1

Discrete-time processes

Our study of stochastic processes, motivated by their use in financial modelling, begins with discrete-time models, including and generalising the models studied in detail in *Discrete Models of Financial Markets* [DMFM], where the typical 'process' was simply a finite sequence of random variables defined on some finite sample space. We generalise this in two directions, by considering a general probability space $(\Omega, \mathcal{F}, P)$ and allowing our processes to be infinite sequences of random variables defined on this space. Again the key concept is that of martingales, and we study the basic properties of discrete martingales in preparation for our later consideration of their continuous-time counterparts. We then briefly consider how another basic class of discrete-time processes, Markov chains, enters into the study of credit ratings, and develop some of their simple properties. Throughout, we will make extensive use of the fundamental properties of probability measures and random variables described in *Probability for Finance* [PF], and we refer the reader to that text for any probabilistic notions not explicitly defined here.

1.1 General definitions

We take a discrete time scale with $n = 0, 1, 2, \ldots$ denoting the number of consecutive steps of fixed length $h > 0$, so time instants are $t = nh \in [0, \infty)$. In contrast to [DMFM], where we had finitely many times, we allow infinitely many steps as a prelude to the continuous case, where $t \in [0, \infty)$ is arbitrary, which we study in the subsequent chapters.

We assume that a probability space $(\Omega, \mathcal{F}, P)$ is available, sufficiently rich to accomodate the various collections of random variables we wish to define. We have to allow infinite Ω to be able to discuss random variables without restricting their values to some finite set. Thus Ω is an arbitrary set, while $\mathcal{F}$ is a σ-field, and P a countably additive probability measure.

From the financial perspective, a random variable is a mathematical object modelling an unknown quantity such as a stock price. A sequence of random variables will correspond to its future evolution with no limiting horizon, as described in the next definition.

Definition 1.1
A **discrete-time stochastic process** is a sequence of random variables, that is an $\mathcal{F}$-measurable function

$$X(n) : \Omega \to \mathbb{R} \text{ for } n \geq 0,$$

and we assume that $X(0)$ is constant.

We employ the notation $X = (X(n))_{n \geq 0}$ but often we refer to 'the process $X(n)$' and alternatively we will use X to denote an arbitrary random variable, thus risking a lack of precision. This allows us, for instance, to indicate the time variable and to keep the presentation brief and free of pure formalism. In the same spirit, we will often drop the expression 'almost surely' after any relation between random variables.

Example 1.2
A classical example of a probability space, which, as it will turn out, is sufficiently rich for all our purposes, is $\Omega = [0, 1]$, $\mathcal{F} = \mathcal{B}([0, 1])$ – Borel sets, $P = m$ – Lebesgue measure (a construction can be found in [PF]).

Example 1.3
Consider a binomial tree, discussed in detail in [DMFM], determined by two single-step returns $D < U$. We define a sequence of returns

$$K(n) : [0, 1] \to \mathbb{R},$$

$n = 1, 2, \ldots,$ by

$$K(n, \omega) = U\mathbf{1}_{A_n}(\omega) + D\mathbf{1}_{[0,1]\setminus A_n}(\omega),$$

$$A_n = \left[0, \frac{1}{2^n}\right) \cup \left[\frac{2}{2^n}, \frac{3}{2^n}\right) \cup \cdots \cup \left[\frac{2^n - 2}{2^n}, \frac{2^n - 1}{2^n}\right),$$

and clearly

$$P(K(n) = U) = P(K(n) = D) = \frac{1}{2}.$$

For instance,

$$K(2) = \begin{cases} U \text{ if } \omega \in [0, \frac{1}{4}) \cup [\frac{1}{2}, \frac{3}{4}), \\ D \text{ if } \omega \in [\frac{1}{4}, \frac{1}{2}) \cup [\frac{3}{4}, 1]. \end{cases}$$

The stock prices are defined in a familiar way by

$$S(n) = S(n - 1)(1 + K(n)),$$

$n = 1, 2, \ldots,$ with $S(0)$ given, deterministic.

Exercise 1.1 Show that for each n, the random variables $K(1), \ldots, K(n)$ are independent.

Exercise 1.2 Redesign the random variables $K(n)$ so that $P(K(n) = U) = p \in (0, 1)$, arbitrary.

Example 1.4
A version of a binomial tree with additive rather than multiplicative changes is called a **symmetric random walk** and is defined by taking $Z(0)$

given and

$$Z(n) = Z(n-1) + L(n),$$
$$L(n) = \pm 1, \text{ each with probability } \frac{1}{2}.$$

The sequence $L(n)$ defining a symmetric random walk can conveniently be regarded as representing a sequence of independent tosses of a fair coin. The outcome of each coin toss might determine whether a gambler gains or loses one unit of currency, so that the random variable $Z(n) = Z(0) + \sum_{i=1}^{n} L(i)$ describes his fortune after n games if he starts with $Z(0)$. Alternatively, it could describe the position on the line reached after n steps by a particle starting at position $Z(0)$ and, at the ith step (for each $i \leq n$), moving one unit to the right if $L(i) = 1$, or to the left if $L(i) = -1$. If the particle moves with constant velocity between the changes of direction, its path can be visualised by joining subsequent points $(n, Z(n))$ in the plane with line segments.

Information given by an initial segment $X(0), \ldots, X(n)$ of the sequence X can be captured by means of a filtration of partitions if the number of possible values of random variables is finite. We exploited this in [DMFM], but here we take a more general approach, replacing partitions by σ-fields, which allows us to consider arbitrary random variables.

Definition 1.5
The **filtration generated** by a discrete-time process $(X(n))_{n\geq 0}$ (also known as its **natural** filtration) is a family of σ-fields

$$\mathcal{F}_n^X = \sigma(\{X(k)^{-1}(B) : B \in \mathcal{B}(\mathbb{R}), k = 0, \ldots, n\}),$$

where for any family of sets $\mathcal{A}$, $\sigma(\mathcal{A})$ is the smallest σ-field containing $\mathcal{A}$, and $\mathcal{B}(\mathbb{R})$ is the σ-field of Borel sets on the real line.

Observe that the same result would be obtained by taking B to run through all intervals, or, indeed, all intervals of the form $(-\infty, a]$ for $a \in \mathbb{R}$. Since all elements of the sequence $X(n)$ are $\mathcal{F}$-measurable, $\mathcal{F}_n^X \subset \mathcal{F}$ for each n. In addition, $X(n)$ is clearly $\mathcal{F}_n^X$-measurable.

Note that, by its definition, the sequence $\mathcal{F}_n^X$ is increasing with respect to set inclusion $\subset$. This motivates introducing a general notion to indicate this

Definition 1.6

A **filtration** is a sequence of σ-fields $\mathcal{F}_n$ such that $\mathcal{F}_n \subset \mathcal{F}$ and $\mathcal{F}_n \subset \mathcal{F}_{n+1}$. A process X is **adapted** if each $X(n)$ is $\mathcal{F}_n$-measurable. If an arbitrary filtration $(\mathcal{F}_n)_{n\geq 0.}$ has been fixed, we call $(\Omega, \mathcal{F}, (\mathcal{F})_{n\geq 0}, P)$ a **filtered probability space**.

Note that for any process its natural filtration is the smallest filtration to which it is adapted.

As we wish $X(0)$ to be constant (almost surely, of course!), we assume that $\mathcal{F}_0$ is **trivial**; that is, it is simply made up of all P-null sets and their complements.

Example 1.7

Consider $K(n)$ as given in Example 1.3. Clearly, for every $n \geq 1$

$$\{K(n)^{-1}(B) : B \in \mathcal{B}(\mathbb{R})\} = \{\emptyset, [0, 1], A_n, [0, 1] \setminus A_n\}$$

and the σ-field $\mathcal{F}_n$ consists of all null sets and all possible unions of intervals of the form $[\frac{i-1}{2^n}, \frac{i}{2^n})$, $i = 1, \dots, 2^n - 1$ and $[\frac{2^n-1}{2^n}, 1]$ (This is an example of a field generated by so-called **atoms**.)

Exercise 1.3 Find the filtration in $\Omega = [0, 1]$ generated by the process $X(n, \omega) = 2\omega \mathbf{1}_{[0,1-\frac{1}{n}]}(\omega)$.

Example 1.8

The previous exercise illustrates the idea of flow of information described by a filtration. With increasing n, the shape of some function defined on Ω (here $\omega \mapsto 2\omega$) is gradually revealed. The values of X, if known, allow us to make a guess about the location of ω. If for instance $X(2, \omega) = \frac{1}{2}$, we know that $\omega = \frac{1}{4}$, but if $X(2, \omega) = 0$, we only know that $\omega \in (\frac{1}{2}, 1]$. Clearly, our information about ω, given the value of X, increases with n.

1.2 Martingales

We prepare ourselves for the estimation of some random variable Y after we observe a process X up to some future time n. The information obtained will be reflected in the σ-field $\mathcal{G} = \mathcal{F}_n$. The idea of the approximation of Y given the information contained in $\mathcal{G}$ is explained in the next definition (see [PF] for more details).

Definition 1.9
The **conditional expectation** of Y given $\mathcal{G}$ is a random variable $\mathbb{E}(Y|\mathcal{G})$, which is:

1. $\mathcal{G}$-measurable,
2. for all $A \in \mathcal{G}$,

$$\int_A \mathbb{E}(Y|\mathcal{G})dP = \int_A YdP.$$

Example 1.10
In $\Omega = [0, 1]$, take the σ-field

$$\mathcal{G} = \left\{B \subset \left[0, \frac{1}{2}\right) : B \in \mathcal{B}(\mathbb{R})\right\} \cup \left\{B \cup \left[\frac{1}{2}, 1\right] : B \subset \left[0, \frac{1}{2}\right), B \in \mathcal{B}(\mathbb{R})\right\}$$

and $Y(\omega) = \omega^2$. Condition 1 imposes the constraint that the conditional expectation be constant on $[\frac{1}{2}, 1]$ but it can be arbitrary on $[0, \frac{1}{2})$. If so, Y will be the best approximation of itself on $[0, \frac{1}{2})$, while the constant c is given by condition 2 with $A = [\frac{1}{2}, 1]$ by solving $cP([\frac{1}{2}, 1]) = \int_{[\frac{1}{2}, 1]} \omega^2 dP$ and so

$$\mathbb{E}(Y|\mathcal{G})(\omega) = \begin{cases} \omega^2 \text{ if } \omega \in [0, \frac{1}{2}), \\ \frac{7}{12} \text{ if } \omega \in [\frac{1}{2}, 1]. \end{cases}$$

With increasing information about the future, as captured by a filtration $\mathcal{F}_n$, the accuracy of prediction improves, which follows from the important **tower property** of conditional expectation

$$\mathbb{E}(Y|\mathcal{F}_n) = \mathbb{E}(\mathbb{E}(Y|\mathcal{F}_{n+1})|\mathcal{F}_n).$$

Writing $M(n) = \mathbb{E}(Y|\mathcal{F}_n)$ above, we have an example of the following notion, crucial in what follows:

Definition 1.11
A process M is a **martingale** with respect to the filtration $\mathcal{F}_n$ if, for all

$n \geq 0$, $\mathbb{E}(|M(n)|) < \infty$, and

$$M(n) = \mathbb{E}(M(n+1)|\mathcal{F}_n).$$

Note that a martingale is $\mathcal{F}_n$-adapted, by the definition of conditional expectation.

Exercise 1.4 Working on $\Omega = [0, 1]$, find (by means of concrete formulae and sketching the graphs) the martingale $\mathbb{E}(Y|\mathcal{F}_n)$, where $Y(\omega) = \omega^2$ and $\mathcal{F}_n$ is generated by $X(n, \omega) = 2\omega\mathbf{1}_{[0,1-\frac{1}{n}]}(\omega)$ (see Exercise 1.3).

Exercise 1.5 Show that the expectation of a martingale is constant in time. Find an example showing that constant expectation does not imply the martingale property.

Exercise 1.6 Show that martingale property is preserved under linear combinations with constant coefficients and adding a constant.

Exercise 1.7 Prove that if M is a martingale, then for $m < n$

$$M(m) = \mathbb{E}(M(n)|\mathcal{F}_m).$$

Another classical example of a martingale is the random walk $Z(n) = Z(n-1)+L(n)$, with filtration generated by $L(n)$. The proof of the martingale property of a random walk is exactly the same as that of the general result below.

Proposition 1.12
The sequence obtained by sums of independent random variables with zero mean is a martingale with respect to the filtration it generates.

Proof Assume that $L(n)$ is an arbitrary sequence of independent random variables with $\mathbb{E}(L(n)) = 0$, and write

$$Z(n) = Z(0) + \sum_{j=1}^{n} L(j),$$
$$\mathcal{F}_n = \sigma(Z(k) : 0 \leq k \leq n).$$

The properties of the conditional expectation immediately give the result:

$$
\begin{aligned}
\mathbb{E}(Z(n+1)|\mathcal{F}_n) &= \mathbb{E}(Z(n)+L(n+1)|\mathcal{F}_n) \quad \text{(definition of } Z\text{)}\\
&= \mathbb{E}(Z(n)|\mathcal{F}_n)+\mathbb{E}(L(n+1)|\mathcal{F}_n) \quad \text{(linearity)}\\
&= Z(n)+\mathbb{E}(L(n+1)) \quad \text{(definition of } \mathcal{F}_n\text{, independence)}\\
&= Z(n).
\end{aligned}
$$

□

We now develop a method of producing new martingales from a given one, which was already exploited in [DMFM], where we discussed the value process of a trading strategy. The result we wish to highlight is sometimes called a theorem on discrete stochastic integration. As we shall see later, the name is well deserved. To state this theorem, we need one more definition, which codifies an important property of trading strategies: recall that, for $n > 0$, an investor's decision to hold $x(n)$ units of stock throughout the period between trading dates $n-1$ and n was based on his knowledge of the behaviour of the stock price up to time $n-1$, so that the random variable $x(n)$ is $\mathcal{F}_{n-1}$-measurable. Such a process $x(n)$ was said to be predictable, and we use this idea for a general definition:

Definition 1.13
A process $X = (X(n))_{n\geq 1}$ is **predictable** relative to a given filtration $(\mathcal{F}_n)_{n\geq 0}$ if for every $n \geq 1$, $X(n)$ is $\mathcal{F}_{n-1}$-measurable.

Note that the sequence starts at $n = 1$, so there is no $X(0)$. Recall also that we take the σ-field $\mathcal{F}_0$ to be trivial. Thus if X is predictable, $X(1)$ is a constant function.

As we saw for trading strategies, predictability means that the variable $X(n)$ is 'known' by time $n-1$, so we can 'predict' future values of X one step ahead. This property is incompatible with the martingale property unless the process is constant.

Proposition 1.14
A predictable martingale is constant.

Proof By definition of martingale $M(n-1) = \mathbb{E}(M(n)|\mathcal{F}_{n-1})$, which equals $M(n)$ if M is predictable. □

We return to our stock price example. Having bought $x(n)$ shares at time $n-1$, we are of course interested in our gains. When Ω is finite and $P(\omega) > 0$ for each ω in Ω we know that, under the No Arbitrage Principle, the pricing model admits a risk-neutral probability Q, with $Q(\omega) > 0$ for each

ω (as discussed in [DMFM]). This means that, assuming that investment in a risk-free asset attracts a constant return $R \geq 0$, the discounted stock prices $\tilde{S}(n) = (1 + R)^{-n} S(n)$ follow a martingale under Q. Working with discounted values is a sound concept since the alternative risk-free investment provides a natural benchmark for returns. Our discounted gain (or loss) at the nth step will be $x(n)[\tilde{S}(n) - \tilde{S}(n-1)]$. It is natural to consider the gains accumulated from time zero:

$$G(n) = V(0) + \sum_{i=1}^{n} x(i)[\tilde{S}(i) - \tilde{S}(i-1)].$$

As we will see, the resulting process, expressing the discounted values, remains a martingale. In other words, the fairness of the discounted stock prices expressed by means of martingale property – the 'best guess' of future prices is the current price – is preserved by the strategy, so this market cannot be 'second-guessed' legally to ensure a profit by making cleverly chosen purchases and sales of the stock.

Theorem 1.15
Let M be a martingale and H a predictable process. If H is bounded, or if both H and M are square-integrable, then $X(0) = 0$,

$$X(n) = \sum_{j=1}^{n} H(j)[M(j) - M(j-1)] \text{ for } n > 0,$$

defines a martingale.

Proof This is proved in [DMFM] in a multi-dimensional setting for a finite Ω. The only change needed now is to observe that the conditions imposed here ensure the integrability of $X(n)$, trivially in the first case and by the Cauchy–Schwarz inequality in the second. You should fill in the details by writing out the easy proof of this result yourself, recalling that $H(j)$ is $\mathcal{F}_{j-1}$-measurable). □

The next exercise provides a converse to the theorem in an important special case.

Exercise 1.8 Let M be a martingale with respect to the filtration generated by $L(n)$ (as defined for a random walk), and assume for simplicity $M(0) = 0$. Show that there exists a predictable process H such that $M(n) = \sum_{i=1}^{n} H(i)L(i)$ (that is $M(n) = \sum_{i=1}^{n} H(i)[Z(i) - Z(i-1)]$, where

$Z(i) = \sum_{j=1}^{i} L(j)$. (We are justified in calling this result a representation theorem: each martingale is a discrete stochastic integral).

Adding a constant to the sum on the right in the above theorem preserves the martingale property – unlike non-linear operations, which typically destroy it.

Exercise 1.9 Show that the process $Z^2(n)$, the square of a random walk, is not a martingale, by checking that $\mathbb{E}(Z^2(n+1)|\mathcal{F}_n) = Z^2(n)+1$.

Recall the **Jensen inequality**: If $\varphi : \mathbb{R} \to \mathbb{R}$ is a convex function and $\varphi(X) \in L^1(P)$, then we have

$$\mathbb{E}(\varphi(X)|\mathcal{G}) \geq \varphi(\mathbb{E}(X|\mathcal{G})).$$

Applied to $\varphi(x) = x^2$, $X = M(n+1)$, $\mathcal{G} = \mathcal{F}_n$, this gives the following property of the square of a martingale

$$\mathbb{E}(M^2(n+1)|\mathcal{F}_n) \geq (\mathbb{E}(M(n+1)|\mathcal{F}_n))^2 = M^2(n).$$

This leads to some useful general notions.

Definition 1.16
A process $X(n)$-adapted to a filtration $\mathcal{F}_n$ with all $X(n)$ integrable, is

1. a **submartingale** if $\mathbb{E}(X(n+1)|\mathcal{F}_n) \geq X(n)$,
2. a **supermartingale** if $\mathbb{E}(X(n+1)|\mathcal{F}_n) \leq X(n)$.

Clearly, a martingale is a sub- and supermartingale. A process which is both, sub- and supermartingale, is a martingale. The above application of Jensen's inequality shows that the square of a martingale is a submartingale.

Exercise 1.10 Show that if X is a submartingale, then its expectations increase with n:

$$\mathbb{E}(X(0)) \leq \mathbb{E}(X(1)) \leq \mathbb{E}(X(2)) \leq \cdots,$$