = \overline{\mathbb{E}}[G_\pi \mid \mathcal{F}_t]$. $\pi \in \mathcal{T}_t \subset \mathcal{T}_s$. Then

$$\overline{\mathbb{E}}[V_t \mid \mathcal{F}_s] = \overline{\mathbb{E}}[G_\pi \mid \mathcal{F}_s] \le \sup_{\tau \in \mathcal{T}_s} \mathbb{E}[G_\tau \mid \mathcal{F}_s] = V_s.$$

□

Proposition 22.3. *If Y_t is a supermartingale with $Y_t \geq G_t$ for all t, then $Y_t \geq V_t$.*

Proof. If $\tau \in \mathcal{T}_t$, then since Y_t is a supermartingale, we have

$$\overline{\mathbb{E}}[Y_\tau \mid \mathcal{F}_t] \leq Y_t.$$

So

$$V_t = \sup_{\tau \in \mathcal{T}_t} \overline{\mathbb{E}}[G_\tau \mid \mathcal{F}_t] \leq \sup_{\tau \in \mathcal{T}_t} \overline{\mathbb{E}}[Y_\tau \mid \mathcal{F}_t] \leq Y_t.$$

□

What we have shown is that W_t is equal to V_t. It remains to show that $\overline{\tau}$ is optimal. There may in fact be more than one optimal time, but in any case $\overline{\tau}$ is one of them. Recall we have $\mathcal{F}_0$ is the σ-field generated by S_0, and hence consists of only $\emptyset$ and Ω.

Proposiiton 22.4. *$\overline{\tau}$ is an optimal stopping time.*

Proof. Since $\mathcal{F}_0$ is trivial, $V_0 = \sup_{\tau \in \mathcal{T}_0} \overline{\mathbb{E}}[G_\tau \mid \mathcal{F}_0] = \sup_\tau \overline{\mathbb{E}}[G_\tau]$. Let σ be a stopping time where the supremum is attained. Then

$$V_0 \geq \overline{\mathbb{E}}[V_\sigma \mid \mathcal{F}_0] = \overline{\mathbb{E}}[V_\sigma] \geq \overline{\mathbb{E}}[G_\sigma] = V_0.$$

Therefore all the inequalities must be equalities. Since $V_\sigma \geq G_\sigma$, we must have $V_\sigma = G_\sigma$. Since $\overline{\tau}$ was the first time that W_t equals G_t and $W_t = V_t$, we see that $\overline{\tau} \leq \sigma$. Then

$$\overline{\mathbb{E}}[G_{\overline{\tau}}] = \overline{\mathbb{E}}[V_{\overline{\tau}}] \geq \overline{E} V_\sigma = \overline{\mathbb{E}} G_\sigma.$$

Therefore the expected value of $G_{\overline{\tau}}$ is as least as large as the expected value of G_σ, and hence $\overline{\tau}$ is also an optimal stopping time. □

The above representation of the optimal stopping problem may seem rather bizarre. However, this procedure gives good usable results for some optimal stopping problems. An example is where G_t is a function of just W_t.

23. Term structure.

We now want to consider the case where the interest rate is nondeterministic, that is, it has a random component. To do so, we take another look at option pricing.

Let $r(t)$ be the (random) interest rate at time t. Let

$$\beta(t) = e^{\int_0^t r(u) du}$$

be the accumulation factor. One dollar at time T will be worth $1/\beta(T)$ in today's dollars.

Let $V = (S_T - K)^+$ be the payoff on the standard European call option at time T with strike price K, where S_t is the stock price. In today's dollars it is worth, as we have seen, $V/\beta(T)$. Therefore the price of the option should be

$$\overline{\mathbb{E}}\Big[\frac{V}{\beta(T)}\Big].$$

We can also get an expression for the value of the option at time t. The payoff, in terms of dollars at time t, should be the payoff at time T discounted by the interest or inflation rate, and so should be

$$e^{-\int_t^T r(u)du}(S_T - K)^+.$$

Therefore the value at time t is

$$\overline{\mathbb{E}}\Big[e^{-\int_t^T r(u)du}(S_T - K)^+ \mid \mathcal{F}_t\Big] = \overline{\mathbb{E}}\Big[\frac{\beta(t)}{\beta(T)}V \mid \mathcal{F}_t\Big] = \beta(t)\overline{\mathbb{E}}\Big[\frac{V}{\beta(T)} \mid \mathcal{F}_t\Big].$$

From now on we assume we have already changed to the risk-neutral measure and we write $\mathbb{P}$ instead of $\overline{\mathbb{P}}$.

A zero coupon bond with maturity date T pays \$1 at time T and nothing before. This is equivalent to an option with payoff value $V = 1$. So its price at time t, as above, should be

$$B(t,T) = \beta(t)\mathbb{E}\Big[\frac{1}{\beta(T)} \mid \mathcal{F}_t\Big] = \mathbb{E}\Big[e^{-\int_t^T r(u)du} \mid \mathcal{F}_t\Big].$$

Let's derive the SDE satisfied by $B(t,T)$. Let $N_t = \mathbb{E}\left[1/\beta(T) \mid \mathcal{F}_t\right]$. This is a martingale. By the martingale representation theorem,

$$N_t = \mathbb{E}\left[1/\beta(T)\right] + \int_0^t H_s dW_s$$

for some adapted integrand H_s. So $B(t,T) = \beta(t)N_t$. Here T is fixed. By Ito's product formula,

$$\begin{aligned} dB(t,T) &= \beta(t)dN_t + N_t d\beta(t) \\ &= \beta(t)H_t dW_t + N_t r(t)\beta(t)dt \\ &= \beta(t)H_t dW_t + B(t,T)r(t)dt, \end{aligned}$$

and we thus have

$$dB(t,T) = \beta(t)H_t dW_t + B(t,T)r(t)dt. \tag{23.1}$$

We now discuss forward rates. If one holds T fixed and graphs $B(t,T)$ as a function of t, the graph will not clearly show the behavior of r. One sometimes specifies interest rates by what are known as forward rates.

Suppose you want to borrow \$1 at time T and repay it with interest at time $T+\varepsilon$. At the present time we are at time $t \le T$. One can accomplish this by buying a zero coupon bond with maturity date T and shorting (i.e., selling) $B(t,T)/B(t,T+\varepsilon)$ zero coupon bonds with maturity date $T+\varepsilon$. The value of the portfolio at time t is

$$B(t,T) - \frac{B(t,T)}{B(t,T+\varepsilon)} B(t,T+\varepsilon) = 0.$$

At time T you receive \$1. At time $T+\varepsilon$ you pay $B(t,T)/B(t,T+\varepsilon)$. The effective rate of interest R over the time period T to $T+\varepsilon$ is

$$e^{\varepsilon R} = \frac{B(t,T)}{B(t,T+\varepsilon)}.$$

Solving for R, we have

$$R = \frac{\log B(t,T) - \log B(t,T+\varepsilon)}{\varepsilon}.$$

We now let $\varepsilon \to 0$. We define the forward rate by

$$f(t,T) = -\frac{\partial}{\partial T} \log B(t,T). \tag{23.2}$$

Sometimes interest rates are specified by giving $f(t,T)$ instead of $B(t,T)$ or $r(t)$.

Let us see how to recover $B(t,T)$ from $f(t,T)$. Integrating, we have

$$\begin{aligned}\int_t^T f(t,u)du &= -\int_t^T \frac{\partial}{\partial u} \log B(t,u)du = -\log B(t,u) \mid_{u=t}^{u=T} \\ &= -\log B(t,T) + \log B(t,t).\end{aligned}$$

Since $B(t,t)$ is the value of a zero coupon bond at time t which expires at time t, it is equal to 1, and its log is 0. Solving for $B(t,T)$, we have

$$B(t,T) = e^{-\int_t^T f(t,u)du}. \tag{23.3}$$

Next, let us show how to recover $r(t)$ from the forward rates. We have

$$B(t,T) = \mathbb{E}\left[e^{-\int_t^T r(u)du} \mid \mathcal{F}_t\right].$$

Differentiating,

$$\frac{\partial}{\partial T} B(t,T) = \mathbb{E}\left[-r(T)e^{-\int_t^T r(u)du} \mid \mathcal{F}_t\right].$$

Evaluating this when $T=t$, we obtain

$$\mathbb{E}\left[-r(t) \mid \mathcal{F}_t\right] = r(t). \tag{23.4}$$

On the other hand, from (23.3) we have

$$\frac{\partial}{\partial T}B(t,T) = -f(t,T)e^{-\int_t^T f(t,u)du}.$$

Setting $T = t$ we obtain $-f(t,t)$. Comparing with (23.4) yields

$$r(t) = f(t,t). \tag{23.5}$$

24. Some interest rate models.

Heath-Jarrow-Morton model

Instead of specifying r, the Heath-Jarrow-Morton model (HJM) specifies the forward rates:

$$df(t,T) = \sigma(t,T)dW_t + \alpha(t,T)dt. \tag{24.1}$$

Let us derive the SDE that $B(t,T)$ satisfies. Let

$$\alpha^*(t,T) = \int_t^T \alpha(t,u)du, \qquad \sigma^*(t,T) = \int_t^T \sigma(t,u)du.$$

Since $B(t,T) = \exp(-\int_t^T f(t,u)du)$, we derive the SDE for B by using Ito's formula with the function e^x and $X_t = -\int_t^T f(t,u)du$. We have

$$\begin{aligned}
dX_t &= f(t,t)dt - \int_t^T df(t,u)du \\
&= r(t)dt - \int_t^T [\alpha(t,u)dt + \sigma(t,u)dW_t]\,du \\
&= r(t)dt - \Big[\int_t^T \alpha(t,u)du\Big]dt - \Big[\int_t^T \sigma(t,u)du\Big]dW_t \\
&= r(t)dt - \alpha^*(t,T)dt - \sigma^*(t,T)dW_t.
\end{aligned}$$

Therefore, using Ito's formula,

$$\begin{aligned}
dB(t,T) &= B(t,T)dX_t + \tfrac{1}{2}B(t,T)(\sigma^*(t,T))^2dt \\
&= B(t,T)\Big[r(t) - \alpha^* + \tfrac{1}{2}(\sigma^*)^2\Big]dt - \sigma^* B(t,T)dW_t.
\end{aligned}$$

From (23.1)

$$dB(t,T) = B(t,T)r(t)dt - \sigma^* B(t,T)dW_t.$$

Comparing, we see that if $\mathbb{P}$ is the risk-neutral measure, we have $\alpha^* = \frac{1}{2}(\sigma^*)^2$.

If $\mathbb{P}$ is not the risk-neutral measure, it is still possible that one exists. Let $\theta(t)$ be a function of t, let $M_t = \exp(-\int_0^t \theta(u)dW_u - \frac{1}{2}\int_0^t \theta(u)^2 du)$ and define $\overline{\mathbb{P}}(A) = \mathbb{E}\,[M_T; A]$ for $A \in \mathcal{F}_T$. By the Girsanov theorem,

$$dB(t,T) = B(t,T)\Big[r(t) - \alpha^* + \tfrac{1}{2}(\sigma^*)^2 + \sigma^*\theta]dt - \sigma^* B(t,T)d\widetilde{W}_t,$$

where $\widetilde{W}_t$ is a Brownian motion under $\overline{\mathbb{P}}$. So we must have

$$\alpha^* = \tfrac{1}{2}(\sigma^*)^2 + \sigma^*\theta.$$

Differentiating with respect to T, we obtain

$$\alpha(t,T) = \sigma(t,T)\sigma^*(t,T) + \sigma(t,T)\theta(t).$$

If we try to solve this equation for θ, there is no reason off-hand that θ depends only on t and not T. However, if θ does not depend on T, $\overline{\mathbb{P}}$ will be the risk-neutral measure.

Hull and White model

In this model, the interest rate r is specified as the solution to the SDE

$$dr(t) = \sigma(t)dW_t + (a(t) - b(t)r(t))dt.$$

Here σ, a, b are deterministic functions. The stochastic integral term introduces randomness, while the $a - br$ term causes a drift toward $a(t)/b(t)$.

This is one of those SDE's that can be solved explicitly. Let $K(t) = \int_0^t b(u)du$. Then

$$\begin{aligned} d\Big[e^{K(t)}r(t)\Big] &= e^{K(t)}r(t)b(t) + e^{K(t)}\Big[a(t) - b(t)r(t)\Big]dt + e^{K(t)}[\sigma(t)dW_t] \\ &= e^{K(t)}a(t)dt + e^{K(t)}[\sigma(t)dW_t]. \end{aligned}$$

Integrating both sides,

$$e^{K(t)}r(t) = r(0) + \int_0^t e^{K(u)}a(u)du + \int_0^t e^{K(u)}\sigma(u)dW_u.$$

Multiplying both sides by $e^{-K(t)}$, we have the explicit solution

$$r(t) = e^{-K(t)}\Big[r(0) + \int_0^t e^{K(u)}a(u)du + \int_0^t e^{K(u)}\sigma(u)dW_u\Big].$$

If $F(u)$ is deterministic, then

$$\int_0^t F(u)dW_u = \lim \sum F(u_i)(W_{u_{i+1}} - W_{u_i}).$$

Linear combinations of Gaussian r.v.'s (Gaussian = normal) are Gaussian, and limits of Gaussian r.v.'s are Gaussian, so we conclude $\int_0^t F(u)dW_u$ is a Gaussian r.v. We see that the mean at time t is

$$\mathbb{E}\, r(t) = e^{-K(t)} \Big[r(0) + \int_0^t e^{K(u)} a(u) du \Big].$$

We know how to calculate the second moment of a stochastic integral, so

$$\operatorname{Var} r(t) = e^{-2K(t)} \int_0^t e^{2K(u)} \sigma(u)^2 du.$$

(One can similarly calculate the covariance of $r(s)$ and $r(t)$.) Limits of linear combinations of Gaussians are Gaussian, so we can calculate the mean and variance of $\int_0^T r(t)dt$ and get an explicit expression for

$$B(0,T) = \mathbb{E}\, e^{-\int_0^T r(u)du}.$$

Cox-Ingersoll-Ross model

One drawback of the Hull and White model is that since $r(t)$ is Gaussian, it can take negative values with positive probability, which doesn't make sense. The Cox-Ingersoll-Ross model avoids this by modeling r by the SDE

$$dr(t) = (a - br(t))dt + \sigma \sqrt{r(t)} dW_t.$$

The difference from the Hull and White model is the square root of r in the stochastic integral term. Provided $a \geq \frac{1}{2}\sigma^2$, it can be shown that $r(t)$ will never hit 0 and will always be positive. Although one cannot solve for r explicitly, one can calculate the distribution of r. It turns out to be related to the square of what are known in probability theory as Bessel processes. (The density of $r(t)$, for example, will be given in terms of Bessel functions.)

25. Foreign exchange.

Suppose we can buy bonds in dollars with constant interest rate r. So the value of the bond at time t is $B_t = B_0 e^{rt}$. Suppose we can buy bonds in pounds sterling with constant rate u. If D_t is the price of the bond, then $D_t = D_0 e^{ut}$.

Let C_t be the exchange rate; at time t one pound is worth C_t dollars. Not surprisingly, one model for exchange rates is to set

$$C_t = C_0 e^{\sigma W_t - \sigma^2 t/2 + \mu t},$$

where W_t is a Brownian motion.

We want to see how to price options in terms of the exchange rate. There are two main differences from the stock case. One is that if we buy pounds, we can invest them in sterling bonds and receive interest. The other is that C_t is not a quantity that one can invest in. One cannot buy an exchange rate.

Let $S_t = C_t D_t$. Start with with $S_0 = C_0 D_0$ dollars at time 0, exchange them for S_0/C_0 pounds, and buy $S_0/(C_0 D_0)$ sterling bonds. At time t redeem the bonds for D_t pounds each and convert back to dollars. So at time t, you have $C_t D_t = S_t$ dollars. Therefore S_t is tradable, even if C_t is not.

The present value of S_t is $e^{-rt} S_t = B_t^{-1} S_t$. Let us set $Z_t = B_t^{-1} S_t$. Substituting for C_t and D_t, we have

$$Z_t = e^{-rt} e^{ut} e^{\sigma W_t - \sigma^2 t/2 + \mu t} = e^{\sigma W_t - \sigma^2 t/2 + (\mu + u - r)t}.$$

Let $M_t = \exp(-(\mu + u - r)W_t - (\mu + u - r)^2 t/2)$ and define $\mathbb{Q}(A) = \mathbb{E}\,[M_t; A]$ if $A \in \mathcal{F}_t$. By what are now standard calculations, Z_t is a martingale under $\mathbb{Q}$.

Now let's price options. For example, one might have a sterling call, which is the right at time T to buy a pound for K dollars. The payoff at time T is $V = (C_T - K)^+$. We use the martingale representation theorem, just as in the Black-Scholes derivation, to write

$$V_T = c + \int_0^T H_s dZ_s.$$

Since Z is a martingale under $\mathbb{Q}$, then $c = \mathbb{E}_{\mathbb{Q}} V$ and this is the value of the call in today's dollars. H gives the hedging strategy: at time s hold H_s units of sterling bonds, and the remainder, namely, $V_s - H_s Z_s$, in dollar bonds.

Are the calculations hard? Not really – the situation looks exactly the same as the case of stock prices, but the interest rate is now $r - u$ instead of r. This makes sense; we don't discount by the rate r, because insteading of holding stocks, which pay no interest, we hold sterling bonds, which pay at interest rate u.

One could also do all these calculations from the point of view of the English investor. He has the alternative of putting his pounds into sterling bonds, or else exchanging them for dollars and investing in dollar bonds. The calculations are very similar.

It turns out that the martingale measure ($\mathbb{Q}$) will be different for the American and the English investor. Nevertheless, they will come up with the identical value for an option and will have the identical hedging strategy.

26. Dividends.

How do things change if a stock issues dividends? Of course, typically a stock issues dividends at fixed discrete times, but for the sake of our continuous model, let us assume that the stock issues dividends at a continuous rate.

There are two possibilities. One possibility is that the stock issues dividends at a fixed rate, regardless of what the stock price is. This case is very similar to the foreign exchange discussion: like the investor who exchanges his dollars for pounds and invests in sterling bonds, the security purchased (sterling bonds in the foreign exchange case, stock in the present case) pays interest.

So we will look at the other case, where we suppose that the stock issues dividends at a rate proportional to the stock price. We suppose in a short time interval Δt that the amount of dividends issued is $\delta S_t \Delta t$. Here δ is the dividend rate.

Suppose that as soon as we receive dividends, we immediately buy stock with it. Our net gain in cash is then zero, but the number of shares we hold has increased. If S_t is the stock price, we hold $e^{\delta t}$ times as many shares, so our investment in the stock is now

$$\widetilde{S}_t = e^{\delta t} S_t = e^{\sigma W_t - \sigma^2 t/2 + \mu t + \delta t}.$$

We are back in the standard stock model, except that our mean rate of return has become $\mu + \delta$ instead of μ. The Black-Scholes formulas are much as before.