

Lecture Notes for Finance 1 (and More).

David Lando
Simon Ellersgaard Nielsen
Rolf Poulsen

May 12, 2015

Contents

1	Preface	5
2	Introduction	7
2.1	The role of financial markets	9
2.2	The Kelly criterion: Optimal betting	11
3	Payment Streams under Certainty	15
3.1	Financial markets and arbitrage	15
3.2	Zero coupon bonds and the term structure	19
3.3	Annuities, serial loans and bullet bonds	22
3.4	The net present value (NPV) rule	29
3.4.1	Several projects	33
3.5	Duration, convexity and immunization.	35
3.5.1	Duration with a flat term structure.	35
3.5.2	Relaxing the assumption of a flat term structure.	40
3.6	Two examples to mess with your head	42
3.6.1	The Barbell. Or immunization and why that can't be the whole story	42
3.6.2	Riding the yield curve	44
4	Arbitrage pricing in a one-period model	47
4.1	An appetizer.	48
4.2	The single period model	53
4.3	The economic intuition	59
5	Arbitrage pricing in the multi-period model	69
5.1	An appetizer	69
5.2	Price processes, trading and arbitrage	72
5.3	No arbitrage and price functionals	78
5.4	Conditional expectations and martingales	80
5.5	Equivalent martingale measures	83

5.6	One-period submodels	89
5.7	The Standard Binomial Model	90
6	Option pricing	97
6.1	Terminology	97
6.2	Diagrams, strategies and put-call parity	98
6.3	Restrictions on option prices	103
6.4	Binomial models for stock options	105
6.5	Pricing the European call	107
6.6	Hedging the European call	108
6.7	Recombining tree representation	111
6.8	The binomial model for American puts	114
6.9	Implied volatility	115
6.10	Portfolio insurance, implied volatility and crash fears	117
6.11	Debt and equity as options on firm value	118
7	The Black-Scholes formula	121
7.1	Black-Scholes as a limit of binomial models	121
7.2	The Black-Scholes model	125
7.3	A derivation of the Black-Scholes formula	128
7.3.1	Hedging the call	131
8	Stochastic Interest Rates	133
8.1	Introduction	133
8.2	Constructing an arbitrage free model	134
8.2.1	Constructing a Q -tree/lattice for the short rate that fits the initial term structure	137
8.3	Flat shifts of flat term structures	140
8.4	Forwards and futures	142
8.5	Swap contracts	146
8.6	Flat shifts of flat term structures	149
9	Portfolio Theory	151
9.1	Mathematics of Minimum Variance Portfolios	154
9.1.1	The case with no riskfree asset	154
9.1.2	The case with a riskfree asset	161
9.2	The Capital Asset Pricing Model (CAPM)	164
9.3	Relevant but unstructured remarks on CAPM	169
9.3.1	Systematic and non-systematic risk	169
9.3.2	Problems in testing the CAPM	170
9.3.3	Testing the efficiency of a given portfolio	171

Chapter 1

Preface

These notes are intended for the introductory finance course mathematics-economics program at the University of Copenhagen. They cover (the) basic pillars of finance: (1) analysis of deterministic cash-flows (Chapter 3), (2) mean-variance analysis and the capital asset pricing model (CAPM) (Chapter 9), (3) valuation by absence of arbitrage in multi-period models (Chapters 4-6). (For those with OCD: Chapter 2 is an introduction with two examples — which we will not really return to, Chapter 7 is a brief look at the continuous-time Black-Scholes model and formula, and Chapter 8 analyses stochastic interest rate models.)

The aim is to be mathematically precise without abandoning neither the economic intuition (such intuition is hard word, not just hand-waving) nor the ability to be quantitative (i.e. do calculations with sensible numbers).

Except for the brief introduction to the Black-Scholes model in Chapter 7, the presentation is done through discrete-time models emphasizing definitions and setups that prepare the students for the study of continuous-time models.

The notes are not littered with references to books and research papers. Let's say that is intentional. But let us mention two standard text-books — from which we have learned a lot — that cover roughly the same material: John Hull's *Options, Futures and Other Derivative Securities* (it comes in a new edition roughly every second year) and David Luenberger's lesser known *Investment Science* (whose only edition so far was published in 1997 by Oxford University Press).

Chapter 2

Introduction

A student applying for student loans is investing in his or her human capital. Typically, the income of a student is not large enough to cover living expenses, books etc., but the student is hoping that the education will provide future income which is more than enough to repay the loans. The government subsidizes students because it believes that the future income generated by highly educated people will more than compensate for the costs of subsidy, for example through productivity gains and higher tax revenues.

A first time home buyer is typically not able to pay the price of the new home up front but will have to borrow against future income and using the house as collateral.

A company which sees a profitable investment opportunity may not have sufficient funds to launch the project (buy new machines, hire workers) and will seek to raise capital by issuing stocks and/or borrowing money from a bank.

The student, the home buyer and the company are all in need of money to invest now and are confident that they will earn enough in the future to pay back loans that they might receive.

Conversely, a pension fund receives payments from members and promises to pay a certain pension once members retire.

Insurance companies receive premiums on insurance contracts and delivers a promise of future payments in the events of property damage or other unpleasant events which people wish to insure themselves against.

A new lottery millionaire would typically be interested in investing his or her fortune in some sort of assets (government bonds for example) since this will provide a larger income than merely saving the money in a mattress.

The pension fund, the insurance company and the lottery winner are all looking for profitable ways of placing current income in a way which will provide income in the future.

A key role of financial markets is to find efficient ways of connecting the demand for capital with the supply of capital. The examples above illustrated the need for economic agents to substitute income intertemporally. An equally important role of financial markets is to allow risk averse agents (such as insurance buyers) to share risk.

In understanding the way financial markets allocate capital we must understand the chief mechanism by which it performs this allocation, namely through prices. Prices govern the flow of capital, and in financial markets investors will compare the price of some financial security with its promised future payments. A very important aspect of this comparison is the riskiness of the promised payments. We have an intuitive feeling that it is reasonable for government bonds to give a smaller expected return than stocks in risky companies, simply because the government is less likely to default. But exactly how should the relationship between risk and reward (return on an investment) be in a well functioning market? Trying to answer that question is a central part of this course. The best answers delivered so far are in a set of mathematical models developed over the last 50 years or so. One set of models, with the so-called CAPM which we will meet later as the prime example, consider expected return and variance on return as the natural definitions of reward and risk, respectively and tries to answer how these should be related. Another set of models are based on arbitrage pricing, which is a very powerful application of the simple idea, that two securities which deliver the same payments should have the same price. This is typically illustrated through option pricing models and in the modelling of bond markets, but the methodology actually originated partly in work which tried to answer a somewhat different question, which is an essential part of financial theory as well: How should a firm finance its investments? Should it issue stocks and/or bonds or maybe something completely different? How should it (if at all) distribute dividends among shareholders? The so-called Modigliani-Miller theorems provide a very important starting point for studying these issues which currently are by no means resolved.

A historical survey of how finance theory has evolved will probably be more interesting at the end of the course since we will at that point understand versions of the central models of the theory.

But let us start by considering a classical explanation of the significance of financial markets in a microeconomic setting.

2.1 The role of financial markets

equilibrium
utility function

Consider the definition of a private ownership economy as standard economic textbooks: Assume for simplicity that there is only one good and one firm with production set Y . The i th consumer is characterized by a consumption set X_i , a preference preordering \preceq_i , an endowment ω_i and shares in the firm θ_i . Given a price system p , and given a profit maximizing choice of production y , the firm then has a profit of $\pi(p) = p \cdot y$ and this profit is distributed to shareholders such that the wealth of the i th consumer becomes

$$w_i = p \cdot \omega_i + \theta_i \pi(p) \quad (2.1)$$

The definition of an equilibrium in such an economy then has three seemingly natural requirements: The firm maximizes profits, consumers maximize utility subject to their budget constraint and markets clear, i.e. consumption equals the sum of initial resources and production. But why should the firm maximize its profits? After all, the firm has no utility function, only consumers do. But note that given a price system p , the shareholders of the firm all agree that it is desirable to maximize profits, for the higher profits the larger the consumers wealth, and hence the larger is the set of feasible consumption plans, and hence the larger is the attainable level of utility. In this way the firm's production choice is separated from the shareholders' choice of consumption. There are many ways in which we could imagine shareholders disagreeing over the firm's choice of production. Some examples could include cases where the choice of production influences on the consumption sets of the consumers, or if we relax the assumption of price taking behavior, where the choice of production plan affects the price system and thereby the initial wealth of the shareholders. Let us, by two examples, illustrate in what sense the price system changes the behavior of agents.

Example 1. Consider a single agent who is both a consumer and a producer. The agent has an initial endowment $e_0 > 0$ of the date 0 good and has to divide this endowment between consumption at date 0 and investment in production of a time 1 good. Assume that only non-negative consumption is allowed. Through investment in production, the agent is able to transform an input of i_0 into $f(i_0)$ units of date 1 consumption. The agent has a utility function $U(c_0, c_1)$ which we assume is strictly increasing. The agent's problem is then to maximize utility of consumption, i.e. to maximize $U(c_0, c_1)$ subject to the constraints $c_0 + i_0 \leq e_0$ and $c_1 = f(i_0)$ and we may rewrite this problem as

$$\begin{aligned} \max v(c_0) &\equiv U(c_0, f(e_0 - c_0)) \\ \text{subject to } c_0 &\leq e_0 \end{aligned}$$

If we impose regularity conditions on the functions f and U (for example that they are differentiable and strictly concave and that utility of zero consumption in either period is $-\infty$) then we know that at the maximum c_0^* we will have $0 < c_0^* < e_0$ and $v'(c_0^*) = 0$ i.e.

$$D_1U(c_0^*, f(e_0 - c_0^*)) \cdot 1 - D_2U(c_0^*, f(e_0 - c_0^*))f'(e_0 - c_0^*) = 0$$

where D_1 means differentiation after the first variable. Defining i_0^* as the optimal investment level and $c_1^* = f(e_0 - c_0^*)$, we see that

$$f'(i_0^*) = \frac{D_1U(c_0^*, c_1^*)}{D_2U(c_0^*, c_1^*)}$$

and this condition merely says that the marginal rate of substitution in production is equal to the marginal rate of substitution of consumption.

The key property to note in this example is that what determines the production plan in the absence of prices is the preferences for consumption of the consumer. If two consumers with no access to trade owned shares in the same firm, but had different preferences and identical initial endowments, they would bitterly disagree on the level of the firm's investment.

Example 2. Now consider the setup of the previous example but assume that a price system (p_0, p_1) (whose components are strictly positive) gives the consumer an additional means of transferring date 0 wealth to date 1 consumption. Note that by selling one unit of date 0 consumption the agent acquires $\frac{p_0}{p_1}$ units of date 1 consumption, and we define $1 + r = \frac{p_0}{p_1}$. The initial endowment must now be divided between three parts: consumption at date 0 c_0 , input into production i_0 and s_0 which is sold in the market and whose revenue can be used to purchase date 1 consumption in the market.

With this possibility the agent's problem becomes that of maximizing $U(c_0, c_1)$ subject to the constraints

$$\begin{aligned} c_0 + i_0 + s_0 &\leq e_0 \\ c_1 &\leq f(i_0) + (1 + r)s_0 \end{aligned}$$

and with monotonicity constraints the inequalities may be replaced by equalities. Note that the problem then may be reduced to having two decision variables c_0 and i_0 and maximizing

$$v(c_0, i_0) \equiv U(c_0, f(i_0) + (1 + r)(e_0 - c_0 - i_0)).$$

Again we may impose enough regularity conditions on U (strict concavity, twice differentiability, strong aversion to zero consumption) to ensure that it

attains its maximum in an interior point of the set of feasible pairs (c_0, i_0) Fisher Separation and that at this point the gradient of v is zero, i.e.

$$\begin{aligned} D_1U(c_0^*, c_1^*) \cdot 1 - D_2U(c_0^*, f(i_0^*) + (1+r)(e_0 - c_0^* - i_0^*))(1+r) &= 0 \\ D_2U(c_0^*, f(i_0^*) + (1+r)(e_0 - c_0^* - i_0^*))(f'(i_0^*) - (1+r)) &= 0 \end{aligned}$$

With the assumption of strictly increasing U , the only way the second equality can hold, is if

$$f'(i_0^*) = (1+r)$$

and the first equality holds if

$$\frac{D_1U(c_0^*, c_1^*)}{D_2U(c_0^*, c_1^*)} = (1+r)$$

We observe two significant features:

First, the production decision is independent of the utility function of the agent. Production is chosen to a point where the marginal benefit of investing in production is equal to the 'interest rate' earned in the market. The consumption decision is separate from the production decision and the marginal condition is provided by the market price. In such an environment we have what is known as Fisher Separation where the firm's decision is independent of the shareholder's utility functions. Such a setup rests critically on the assumptions of the perfect competitive markets where there is price taking behavior and a market for both consumption goods at date 0. Whenever we speak of firms having the objective of maximizing shareholders' wealth we are assuming an economy with a setup similar to that of the private ownership economy of which we may think of the second example as a very special case.

Second, the solution to the maximization problem will typically have a higher level of utility for the agent at the optimal point: Simply note that any feasible solution to the first maximization problem is also a solution to the second. This is an improvement which we take as a 'proof' of the significance of the existence of markets. If we consider a private ownership economy equilibrium, the equilibrium price system will see to that consumers and producers coordinate their activities simply by following the price system and they will obtain higher utility than if each individual would act without a price system as in example 1.


2.2 The Kelly criterion: Optimal betting

Suppose a bookie offers us the following classical scenario: for the price of 1 [units of currency] we may enter a game in which we gain $b + 1$ with a

odds
edge

probability p , but lose everything with a probability of $q = 1 - p$. Here, b is a positive constant known as the **odds**. (People familiar with betting will note that these are so-called UK odds, i.e. how odds are quoted in Britain. Often it's most convenient to use decimal odds (why?), but here that would ruin the aesthetics of the resulting formula.)

Graphically, the pay-off structure of the game may be represented as the binomial tree diagram:



Clearly, our expected winnings from entering such a game is

$$\underbrace{p \cdot (b + 1) + q \cdot 0}_{\text{expected pay-off}} - \underbrace{1}_{\text{cost}} = pb + p - 1 = pb - q,$$

where the last equality uses the definition $q = 1 - p$. $pb - q$ is sometimes referred to as your **edge** if greater than zero.

Insofar as we repeat the game in such a manner that we in every round bet the fraction f of our total wealth ($W(t)$ in round t) then we have that

$$W(t + 1) = W(t)X(t + 1),$$

where $t = 0, 1, 2, \dots, T$ and

$$X(t + 1) = \begin{cases} 1 - f + (b + 1)f = 1 + bf, & \text{with probability } p \\ 1 - f, & \text{with probability } q. \end{cases}$$

Using this relation repeatedly we find that

$$W(T) = W(0) \prod_{t=1}^T X(t). \quad (2.2)$$

We can now introduce the logarithmic growth rate $R(T)$ for our wealth:

Kelly criterion

$$\frac{W(T)}{W(0)} = e^{R(T)T} \Leftrightarrow R(T) = \frac{1}{T} \ln \left(\frac{W(T)}{W(0)} \right).$$

Inserting the relation (2.2) and using the law of large numbers we find that

$$R(T) = \frac{1}{T} \sum_{t=1}^T \ln X(t) \xrightarrow{T \rightarrow \infty} E(\ln X),$$

where X encodes the common distribution of each of the $X(t)$ s (which plausibly have been assumed i.i.d.). Upon noticing that the X distribution depends on f we may now get the idea of solving

$$\max_f E(\ln X) := \max_f \{p \ln(1 + bf) + q \ln(1 - f)\},$$

which will give us the maximal growth rate for our personal wealth (asymptotically - but deterministically). Differentiating and equating to zero we find that

$$\left. \frac{\partial E(\ln X)}{\partial f} \right|_{f=f^*} = \frac{pb}{1 + f^*b} - \frac{q}{1 - f^*} = 0,$$

which may be re-arranged to give

$$\boxed{f^* = \frac{bp - q}{b}} \quad (2.3)$$

This growth optimal betting strategy is commonly referred to as the **Kelly bet** or **Kelly criterion** after the physicist John Larry Kelly, Jr. who first derived the formula in 1956.¹ Clearly, there's a trade-off between expected winnings (the numerator) and odds (the denominator). This is hardly surprising: if the odds are high, then conceivably there is also a low probability of winning.

Although the Kelly criterion is a well established result in investment theory, which reportedly has been used by both Warren Buffet and James Harries Simons, it is not altogether uncontroversial. E.g. one might argue that an individual's specific investing constraints may override his or her desire for an optimal growth rate. Kelly himself apparently never used his own criterion to make money.

¹The formula may be recalled using the mnemonic "edge over odds".

Chapter 3

Payment Streams under Certainty

3.1 Financial markets and arbitrage

In this section we consider a very simple setup with no uncertainty. There are three reasons that we do this:

First, the terminology of bond markets is conveniently introduced in this setting, for even if there were uncertainty in our model, bonds would be characterized by having payments whose size at any date are constant and known in advance.

Second, the classical net present value (NPV) rule of capital budgeting is easily understood in this framework.

And finally, the mathematics introduced in this section will be extremely useful in later chapters as well.

A note on notation: If $v \in \mathbb{R}^N$ is a vector the following conventions for “vector positivity” are used:

- $v \geq 0$ (“ v is non-negative”) means that all of v 's coordinates are non-negative. ie. $\forall i: v_i \geq 0$.
- $v > 0$ (“ v is positive”) means that $v \geq 0$ and that at least one coordinate is strictly positive, ie. $\forall i: v_i \geq 0$ and $\exists i: v_i > 0$, or differently that $v \geq 0$ and $v \neq 0$.
- $v \gg 0$ (“ v is strictly positive”) means that every coordinate is strictly positive, $\forall i: v_i > 0$. This (when v is N -dimensional) we will sometimes write as $v \in \mathbb{R}_{++}^N$. (This saves a bit of space, when we want to indicate both strict positivity and the dimension of v .)

financial market
 security price
 system
 payment stream
 portfolio
 short position
 long position
 arbitrage
 opportunity

Throughout we use v^\top to denote the transpose of the vector v . Vectors without the transpose sign are always thought of as column vectors.

We now consider a model for a financial market (sometimes also called a security market or price system; individual components are then referred to as securities) with $T + 1$ dates: $0, 1, \dots, T$ and no uncertainty.

Definition 1. A financial market consists of a pair (π, C) where $\pi \in \mathbb{R}^N$ and C is an $N \times T$ -matrix.

The interpretation is as follows: By paying the price π_i at date 0 one is entitled to a stream of payments (c_{i1}, \dots, c_{iT}) at dates $1, \dots, T$. Negative components are interpreted as amounts that the owner of the security has to pay. There are N different payment streams trading. But these payment streams can be bought or sold in any quantity and they may be combined in portfolios to form new payment streams:

Definition 2. A portfolio θ is an element of \mathbb{R}^N . The payment stream generated by θ is $C^\top \theta \in \mathbb{R}^T$. The price of the portfolio θ at date 0 is $\pi \cdot \theta$ ($= \pi^\top \theta = \theta^\top \pi$).

Note that allowing portfolios to have negative coordinates means that we allow securities to be sold. We often refer to a negative position in a security as a short position and a positive position as a long position. Short positions are not just a convenient mathematical abstraction. For instance when you borrow money to buy a home, you take a short position in bonds.

Before we even think of adopting (π, C) as a model of a security market we want to check that the price system is sensible. If we think of the financial market as part of an equilibrium model in which the agents use the market to transfer wealth between periods, we clearly want a payment stream of $(1, \dots, 1)$ to have a lower price than that of $(2, \dots, 2)$. We also want payment streams that are non-negative at all times to have a non-negative price. More precisely, we want to rule out arbitrage opportunities in the security market model:

Definition 3. A portfolio θ is an arbitrage opportunity (of type 1 or 2) if it satisfies one of the following conditions:

1. $\pi \cdot \theta = 0$ and $C^\top \theta > 0$.
2. $\pi \cdot \theta < 0$ and $C^\top \theta \geq 0$.

Alternatively, we can express this as $(-\pi \cdot \theta, C^\top \theta) > 0$.

In short, an arbitrage (opportunity) is a “a money machine” or “a free lunch”, i.e. something that is really too good to be true. Note that an arbitrage is much better than just the favourable bets that we encountered in Chapter 2. *arbitrage, odds*

Example 3. To illustrate the arguably abstract concept of an arbitrage, consider the following odds (decimal odds to be precise; i.e. betting \$1 and winning gives you \$1 · odds back) that a number internet bookmakers put on the 2004 African Nation’s Cup match between Burkina Faso and Mali.

Bookmaker	Burkina Faso - Mali		
	1 (B F win)	X (draw)	2 (Mali win)
Aebet	5.50	3.10	1.61
Bet-at-home.com	3.65	3.20	1.75
EasyBets	4.20	3.30	1.73
Expekt	4.05	3.15	1.85
InterWetten	3.50	2.80	2.00
MrBookmaker	4.60	3.05	1.73

Now imagine that that we pick the best odds for each outcome and bet \$ $1/5.5 = 0.1818$ on Burkina Faso, $\$1/3.3 = 0.3030$ on a draw, and $\$1/2.0 = 0.5$ on Mali. The total cost of this is \$0.9848. Irrespective of what happens we win \$1. This is an arbitrage. As an exercise, try to formulate this strategy in the (C, π) and θ -formalism from above (and below). (What are we implicitly saying about a risk-free asset/cash?) If our framework were taken completely literally, any two different odds on the same outcome would constitute an arbitrage. That, however, is not the situation in practice. A quoted odds only means that the bookmaker will take your money and pay you back if you win, not the other way round. Or differently put, they are selling prices.

The example above notwithstanding, a prudent financial assumption is that markets do not contain arbitrages.

Definition 4. *The security market is arbitrage-free if it contains no arbitrage opportunities.*

If arbitrages do exist, then we would of love to find them. The way to do that, however, is to study closely the consequences of absence of arbitrage. If they are violated, then there must be arbitrage, and our means of analysis give us constructive ways of find it/them. It turns out that there is a simple characterization of arbitrage-free markets. For that we need a lemma that is very similar to Farkas’ theorem of alternatives, which is often encountered when duality for linear programming is studied.

Stiemke's lemma
discount factors
complete market

Lemma 1. (Stiemke's lemma) Let A be an $n \times m$ -matrix: Then precisely one of the following two statements is true:

1. There exists $x \in \mathbb{R}_{++}^m$ such that $Ax = 0$.
2. There exists $y \in \mathbb{R}^n$ such that $y^\top A > 0$.

We will not prove the lemma here, but rather use it as a steppingstone to our next theorem:

Theorem 1. The security market (π, C) is arbitrage-free if and only if there exists a strictly positive vector $d \in \mathbb{R}_{++}^T$ such that $\pi = Cd$.

In the context of our security market the vector d will be referred to as a vector of discount factors. This use of language will be clear shortly.

Proof. Define the matrix

$$A = \begin{pmatrix} -\pi_1 & c_{11} & c_{12} & \cdots & c_{1T} \\ -\pi_2 & c_{21} & c_{22} & \cdots & c_{2T} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ -\pi_N & c_{N1} & c_{N2} & \cdots & c_{NT} \end{pmatrix}$$

First, note that the existence of $x \in \mathbb{R}_{++}^{T+1}$ such that $Ax = 0$ is equivalent to the existence of a vector of discount factors since we may define

$$d_i = \frac{x_i}{x_0} \quad i = 1, \dots, T.$$

Hence if the first condition of Stiemke's lemma is satisfied, a vector d exists such that $\pi = Cd$. The second condition corresponds to the existence of an arbitrage opportunity: If $y^\top A > 0$ then we have either

$$(y^\top A)_1 > 0 \text{ and } (y^\top A)_i \geq 0 \quad i = 1, \dots, T + 1$$

or

$$(y^\top A)_1 = 0, y^\top A \geq 0 \text{ and } (y^\top A)_i > 0 \quad \text{some } i \in \{2, \dots, T + 1\}$$

and this is precisely the condition for the existence of an arbitrage opportunity. Now use Stiemke's lemma. ■

Another important concept is market completeness (in Danish: *Komplethed* or *fuldstændighed*).

Definition 5. *The security market is complete if for every $y \in \mathbb{R}^T$ there exists a $\theta \in \mathbb{R}^N$ such that $C^\top \theta = y$.*

zero coupon bond,
ZCB
discount factors

In linear algebra terms this means that the rows of C span \mathbb{R}^T , which can only happen if $N \geq T$, and in our interpretation it means that any desired payment stream can be generated by an appropriate choice of portfolio.

Theorem 2. *Assume that (π, C) is arbitrage-free. Then the market is complete if and only if there is a unique vector of discount factors.*

Proof. Since the market is arbitrage-free we know that there exists $d \gg 0$ such that $\pi = Cd$. Now if the model is complete then \mathbb{R}^T is spanned by the columns of C^\top , ie. the rows of C of which there are N . This means that C has T linearly independent rows, and from basic linear algebra (look around where *rank* is defined) it also has T linearly independent columns, which is to say that all the columns are independent. They therefore form a basis for a T -dimensional linear subspace of \mathbb{R}^N (remember we must have $N \geq T$ to have completeness), ie. any vector in this subspace has unique representation in terms of the basis-vectors. Put differently, the equation $Cx = y$ has at most one solution. And in case where $y = \pi$, we know there is one by absence of arbitrage. For the other direction assume that the model is incomplete. Then the columns of C are linearly dependent, and that means that there exists a vector $\tilde{d} \neq 0$ such that $0 = C\tilde{d}$. Since $d \gg 0$, we may choose $\epsilon > 0$ such that $d + \epsilon\tilde{d} \gg 0$. Clearly, this produces a vector of discount factors different from d . ■

3.2 Zero coupon bonds and the term structure

Assume throughout this section that the model (π, C) is complete and arbitrage-free and let $d^\top = (d_1, \dots, d_T)$ be the unique vector of discount factors. Since there must be at least T securities to have a complete model, C must have at least T rows. On the other hand if C has exactly T linearly independent rows, then adding other securities to C will not add any more possibilities of wealth transfer to the market. Hence we can assume that C is an invertible $T \times T$ matrix.

Definition 6. *The payment stream of a zero coupon bond with maturity t is given by the t 'th unit vector e_t of \mathbb{R}^T .*

Next we see why the words discount factors were chosen:

Proposition 1. *The price of a zero coupon bond with maturity t is d_t .*

forward rates
short rate

Proof. Let θ_t be the portfolio such that $C^\top \theta_t = e_t$. Then

$$\pi^\top \theta_t = (Cd)^\top \theta_t = d^\top C^\top \theta_t = d^\top e_t = d_t$$

■

Note from the definition of d that we get the value of a stream of payments c by computing $\sum_{t=1}^T c_t d_t$. In other words, the value of a stream of payments is obtained by discounting back the individual components. There is nothing in our definition of d which prevents $d_s > d_t$ even when $s > t$, but in the models we will consider this will not be relevant: It is safe to think of d_t as decreasing in t corresponding to the idea that the longer the maturity of a zero coupon bond, the smaller is its value at time 0.

From the discount factors we may derive/define various types of interest rates which are essential in the study of bond markets.:

Definition 7. (*Short and forward rates.*) *The short rate at date 0 is given by*

$$r_0 = \frac{1}{d_1} - 1.$$

The (one-period) time t - forward rate at date 0, is equal to

$$f(0, t) = \frac{d_t}{d_{t+1}} - 1,$$

where $d_0 = 1$ by convention.

The interpretation of the short rate should be straightforward: Buying $\frac{1}{d_1}$ units of a maturity 1 zero coupon bond costs $\frac{1}{d_1} d_1 = 1$ at date 0 and gives a payment at date 1 of $\frac{1}{d_1} = 1 + r_0$. The forward rate tells us the rate at which we may agree at date 0 to borrow (or lend) between dates t and $t + 1$. To see this, consider the following strategy at time 0 :

- Sell 1 zero coupon bond with maturity t .
- Buy $\frac{d_t}{d_{t+1}}$ zero coupon bonds with maturity $t + 1$.

Note that the amount raised by selling precisely matches the amount used for buying and hence the cash flow from this strategy at time 0 is 0. Now consider what happens if the positions are held to the maturity date of the bonds: At date t the cash flow is then -1 and at date $t + 1$ the cash flow is $\frac{d_t}{d_{t+1}} = 1 + f(0, t)$.

Definition 8. *The yield (or yield to maturity) at time 0 of a zero coupon bond with maturity t is given as*

$$y(0, t) = \left(\frac{1}{d_t} \right)^{\frac{1}{t}} - 1.$$

*term structure of
interest rates
yield to maturity
compounding
periods*

Note that

$$d_t(1 + y(0, t))^t = 1.$$

and that one may therefore think of the yield as an 'average interest rate' earned on a zero coupon bond. In fact, the yield is a geometric average of forward rates:

$$1 + y(0, t) = ((1 + f(0, 0)) \cdots (1 + f(0, t - 1)))^{\frac{1}{t}}$$

Definition 9. *The term structure of interest rates (or the yield curve) at date 0 is given by $(y(0, 1), \dots, y(0, T))$.*

Note that if we have any one of the vector of yields, the vector of forward rates and the vector of discount factors, we may determine the other two. Therefore we could equally well define a term structure of forward rates and a term structure of discount factors. In these notes unless otherwise stated, we think of the term structure of interest rates as the yields of zero coupon bonds as a function of time to maturity. It is important to note that the term structure of interest rate depicts yields of zero coupon bonds. We do however also speak of yields on securities with general positive payment streams:

Definition 10. *The yield (or yield to maturity) of a security $c^\top = (c_1, \dots, c_T)$ with $c > 0$ and price π is the unique solution $y > -1$ of the equation*

$$\pi = \sum_{i=1}^T \frac{c_i}{(1 + y)^i}.$$

Example 4 (Compounding Periods). In most of the analysis in this chapter the time is "stylized"; it is measured in some unit (which we think of and refer to as "years") and cash-flows occur at dates $\{0, 1, 2, \dots, T\}$. But it is often convenient (and not hard) to work with dates that are not integer multiples of the fundamental time-unit. We quote interest rates in units of years⁻¹ ("per year"), but to any interest rate there should be a number, m , associated stating how often the interest is compounded. By this we mean the following: If you invest 1 \$ for n years at the m -compounded rate r_m you end up with

$$\left(1 + \frac{r_m}{m} \right)^{mn}. \quad (3.1)$$

continuously
compounded
interest rate
annuity
serial loan
bullet bond

The standard example: If you borrow 1\$ in the bank, a 12% interest rate means they will add 1% to you debt each month (i.e. $m = 12$) and you will end up paying back 1.1268 \$ after a year, while if you make a deposit, they will add 12% after a year (i.e. $m = 1$) and you will of course get 1.12\$ back after one year. If we keep r_m and n fixed in (3.1) (and then drop the m -subscript) and let m tend to infinity, it is well known that we get:

$$\lim_{m \rightarrow \infty} \left(1 + \frac{r}{m}\right)^{mn} = e^{nr},$$

and in this case we will call r the continuously compounded interest rate. In other words: If you invest 1 \$ and the continuously compounded rate r_c for a period of length t , you will get back e^{tr_c} . Note also that a continuously compounded rate r_c can be used to find (uniquely for any m) r_m such that 1 \$ invested at m -compounding corresponds to 1 \$ invested at continuous compounding, i.e.

$$\left(1 + \frac{r_m}{m}\right)^m = e^{r_c}.$$

This means that in order to avoid confusion – even in discrete models – there is much to be said in favor of quoting interest rates on a continuously compounded basis. But then again, in the highly stylized discrete models it would be pretty artificial, so we will not do it (rather it will always be $m = 1$).

And then a final piece of advice: Whenever you do calculations be careful always to plug in interest rates as decimal numbers, not as percentages. There is a large difference between $e^{0.12}$ and e^{12} , much larger than what can be recovered by dividing the end result by 100.

3.3 Annuities, serial loans and bullet bonds

Typically, zero-coupon bonds do not trade in financial markets and one therefore has to deduce prices of zero-coupon bonds from other types of bonds trading in the market. Three of the most common types of bonds which do trade in most bond markets are annuities, serial loans and bullet bonds. (In literature relating to the American market, “bond” is usually understood to mean “bullet bond with 2 yearly payments”. Further, “bills” are term short bonds, annuities explicitly referred to as such, and serial loans rare.) We now show how knowing to which of these three types a bond belongs and knowing three characteristics, namely the maturity, the principal and the coupon rate, will enable us to determine the bond’s cash flow completely.

Let the principal or face value of the bond be denoted F . Payments on the bond start at date 1 and continue to the time of the bond’s maturity, which

we denote τ . The payments are denoted c_t . We think of the principal of a *annuity* bond with coupon rate R and payments c_1, \dots, c_τ as satisfying the following difference equation:

$$p_t = (1 + R)p_{t-1} - c_t \quad t = 1, \dots, \tau, \quad (3.2)$$

with the boundary conditions $p_0 = F$ and $p_\tau = 0$.

Think of p_t as the remaining principal right after a payment at date t has been made. For accounting and tax purposes and also as a helpful tool in designing particular types of bonds, it is useful to split payments into a part which serves as reduction of principal and one part which is seen as an interest payment. We define the reduction in principal at date t as

$$\delta_t = p_{t-1} - p_t$$

and the interest payment as

$$i_t = Rp_{t-1} = c_t - \delta_t.$$

Definition 11. *An annuity with maturity τ , principal F and coupon rate R is a bond whose payments are constant between dates 1 and τ , and whose principal evolves according to Equation (3.2).*

With constant payments we can use (3.2) repeatedly to write the remaining principal at time t as

$$p_t = (1 + R)^t F - c \sum_{j=0}^{t-1} (1 + R)^j \quad \text{for } t = 1, 2, \dots, \tau.$$

To satisfy the boundary condition $p_\tau = 0$ we must therefore have

$$F - c \sum_{j=0}^{\tau-1} (1 + R)^{j-\tau} = 0,$$

so by using the well-known formula $\sum_{i=0}^{n-1} x^i = (x^n - 1)/(x - 1)$ for the summation of a geometric series, we get

$$c = F \left(\sum_{j=0}^{\tau-1} (1 + R)^{j-\tau} \right)^{-1} = F \frac{R(1 + R)^\tau}{(1 + R)^\tau - 1} = F \frac{R}{1 - (1 + R)^{-\tau}}.$$

Note that the size of the payment is homogeneous (of degree 1) in the principal, so it's usually enough to look at the $F = 1$. (This rather trivial observation can in fact be extremely useful in a dynamic context.) It is common

alfahage; \$“alpha
 n“rceil R \$
 bullet bond
 serial loan

to use the shorthand notation

$$\alpha_{n|R} = (\text{“Alfahage”}) = \frac{(1+R)^n - 1}{R(1+R)^n}.$$

Having found what the size of the payment must be we may derive the interest and the deduction of principal as well: Let us calculate the size of the payments and see how they split into deduction of principal and interest payments. First, we derive an expression for the remaining principal:

$$\begin{aligned} p_t &= (1+R)^t F - \frac{F}{\alpha_{\tau|R}} \sum_{j=0}^{t-1} (1+R)^j \\ &= \frac{F}{\alpha_{\tau|R}} \left((1+R)^t \alpha_{\tau|R} - \frac{(1+R)^t - 1}{R} \right) \\ &= \frac{F}{\alpha_{\tau|R}} \left(\frac{(1+R)^\tau - 1}{R(1+R)^{\tau-t}} - \frac{(1+R)^t - 1}{R(1+R)^{\tau-t}} \right) \\ &= \frac{F}{\alpha_{\tau|R}} \alpha_{\tau-t|R}. \end{aligned}$$

This gives us the interest payment and the deduction immediately for the annuity:

$$\begin{aligned} i_t &= R \frac{F}{\alpha_{\tau|R}} \alpha_{\tau-t+1|R} \\ \delta_t &= \frac{F}{\alpha_{\tau|R}} (1 - R \alpha_{\tau-t+1|R}). \end{aligned}$$

In the definition of an annuity, the size of the payments is implicitly defined. The definitions of bullets and serials are more direct.

Definition 12. A bullet bond¹ with maturity τ , principal F and coupon rate R is characterized by having $i_t = c_t$ for $t = 1, \dots, \tau - 1$ and $c_\tau = (1+R)F$.

The fact that we have no reduction in principal before τ forces us to have $c_t = RF$ for all $t < \tau$.

Definition 13. A serial loan or bond with maturity τ , principal F and coupon rate R is characterized by having δ_t , constant for all $t = 1, \dots, \tau$.

¹In Danish: Et stående lån

Since the deduction in principal is constant every period and we must have $p_\tau = 0$, it is clear that $\delta_t = \frac{F}{\tau}$ for $t = 1, \dots, \tau$. From this it is straightforward to calculate the interest using $i_t = Rp_{t-1}$.

We summarize the characteristics of the three types of bonds in the table below:

	payment	interest	deduction of principal
Annuity	$F\alpha_{\tau R}^{-1}$	$R\frac{F}{\alpha_{\tau R}}\alpha_{\tau-t+1 R}$	$\frac{F}{\alpha_{\tau R}}(1 - R\alpha_{\tau-t+1 R})$
Bullet	RF for $t < \tau$ $(1 + R)F$ for $t = \tau$	RF	0 for $t < \tau$ F for $t = \tau$
Serial	$\frac{F}{\tau} + R\left(F - \frac{t-1}{\tau}F\right)$	$R\left(F - \frac{t-1}{\tau}F\right)$	$\frac{F}{\tau}$

Example 5 (A Simple Bond Market). Consider the following bond market where time is measured in years and where payments are made at dates $\{0, 1, \dots, 4\}$:

Bond (i)	Coupon rate (R_i)	Price at time 0 ($\pi_i(0)$)
1 yr bullet	5	100.00
2 yr bullet	5	99.10
3 yr annuity	6	100.65
4 yr serial	7	102.38

We are interested in finding the zero-coupon prices/yields in this market. First we have to determine the payment streams of the bonds that are traded (the C -matrix). Since $\alpha_{3|6} = 2.6730$ we find that

$$C = \begin{bmatrix} 105 & 0 & 0 & 0 \\ 5 & 105 & 0 & 0 \\ 37.41 & 37.41 & 37.41 & 0 \\ 32 & 30.25 & 28.5 & 26.75 \end{bmatrix}$$

Clearly this matrix is invertible so $e_t = C^\top \theta_t$ has a unique solution for all $t \in \{1, \dots, 4\}$ (namely $\theta_t = (C^\top)^{-1}e_t$). If the resulting t -zero-coupon bond prices, $d_t(0) = \pi(0) \cdot \theta_t$, are strictly positive then there is no arbitrage. Performing the inversion and the matrix multiplications we find that

$$(d_1(0), d_2(0), d_3(0), d_4(0))^\top = (0.952381, 0.898458, 0.839618, 0.7774332),$$

or alternatively the following zero-coupon yields

$$100 * (y(0, 1), y(0, 2), y(0, 3), y(0, 4))^\top = (5.00, 5.50, 6.00, 6.50).$$

clean price
dirty price

Now suppose that somebody introduces a 4 yr annuity with a coupon rate of 5 % . Since $\alpha_{4|5} = 3.5459$ this bond has a unique arbitrage-free price of

$$\pi_5(0) = \frac{100}{3.5459} (0.952381 + 0.898458 + 0.839618 + 0.7774332) = 97.80.$$

Notice that bond prices are always quoted per 100 *units* (e.g. \$ or DKK) of principal. This means that if we assume the yield curve is the same at time 1 the price of the serial bond would be quoted as

$$\pi_4(1) = \frac{d_{1:3}(0) \cdot C_{4,2:4}}{0.75} = \frac{76.87536}{0.75} = 102.50$$

(where $d_{1:3}(0)$ means the first 3 entries of $d(0)$ and $C_{4,2:4}$ means the entries 2 to 4 in row 4 of C).

Example 6 (Reading the financial pages). This example gives concrete calculations for a specific Danish Government bond traded at the Copenhagen Stock Exchange(CSX): A bullet bond with a 4 % coupon rate and yearly coupon payments that matures on January 1 2010. Around February 1 2005 you could read the following on the CSX homepage or on the financial pages of decent newspapers

Bond type	Current date	Maturity date	Price	Yield
4% bullet	February 1 2005	January 1 2010	104.02	3.10 %

Let us see how the yield was calculated. First, we need to set up the cash-flow stream that results from buying the bond. The first cash-flow, π in the sense of Definition 8 would take place today. (Actually it wouldn't, even these days trades take a couple of day to be in effect; *valør* in Danish. We don't care here.) And how large is it? By convention, and reasonably so, the buyer has to pay the price (104.02; this is called the *clean price*) plus compensate the seller of the bond for the accrued interest over the period from January 1 to February 1, ie. for 1 month, which we take to mean 1/12 of a year. (This is not as trivial as it seems. In practice there are a lot of finer - and extremely boring - points about how days are counted and fractions calculated. Suffice it to say that mostly actual days are used in Denmark.) By definition the buyer has to pay accrued interest of "coupon \times year-fraction", ie. $4 \times 1/12 = 0.333$, so the total payment (called the *dirty price*) is $\pi = 104.35$. So now

we can write down the cash-flows and verify the yield calculation:

Date	t_k	Cash-flow (c_k)	$d_k = (1 + 0.0310)^{-t_k}$	PV = $d_k * c_k$
Feb. 1 2005	0	- 104.35	1	
Jan. 1 2006	$\frac{11}{12}$	4	0.9724	3.890
Jan. 1 2007	$1\frac{11}{12}$	4	0.9432	3.772
Jan. 1 2008	$2\frac{11}{12}$	4	0.9148	3.660
Jan. 1 2009	$3\frac{11}{12}$	4	0.8873	3.549
Jan. 1 2010	$4\frac{11}{12}$	104	0.8606	89.505
				SUM = 104.38

(The match, 104.35 vs. 104.38 isn't perfect. But to 3 significant digits 0.0310 is the best solution, and anything else can be attributed to out rough approach to exact dates.)

Example 7 (Finding the yield curve). In early February you could find prices 4%-coupon rate bullet bonds with a range of different maturities (all maturities fall on January firsts):

Maturity year	2006	2007	2008	2009	2010
Clean price	101.46	102.69	103.43	103.88	104.02
Maturity year	2011	2012	2013	2014	2015
Clean price	103.80	103.50	103.12	102.45	102.08

These bonds (with names like 4%10DsINKx) are used for the construction of private home-owners variable/floating rate loans such as “FlexLån”. (Hey! How does the interest rate get floating? Well, it does if you (completely) refinance your 30-year loan every year or every 5 years with shorter maturity bonds.) In many practical contexts these are not the right bonds to use; yield curves “should” be inferred from government bonds. (Of course this statement makes no sense within our modelling framework.)

Dirty prices, these play the role of π , are found as in Example 6, and the (10 by 10) C -matrix has the form

$$C_{i,j} = \begin{cases} 4 & \text{if } j < i \\ 104 & \text{if } j = i \\ 0 & \text{if } j > i \end{cases}$$

The system $Cd = \pi$ has the positive (\sim no arbitrage) unique (\sim completeness) solution

$$d = (0.9788, 0.9530, 0.9234, 0.8922, 0.8593, 0.8241, 0.7895, 0.7555, 0.7200, 0.6888)^\top.$$




Figure 3.1: The term structure of interest rates in Denmark, February 2005. The o's are the points we have actually calculated, the rest is just linear interpolation.

and that corresponds to these (yearly compounded) zero coupon yields:

Maturity	0.92	1.92	2.92	3.92	4.92	5.92	6.92	7.92	8.92	9.92
ZC yield in %	2.37	2.55	2.77	2.95	3.13	3.32	3.48	3.61	3.75	3.83

as depicted in Figure 3.1. Estimating yield curves (also known determining discount factors) is a very important, though not particularly glamorous, task in the financial sector. Two things that make it challenging are (1) there are more relevant payments dates than there bonds, (2) following the 2007-8-9 financial crisis/turmoil credit/default/bankruptcy risk can be/is being ignored less.

Example 8. Saving for retirement. Annuity type calculations are very useful for pension savings calculations. Suppose that a newly (i.e. at time 0) graduated person saves the fraction x of his or her salary (assumed fixed) every year for retirement, and that the pension savings have a yearly (deterministic) rate of return of r . We assume that payments are made at times $1, 2, \dots, T$ where the person retires. The pension is then paid out in yearly installments of y (which can be interpreted as a fraction of the pre-retirement salary) at times $T + 1, T + 2, \dots, T + \tau$. Using the geometric summation

	A	B	C	D	E	F	G	H
1	Savings as fraction of salary (x)	0.180328						
2	Rate of return on savings (r)	0.025		VPI(T)	9.905101		-1E-06	
3	Years until retirement (T)	35		VPO(T)	9.905102			
4	Years in retirement (tau)	15						
5	Pension payout as fraction of salary (y)	0.8						
6								
7								
8								
9								
10								
11								

internal rate of
return

Figure 3.2: Pension savings calculations done in Excel.

formula $\sum_{j=0}^{n-1} z^j = \frac{z^n - 1}{z - 1}$ we can find the value at time T of the money that has been paid in to be

$$VPI(T) := x + x(1+r) + \dots + x(1+r)^{T-1} = x \frac{(1+r)^T - 1}{r}.$$

Similarly, the time- T value of the pension payouts is

$$VPO(T) := \frac{y}{1+r} + \dots + \frac{y}{(1+r)^\tau} = y \frac{1 - (1+r)^{-\tau}}{r}.$$

Various strategies for saving for retirement can then be analyzed by studying $VPI(T) = VPO(T)$. Or more concretely, by fixing all-but-one input parameters and solving (possibly numerically) for the remaining one. A spreadsheet is very well-suited for this. Figure 3.2 shows a numerical example. The 18% savings rate is typical for Denmark. The 2.5% rate of return (which we should think of as being in excess of inflation, that we ignore here) is on the conservative side. Or rather it should be, but is actually above what most pension companies will promise you these days.

3.4 The net present value (NPV) rule

Similarly to yield, we can define the so-called internal rate of return (IRR) for an arbitrary cash flow stream, i.e. on securities which may have negative cash flows as well:

Definition 14. An internal rate of return of a security (c_1, \dots, c_T) with price $\pi \neq 0$ is a solution (with $y > -1$) of the equation

$$\pi = \sum_{i=1}^T \frac{c_i}{(1+y)^i}.$$

net present value,
NPV
NPV criterion

Hence the definitions of yield and internal rate of return are identical for positive cash flows. It is easy to see that for securities whose future payments are both positive and negative we may have several IRRs. This is one reason that one should be very careful interpreting and using this measure at all when comparing cash flows. We will see below that there are even more serious reasons. When judging whether a certain cash flow is 'attractive' the correct measure to use is net present value:

Definition 15. *The PV and NPV of security (c_1, \dots, c_T) with price c_0 given a term structure $(y(0, 1), \dots, y(0, T))$ are defined as*

$$PV(c) = \sum_{i=1}^T \frac{c_i}{(1 + y(0, i))^i}$$

$$NPV(c) = \sum_{i=1}^T \frac{c_i}{(1 + y(0, i))^i} - c_0$$

Next, we will see how these concepts are used in deciding how to invest under certainty.

Assume throughout this section that we have a complete security market as defined in the previous section. Hence a unique discount function d is given as well as the associated concepts of interest rates and yields. We let y denote the term structure of interest rates and use the short hand notation y_i for $y(0, i)$.

In capital budgeting we analyze how firms should invest in projects whose payoffs are represented by cash flows. Whereas we assumed in the security market model that a given security could be bought or sold in any quantity desired, we will use the term project more restrictively: We will say that the project is scalable by a factor $\lambda \neq 1$ if it is possible to start a project which produces the cash flow λc by paying λc_0 initially. A project is not scalable unless we state this explicitly and we will not consider any negative scaling.

In a complete financial market an investor who needs to decide on only one project faces a very simple decision: Accept the project if and only if it has positive NPV. We will see why this is shortly. Accepting this fact we will see examples of some other criteria which are generally inconsistent with the NPV criterion. We will also note that when a collection of projects are available capital budgeting becomes a problem of maximizing NPV over the range of available projects. The complexity of the problem arises from the constraints that we impose on the projects. The available projects may be non-scalable or scalable up to a certain point, they may be mutually exclusive (i.e. starting one project excludes starting another), we may impose

restrictions on the initial outlay that we will allow the investor to make (representing limited access to borrowing in the financial market), we may assume that a project may be repeated once it is finished and so on. In all cases our objective is simple: Maximize NPV.

First, let us note why looking at NPV is a sensible thing to do:

Proposition 2. *Given a cash flow $c = (c_1, \dots, c_T)$ and given c_0 such that $NPV(c_0; c) < 0$. Then there exists a portfolio θ of securities whose price is c_0 and whose payoff satisfies*

$$C^T \theta > \begin{pmatrix} c_1 \\ \vdots \\ c_T \end{pmatrix}.$$

Conversely, if $NPV(c_0; c) > 0$, then every θ with $C^T \theta = c$ satisfies $\pi^T \theta > c_0$.

Proof. Since the security market is complete, there exists a portfolio θ^c such that $C^T \theta^c = c$. Now $\pi^T \theta^c < c_0$ (why?), hence we may form a new portfolio by investing the amount $c_0 - \pi^T \theta^c$ in some zero coupon bond (e_1 , say) and also invest in θ^c . This generates a stream of payments equal to $C^T \theta^c + \frac{(c_0 - \pi^T \theta^c)}{d_1} e_1 > c$ and the cost is c_0 by construction.

The second part is left as an exercise. ■

The interpretation of this lemma is the following: One should never accept a project with negative NPV since a strictly larger cash flow can be obtained at the same initial cost by trading in the capital market. On the other hand, a positive NPV project generates a cash flow at a lower cost than the cost of generating the same cash flow in the capital market. It might seem that this generates an arbitrage opportunity since we could buy the project and sell the corresponding future cash flow in the capital market generating a profit at time 0. However, we insist on relating the term arbitrage to the capital market only. Projects should be thought of as 'endowments': Firms have an available range of projects. By choosing the right projects the firms maximize the value of these 'endowments'.

Some times when performing NPV-calculations, we assume that 'the term structure is flat'. What this means is that the discount function has the particularly simple form

$$d_t = \frac{1}{(1+r)^t}$$

for some constant r , which we will usually assume to be non-negative, although our model only guarantees that $r > -1$ in an arbitrage-free market. A flat term structure is very rarely observed in practice - a typical real world

Gordon's growth
formula
capital budgeting

term structure will be upward sloping: Yields on long maturity zero coupon bonds will be greater than yields on short bonds. Reasons for this will be discussed once we model the term structure and its evolution over time - a task which requires the introduction of uncertainty to be of any interest. When the term structure is flat then evaluating the NPV of a project having a constant cash flow is easily done by summing the geometric series. The present value of n payments starting at date 1, ending at date n each of size c , is

$$\sum_{i=1}^n cd^i = cd \sum_{i=0}^{n-1} d^i = cd \frac{1-d^n}{1-d}, \quad d \neq 1$$

Another classical formula concerns the present value of a geometrically growing payment stream $(c, c(1+g), \dots, c(1+g)^{n-1})$ as

$$\begin{aligned} & \sum_{i=1}^n c \frac{(1+g)^{i-1}}{(1+r)^i} \\ &= \frac{c}{1+r} \sum_{i=0}^{n-1} \frac{(1+g)^i}{(1+r)^i} \\ &= \frac{c}{r-g} \left(1 - \left(\frac{1+g}{1+r} \right)^n \right). \end{aligned}$$

Although we have not taken into account the possibility of infinite payment streams, we note for future reference, that for $0 \leq g < r$ we have what is known as Gordon's growth formula:

$$\sum_{i=1}^{\infty} \frac{c(1+g)^{i-1}}{(1+r)^i} = \frac{c}{r-g}.$$

Example 9. *Some seemingly sensible rules that are inconsistent with the NPV rule.* Corresponding to our definition of internal rate of return in Chapter 3, we define an internal rate of return on a project c with initial cost $c_0 > 0$, denoted $IRR(c_0; c)$, as a solution to the equation

$$c_0 = \sum_{i=1}^T \frac{c_i}{(1+x)^i}, \quad x > -1$$

As we have noted earlier such a solution need not be unique unless $c > 0$ and $c_0 > 0$. Note that an internal rate of return is defined without referring to the underlying term structure. The internal rate of return describes the level of a flat term structure at which the NPV of the project is 0. The idea behind its use in capital budgeting would then be to say that the higher the

level of the interest rate, the better the project (and some sort of comparison with the existing term structure would then be appropriate when deciding whether to accept the project at all). But as we will see in the following example, *IRR* and *NPV* may disagree on which project is better: Consider the projects shown in the table below (whose last column shows a discount function d):

date	proj 1	proj 2	d
0	-100	-100	1
1	50	50	0.95
2	5	80	0.85
3	90	4	0.75
IRR	0.184	0.197	-
NPV	19.3	18.5	-

Project 2 has a higher IRR than project 1, but 1 has a larger NPV than 2. Using the same argument as in the previous section it is easy to check, that even if a cash flow similar to that of project 2 is desired by an investor, he would be better off investing in project 1 and then reforming the flow of payments using the capital market. Another problem with trying to use IRR as a decision variable arises when the IRR is not uniquely defined - something which typically happens when the cash flows exhibit sign changes. Which IRR should we then choose? One might also contemplate using the payback method and count the number of years it takes to recover the initial cash outlay - possibly after discounting appropriately the future cash flows. Project 2 in the table has a payback of 2 years whereas project 1 has a payback of three years. The example above therefore also shows that choosing projects with the shortest payback time may be inconsistent with the NPV method.

3.4.1 Several projects

Consider someone with $c_0 > 0$ available at date 0 who wishes to allocate this capital over the $T + 1$ dates, and who considers a project c with initial cost c_0 . We have seen that precisely when $NPV(c_0; c) > 0$ this person will be able to obtain better cash flows by adopting c and trading in the capital market than by trading in the capital market alone.

When there are several projects available the situation really does not change much: Think of the i 'th project (p_0^i, p) as an element of a set $P_i \subset \mathbb{R}^{T+1}$. Assume that $0 \in P_i$ all i representing the choice of not starting the i 'th project. For a non-scalable project this set will consist of one point in addition to 0.

Given a collection of projects represented by $(P_i)_{i \in I}$. Situations where there is a limited amount of money to invest at the beginning (and borrowing is not permitted), where projects are mutually exclusive etc. may then be described abstractly by the requirement that the collection of selected projects $(p_0^i, p^i)_{i \in I}$ are chosen from a feasible subset P of the Cartesian product $\times_{i \in I} P_i$. The NPV of the chosen collection of projects is then just the sum of the NPVs of the individual projects and this in turn may be written as the NPV of the sum of the projects:

$$\sum_{i \in I} NPV(p_0^i; p^i) = NPV \left(\sum_{i \in I} (p_0^i, p^i) \right).$$

Hence we may think of the chosen collection of projects as producing one project and we can use the result of the previous section to note that clearly an investor should choose a project giving the highest NPV. In practice, the maximization over feasible “artificial” may not be easy at all.

Let us look at an example.

Example 10. (*Adapted from Luenberger (1997)*) A company need a certain type of machine to produce *widgets*. The machine costs \$10,000 (say, to paid at time 0) to buy and its yearly maintenance costs grow linearly; \$2,000 in year 1 (say, paid at time 1), 3,000 in year 2, and so on. At any time the company can buy a new machine (suppose the old one has 0 scrap value). The new machine has (even in nominal terms) the same price and cost profile. The yield curve is flat at 10%. How often should the machine be changed? Let’s analyze: Changing every year gives the cash flow stream $(-10, -2, 0, \dots) + (0, -10, -2, 0, \dots) + (0, 0, -10, -2, 0, \dots) + \dots$. The present value of this (up to change of sign and division by 1,000) must (as everything starts over at time time 1) solve

$$PV = 10 + 2/1.1 + PV/1.1 \Rightarrow PV = 130.$$

(Pedants should verify that the infinite sum we work with here are sufficiently convergent for such manipulation to be allowed.) If instead we change every k years and denote the present value of the all payments by $PV_{k;total}$ then

$$PV_{k;total} = PV_{k;1\ cycle} + \left(\frac{1}{1.1} \right)^k PV_{k;total},$$

where $PV_{k;1\ cycle} = 10 + \sum_{j=1}^k \frac{j+1}{1.1^j}$. The task is to chose k , such that the total PV minimized (yes, minimized; we changed the sign). This is done

numerically:

k	$PV_{k;total}$
1	130.00
2	82.38
3	69.58
4	65.36
5	64.48
6	65.20
7	66.76
8	68.79
9	71.09
10	73.53

So changing after five years is optimal.

The moral of this section is simple: Given a perfect capital market, investors who are offered projects should simply maximize NPV. This is merely an equivalent way of saying that profit maximization with respect to the existing price system (as represented by the term structure) is the appropriate strategy when a perfect capital market exists. The technical difficulties arise from the constraints that we impose on the projects and these constraints easily lead to linear programming problems, integer programming problems or even non-linear optimization problems.

However, real world projects typically do not generate cash flows which are known in advance. Real world projects involve risk and uncertainty and therefore capital budgeting under certainty is really not sophisticated enough for a manager deciding which projects to undertake. A key objective of this course is to try and model uncertainty and to construct models of how risky cash flows are priced. This will give us definitions of NPV which work for uncertain cash flows as well.

3.5 Duration, convexity and immunization.

3.5.1 Duration with a flat term structure.

In this chapter we introduce the notions of duration and convexity which are often used in practical bond risk management and asset/liability management. It is worth stressing that when we introduce dynamic models of the term structure of interest rates in a world with uncertainty, we obtain much more sophisticated methods for measuring and controlling interest rate risk than the ones presented in this section.

*duration,
Macaulay*

Consider an arbitrage-free and complete financial market where the discount function $d = (d_1, \dots, d_T)$ satisfies

$$d_i = \frac{1}{(1+r)^i} \text{ for } i = 1, \dots, T.$$

This corresponds to the assumption of a flat term structure. We stress that this assumption is rarely satisfied in practice but we will see how to relax this assumption.

What we are about to investigate are changes in present values as a function of changes in r . We will speak freely of 'interest changes' occurring even though strictly speaking, we still do not have uncertainty in our model.

With a flat term structure, the present value of a payment stream $c = (c_1, \dots, c_T)$ is given by

$$PV(c; r) = \sum_{t=1}^T \frac{c_t}{(1+r)^t}$$

We have now included the dependence on r explicitly in our notation since what we are about to model are essentially derivatives of $PV(c; r)$ with respect to r .

Definition 16. *Let c be a non-negative payment stream. The Macaulay duration $D(c; r)$ of c is given by*

$$\begin{aligned} D(c; r) &= \left(-\frac{\partial}{\partial r} PV(c; r) \right) \frac{1+r}{PV(c; r)} \\ &= \frac{1}{PV(c; r)} \sum_{t=1}^T t \frac{c_t}{(1+r)^t} \end{aligned} \quad (3.3)$$

The Macaulay duration and is the classical one (many more advanced durations have been proposed in the literature). Note that rather than saying it is based on a flat term structure, we could refer to it as being based on the yield of the bond (or portfolio).

If we define

$$w_t = \frac{c_t}{(1+r)^t} \frac{1}{PV(c; r)}, \quad (3.4)$$

then we have $\sum_{t=1}^T w_t = 1$, hence

$$D(c; r) = \sum_{t=1}^T t w_t.$$

This shows that duration has a dual interpretation. On the one hand it is *convexity* (by our definition) a price sensitivity (a sign changed elasticity, to be precise) to interest rate changes. On the other hand it is (as it turns out from the math above) a value-weighted average of payment dates.

Definition 17. *The convexity of c is given by*

$$K(c; r) = \sum_{t=1}^T t^2 w_t. \quad (3.5)$$

where w_t is given by (3.4).

Let us try to interpret D and K by computing the first and second derivatives² of $PV(c; r)$ with respect to r .

$$\begin{aligned} PV'(c; r) &= - \sum_{t=1}^T t c_t \frac{1}{(1+r)^{t+1}} \\ &= - \frac{1}{1+r} \sum_{t=1}^T t c_t \frac{1}{(1+r)^t} \\ PV''(c; r) &= \sum_{t=1}^T t(t+1) \frac{c_t}{(1+r)^{t+2}} \\ &= \frac{1}{(1+r)^2} \left[\sum_{t=1}^T t^2 c_t \frac{1}{(1+r)^t} + \sum_{t=1}^T t c_t \frac{1}{(1+r)^t} \right] \end{aligned}$$

Now consider the relative change in $PV(c; r)$ when r changes to $r + \Delta r$, i.e.

$$\frac{PV(c; r + \Delta r) - PV(c; r)}{PV(c; r)}$$

By considering a second order Taylor expansion of the numerator, we obtain

$$\begin{aligned} \frac{PV(c; r + \Delta r) - PV(c; r)}{PV(c; r)} &\approx \frac{PV'(c; r)\Delta r + \frac{1}{2}PV''(c; r)(\Delta r)^2}{PV(c; r)} \\ &= -D \frac{\Delta r}{(1+r)} + \frac{1}{2}(K + D) \left(\frac{\Delta r}{1+r} \right)^2 \end{aligned}$$

Hence D and K can be used to approximate the relative change in $PV(c; r)$ as a function of the relative change in r (or more precisely, relative changes in $1 + r$, since $\frac{\Delta(1+r)}{1+r} = \frac{\Delta r}{1+r}$).

²From now on we write $PV'(c; r)$ and $PV''(c; r)$ instead of $\frac{\partial}{\partial r}PV(c; r)$ resp. $\frac{\partial^2}{\partial r^2}PV(c; r)$

duration,
modified

Sometimes one finds the expression modified duration defined by

$$MD(c; r) = \frac{D}{1+r}$$

and using this in a first order approximation, we get the relative change in $PV(c; r)$ expressed by $-MD(c; r)\Delta r$, which is a function of Δr itself. The interpretation of D as a price elasticity gives us no reasonable explanation of the word 'duration', which certainly leads one to think of quantity measured in units of time. If we use the definition of w_t we have the following simple expression for the duration:

$$D(c; r) = \sum_{t=1}^T t w_t.$$

Notice that w_t expresses the present value of c_t divided by the total present value, i.e. w_t expresses the weight by which c_t is contributing to the total present value. Since $\sum_{t=1}^T w_t = 1$ we see that $D(c; r)$ may be interpreted as a 'mean waiting time'. The payment which occurs at time t is weighted by w_t .

Example 11. For the bullet bond in Example 6 the present value of the payment stream is 104.35 and $y = 0.0310$, so therefore the Macaulay duration is

$$\frac{\sum_{k=1}^4 t_k c_k (1+y)^{-t_k}}{PV} = \frac{475.43}{104.35} = 4.556$$

while the convexity is

$$\frac{\sum_{k=1}^4 t_k^2 c_k (1+y)^{-t_k}}{PV} = \frac{2266.35}{104.35} = 21.72,$$

and the following table shows the the exact and approximated relative changes in present value when the yield changes:

Yield	Δ yield	Exact rel. (%) PV-change	First order approximation	Second order approximation
0.021	-0.010	4.57	4.42	4.54
0.026	-0.005	2.27	2.21	2.24
0.031	0	0	0	0
0.036	0.005	-2.15	-2.21	-2.18
0.041	0.010	-4.27	-4.42	- 4.30

Notice that since PV is a decreasing, convex function of y we know that the first order approximation will underestimate the effect of decreasing y (and overestimate the effect of increasing it).

Notice that for a zero coupon bond with time to maturity t the duration is t . For other kinds of bonds with time to maturity t , the duration is less than t . Furthermore, note that investing in a zero coupon bond with yield to maturity r and holding the bond to expiration guarantees the owner an annual return of r between time 0 and time t . This is not true of a bond with maturity t which pays coupons before t . For such a bond the duration has an interpretation as the length of time for which the bond can ensure an annual return of r :

Let $FV(c; r, H)$ denote the (future) value of the payment stream c at time H if the interest rate is fixed at level r . Then

$$\begin{aligned} FV(c; r, H) &= (1+r)^H PV(c; r) \\ &= \sum_{t=1}^{H-1} c_t (1+r)^{H-t} + c_H + \sum_{t=H+1}^T c_t \frac{1}{(1+r)^{t-H}} \end{aligned}$$

Consider a change in r which occurs an instant after time 0. How would such a change affect $FV(c; r, H)$? There are two effects with opposite directions which influence the future value: Assume that r decreases. Then the first sum in the expression for $FV(c; r, H)$ will decrease. This decrease can be seen as caused by reinvestment risk: The coupons received up to time H will have to be reinvested at a lower level of interest rates. The last sum will increase when r decreases. This is due to price risk : As interest rates fall the value of the remaining payments after H will be higher since they have to be discounted by a smaller factor. Only c_H is unchanged.

The natural question to ask then is for which H these two effects cancel each other. At such a time point we must have $\frac{\partial}{\partial r} FV(c; r, H) = 0$ since an infinitesimal change in r should have no effect on the future value. Now,

$$\begin{aligned} \frac{\partial}{\partial r} FV(c; r, H) &= \frac{\partial}{\partial r} [(1+r)^H PV(c; r)] \\ &= H(1+r)^{H-1} PV(c; r) + (1+r)^H PV'(c; r) \end{aligned}$$

Setting this expression equal to 0 gives us

$$\begin{aligned} H &= \frac{-PV'(c; r)}{PV(c; r)} (1+r) \\ \text{i.e. } H &= D(c; r) \end{aligned}$$

Furthermore, at $H = D(c; r)$, we have $\frac{\partial^2}{\partial r^2} FV(c; r, H) > 0$. This you can check by computing $\frac{\partial^2}{\partial r^2} ((1+r)^H PV(c; r))$, reexpressing in terms D and K ,

immunization

and using the fact that $K > D^2$. Hence, at $H = D(c; r)$, $FV(c; r, H)$ will have a minimum in r . We say that $FV(c; r, H)$ is immunized towards changes in r , but we have to interpret this expression with caution: The only way a bond really can be immunized towards changes in the interest rate r between time 0 and the investment horizon t is by buying zero coupon bonds with maturity t . Whenever we buy a coupon bond at time 0 with duration t , then to a first order approximation, an interest change immediately after time 0, will leave the future value at time t unchanged. However, as date 1 is reached (say) it will not be the case that the duration of the coupon bond has decreased to $t - 1$. As time passes, it is generally necessary to adjust bond portfolios to maintain a fixed investment horizon, even if r is unchanged. This is true even in the case of certainty.

Later when we introduce dynamic hedging strategies we will see how a portfolio of bonds can be dynamically managed so as to truly immunize the return.

3.5.2 Relaxing the assumption of a flat term structure.

What we have considered above were parallel changes in a flat term structure. Since we rarely observe this in practice, it is natural to try and generalize the analysis to different shapes of the term structure. Consider a family of structures given by a function r of two variables, t and x . Holding x fixed gives a term structure $r(\cdot, x)$.

For example, given a current term structure (y_1, \dots, y_T) we could have $r(t, x) = y_t + x$ in which case changes in x correspond to additive changes in the current term structure (the one corresponding to $x = 0$). Or we could have $1 + r(t, x) = (1 + y_t)x$, in which case changes in x would produce multiplicative changes in the current (obtained by letting $x = 1$) term structure.

Now let us compute changes in present values as x changes:

$$\frac{\partial PV}{\partial x} = - \sum_{t=1}^T t c_t \frac{1}{(1 + r(t, x))^{t+1}} \frac{\partial r(t, x)}{\partial x}$$

which gives us

$$\frac{\partial PV}{\partial x} \frac{1}{PV} = - \sum_{t=1}^T \frac{t w_t}{1 + r(t, x)} \frac{\partial r(t, x)}{\partial x}$$

where

$$w_t = \frac{c_t}{(1 + r(t, x))^t} \frac{1}{PV}$$

We want to try and generalize the 'investment horizon' interpretation of duration, and hence calculate the future value of the payment stream at

time H and differentiate with respect to x . Assume that the current term structure is $r(\cdot, x_0)$.

duration
duration,
Fisher-Weil

$$FV(c; r(H, x_0), H) = (1 + r(H, x_0))^H PV(c; r(t, x_0))$$

Differentiating we get

$$\begin{aligned} \frac{\partial}{\partial x} FV(c; r(H, x), H) &= (1 + r(H, x))^H \frac{\partial PV}{\partial x} \\ &\quad + H(1 + r(H, x))^{H-1} \frac{\partial r(H, x)}{\partial x} PV(c; r(t, x)) \end{aligned}$$

Evaluate this derivative at $x = x_0$ and set it equal to 0:

$$\left. \frac{\partial PV}{\partial x} \right|_{x=x_0} \frac{1}{PV} = -H \left. \frac{\partial r(H, x)}{\partial x} \right|_{x=x_0} (1 + r(H, x_0))^{-1}$$

and hence we could define the duration corresponding to the given parametrization as the value D for which

$$\left. \frac{\partial PV}{\partial x} \right|_{x=x_0} \frac{1}{PV} = -D \left. \frac{\partial r(D, x)}{\partial x} \right|_{x=x_0} (1 + r(D, x_0))^{-1}.$$

The additive case would correspond to

$$\left. \frac{\partial r(D, x)}{\partial x} \right|_{x=0} = 1,$$

and the multiplicative case to

$$\left. \frac{\partial r(D, x)}{\partial x} \right|_{x=1} = 1 + y_D.$$

The multiplicative case is by far the most common one. So common that many sources other choices. This has name Fisher-Weil duration),

$$D_{FW} = -\frac{\partial PV}{\partial x} \frac{1}{PV} = \sum_{t=1}^T tw_t.$$

Given the value-weighted average of payment dates interpretation of (Macaulay) duration, this is exactly what we would conjecture a duration measure based on a non-flat term structure to look like. But by going through this analysis, we maintaining the connection to yield curve shifts.

barbell strategy

A slightly different path to Fisher-Weil duration is this: Let us write the discount function as

$$d(t; x) = \exp(-t \times (z_0(t) + x)),$$

where we suppose $x = 0$ gives us today's zero coupon yield curve. The variable x then creates parallel (additive) shifts in the continuously compounded zero coupon rates. We have $PV(c, x) = \sum_t c_t d(t; x)$ and from this

$$-\frac{1}{PV(c, x)} \frac{\partial PV(c, x)}{\partial x} \Big|_{x=0} = D_{FW}(c)$$

So Fisher-Weil duration is sensitivity to additive shifts in continuously compounded rates.

Example 12 (Macaulay vs. Fisher-Weil). Consider again the small bond market from Example 5. We have already found the zero-coupon yields in the market, and find that the Fisher-Weil duration of the 4 yr serial bond is

$$\frac{1}{102.38} \left(\frac{32}{1.0500} + \frac{2 * 30.25}{1.0550^2} + \frac{3 * 28.5}{1.0600^3} + \frac{4 * 26.75}{1.0650^4} \right) = 2.342,$$

and the following table gives the yields, Macaulay durations based on yields and Fisher-Weil durations for all the coupon bonds:

Bond	Yield ()	M-duration	FW-duration
1 yr bullet	5	1	1
2 yr bullet	5.49	1.952	1.952
3 yr annuity	5.65	1.963	1.958
4 yr serial	5.93	2.354	2.342

So not much difference.

Similarly, the Fisher-Weil duration of the bullet bond from Examples 6, 7 and 11 is 4.552, whereas its Macaulay duration was 4.556.

3.6 Two examples to mess with your head

3.6.1 The Barbell. Or immunization and why that can't be the whole story

We finish this chapter with an example (with something usually referred to as a barbell strategy) which is intended to cause some concern. Some of the claims are for you to check!

A financial institution issues 100 million \$ worth of 10 year bullet bonds with time to maturity 10 years and a coupon rate of 7 percent. Assume that the term structure is flat at $r = 7$ percent. The revenue (of 100 million \$) is used to purchase 10-and 20 year annuities also with coupon rates of 7%. The numbers of the 10 and 20 year annuities purchased are chosen in such a way that the duration of the issued bullet bond matches that of the portfolio of annuities. Now there are three facts you need to know at this stage. Letting T denote time to maturity, r the level of the term structure and γ the coupon rate, we have that the duration of an annuity is given by

$$D_{ann} = \frac{1+r}{r} - \frac{T}{(1+r)^T - 1}.$$

Note that since payments on an annuity are equal in all periods we need not know the size of the payments to calculate the duration.

The duration of a bullet bond is

$$D_{bullet} = \frac{1+r}{r} - \frac{1+r - T(r-R)}{R((1+r)^T - 1) + r}$$

which of course simplifies when $r = R$.

The third fact you need to check is that if a portfolio consists of two securities whose values are P_1 and P_2 respectively, then the duration of the portfolio $P_1 + P_2$ is given as

$$D(P_1 + P_2) = \frac{P_1}{P_1 + P_2} D(P_1) + \frac{P_2}{P_1 + P_2} D(P_2).$$

Using these three facts you will note that a portfolio consisting of 23.77 million dollars worth of the 10-year annuity and 76.23 million dollars worth of the 20-year annuity will produce a portfolio whose duration exactly matches that of the issued bullet bond. By construction the present value of the two annuities equals that of the bullet bond. The present value of the whole transaction in other words is 0 at an interest level of 7 percent. However, for all other levels of the interest rate, the present value is strictly positive! In other words, any change away from 7 percent will produce a profit to the financial institution.

What this example shows is that our fundamentally deterministic framework is not good enough to deal with uncertainty, with changes; from seemingly sensible assumptions we get out paradoxical results. (After treating arbitrage-free multi-period stochastic model, we will show in Section 8.6 that we can't have only flat yield curves in an arbitrage-free model.)

Riding the yield
curve

3.6.2 Riding the yield curve

For some notation. Slightly cumbersome, but we need it. For a calendar date t and a time to maturity τ , let $y(t, \tau)$ be the continuously compounded zero-coupon rate for time to maturity τ , i.e. the (annualized) rate of return we get by investing (at time t) in zero-coupon bonds maturing at time $t + \tau$ and holding them until they mature. The mapping

$$\tau \mapsto y(t, \tau)$$

we call the (time- t zero-coupon) yield curve. In terms of zero-coupon bond prices ($P(t, T)$, first argument current time t , second argument the maturity date T) we have

$$P(t, T) = e^{-(T-t)y(t, T-t)}.$$

Suppose we want to invest at time 0, look one year ahead, and have at our disposal zero-coupon bonds for all possible maturity dates. The rate of return we get from investing in a maturity date T zero-coupon bond is

$$i_1(T) := \frac{P(t, T) - P(0, T)}{P(0, T)} = \frac{P(1, T)}{P(0, T)} - 1.$$

(The notation “:=” means “equal to by definition”, with the term nearest the : being defined.) A natural question is: Can we maximize this rate of return by choosing an appropriate T . The short answer is no, not without being able to look into the future. We do not know until time 1 what $P(1, T)$ turns out to be. So we need to make further modelling assumptions. A natural first step is the hypothesis,

H_0 : The time 1 yield curve will be the same as the time 0 yield curve.

At time 1, a maturity date T zero-coupon bond has time t to maturity $T - 1$, so under the H_0 -hypothesis we have

$$P(1, T) = e^{-(T-1)y(0, T-1)} = P(0, T - 1).$$

Thus the rate of return becomes known. It is in fact our old friend the forward rate,

$$i_1(T) = \frac{P(0, T - 1)}{P(0, T)} - 1 = f(0, T - 1).$$

So to maximize we should invest according to the highest forward rate. (Or more precisely: Find the time to maturity for the maximal forward rate and then invest in zero-coupon bonds with 1 year more to maturity than that.)

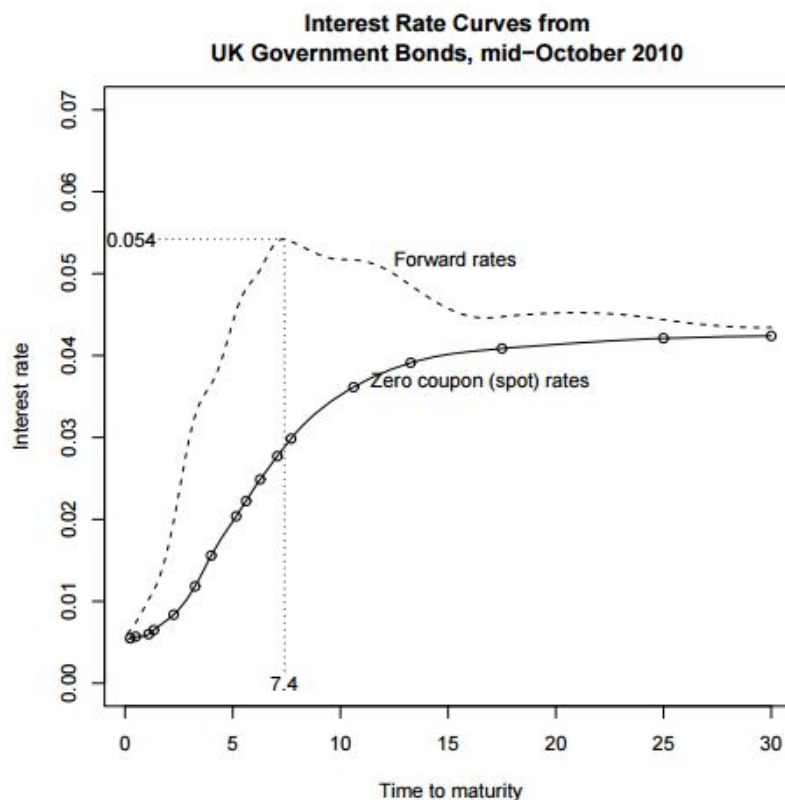


Figure 3.3: UK zero-coupon yield and forward rate curves from mid-October 2010.

And that forward rate will then be our return. This strategy is called “riding the yield curve”. Note that it can be carried out whether H_0 holds or not but only in the former case are we sure what our return will be. In words, this strategy says that to maximize investment returns, go not where the yield curve is at its highest, but where it is at its steepest. If the yield curve is “truly curved” then this can have surprising effects.

A good example is provided by the UK yield curve from mid-October 2010. The zero-coupon and (1-year-ahead) forward curves are shown in Figure 3.3. The circles are (more or less) observed points on the zero-coupon curve. The smooth curve was fitted through them with an interpolation technique called cubic spline. The smooth curve was then used to calculate forward rates. (Notice how the forward rate curve is considerable less

smooth than the zero-coupon rate curve. How to deal with this is the focus of numerous research articles.) We see that the forward rate curve attains its maximum around 7.5 years, and that the maximal value (5.4% continuously compounded) is considerably above any zero-coupon rate; that curve has its maximum at 4.2% for 30-year maturities. Could this trick be repeated over and over for, say, 30 years 1 would grow to $e^{30 \cdot 0.054} = 5.08$, while investing in maturity-date 30 zero coupon bonds and holding pays back only 3.57. That difference should make any pension fund manager sit up and take notice.

Magic? Alchemy? Or: Is yield curve riding really the free lunch that it seems to be? Of course not. First, it's risky. We only get the return we think if hypothesis H_0 holds, i.e. if the yield curve does not move. And that is a big if. And the longer the maturity of the bond we have invested in, the greater the sensitivity. Second, we may reverse the question and ask: How should the yield curve move for all 1-year returns to be the same? It turns out that if future zero-coupon spot rates are realized at the current forward rates then no gains can be achieved by short term riding or rolling. (It should be added: "for appropriately matched times to maturity and years-ahead". To make the statement precise we need three time indices, so we'll spare the reader.) This then leads to a counter-argument called the unbiased expectations hypothesis,

H_1 : Forward rates are expected future zero-coupon (spot) rates.

So which hypothesis is it then? Well, the truth (if such a thing exists at all) is somewhere in between. First, these are technical arguments (to do with expectations of non-linear functions, Jensen's Inequality, and absence of arbitrage) against the unbiased expectations hypothesis. But the main reason is risk-aversion: If all bonds give the same expected return, then why invest in risky ones at all? Thus prices of long-term bonds will be "low" and one is rewarded for taking the riskier positions such as riding the yield curve. (But, arguably, riding the yield curve gives "double exposure: It's risky and to the extent that it's not, movements will go against you!")

Chapter 4

Arbitrage pricing in a one-period model

One of the biggest success stories of financial economics is the *Black-Scholes model of option pricing*. But even though the formula itself is easy to use, a rigorous presentation of how it comes about requires some fairly sophisticated mathematics. Fortunately, the so-called binomial model of option pricing offers a much simpler framework and gives almost the same level of understanding of the way option pricing works. Furthermore, the binomial model turns out to be very easy to generalize (to so-called multinomial models) and more importantly to use for pricing other derivative securities (i.e. different contract types or different underlying securities) where an extension of the Black-Scholes framework would often turn out to be difficult. The flexibility of binomial models is the main reason why these models are used daily in trading all over the world.

Binomial models are often presented separately for each application. For example, one often sees the "classical" binomial model for pricing options on stocks presented separately from binomial term structure models and pricing of bond options etc.

The aim of this chapter is to present the underlying theory at a level of abstraction which is high enough to understand all binomial/multinomial approaches to the pricing of derivative securities as special cases of one model.

Apart from the obvious savings in allocation of brain RAM that this provides, it is also the goal to provide the reader with a language and framework which will make the transition to continuous-time models in future courses much easier.

binomial model,
one-period
stock
bank account
call option,
European


4.1 An appetizer.

Before we introduce our model of a financial market with uncertainty formally, we present a little appetizer. Despite its simplicity it contains most of the insights that we are about to get in this chapter.

Consider a one-period model with two states of nature, ω_1 and ω_2 . At time $t = 0$ nothing is known about the time state, at time $t = 1$ the state is revealed. State ω_1 occurs with probability p . Two securities are traded:

- A stock which costs S at time 0 and is worth uS at time 1 in one state and dS in the other.¹
- A money market account or bank account which costs 1 at time 0 and is worth R at time 1 regardless of the state.

Assume $0 < d < R < u$. (This condition will be explained later). We summarize the setup with a graph:



Now assume that we introduce into the economy a *European call option*² on the stock with exercise (or strike) price K and expiry (sometimes called maturity, although this is primarily used for bonds) 1. At time 1 the value of this call is equal to (where the notation $[y]^+$ (or sometimes $(y)^+$) means $\max(y, 0)$)

$$C_1(\omega) = \begin{cases} [uS - K]^+ & \text{if } \omega = \omega_1 \\ [dS - K]^+ & \text{if } \omega = \omega_2 \end{cases}$$

We will discuss options in more detail later. For now, note that it can be thought of as a contract giving the owner the right but not the obligation to buy the stock at time 1 for K .

¹We use *stock* as a generic term for a risky asset whose stochastic price behaviour we take as given. Words such as “share” or “equity” are largely synonymous.

²The term “European” is used for historical reasons; it has no particular meaning today, and it is doubtful if it ever had.

To simplify notation, let $C_u = C_1(\omega_1)$ and $C_d = C_1(\omega_2)$. The question is: replication
 What should the price of this call option be at time 0? A simple replication argument will give the answer: Let us try to form a portfolio at time 0 using only the stock and the money market account which gives the same payoff as the call at time 1 regardless of which state occurs. Let (a, b) denote, respectively, the number of stocks and units of the money market account held at time 0. If the payoff at time 1 has to match that of the call, we must have

$$a(uS) + bR = C_u$$

$$a(dS) + bR = C_d$$

Subtracting the second equation from the first we get

$$a(u - d)S = C_u - C_d$$

i.e.

$$a = \frac{C_u - C_d}{S(u - d)}$$

and algebra gives us

$$b = \frac{1}{R} \frac{uC_d - dC_u}{(u - d)}$$

where we have used our assumption that $u > d$. The cost of forming the portfolio (a, b) at time 0 is

$$\begin{aligned} \frac{(C_u - C_d)}{S(u - d)} S + \frac{1}{R} \frac{uC_d - dC_u}{(u - d)} \cdot 1 &= \frac{R(C_u - C_d)}{R(u - d)} + \frac{1}{R} \frac{uC_d - dC_u}{(u - d)} \\ &= \frac{1}{R} \left[\frac{R - d}{u - d} C_u + \frac{u - R}{u - d} C_d \right]. \end{aligned}$$

We will formulate below exactly how to define the notion of no arbitrage when there is uncertainty, but it should be clear that the argument we have just given shows why the call option must have the price

$$C_0 = \frac{1}{R} \left[\frac{R - d}{u - d} C_u + \frac{u - R}{u - d} C_d \right]$$

Rewriting this expression we get

$$C_0 = \left(\frac{R - d}{u - d} \right) \frac{C_u}{R} + \left(\frac{u - R}{u - d} \right) \frac{C_d}{R}$$

and if we let

$$q = \frac{R - d}{u - d}$$

we get

$$C_0 = q \frac{C_u}{R} + (1 - q) \frac{C_d}{R}.$$

If the price were lower, one could buy the call and sell the portfolio (a, b) , receive cash now as a consequence and have no future obligations except to exercise the call if necessary.

Remark 1. Notice that the hedge position in the stock (a) is positive, whilst the hedge position in the bank (b) is negative (i.e. we borrow money). This is quite apparent from the condition $d < R < u$: there are three different scenarios which arise: (I) $C_u > 0$ and $C_d > 0$, (II) $C_u > 0$ and $C_d = 0$, and the trivial case (III) $C_u = C_d = 0$. Clearly, nobody would bother hedging (III). However, (I) and (II) are readily shown to have $a > 0$ and $b < 0$.

We are now able to draw the following highly significant conclusions about the valuation of risky securities:

- The physical probability p plays **NO** role in the expression for C_0 . The fair price of the option is

$$\text{NOT: } C_0 = E[C_1(\omega)] = pC_u + (1 - p)C_d,$$

as one would perhaps initially conjecture.

- Rather, the fair price is given by the expression

$$C_0 = q \frac{C_u}{R} + (1 - q) \frac{C_d}{R},$$

where $q = \frac{R-d}{u-d}$ and $1 - q = \frac{u-R}{u-d}$.

- q and $1 - q$ formally satisfy the requirements of being *probability weights* (i.e. $\{q, 1 - q\} \in (0, 1)$ and $q + (1 - q) = 1$) as per the criterion $d < R < u$. To see this, consider weight q : since $R > d$ and $u > d$ it immediately follows that $q > 0$. On the other hand, suppose $q \geq 1$: then $R - d \geq u - d \Leftrightarrow R \geq u$ which contradicts $u > R$. So $q < 1$.
- Hence we may equivalently write the valuation formula as the Q -expectation of the discounted payoff

$$C_0 = E^Q \left[\frac{C_1(\omega)}{R} \right],$$


where Q is defined such that $Q(\omega_1) = q$, $Q(\omega_2) = 1 - q$, etc.

- The method of pricing the call really did not use the fact that C_u and C_d were call-values. Any security, V , with a time 1 value depending on ω_1 and ω_2 could have been priced according to the key valuation formula numeraire asset
risk neutral
measure
equivalent
measure

$$V_0 = E^Q \left[\frac{V_1(\omega)}{R} \right] = q \frac{V_u}{R} + (1 - q) \frac{V_d}{R}, \quad (4.1)$$

where $V_u = V_1(\omega_1)$ and $V_d = V_1(\omega_2)$ and q is as given above.

More prosaically we might restate the implication of (4.1) as follows: suppose we use the money market account B as our *numeraire asset* - i.e. suppose we measure the value of the security V in terms of how many units of B it corresponds to. Then the fair price of V at time zero [measured in units of the initial money market account (=1)] is related to the terminal price of V [measured in units of the terminal money market account (now = R)] *as though* the up state occurs with a probability of q and the down state occurs with a probability of $1 - q$. Since the numeraire asset B is manifestly deterministic and therefore void of any financial risk, it is customary to refer to Q as the *risk neutral measure*.



Now this exposition is bound to raise some questions: what (if anything) is so special about the money market account? Couldn't we have priced the security in terms of some other numeraire asset (say, the stock price process S)? The short answer is that there is nothing per se which singles out the money market account as the preferred numeraire. In fact, provided that we perform an *equivalent change of probability measure*, we can gracefully move to whichever numeraire asset tickles our fancy.³ To see how this works out, consider changing the measure in (4.1) from Q to some $Q' := \xi^{-1}Q$ where ξ is a non-negative random variable

³Two probability measures Q and Q' are said to be equivalent provided that they agree on which events have probability zero. We denote this property by $Q \sim Q'$.

$$\begin{aligned}
V_0 &= E^Q[R^{-1}V_1(\omega)] = R^{-1}(Q(\omega_1)V_1(\omega_1) + Q(\omega_2)V_1(\omega_2)) \\
&= R^{-1}\xi(Q'(\omega_1)V_1(\omega_1) + Q'(\omega_2)V_1(\omega_2)) \\
&= E^{Q'}[R^{-1}V_1(\omega)\xi].
\end{aligned}$$

In particular, suppose we specify ξ such that


$$\xi = \frac{RS}{S_1(\omega)},$$

then

$$V_0 = SE^{Q'}\left[\frac{V_1(\omega)}{S_1(\omega)}\right] = q'\frac{V_u}{u} + (1 - q')\frac{V_d}{d}, \quad (4.2)$$

where $q' = Q'(\omega_1) = \xi^{-1}Q(\omega_1) = (RS/(uS))^{-1}q = uq/R$. Make sure that you see that (4.2) is consistent with formula (4.1).

We may thus repeat the conclusion above in an analogous manner: suppose we use the stock price S as our *numeraire asset* - i.e. suppose we measure the value of the security V terms of how many units of S it corresponds to. Then the fair price of V at time zero [measured in units of initial stock ($=S$)] is related to the terminal price of V [measured in units of terminal stock ($= uS$ or $= dS$)] *as though* the up state occurs with a probability of q' and the down state occurs with a probability of $1 - q'$.



Whilst this numeraire-invariance (modulo a change a measure) of the valuation formula for risky securities is a neat theoretical result, one must inevitably wonder whether the result carries any practical implications. What could possibly warrant a preference for formula (4.2) over (4.1)? In discrete time “not too much” is generally the answer. However, upon moving to continuous time finance, an apt choice of numeraire can have a profound impact on our quest for a closed form option pricing formula: a seemingly impenetrable valuation exercise under one probability measure, may decompose into

a few lines of routine calculations under another. Indeed the risk neutral measure is not always to be preferred.

Example 13. Suppose we try to value a security which pays out whatever is the stock price according to the framework above. Obviously, for our model to be consistent with the absence of arbitrage, the time zero value should be equal to S . Let's check this. From (4.1):

$$\begin{aligned} V_0 = E^Q \left[\frac{S_1(\omega)}{R} \right] &= \left(\frac{R-d}{u-d} \right) \frac{1}{R} (uS) + \left(\frac{u-R}{u-d} \right) \frac{1}{R} (dS) \\ &= \frac{1}{(u-d)R} (RuS - duS + udS - RdS) \\ &= S. \end{aligned}$$

Equivalently, if we use equation (4.2) we find

$$V_0 = SE^{Q'} \left[\frac{S_1(\omega)}{S_1(\omega)} \right] = S(q' + (1 - q')) = S,$$

as desired.

Example 14. Finally, let us show that the call option price C_0 is increasing in the interest rate R . This is quite apparent upon remembering that

$$C_0 = aS + b,$$

where

$$a = \frac{C_u - C_d}{S(u-d)}, \quad \text{and} \quad b = \frac{1}{R} \frac{uC_d - dC_u}{(u-d)}.$$

It follows that as R increases R^{-1} decreases whence b decreases. But (as a bit of algebra shows; this relies explicitly on the call-option's payoff structure), $b < 0$ so C_0 increases. Simply put: a call option increases with the interest rate because borrowing becomes more expensive.

4.2 The single period model

The mathematics used when considering a one-period financial market with uncertainty is exactly the same as that used to describe the bond market in a multiperiod model with certainty: Just replace dates by states.

Given two time points $t = 0$ and $t = 1$ and a finite state space

$$\Omega = \{\omega_1, \dots, \omega_S\}.$$

probability
measure
financial market
security price
system
portfolio
arbitrage
opportunity
state-price vector
Arrow-Debreu
prices
riskfree asset

Whenever we have a probability measure P (or Q) we write p_i (or q_i) instead of $P(\{\omega_i\})$ (or $Q(\{\omega_i\})$).

A financial market or a security price system consists of a vector $\pi \in \mathbb{R}^N$ and an $N \times S$ matrix D where we interpret the i 'th row (d_{i1}, \dots, d_{iS}) of D as the payoff at time 1 of the i 'th security in states $1, \dots, S$, respectively. The price at time 0 of the i 'th security is π_i . A portfolio is a vector $\theta \in \mathbb{R}^N$ whose coordinates represent the number of securities bought at time 0. The price of the portfolio θ bought at time 0 is $\pi \cdot \theta$.

Definition 18. *An arbitrage in the security price system (π, D) is a portfolio θ which satisfies either*

$$\pi \cdot \theta \leq 0 \in \mathbb{R} \text{ and } D^\top \theta > 0 \in \mathbb{R}^S$$

or

$$\pi \cdot \theta < 0 \in \mathbb{R} \text{ and } D^\top \theta \geq 0 \in \mathbb{R}^S$$

A security price system (π, D) for which no arbitrage exists is called *arbitrage-free*.

Remark 2. *Our conventions when using inequalities on a vector in \mathbb{R}^k are the same as described in Chapter 3.*

When a market is arbitrage-free we want a vector of prices of 'elementary securities' - just as we had a vector of discount factors in Chapter 3.

Definition 19. $\psi \in \mathbb{R}_{++}^S$ (i.e. $\psi \gg 0$) is said to be a *state-price vector* for the system (π, D) if it satisfies

$$\pi = D\psi$$

Sometimes ψ is called a *state-price density*, or its elements referred to as *Arrow-Debreu-prices* and the term *Arrow-Debreu* attached to the elementary securities. Clearly, we have already proved the following in Chapter 3:

Proposition 3. *A security price system is arbitrage-free if and only if there exists a state-price vector.*

Unlike the model we considered in Chapter 3, the security which pays 1 in every state plays a special role here. If it exists, it allows us to speak of an 'interest rate':

Definition 20. *The system (π, D) contains a riskfree asset if there exists a linear combination of the rows of D which gives us $(1, \dots, 1) \in \mathbb{R}^S$.*

In an arbitrage-free system the price of the riskless asset d_0 is called the discount factor and $R_0 \equiv \frac{1}{d_0}$ is the return on the riskfree asset. Note that when a riskfree asset exists, and the price of obtaining it is d_0 , we have

$$d_0 = \theta_0^\top \pi = \theta_0^\top D\psi = \psi_1 + \cdots + \psi_S$$

where θ_0 is the portfolio that constructs the riskfree asset.

Now define

$$q_i = \frac{\psi_i}{d_0}, i = 1, \dots, S$$

Clearly, $q_i > 0$ and $\sum_{i=1}^S q_i = 1$, so we may interpret the q_i 's as probabilities. We may now rewrite the identity (assuming no arbitrage) $\pi = D\psi$ as follows:

$$\pi = d_0 Dq = \frac{1}{R_0} Dq, \text{ where } q = (q_1, \dots, q_S)^\top$$

If we read this coordinate by coordinate it says that

$$\pi_i = \frac{1}{R_0} (q_1 d_{i1} + \dots + q_S d_{iS})$$

which is the discounted expected value using q of the i^{th} security's payoff at time 1. Note that since R_0 is a constant we may as well say "expected discounted ...".

We assume throughout the rest of this section that a riskfree asset exists.

Definition 21. A security $c = (c_1, \dots, c_S)$ is redundant given the security price system (π, D) if there exists a portfolio θ_c such that $D^\top \theta_c = c$.

Proposition 4. Let an arbitrage-free system (π, D) and a redundant security c be given. The augmented system $(\hat{\pi}, \hat{D})$ obtained by adding π_c to the vector π and $c \in \mathbb{R}^S$ as a row of D is arbitrage-free if and only if

$$\pi_c = \frac{1}{R_0} (q_1 c_1 + \dots + q_S c_S) \equiv \psi_1 c_1 + \dots + \psi_S c_S.$$

Proof. Assume $\pi_c < \psi_1 c_1 + \dots + \psi_S c_S$. Buy the security c and sell the portfolio θ_c . The price of θ_c is by assumption higher than π_c , so we receive a positive cash-flow now. The cash-flow at time 1 is 0. Hence there is an arbitrage opportunity. If $\pi_c > \psi_1 c_1 + \dots + \psi_S c_S$ reverse the strategy. ■

Definition 22. The market is complete if for every $y \in \mathbb{R}^S$ there exists a $\theta \in \mathbb{R}^N$ such that

$$D^\top \theta = y \tag{4.3}$$

i.e. if the rows of D (the columns of D^\top) span \mathbb{R}^S .

contingent claim
risk-neutral
pricing

Proposition 5. *If the market is complete and arbitrage-free, there exists precisely one state-price vector ψ .*

The proof is exactly as in Chapter 3 and we are ready to price new securities in the financial market; also known as pricing of *contingent claim*.⁴

Here is how it is done in a one-period model: Construct a set of securities (the D -matrix,) and a set of prices. Make sure that (π, D) is arbitrage-free. Make sure that either

- (a) The model is complete, i.e. there are as many linearly independent securities as there are states.

Or

- (b) The contingent claim we wish to price is redundant given (π, D) .

Clearly, (a) implies (b) but not vice versa. (a) is almost always what is done in practice. Given a contingent claim $c = (c_1, \dots, c_S)$. Now compute the price of the contingent claim as

$$\pi(c) = \frac{1}{R_0} E^q(c) \equiv \frac{1}{R_0} \sum_{i=1}^S q_i c_i, \quad (4.4)$$

where $q_i = \frac{\psi_i}{d_0} \equiv R_0 \psi_i$. Again, the method in Equation (4.4) (and the generalizations of it we'll meet in Chapters 5 and 6) is referred to as *risk-neutral pricing*. Arbitrage-free prices are calculated as discounted expected values (with some new or artificial probabilities, the q s), i.e. as if agents were risk-neutral. But the “as if” is important to note:

No assumption of actual agent risk-neutrality is used to derive the risk neutral pricing formula (4.4) - just that they prefer more to less (see the next subsection). As a catch-phrase: “Risk-neutral pricing does not assume risk-neutrality”.

Measure Q might therefore be construed as a mathematical convenience tool, which allows us to do arbitrage free valuation. Of course, the existence of Q in turn depends on whether the financial market (we, the agents)

⁴A contingent claim just a random variable describing pay-offs; the pay-off is *contingent* on ω . The term (financial) derivative (asset, contract, or security) is largely synonymous, except that we are usually more specific about the pay-off being contingent on another financial asset such as a stock. We say option, even when we more specific pay-off structures in mind.

have valued existing assets consistently (i.e. without arbitrage). Indeed, its uniqueness (our ability to extract just one arbitrage free price) depends on whether the market is complete. But in the real world this assumption is obviously extremely hard to check, which has brought some skeptics to voice their dissatisfaction with the risk-neutral pricing enterprise.

Let us return to our example in the beginning of this chapter: The security price system is

$$(\pi, D) = \left(\left(\begin{array}{c} 1 \\ S \end{array} \right), \left(\begin{array}{cc} R & R \\ uS & dS \end{array} \right) \right).$$

For this to be arbitrage-free, proposition (3) tells us that there must be a solution (ψ_1, ψ_2) with $\psi_1 > 0$ and $\psi_2 > 0$ to the equation

$$\left(\begin{array}{c} 1 \\ S \end{array} \right) = \left(\begin{array}{cc} R & R \\ uS & dS \end{array} \right) \left(\begin{array}{c} \psi_1 \\ \psi_2 \end{array} \right).$$

$u \neq d$ ensures that the matrix D has full rank. $u > d$ can be assumed without loss of generality. We find

$$\psi_1 = \frac{R-d}{R(u-d)}, \quad \text{and} \quad \psi_2 = \frac{u-R}{R(u-d)},$$

and note that the solution is strictly positive precisely when $u > R > d$ (given our assumption that $u > d > 0$). The risk-free asset has a rate of return of $R-1$, and

$$q_1 = R\psi_1 = \frac{R-d}{u-d}, \quad \text{and} \quad q_2 = R\psi_2 = \frac{u-R}{u-d},$$

are the probabilities defining the measure q which can be used for pricing. Note that the market is complete, and this explains why we could use the procedure in the previous example to say what the correct price at time 0 of any claim (c_1, c_2) should be.

Example 15. Consider an arbitrage free market comprised three securities all valued at 2 [units of currency] with associated pay-offs $\begin{pmatrix} 2 \\ 3 \end{pmatrix}$, $\begin{pmatrix} 1 \\ 5 \end{pmatrix}$, and $\begin{pmatrix} 3 \\ 1 \end{pmatrix}$. Is the market complete? Is it arbitrage free? Suppose we introduce a fourth security with pay-off $\begin{pmatrix} 0 \\ 10 \end{pmatrix}$. What is its fair price?

The security price system is

$$(\pi, D) = \left(\left(\begin{array}{c} 2 \\ 2 \\ 2 \end{array} \right), \left(\begin{array}{cc} 2 & 3 \\ 1 & 5 \\ 3 & 1 \end{array} \right) \right).$$

Evidently the market is complete as we can form a basis in \mathbb{R}^2 from the existing securities (they do not all lie on the same line). Indeed, one security is redundant, which can be seen by noting that $2\begin{pmatrix} 2 \\ 3 \end{pmatrix} - \begin{pmatrix} 1 \\ 5 \end{pmatrix} = \begin{pmatrix} 3 \\ 1 \end{pmatrix}$. To check for the absence of arbitrage, is equivalent to checking whether there exists a strictly positive vector $\psi = (\psi_1, \psi_2)^\top$ such that $\pi = D\psi$. To this end we notice that D qua its dimensionality is non-invertible. However, upon multiplying $\pi = D\psi$ through by D^\top we have a system which is solvable. Indeed,

$$\psi = (D^\top D)^{-1} D^\top \pi \in \mathbb{R}_{++}^2,$$

so the system is arbitrage free (the reader should check that this is in fact the case). Finally, let's put a fair price on the (redundant) security $y = \begin{pmatrix} 0 \\ 10 \end{pmatrix}$. Now we might try to do this by solving equation (4.3) as $\theta = (DD^\top)^{-1} D y$, however, the matrix DD^\top is singular. Instead, we pick an invertible sub-matrix D and perform the valuation accordingly (specifically, we pick two (independent) assets and solve the problem). E.g. using assets 1 & 2 we find that

$$\begin{pmatrix} 2 & 3 \\ 1 & 5 \end{pmatrix} \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} 0 \\ 10 \end{pmatrix} \Leftrightarrow \begin{pmatrix} \theta_1 \\ \theta_2 \end{pmatrix} = \begin{pmatrix} -\frac{10}{7} \\ \frac{20}{7} \end{pmatrix}.$$

Hence the no-arbitrage price is $-2 \cdot \frac{10}{7} + 2 \cdot \frac{20}{7} = \frac{20}{7}$. The important point is that we arrive at this price irrespective of which replicating securities we choose. Thus, the reader might like to verify that the security pairs $\{\begin{pmatrix} 2 \\ 3 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \end{pmatrix}\}$ and $\{\begin{pmatrix} 1 \\ 5 \end{pmatrix}, \begin{pmatrix} 3 \\ 1 \end{pmatrix}\}$ both entail the same price for $\begin{pmatrix} 0 \\ 10 \end{pmatrix}$.

Example 16. Consider the following curious set-up: suppose there are three securities on the market with $t = 0$ prices 1, 1 and γ [\mathcal{L}], where γ is a positive constant. At time $t = 1$ their pay-offs are determined based on the local temperature (T) in London: if $T \geq 20^\circ$ the securities respectively pay out 1, 2 and 1 [\mathcal{L}]. If $20^\circ > T \geq 15^\circ$ they pay out 1, 1 and γ [\mathcal{L}]. Finally, if $T < 15^\circ$ the securities pay out 1, 0 and 1 [\mathcal{L}]. Is this market arbitrage free?

The security price system is

$$(\pi, D) = \left(\begin{pmatrix} 1 \\ 1 \\ \gamma \end{pmatrix}, \begin{pmatrix} 1 & 1 & 1 \\ 2 & 1 & 0 \\ 1 & \gamma & 1 \end{pmatrix} \right). \quad (4.5)$$

Note that the first security is risk free. To check whether the system is arbitrage free, we must establish whether there exists a strictly positive state price vector $\psi = (\psi_1, \psi_2, \psi_3)^\top$ such that $\pi = D\psi$. To this end, consider the inverse matrix

$$D^{-1} = \frac{1}{2(\gamma - 1)} \begin{pmatrix} 1 & \gamma - 1 & -1 \\ -2 & 0 & 2 \\ 2\gamma - 1 & 1 - \gamma & -1 \end{pmatrix}.$$

This is well-defined if and only if $\gamma \neq 1$. Assuming this to be the case, we readily find that $\psi = D^{-1}\pi = (0, 1, 0)^\top$ which clearly does **not** meet the requirement of strict positivity. $\gamma \neq 1$ thus corresponds to a complete market with arbitrage opportunities. What about the case $\gamma = 1$? From (4.5) we see that the market becomes incomplete: specifically, assets one and three are now identical (both in price and in pay-off). Any residual arbitrage should therefore be between the risk free asset and asset 2. To check if arbitrage obtains, let us look for a vector θ such that $\pi \cdot \theta < 0$ and $D^\top \theta \geq 0$. Without loss of generality set $\theta_3 = 0$, then we may recast this problem as the linear programme

$$\begin{array}{ll} \min & \theta_1 + \theta_2 \\ \text{s.t.} & \theta_1 + 2\theta_2 \geq 0, \\ & \theta_1 + \theta_2 \geq 0, \\ & \theta_1 \geq 0. \end{array}$$

Clearly, the second constraint is incompatible with a situation in which $\min \theta_1 + \theta_2 < 0$. Hence, $\gamma = 1$ corresponds to an incomplete market without arbitrage opportunities.

4.3 The economic intuition

At first, it may seem surprising that the “objective” probability p does not enter into the expression for the option price. Even if the the probability is 0.99 making the probability of the option paying out a positive amount very large, it does not alter the option’s price at time 0. Looking at this problem from a mathematical viewpoint, one can just say that this is a consequence of the linear algebra of the problem: The cost of forming a replicating strategy does not depend on the probability measure and therefore it does not enter into the contract. But this argument will not (and should not) convince a person who is worried by the economic interpretation of a model. Addressing the problem from a purely mathematical angle leaves some very important economic intuition behind. We will try in this section to get the economic intuition behind this ‘invariance’ to the choice of p straight. This will provide an opportunity to outline how the financial markets studied in this course fit in with a broader microeconomic analysis.

utility function

Before the more formal approach, here is the story in words: If we argue (erroneously) that changing p ought to change the option's price at time 0, the same argument should also lead to a suggested change in S_0 . But the experiment involving a change in p is an experiment which holds S_0 fixed. The given price of the stock is supposed to represent a 'sensible' model of the market. If we change p without changing S_0 we are implicitly changing our description of the underlying economy. An economy in which the probability of an up jump p is increased to 0.99 while the initial stock price remains fixed must be a description of an economy in which payoff in the upstate has lost value relative to a payoff in the downstate. These two opposite effects precisely offset each other when pricing the option.

The economic model we have in mind when studying the financial market is one in which utility is a function of wealth in each state and wealth is measured by a scalar (kroner, dollars, ...). Think of the financial market as a way of transferring money between different time periods and different states. A real economy would have a (spot) market for real goods also (food, houses, TV-sets, ...) and perhaps agents would have known endowments of real goods in each state at each time. If the spot prices of real goods which are realized in each state at each future point in time are known at time 0, then we may as well express the initial endowment in terms of wealth in each state. Similarly, the optimal consumption plan is associated with a precise transfer of wealth between states which allows one to realize the consumption plan. So even if utility is typically a function of the real goods (most people like money because of the things it allows them to buy), we can formulate the utility as a function of the wealth available in each state.

Consider⁵ an agent who has an endowment $e = (e_1, \dots, e_S) \in \mathbb{R}_+^S$. This vector represents the random wealth that the agent will have at time 1. The agent has a utility function $U : \mathbb{R}_+^S \rightarrow \mathbb{R}$ which we assume to be concave, differentiable and strictly increasing in each coordinate. Given a financial market represented by the pair (π, D) , the agent's problem is

$$\begin{aligned} \max_{\theta} \quad & U(e + D^\top \theta) \\ \text{s.t.} \quad & \pi^\top \theta \leq 0. \end{aligned} \tag{4.6}$$

If we assume that there exists a security with a non-negative payoff which is strictly positive in at least one state, then because the utility function is increasing we can replace the inequality in the constraint by an equality.

⁵This closely follows Darrell Duffie: *Dynamic Asset Pricing Theory*. Princeton University Press. 1996

And then the interpretation is simply that the agent sells endowment in some states to obtain more in other states. But no cash changes hands at time 0. Note that while utility is defined over all (non-negative) consumption vectors, it is the rank of D which decides in which directions the consumer can move away from the initial endowment.

Proposition 6. *If there exists a portfolio θ_0 with $D^\top \theta_0 > 0$ then the agent can find a solution to the maximization problem if and only if (π, D) is arbitrage-free.*

Proof. **The “only if” part:** We want to show that if (4.6) admits an optimal solution, then there is no arbitrage - or, equivalently, if there is an arbitrage, then there is no solution to (4.6). Suppose there is an arbitrage portfolio $\tilde{\theta}$ and that $c^* = e + D^\top \theta^*$ is an optimal solution to (4.6). Let the arbitrage be of the first kind, i.e. $\pi^\top \tilde{\theta} \leq 0$ and $D^\top \tilde{\theta} > 0$, then the agent would be better off investing in the arbitrage portfolio. Specifically, there must exist a non-negative α for which $D^\top(\alpha \tilde{\theta}) > D^\top \theta^*$ and since U is strictly increasing $U(e + D^\top(\alpha \tilde{\theta})) > U(e + D^\top \theta^*)$. This contradicts the assumption that θ^* is optimal. Now suppose the arbitrage is of the second kind, specifically the case where $\pi^\top \tilde{\theta} < 0$ and $D^\top \tilde{\theta} = 0$. Then the agent may invest the proceeds from the arbitrage portfolio into the portfolio θ_0 for which $D^\top \theta_0 > 0$ (the assumption that such a portfolio exists is a very mild condition). Once again, this will allow us to contradict the assumption that θ^* is the optimal portfolio.

The “if” part: We will now show that in the absence of arbitrage, there exists a solution to (4.6). To this end, it would be convenient to use the *extreme value theorem* which establishes that a continuous real-valued function on a nonempty compact space is bounded above and attains its supremum. That $X = \{e + D^\top \theta \in \mathbb{R}_+^S \mid \theta \in \mathbb{R}^N, \pi^\top \theta \leq 0\}$ constitutes a non-empty compact (convex) space is readily demonstrated: by assumption it is not empty, and closure follows from the “ \leq ”. Convexity is likewise trivial: if θ_1 and θ_2 are two arbitrary portfolios which both satisfy the conditions of X , then so does the portfolio $\lambda \theta_1 + (1 - \lambda) \theta_2 \forall \lambda \in (0, 1)$. Finally, we can argue for boundedness by contradiction: suppose the convex space X is unbounded, then each $c \in X$ has an associated *ray* i.e. an infinite straight line which can be traversed without leaving X .⁶ However, such a ray corresponds to the existence of an arbitrage (why?), which contradicts our assumptions. Hence, X must be non-empty and compact and the extreme value theorem entails that (4.6) has a solution. \square

⁶The precise statement is: $\forall c \in X \exists h \in \mathbb{R}^N$ such that $h \neq 0$ and $\ell = \{x \in \mathbb{R}^N \mid x = c + th, t \geq 0\} \subset X$.

state-price vector

The important insight is the following (see Proposition 1C in Duffie (1996)):

Proposition 7. *Assume that there exists a portfolio θ_0 with $D^\top \theta_0 > 0$. If there exists a solution θ^* to (4.6) and the associated optimal consumption is given by $c^* := e + D^\top \theta^* \gg 0$, then the gradient $\nabla U(c^*)$ (thought of as a column vector) is proportional to a state-price vector. The constant of proportionality is positive.*

Proof. Since c^* is strictly positive, then for any portfolio θ there exists some $k(\theta)$ such that $c^* + \alpha D^\top \theta \geq 0$ for all α in $[-k(\theta), k(\theta)]$. Define

$$g_\theta : [-k(\theta), k(\theta)] \rightarrow \mathbb{R}$$

as

$$g_\theta(\alpha) = U(c^* + \alpha D^\top \theta)$$

Now consider a θ with $\pi^\top \theta = 0$. Since c^* is optimal, g_θ must be maximized at $\alpha = 0$ and due to our differentiability assumptions we must have

$$g'_\theta(0) = (\nabla U(c^*))^\top D^\top \theta = 0.$$

We can conclude that any θ with $\pi^\top \theta = 0$ satisfies $(\nabla U(c^*))^\top D^\top \theta = 0$. Transposing the last expression, we may also write $\theta^\top D \nabla U(c^*) = 0$. In words, *any* vector which is orthogonal to π is also orthogonal to $D \nabla U(c^*)$. This means that $\mu \pi = D \nabla U(c^*)$ for some μ showing that $\nabla U(c^*)$ is proportional to a state-price vector. Choosing a θ_0 with $D^\top \theta_0 > 0$ we know from no arbitrage that $\pi^\top \theta_0 > 0$ and from the assumption that the utility function is strictly increasing, we have $\nabla U(c^*) D^\top \theta_0 > 0$. Hence μ must be positive. \square

Remark 3. *To understand the implications of this result we turn to the special case where the utility function has an expected utility representation, i.e. where we have a set of probabilities (p_1, \dots, p_S) and a function u such that*

$$U(c) = \sum_{i=1}^S p_i u(c_i).$$

In this special case we note that the coordinates of the state-price vector satisfy

$$\psi_i = \lambda p_i u'(c_i^*), \quad i = 1, \dots, S. \quad (4.7)$$

where λ is some constant of proportionality. Suppose the market is complete. Then proposition 7 effectively reduces the optimal portfolio problem (4.6) to a one-dimensional problem. Specifically, the lefthand side of (4.7) is determined from market data independently of the agent. So we divide λ and p_i

over and take $(u')^{-1}$ (u is concave and smooth, so u' is a continuous and decreasing function, hence it has an inverse) thus determining c_i^* . (At least up to knowledge of the scalar λ , which in practical cases would be determined from the agent's budget constraint; λ has to be such that the time zero price of the agent's consumption is no more than what he has to spend.) This is Pliska's martingale method of stochastic control.

So what if the market is incomplete? In this case we may construe proposition 7 in reverse order as a way to pin down a "reasonable" martingale measure. Specifically, some understanding of the righthand side, will lead to a concretisation of the (underdetermined) ψ .

Now we can state the economic intuition behind the option example as follows (and it is best to think of a complete market to avoid ambiguities in the interpretation): Given the complete market (π, D) we can find a unique state price vector ψ . This state price vector does not depend on p . Thus if we change p and we are thinking of some agent out there 'justifying' our assumptions on prices of traded securities, it must be the case that the agent has different marginal utilities associated with optimal consumption in each state. The difference must offset the change in p in such a way that (4.7) still holds. We can think of this change in marginal utility as happening in two ways: One way is to change utility functions altogether. Then starting with the same endowment the new utility functions would offset the change in probabilities so that the equality still holds. Another way to think of state prices as being fixed with new probabilities but utility functions unchanged, is to think of a different value of the initial endowment. If the endowment is made very large in one state and very small in the other, then this will offset the large change in probabilities of the two states. The analysis of the single agent can be carried over to an economy with many agents with suitable technical assumptions. Things become particularly easy when the equilibrium can be analyzed by considering the utility of a single, 'representative' agent, whose endowment is the sum of all the agents' endowments. An equilibrium then occurs only if this representative investor has the initial endowment as the solution to the utility maximization problem and hence does not need to trade in the market with the given prices. In this case the aggregate endowment plays a crucial role. Increasing the probability of a state while holding prices and the utility function of the representative investor constant must imply that the aggregate endowment is different with more endowment (low marginal utility) in the states with high probability and low endowment (high marginal utility) in the states with low probability. This intuition is very important when we discuss the Capital Asset Pricing Model later in the course.

A market where we are able to separate out the financial decisions as above is the one we will have in our mind throughout this course. But do keep in mind that this leaves out many interesting issues in the interaction between real markets and financial markets. For example, it is easy to imagine that an incomplete financial market (i.e. one which does not allow any distribution of wealth across states and time periods) makes it impossible for agents to realize consumption plans that they would find optimal in a complete market. This in turn may change equilibrium prices on real markets because it changes investment behavior. For example, returning to the house market, the fact that financial markets allows young agents to borrow against future income, makes it possible for more consumers to buy a house early in their lives. If all of a sudden we removed the possibility of borrowing we could imagine that house prices would drop significantly, since the demand would suddenly decrease.

Example 17. *Notice:* This exercise requires programming. The reader is strongly encouraged to have a personal go at the questions before consulting the answers.

Consider an investor who can choose between a (risky) stock and a (risk-free) bond investment, with the aim of maximising his expected terminal wealth. Specifically, he wants to solve the following optimisation problem

$$\begin{aligned} \max_{x_S, x_b} E(u(W(1))) \quad \text{s.t.} \quad & x_S S(0) + x_b \leq W(0), \\ & x_S S(1) + x_b(1+r) = W(1), \\ & x_S, x_b \geq 0, \quad (\text{i.e. no short selling}) \end{aligned} \tag{4.8}$$

where $W(1)$ denotes his (stochastic) time-1 wealth and his criterion or utility function has the form

$$u(x) = \frac{x^\gamma - 1}{\gamma},$$

for some γ , that is then one minus investor's (constant) relative risk-aversion. To get the ball rolling, let us look at a two-state model for $S(1)$

$$S(1) = \begin{cases} uS(0) & \text{with probability } p, \\ dS(0) & \text{with probability } 1 - p, \end{cases}$$

As default, assume $W(0) = 100$, $S(0) = 100$, $u = 1.25$, $d = 0.95$, $r = 0.095$, $p = 0.5$ and $\gamma = 0.5$.

(1) Formulate and solve the problem using a program, you find suitable (R, Excel, Maple, Matlab, GAMS, ...), i.e. find the optimal portfolio (x_S, x_b) . (2) How does this fraction allocated to the risky asset depend (or not depend) on $W(0)$? (3) What would happen with quadratic utility? (4) What happens to the optimal portfolio when you vary γ ?

1. Since the utility function is non-linear, so is the optimisation problem. Entering the following command in Maple

$$NLP\text{Solve}\left(0.5 \cdot \frac{(x_S \cdot 1.25 \cdot 100 + x_b \cdot (1 + 0.095))^{0.5} - 1}{0.5} + \right. \\ \left. 0.5 \cdot \frac{(x_S \cdot 0.95 \cdot 100 + x_b \cdot (1 + 0.095))^{0.5} - 1}{0.5}, \right. \\ \left. \{100 \cdot x_S + x_b \leq 100\}, \text{assume} = \text{nonnegative, maximize} \right)$$

one finds that the optimal investment for the stock is $x_S = .487$, whilst the optimal bond investment is $x_b = 51.3$ (i.e. we place the fraction of our initial wealth $x_b/W(0) = 0.513$ in the bond).

2. Although the numerical results for x_S and x_b change for different choices of $W(0)$, the fraction of the investors initial wealth that go into the stock and the bond remain unchanged. I.e. the weights

$$(w_S, w_b) := (x_S S(0)/W(0), x_b/W(0)),$$

are invariant to the investors initial wealth level.

3. For $\gamma = 2$ the investor's utility function is convex, wherefore he is risk loving rather than risk averse ($\gamma < 1$). This means the the investor allocates all of his initial resources to the stock ($x_S = 1, x_b = 0$).
4. The higher the risk aversion (effectively codified by $-\gamma$) the lower is the investors appetite for the stock. E.g. for $\gamma = 0.3$ we find that $(x_S, x_b) = (0.348, 65.2)$. Conversely, for $\gamma = 0.7$ one finds that $(x_S, x_b) = (0.811, 18.9)$.

Suppose now that $\gamma = 1$. (5) Why is that case somewhat special? (6) What happens when you increase r ? Or more accurately: When does the solution change? And on a possibly related note: What is the expected rate of return on the stock $E(R) := E((S(1) - S(0))/S(0))$? (7) What happens if the short sales constraints are removed?

5. When $\gamma = 1$ the investor's utility function becomes linear (we say he is risk neutral), whence programme (4.8) reduces to a linear optimisation problem. It is at this value for γ that the investor shifts his entire wealth to the stock market.
6. Upon increasing r to some value > 0.1 we find that the investor reallocates his entire wealth to the bond. The reason for this is simply that the *guaranteed* return from the bond now overshadows the *expected* return on the stock: $E(R) = p \cdot u + (1 - p) \cdot d - 1 = 0.1$. As the investor is not risk loving he has thus been completely disincentivised from investing in the risky asset.
7. Removing the short selling restrictions will result in unbounded short selling of the asset with the lowest (expected) return.

Go back to the default settings, but let's play around with the risk-free rate. **(8)** Is it correct that risk-averse investors always "diversify" i.e. invest in both the stock and the bond?

8. No. Consider e.g. the extreme example of a negative interest rate. Clearly, even though the investor is risk averse, he would prefer to have an expected positive return over a guaranteed negative return. The catchy way to put this is: "he's risk averse - not stupid".

Still default settings. **(9)** Now what if the short sales constraints are removed? **(10)** How does the solution look (or: the solver behave) for $r = 0.05$, $r = 0$, $r = -0.1$, $r = 0.2$, and $r = 0.3$?

9. Removing the restriction on short selling, will not perturb the values of the original optimal investment $(x_S, x_b) = (0.487, 51.3)$.
10. For $r = -0.1$ and $r = 0.3$ there is arbitrage on the market (the reader should verify that the associated state price vectors are not strictly positive). This incentivises the investor to perform unlimited short selling of the under-performing asset (respectively the bond and the stock) for an unlimited profit (i.e. no maximum in programme (4.8) can be found). For the remaining no-arbitrage values of r the situation is exactly as you would expect: the higher the risk free return, the keener is the investor on allocating his funds to the bond. Maple returns the following values:

r	0	0.05	0.2
x_S	16	5.25	-19.2
x_b	-1500	-425	2020

This can be seen a numerical illustration of Proposition 6.⁷ Only when the market is arbitrage free can the investor find an optimum to his utility maximisation problem.

⁷Purists can rightly claim than the utility function here does not satisfy the assumptions for said Proposition. Specifically, for $\gamma \leq 0$ the utility function is non-defined (non-finite) in $x = 0$. However, for our choice of $\gamma = 0.5$ this is obviously not a problem.

Chapter 5

Arbitrage pricing in the multi-period model

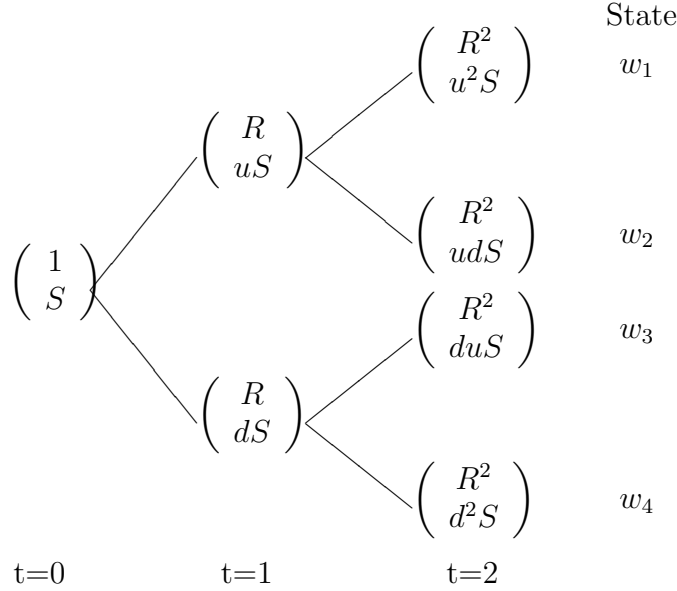
5.1 An appetizer

It is fair to argue that to get realism in a model with finite state space we need the number of states to be large. After all, why would the stock take on only two possible values at the expiration date of the option? On the other hand, we know from the previous section that in a model with many states we need many securities to have completeness, which (in arbitrage-free models) is a requirement for pricing every claim. And if we want to price an option using only the underlying stock and a money market account, we only have two securities to work with. Fortunately, there is a clever way out of this.

Assume that over a short time interval the stock can only move to two different values and split up the time interval between 0 and T (the expiry date of an option) into small intervals in which the stock can be traded. Then it turns out that we can have both completeness and therefore arbitrage pricing even if the number of securities is much smaller than the number of states. Again, before we go into the mathematics, we give an example to help with the intuition.

Assume that $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$ and that there are three dates: $t \in \{0, 1, 2\}$. We specify the behavior of the stock and the money market account as follows: Assume that $0 < d < R < u$ and that $S > 0$. Consider the following graph:

call option,
European



At time 0 the stock price is S , the bank account is worth 1. At time 1, if the state of the world is ω_1 or ω_2 , the prices are uS and R , respectively, whereas if the true state is ω_3 or ω_4 , the prices are dS and R . And finally, at time $t = 2$, the prices of the two instruments are as shown in the figure above. Note that $\omega \in \Omega$ describes a whole "sample path" of the stock price process and the bank account, i.e. it tells us not only the final time 2 value, but the entire history of values up to time 2.

Now suppose that we are interested in the price of a European call option on the stock with exercise price K and expiry $T = 2$. At time 2, we know it is worth

$$C_2(\omega) = [S_2(\omega) - K]^+$$

where $S_2(\omega)$ is the value of the stock at time 2 if the true state is ω .

At time 1, if we are in state ω_1 or ω_2 , the money market account is worth R and the stock is worth uS , and we know that there are only two possible time 2 values, namely (R^2, u^2S) or (R^2, duS) . But then we can use the argument of the one period example to see that at time 1 in state ω_1 or ω_2 we can replicate the calls payoff by choosing a suitable portfolio of stock and money market account: Simply solve the system:

$$au^2S + bR^2 = [u^2S - K]^+ \equiv C_{uu}$$

$$aduS + bR^2 = [duS - K]^+ \equiv C_{du}$$

for (a, b) and compute the price of forming the portfolio at time 1. We find

$$a = \frac{C_{uu} - C_{du}}{uS(u-d)}, \quad b = \frac{uC_{du} - dC_{uu}}{(u-d)R^2}.$$

The price of this portfolio is

$$\begin{aligned} auS + bR &= \frac{R}{R} \frac{(C_{uu} - C_{du})}{(u-d)} + \frac{uC_{du} - dC_{uu}}{(u-d)R} \\ &= \frac{1}{R} \left[\frac{(R-d)}{(u-d)} C_{uu} + \frac{(u-R)}{(u-d)} C_{ud} \right] =: C_u \end{aligned}$$

This is clearly what the call is worth at time $t = 1$ if we are in ω_1 or ω_2 , i.e. if the stock is worth uS at time 1. Similarly, we may define $C_{ud} := [udS - K]^+$ (which is equal to C_{du}) and $C_{dd} = [d^2S - K]^+$. And now we use the exact same argument to see that if we are in state ω_3 or ω_4 , i.e. if the stock is worth dS at time 1, then at time 1 the call should be worth C_d where

$$C_d := \frac{1}{R} \left[\frac{(R-d)}{(u-d)} C_{ud} + \frac{(u-R)}{(u-d)} C_{dd} \right].$$

Now we know what the call is worth at time 1 depending on which state we are in: If we are in a state where the stock is worth uS , the call is worth C_u and if the stock is worth dS , the call is worth C_d .

Looking at time 0 now, we know that all we need at time 1 to be able to "create the call", is to have C_u when the stock goes up to uS and C_d when it goes down. But that we can accomplish again by using the one-period example: The cost of getting $\begin{pmatrix} C_u \\ C_d \end{pmatrix}$ is

$$C_0 := \frac{1}{R} \left[\frac{(R-d)}{(u-d)} C_u + \frac{(u-R)}{(u-d)} C_d \right].$$

If we let $q = \frac{R-d}{u-d}$ and if we insert the expressions for C_u and C_d , noting that $C_{ud} = C_{du}$, we find that

$$C_0 = \frac{1}{R^2} [q^2 C_{uu} + 2q(1-q) C_{ud} + (1-q)^2 C_{dd}]$$

which the reader will recognize as a discounted expected value, just as in the one period example. (Note that the representation as an expected value does not hinge on $C_{ud} = C_{du}$.)

The important thing to understand in this example is the following: Starting out with the amount C_0 , an investor is able to form a portfolio in the stock and the money market account which produces the payoffs C_u or C_d

replication
 hedging
 derivative security
 information
 filtration
 partition

at time 1 depending on where the stock goes. Now without any additional costs, the investor can rearrange his/her portfolio at time 1, such that at time 2, the payoff will match that of the option. Therefore, at time 0 the price of the option must be C_0 .

This dynamic replication or hedging argument (If we create the pay-off of the call, we typically say that we replicate, when we create *minus* the call pay-off, we say hedge.) is the key to pricing derivative securities (another word for contingent claims) in discrete-time, finite state space models. We now want to understand the mathematics behind this example.

5.2 Price processes, trading and arbitrage

Given a probability space (Ω, \mathcal{F}, P) with Ω finite, let $\mathcal{F} := 2^\Omega$ (i.e. the set of all subsets of Ω) and assume that $P(\omega) > 0$ for all $\omega \in \Omega$. Also assume that there are $T + 1$ dates, starting at date 0, ending at date T . To formalize how information is revealed through time, we introduce the notion of a filtration:

Definition 23. A filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t=0}^T$ is an increasing sequence of σ -algebras contained in \mathcal{F} : $\mathcal{F}_0 \subseteq \mathcal{F}_1 \subseteq \dots \subseteq \mathcal{F}_T$.


We will always assume that $\mathcal{F}_0 = \{\emptyset, \Omega\}$ and $\mathcal{F}_T = \mathcal{F}$. Since Ω is finite, it will be easy to think of the σ -algebras in terms of partitions:

Definition 24. A partition \mathcal{P}_t of Ω is a collection of non-empty subsets of Ω such that

- $\bigcup_{P_i \in \mathcal{P}_t} P_i = \Omega$
- $P_i \cap P_j = \emptyset$ whenever $i \neq j, P_i, P_j \in \mathcal{P}_t$.

Because Ω is finite, there is a one-to-one correspondence between partitions and σ -algebras: The elements of \mathcal{P}_t corresponds to the atoms of \mathcal{F}_t .

The concepts we have just defined are well illustrated in an event-tree:



The event tree illustrates the way in which we imagine information about the true state being revealed over time. At time $t = 1$, for example, we may find ourselves in one of two nodes: ξ_{11} or ξ_{12} . If we are in the node ξ_{11} , we know that the true state is in the set $\{\omega_1, \omega_2, \dots, \omega_5\}$, but we have no more knowledge than that. In ξ_{12} , we know (only) that $\omega \in \{\omega_6, \omega_7, \dots, \omega_9\}$. At time $t = 2$ we have more detailed knowledge, as represented by the partition \mathcal{P}_2 . Elements of the partition \mathcal{P}_t are events which we can decide as having occurred or not occurred at time t , regardless of what the true ω is. At time 1, we will always know whether $\{\omega_1, \omega_2, \dots, \omega_5\}$ has occurred or not, regardless of the true ω . If we are at node ξ_{12} , we would be able to rule out

measurable
adapted process
dividend process

the event $\{\omega_1, \omega_2\}$ also at time 1, but if we are at node ξ_{11} , we will not be able to decide whether this event has occurred or not. Hence $\{\omega_1, \omega_2\}$ is not a member of the partition.

Make sure you understand the following:

Remark 4. A random variable defined on (Ω, \mathcal{F}, P) is measurable with respect to \mathcal{F}_t precisely when it is constant on each member of \mathcal{P}_t .

A stochastic process $X := (X_t)_{t=0, \dots, T}$ is a sequence of random variables X_0, X_1, \dots, X_T . The process is adapted to the filtration \mathbb{F} , if X_t is \mathcal{F}_t -measurable (which we will often write: $X_t \in \mathcal{F}_t$) for $t = 0, \dots, T$. Returning to the event tree setup, it must be the case, for example, that $X_1(\omega_1) = X_1(\omega_5)$ if X is adapted, but we may have $X_1(\omega_1) \neq X_1(\omega_6)$.

Given an event tree, it is easy to construct adapted processes: Just assign the values of the process using the nodes of the tree. For example, at time 1, there are two nodes ξ_{11} and ξ_{12} . You can choose one value for X_1 in ξ_{11} and another in ξ_{12} . The value chosen in ξ_{11} will correspond to the value of X_1 on the set $\{\omega_1, \omega_2, \dots, \omega_5\}$, the value chosen in ξ_{12} will correspond to the common value of X_1 on the set $\{\omega_6, \dots, \omega_9\}$. When X_t is constant on an event A_t we will sometimes write $X_t(A_t)$ for this value. At time 2 there are five different values possible for X_2 . The value chosen in the top node is the value of X_2 on the set $\{\omega_1, \omega_2\}$.

As we have just seen it is convenient to speak in terms of the event tree associated with the filtration. From now on we will refer to the event tree as the graph Ξ and use ξ to refer to the individual nodes. The notation $p(\xi)$ will denote the probability of the event associated with ξ ; for example $P(\xi_{11}) = P(\{\omega_1, \omega_2, \dots, \omega_5\})$. This graph Ξ will also allow us to identify adapted processes with vectors in \mathbb{R}^Ξ . The following inner products on the space of adapted processes will become useful later: Let X, Y be adapted processes and define

$$\begin{aligned} \sum_{\xi \in \Xi} X(\xi)Y(\xi) &\equiv \sum_{\{(t, A_u): A_u \in \mathcal{P}_t, 0 \leq t \leq T\}} X_t(A_u)Y_t(A_u) \\ E \sum_{\xi \in \Xi} X(\xi)Y(\xi) &\equiv \sum_{\xi \in \Xi} P(\xi)X(\xi)Y(\xi) \\ &\equiv \sum_{\{(t, A_u): A_u \in \mathcal{P}_t, 0 \leq t \leq T\}} P(A_u)X_t(A_u)Y_t(A_u) \end{aligned}$$

Now we are ready to model financial markets in multi-period models. Given is a vector of adapted dividend processes

$$\delta = (\delta^1, \dots, \delta^N)$$

and a vector of adapted security price processes

$$S = (S^1, \dots, S^N).$$

price process
 locally riskfree
 asset
 short rate process
 locally riskfree
 asset

The interpretation is as follows: $S_t^i(\omega)$ is the price of security i at time t if the state is ω . Buying the i 'th security at time t ensures the buyer (and obligates the seller to deliver) the remaining dividends $\delta_{t+1}^i, \delta_{t+2}^i, \dots, \delta_T^i$.¹ Hence the security price process is to be interpreted as an ex-dividend price process and in particular we should think of S_T as 0. In all models considered in these notes we will also assume that there is a bank account which provides locally riskfree borrowing and lending. This is modeled as follows: Given an adapted process - the short rate process

$$\rho = (\rho_0, \rho_1, \dots, \rho_{T-1}).$$

To make the math work, all we need to assume about this process is that it is strictly greater than -1 at all times and in all states, but for modelling purposes it is desirable to have it non-negative. Now we may define the money market account as follows:

Definition 25. *The bank account has the security price process*

$$\begin{aligned} S_t^0 &= 1, & t = 0, 1, \dots, T - 1 \\ S_T^0 &= 0. \end{aligned}$$

and the dividend process

$$\begin{aligned} \delta_t^0(\omega) &= \rho_{t-1}(\omega) \text{ for all } \omega \text{ and } t = 1, \dots, T - 1, \\ \delta_T^0(\omega) &= 1 + \rho_{T-1}(\omega). \end{aligned}$$

This means that if you buy one unit of the money market account at time t you will receive a dividend of ρ_t at time $t + 1$. Since ρ_t is known already at time t , the dividend received on the money market account in the next period $t + 1$ is known at time t . Since the price is also known to be 1 you know that placing 1 in the money market account at time t , and selling the asset at time $t + 1$ will give you $1 + \rho_t$. This is why we refer to this asset as a locally riskfree asset. You may of course also choose to keep the money in the bank account and receive the stream of dividends. Reinvesting the dividends in the money market account will make this account grow according to the process R defined as

$$R_t = (1 + \rho_0) \cdots (1 + \rho_{t-1}).$$

¹We will follow the tradition of probability theory and often suppress the ω in the notation.

trading strategy,
 \mathcal{L}^{ϕ}
 dividend stream;
 $\Delta \hat{\phi}$
 arbitrage
 opportunity
 self-financing
 trading strategy

We will need this process to discount cash flows between arbitrary periods and therefore introduce the following notation:

$$R_{s,t} \equiv (1 + \rho_s) \cdots (1 + \rho_{t-1}).$$

Definition 26. A trading strategy is an adapted process

$$\phi = (\phi_t^0, \dots, \phi_t^N)_{t=0, \dots, T-1}.$$

and the interpretation is that $\phi_t^i(\omega)$ is the number of the i 'th security held at time t if the state is ω . The requirement that the trading strategy is adapted is very important. It represents the idea that the strategy should not be able to see into the future. Returning again to the event tree, when standing in node ξ_{11} , a trading strategy can base the number of securities on the fact that we are in ξ_{11} (and not in ξ_{12}), but not on whether the true state is ω_1 or ω_2 .

The dividend stream generated by the trading strategy ϕ is denoted δ^ϕ and it is defined as

$$\delta_0^\phi = -\phi_0 \cdot S_0, \delta_t^\phi = \phi_{t-1} \cdot (S_t + \delta_t) - \phi_t \cdot S_t \text{ for } t = 1, \dots, T.$$

Definition 27. An arbitrage is a trading strategy for which δ_t^ϕ is a positive process, i.e. always nonnegative and $\delta_t^\phi(\omega) > 0$ for some t and ω . The model is said to be arbitrage-free if it contains no arbitrage opportunities.

In words, there is arbitrage if we can adopt a trading strategy which at no point in time requires us to pay anything but which at some time in some state gives us a strictly positive payout. Note that since we have included the initial payout as part of the dividend stream generated by a trading strategy, we can capture the definition of arbitrage in this one statement. This one statement captures arbitrage both in the sense of receiving money now with no future obligations and in the sense of paying nothing now but receiving something later.

Definition 28. A trading strategy ϕ is self-financing if it satisfies

$$\phi_{t-1} \cdot (S_t + \delta_t) = \phi_t \cdot S_t \quad \text{for } t = 1, \dots, T.$$

The interpretation is as follows: Think of forming a portfolio ϕ_{t-1} at time $t - 1$. Now as we reach time t , the value of this portfolio is equal to $\phi_{t-1} \cdot (S_t + \delta_t)$, and for a self-financing trading strategy, this is precisely the amount of money which can be used in forming a new portfolio at time t . We will let Φ denote the set of self-financing trading strategies.

Example 18. A classical exercise in *filtrations and measurability* runs along the following lines: Suppose we have a fair coin which we flip twice. The outcome ω of this experiment is the sequence $\omega = \omega_1\omega_2 \in \Omega$ where each individual coin flip ω_i ($i = 1, 2$) can come out either heads (H) or tails (T). We also imagine that we monitor the evolution of the experiment: first at time $t = 1$ after the coin is flipped for the first time, and subsequently at time $t = 2$ when the coin has been flipped again. Question: what is the associated probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t=0,1,2}$ for this experiment?

Clearly, the experiment has a totality of $2^2 = 4$ possible outcomes, which we represent by the sample space $\Omega = \{HH, HT, TH, TT\}$. The associated event space is the power set 2^Ω :

$$\begin{aligned} \mathcal{F} = \{ & \emptyset, HH, HT, TH, TT, HH \cup HT, HH \cup TH, HH \cup TT, \\ & HT \cup TH, HT \cup TT, TH \cup TT, HH \cup HT \cup TH, \\ & HH \cup HT \cup TT, HH \cup TH \cup TT, HT \cup TH \cup TT, \Omega \}. \end{aligned} \quad (5.1)$$

You should check for yourself that \mathcal{F} satisfies the σ -algebra properties of closure under complementation and countable unions. Notice that the cardinality of the filtration is $|\mathcal{F}| = 16$ or, equivalently, $|\mathcal{F}| = 2^{|\Omega|} = 2^4$ (this explains the notation for the power set). Finally, the real world probability measure \mathbb{P} specifies the probability of every event in \mathcal{F} : e.g. since the coin is fair, we have for each of the elementary events $\omega \in \{HH, HT, TH, TT\}$ that $\mathbb{P}(\omega) = \frac{1}{2} \times \frac{1}{2} = \frac{1}{4}$. Other probabilities follow from the axioms of probability: e.g. $\mathbb{P}(\Omega) = 1$, $\mathbb{P}(\emptyset) = 0$, $\mathbb{P}(HH \cup HT) = \mathbb{P}(HH) + \mathbb{P}(HT) = \frac{1}{4} + \frac{1}{4} = \frac{1}{2}$ and so forth.

As for the filtration $\mathbb{F} = \{\mathcal{F}_t\}_{t=0,1,2}$, the three σ -algebras \mathcal{F}_0 , \mathcal{F}_1 and \mathcal{F}_2 effectively encode the information available to us at times $t = 0$, $t = 1$ and $t = 2$. Clearly, at $t = 0$, before any observation is made, we can only deduce the trivial events *something happened* or *nothing happened*, whence $\mathcal{F}_0 = \{\emptyset, \Omega\}$. At $t = 1$ the outcome of the first coin flip has been revealed (either $\omega_1 = H$ occurred, or $\omega_1 = T$ occurred) while ω_2 remains undisclosed. Hence, $\mathcal{F}_1 = \{\emptyset, HH \cup HT, TT \cup TH, \Omega\}$ where $HH \cup HT$ and $TT \cup TH$ are the *atoms* of \mathcal{F}_1 . Finally, at $t = 2$ the outcome of the second coin flip has been revealed, thus resolving any ambiguity about the experiment. The σ -algebra of identifiable events is therefore $\mathcal{F}_2 = \mathcal{F}$, where \mathcal{F} is given by (5.1).

Now suppose we (costlessly) enter a game which pays out \$1 every time the coin comes out heads, but deducts \$1 every time the coin comes out

pricing
functional

tails. Our cumulative gains are thus represented by the stochastic process $G_t : \Omega \times \{0, 1, 2\} \mapsto \mathbb{R}$ where $G_0 = 0$ and

$$G_1(\omega) = \begin{cases} +1, & \text{if } \omega \in \{HH, HT\} \\ -1, & \text{if } \omega \in \{TH, TT\} \end{cases} \quad G_2(\omega) = \begin{cases} +2, & \text{if } \omega = HH \\ 0, & \text{if } \omega \in \{HT, TH\} \\ -2, & \text{if } \omega = TT \end{cases}$$

What is the smallest σ -algebras \mathcal{F}_{G_1} and \mathcal{F}_{G_2} generated by the random variables G_1 and G_2 ? I.e. what are the sets of possible outcomes that can be deduced solely by monitoring our cumulative gains and not the actual coin flips? With respect to which of the six σ -algebras $\mathcal{F}, \mathcal{F}_0, \mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_{G_1}, \mathcal{F}_{G_2}$ are the random variables G_1 and G_2 measurable?

It is quite clear that $\mathcal{F}_{G_1} = \mathcal{F}_1$: there is no information difference between calling out $+1/-1$ or calling out H/T after the first coin flip. On the other hand, $\mathcal{F}_{G_2} \neq \mathcal{F}_2$: clearly, G_2 encodes *less* information since the outcome 0 tells us nothing about whether $\omega = HT$ or $\omega = TH$ occurred. The correct σ -algebra is readily shown to be

$$\mathcal{F}_{G_2} = \{\emptyset, HH, TT, HT \cup TH, HH \cup TT, HH \cup HT \cup TH, TT \cup HT \cup TH, \Omega\}.$$

To determine the measurability of a random variable X with respect to a given σ -algebra \mathcal{F} , we must check that the σ -algebra generated by X , \mathcal{F}_X , is a subset of \mathcal{F} . Since $\mathcal{F}_{G_1} \subseteq \mathcal{F}, \mathcal{F}_1, \mathcal{F}_2, \mathcal{F}_{G_1}$ the random variable G_1 is measurable with respect to those filtrations. Analogously, you should check that G_2 is measurable with respect to $\mathcal{F}, \mathcal{F}_2$ and \mathcal{F}_{G_2} .

5.3 No arbitrage and price functionals

We have seen in the one period model that there is equivalence between the existence of a state price vector and absence of arbitrage. In this section we show the multi-period analogue of this theorem.

The goal of this section is to prove the existence of the multi-period analogue of state-price vectors in the one-period model. Let \mathbb{L} denote the set of adapted processes on the given filtration.

Definition 29. A pricing functional F is a linear functional

$$F : \mathbb{L} \rightarrow \mathbb{R}$$

which is strictly positive, i.e.

$$\begin{aligned} F(X) &\geq 0 \text{ for } X \geq 0 \\ F(X) &> 0 \text{ for } X > 0. \end{aligned}$$

Riesz' representation theorem
separating hyperplane

Definition 30. A pricing functional F is consistent with security prices if

$$F(\delta^\phi) = 0 \text{ for all trading strategies } \phi.$$

Note that if there exists a consistent pricing functional we may arbitrarily assume that the value of the process $1_{\{t=0\}}$ (i.e. the process which is 1 at time 0 and 0 thereafter) is 1. By Riesz' representation theorem we can represent the functional F as

$$F(X) = \sum_{\xi \in \Xi} X(\xi) f(\xi)$$

With the convention $F(1_{\{t=0\}}) = 1$, we then note that if there exists a trading strategy ϕ which is initiated at time 0 and which only pays a dividend of 1 in the node ξ , then

$$\phi_0 \cdot S_0 = f(\xi).$$

Hence $f(\xi)$ is the price at time 0 of having a payout of 1 in the node ξ .

Proposition 8. The model (δ, S) is arbitrage-free if and only if there exists a consistent pricing functional.

Proof. First, assume that there exists a consistent pricing functional F . Any dividend stream δ^ϕ generated by a trading strategy which is positive must have $F(\delta^\phi) > 0$ but this contradicts consistency. Hence there is no arbitrage. The other direction requires more work:

Define the sets

$$\begin{aligned} \mathbb{L}^1 &= \left\{ X \in \mathbb{L} \mid X > 0 \text{ and } \sum_{\xi \in \Xi} X(\xi) = 1 \right\} \\ \mathbb{L}^0 &= \{ \delta^\phi \in \mathbb{L} \mid \phi \text{ trading strategy} \} \end{aligned}$$

and think of both sets as subsets of \mathbb{R}^Ξ . Note that \mathbb{L}^1 is convex and compact and that \mathbb{L}^0 is a linear subspace, hence closed and convex. By the no arbitrage assumption the two sets are disjoint. Therefore, there exists a separating hyperplane $H(f; \alpha) := \{x \in \mathbb{R}^\Xi : f \cdot x = \alpha\}$ which separates the two sets strictly and we may choose the direction of f such that $f \cdot x \leq \alpha$ for $x \in \mathbb{L}^0$. Since \mathbb{L}^0 is a linear subspace we must have $f \cdot x = 0$ for $x \in \mathbb{L}^0$ (why?).

complete market

Strict separation then gives us that $f \cdot x > 0$ for $x \in \mathbb{L}^1$, and that in turn implies $f \gg 0$ (why?). Hence the functional

$$F(X) = \sum_{\xi \in \Xi} f(\xi) X(\xi)$$

is consistent. ■

By using the same geometric intuition as in Chapter 2, we note that there is a connection between completeness of the market and uniqueness of the consistent price functional:

Definition 31. *The security model is complete if for every $X \in \mathbb{L}$ there exists a trading strategy ϕ such that $\delta_t^\phi = X_t$ for $t \geq 1$.*

If the model is complete and arbitrage-free, there can only be one consistent price functional (up to multiplication by a scalar). To see this, assume that if we have two consistent price functionals F, G both normed to have $F(1_{\{t=0\}}) = G(1_{\{t=0\}}) = 1$. Then for any trading strategy ϕ we have

$$\begin{aligned} 0 &= -\phi_0 \cdot S_0 + F(1_{\{t>0\}} \delta^\phi) \\ &= -\phi_0 \cdot S_0 + G(1_{\{t>0\}} \delta^\phi) \end{aligned}$$

hence F and G agree on all processes of the form $1_{\{t>0\}} \delta^\phi$. But they also agree on $1_{\{t=0\}}$ and therefore they are the same since by the assumption of completeness every adapted process can be obtained as a linear combination of these processes.

Given a security price system (π, D) , the converse is shown in a way very similar to the one-period case. Assume the market is arbitrage-free and incomplete. Then there exists a process π in \mathbb{L} , whose restriction to time $t \geq 1$ is orthogonal to any dividend process generated by a trading strategy. By letting $\pi_0 = 0$ and choosing a sufficiently small $\varepsilon > 0$, the functional defined by

$$(F + \varepsilon\pi)(\delta^\phi) = \sum_{\xi \in \Xi} (f(\xi) + \varepsilon\pi(\xi)) \delta^\phi(\xi)$$

is consistent. Hence we have shown:

Proposition 9. *If the market is arbitrage-free, then the model is complete if and only if the consistent price functional is unique.*

5.4 Conditional expectations and martingales

Consistent price systems turn out to be less interesting for computation when we look at more general models, and they do not really explain the strange

probability measure q which we saw earlier. We are about to remedy both problems, but first we need to make sure that we can handle conditional expectations in our models and that we have a few useful computational rules at our disposal.

*conditional
expectation*
Iterated
expectations

Definition 32. *The conditional expectation of an \mathcal{F}_u -measurable random variable X_u given \mathcal{F}_t , where $\mathcal{F}_t \subseteq \mathcal{F}_u$, is given by*

$$E(X_u | \mathcal{F}_t)(\omega) = \frac{1}{P(A_t)} \sum_{A_v \in \mathcal{P}_t: A_v \subseteq A_t} P(A_v) X_u(A_v) \text{ for } \omega \in A_t$$

where we have written $X_u(A_v)$ for the value of $X_u(\omega)$ on the set A_v and where $A_t \in \mathcal{P}_t$.

We will illustrate this definition in the exercises. Note that we obtain an \mathcal{F}_t -measurable random variable since it is constant over elements of the partition \mathcal{P}_t . The definition above does not work when the probability space becomes uncountable. Then one has to adopt a different definition which we give here and which the reader may check is satisfied by the random variable given above in the case of finite sample space:

Definition 33. *The conditional expectation of an \mathcal{F}_u -measurable random variable X_u given \mathcal{F}_t is a random variable $E(X_u | \mathcal{F}_t)$ which is \mathcal{F}_t -measurable and satisfies*

$$\int_{A_t} E(X_u | \mathcal{F}_t) dP = \int_{A_t} X_u dP$$

for all $A_t \in \mathcal{F}_t$.

It is easy to see that the conditional expectation is linear, i.e. if $X_u, Y_u \in \mathcal{F}_u$ and $a, b \in \mathbb{R}$, then

$$E(aX_u + bY_u | \mathcal{F}_t) = aE(X_u | \mathcal{F}_t) + bE(Y_u | \mathcal{F}_t).$$

We will also need the following computational rules (all of which can be derived by elementary methods from the definition) for conditional expectations:

$$E(E(X_u | \mathcal{F}_t)) = EX_u \tag{5.2}$$

$$E(Z_t X_u | \mathcal{F}_t) = Z_t E(X_u | \mathcal{F}_t) \text{ whenever } Z_t \in \mathcal{F}_t \tag{5.3}$$

$$E(E(X_u | \mathcal{F}_t) | \mathcal{F}_s) = E(X_u | \mathcal{F}_s) \text{ whenever } s \leq t \leq u \tag{5.4}$$

Equation (5.4) is called (the rule of) iterated expectations or the tower law. It is very useful. (But the so-called useful rule is some different.) Using

martingale

Equation (5.3) is sometimes referred to as “taking out what is known”. A consequence of (5.3) obtained by letting $X_u = 1$, is that

$$E(Z_t | \mathcal{F}_t) = Z_t \quad \text{whenever} \quad Z_t \in \mathcal{F}_t. \quad (5.5)$$

Another fact that we will often need: If a random variable Y is independent of the σ -algebra \mathcal{F} (which means exactly what you think it means), then conditional expectation reduces to ordinary expectation, $E(Y|\mathcal{F}) = E(Y)$.

Example 19. The conditional expectation can be interpreted as “our best estimate given the available information”. Or expressed mathematically: For any random variable X and any σ -algebra \mathcal{F} we have that

$$E(X|\mathcal{F}) = \arg \min_{Z: \mathcal{F}\text{-measurable}} E((X - Z)^2).$$

Mathematicians would refer to this as a projection property. To prove it let us first note that for any \mathcal{F} -measurable random variable Y we have that

$$\begin{aligned} E(Y(X - E(X|\mathcal{F}))) &= E(E(Y(X - E(X|\mathcal{F})))|\mathcal{F}) \\ &= E(YE(X - E(X|\mathcal{F}))|\mathcal{F}) \\ &= E(Y(E(X|\mathcal{F}) - E(X|\mathcal{F}))) = 0, \end{aligned}$$

where the first equality comes from iterated expectations and the second from taking out the known Z . Now let us write

$$\begin{aligned} E((X - Z)^2) &= E(((X - E(X|\mathcal{F})) + (E(X|\mathcal{F}) - Z))^2) \\ &= E((X - E(X|\mathcal{F}))^2) \\ &\quad + 2E((X - E(X|\mathcal{F}))(E(X|\mathcal{F}) - Z)) \\ &\quad + E((E(X|\mathcal{F}) - Z)^2). \end{aligned}$$

The first term on the (last) right hand side does not depend on Z . By using $E(X|\mathcal{F}) - Z$, which is \mathcal{F} -measurable, in the role of Y above, we see that the second term is 0. The third term is positive, but by choosing $Z = E(X|\mathcal{F})$, which is allowed in the minimization problem, we can make it 0, which is as small as it can possibly get. Thus the desired result follows.

Now we can state the important definition:

Definition 34. A stochastic process X is a martingale with respect to the filtration \mathbb{F} if it satisfies

$$E(X_t | \mathcal{F}_{t-1}) = X_{t-1} \quad \text{all } t = 1, \dots, T.$$

You can try out the definition immediately by showing:

Lemma 2. *A stochastic process defined as*

$$X_t = E(X | \mathcal{F}_t) \quad t = 0, 1, \dots, T$$

where $X \in \mathcal{F}_T$, is a martingale.

Let $E^P(Y; A) \equiv \int_A Y dP$ for any random variable Y and $A \in \mathcal{F}$. Using this notation and the definition (33) of a martingale, this lemma says that

$$E(X; A) = E(X_t; A) \quad \text{for all } t \text{ and } A \in \mathcal{F}_t$$

When there can be no confusion about the underlying filtration we will often write $E_t(X)$ instead of $E(X | \mathcal{F}_t)$.

Two probability measures are said to be equivalent when they assign zero probability to the same sets and since we have assumed that $P(\omega) > 0$ for all ω , the measures equivalent to P will be the ones which assign strictly positive probability to all events.

We will need a way to translate conditional expectations under one measure to conditional expectations under an equivalent measure. To do this we need the density process:

Definition 35. *Let the density (or likelihood) process Z be defined as*

$$Z_T(\omega) = \frac{Q(\omega)}{P(\omega)}$$

and

$$Z_t = E^P(Z_T | \mathcal{F}_t) \quad t = 0, 1, \dots, T.$$

We will need (but will not prove) the following result of called the *Abstract Bayes Formula*.

Proposition 10. *Let X be a random variable on (Ω, \mathcal{F}) . Then*

$$E^Q(X | \mathcal{F}_t) = \frac{1}{Z_t} E^P(X Z_T | \mathcal{F}_t).$$

5.5 Equivalent martingale measures

In this section we state and prove what is sometimes known as the fundamental theorems of asset pricing. These theorems will explain the mysterious q -probabilities which arose earlier and they will provide an indispensable tool for constructing arbitrage-free models and pricing contingent claims in these models.

equivalent
probability
measures
abstract Bayes
formula

price process
dividend process
equivalent
martingale
measure

We maintain the setup with a bank account generated by the short rate process ρ and N securities with price- and dividend processes $S = (S^1, \dots, S^N)$, $\delta = (\delta^1, \dots, \delta^N)$. Define the corresponding discounted processes $\tilde{S}, \tilde{\delta}$ by defining for each $i = 1, \dots, N$

local
characterization
of martingale
measure

$$\begin{aligned}\tilde{S}_t^i &= \frac{S_t^i}{R_{0,t}} & t = 0, \dots, T, \\ \tilde{\delta}_t^i &= \frac{\delta_t^i}{R_{0,t}} & t = 1, \dots, T.\end{aligned}$$

Definition 36. A probability measure Q on \mathcal{F} is an equivalent martingale measure (EMM) if $Q(\omega) > 0$ all ω and for all $i = 1, \dots, N$

$$\tilde{S}_t^i = E_t^Q \left(\sum_{j=t+1}^T \tilde{\delta}_j^i \right) \quad t = 0, \dots, T-1. \quad (5.6)$$

The term martingale measure has the following explanation: Given a (one-dimensional) security price process S whose underlying dividend process only pays dividend δ_T at time T . Then the existence of an EMM gives us that

$$\tilde{S}_t = E_t^Q \left(\tilde{\delta}_T \right) \quad t = 0, \dots, T-1.$$

and therefore Lemma 2 tell us that the discounted price process $(\tilde{S}_0, \tilde{S}_1, \dots, \tilde{S}_{T-1}, \tilde{\delta}_T)$ is a martingale, which we may more tellingly write as

$$S_t = E_t^Q \left(\frac{S_{t+1}}{1 + \rho_t} \right).$$

We can rewrite the definition of an equivalent martingale measure into the follow local characterization:

Theorem 3. A measure Q is an equivalent martingale measure if and only if the following holds

$$S_t^i = E_t^Q \left(\frac{S_{t+1}^i + \delta_{t+1}^i}{1 + \rho_t} \right) \text{ for all } i \text{ and } t \text{ (and } \omega \text{)}.$$

Proof. (Short form. The reader is encouraged to “cross the dot the i’s and cross the t’s” him- or herself.) Rewrite Equation (5.6) as

$$\frac{S_t^i}{R_{0,t}} = E_t^Q \left(\frac{\delta_{t+1}^i}{R_{0,t+1}} + \sum_{j=t+2}^T \tilde{\delta}_j^i \right).$$

Now use linearity of conditional expectation, use iterated expectations to write “ $E_t^Q = E_t^Q E_{t+1}^Q$ ”, use the definition of R (a couple of times), and finally use Equation (5.6) for $t + 1$. ■

fundamental
theorem of asset
pricing, first

We are now ready to formulate and prove what is sometimes known as the first fundamental theorem of asset pricing in a version with discrete time and finite state space:

Theorem 4. *In our security market model the following statements are equivalent:*

1. *There are no arbitrage opportunities.*
2. *There exists an equivalent martingale measure.*

Proof. We have already seen that no arbitrage is equivalent to the existence of a consistent price functional F . Therefore, what we show in the following is that there is a one-to-one correspondence between consistent price functionals (up to multiplication by a positive scalar) and equivalent martingale measures. We will need the following notation for the restriction of F to an \mathcal{F}_t -measurable random variable: Let δ^X be a dividend process whose only payout is X at time t . Define

$$F_t(X) = F(\delta^X).$$

If we assume (as we do from now on) that $F_0(1) = 1$, we may think of $F_t(1_A)$ as the price at time 0 of a claim (if it trades) paying off 1 at time t if $\omega \in A$. Note that since we have assumed the existence of a money market account, we have

$$F_T(R_{0,T}) = 1 \tag{5.7}$$

First, assume there is no arbitrage and let F be a consistent price functional. Our candidate as equivalent martingale measure is defined as follows:

$$Q(A) = F_T(1_A R_{0,T}) \quad A \in \mathcal{F} \equiv \mathcal{F}_T. \tag{5.8}$$

By the strict positivity, linearity and (5.7) we see that Q is a probability measure which is strictly positive on all ω . We may write (5.8) as

$$E^Q 1_A = F_T(1_A R_{0,T}) \quad A \in \mathcal{F} \equiv \mathcal{F}_T$$

and by writing a random variable X as a sum of constants times indicator functions, we note that

$$E^Q(X) = F_T(X R_{0,T}) \tag{5.9}$$

Now we want to check the condition (5.6). By definition (33) this is equivalent to showing that for every security we have

$$E^Q(1_A \tilde{S}_t^i) = E^Q \left(1_A \sum_{j=t+1}^T \tilde{\delta}_j^i \right) \quad t = 1, \dots, T. \quad (5.10)$$

Consider for given $A \in \mathcal{F}_t$ the following trading strategy ϕ :

- Buy one unit of stock i at time 0 (this costs S_0^i). Invest all dividends before time t in the money market account and keep them there at least until time t .
- At time t , if $\omega \in A$ (and this we know at time t since $A \in \mathcal{F}_t$) sell the security and invest the proceeds in the money market account, i.e. buy S_t^i units of the 0'th security and roll over the money until time T .
- If $\omega \notin A$, then hold the i 'th security to time T .

This strategy clearly only requires an initial payment of S_0^i . The dividend process generated by this strategy is non-zero only at time 0 and at time T . At time T the dividend is

$$\begin{aligned} \delta_T^\phi &= 1_A R_{t,T} \left(S_t^i + \sum_{j=1}^t \delta_j^i R_{j,t} \right) + 1_{A^c} \sum_{j=1}^T \delta_j^i R_{j,T} \\ &= 1_A R_{0,T} \left(\tilde{S}_t^i + \sum_{j=1}^t \tilde{\delta}_j^i \right) + 1_{A^c} \sum_{j=1}^T \delta_j^i R_{j,T} \end{aligned}$$

One could also choose to just buy the i 'th security and then roll over the dividends to time T . Call this strategy ψ . This would generate a terminal dividend which we may write in a complicated but useful way as

$$\begin{aligned} \delta_T^\psi &= 1_A \sum_{j=1}^T \delta_j^i R_{j,T} + 1_{A^c} \sum_{j=1}^T \delta_j^i R_{j,T} \\ &= 1_A R_{0,T} \sum_{j=1}^T \tilde{\delta}_j^i + 1_{A^c} \sum_{j=1}^T \delta_j^i R_{j,T} \end{aligned}$$

The dividend stream of both strategies at time 0 is $-S_0^i$. We therefore have

$$F_T(\delta_T^\phi) = F_T(\delta_T^\psi)$$

which in turn implies

$$F_T \left(1_A R_{0,T} \left(\tilde{S}_t^i + \sum_{j=1}^t \tilde{\delta}_j^i \right) \right) = F_T \left(1_A R_{0,T} \sum_{j=1}^T \tilde{\delta}_j^i \right)$$

i.e.

$$F_T \left(1_A R_{0,T} \tilde{S}_t^i \right) = F_T \left(1_A R_{0,T} \sum_{j=t+1}^T \tilde{\delta}_j^i \right).$$

Now use (5.9) to conclude that

$$E^Q(1_A \tilde{S}_t^i) = E^Q \left(1_A \sum_{j=t+1}^T \tilde{\delta}_j^i \right)$$

and that is what we needed to show. Q is an equivalent martingale measure.

Now assume that Q is an equivalent martingale measure. Define for an arbitrary dividend process δ

$$F(\delta) = E^Q \sum_{j=0}^T \tilde{\delta}_j$$

Clearly, F is linear and strictly positive. Now consider the dividend process δ^ϕ generated by some trading strategy ϕ . To show consistency we need to show that

$$\phi_0 \cdot \tilde{S}_0 = E^Q \sum_{j=1}^T \tilde{\delta}_j^\phi.$$

Notice that we know that for individual securities we have

$$\tilde{S}_0^i = E^Q \sum_{j=1}^T \tilde{\delta}_j^i.$$

We only need to extend that to portfolios. We do some calculations (where

fundamental
theorem of asset
pricing, second

we make good use of the iterated expectations rule $E^Q E_j^Q = E^Q E_{j-1}^Q$)

$$\begin{aligned}
E^Q \sum_{j=1}^T \tilde{\delta}_j^\phi &= E^Q \left(\sum_{j=1}^T \phi_{j-1} \cdot (\tilde{S}_j + \tilde{\delta}_j) - \phi_j \cdot \tilde{S}_j \right) \\
&= E^Q \left(\sum_{j=1}^T \phi_{j-1} \cdot \left(E_j^Q \left(\sum_{k=j}^T \tilde{\delta}_k \right) \right) - \phi_j \cdot E_j^Q \left(\sum_{k=j+1}^T \tilde{\delta}_k \right) \right) \\
&= E^Q \left(\sum_{j=1}^T \phi_{j-1} \cdot \left(E_{j-1}^Q \left(\sum_{k=j}^T \tilde{\delta}_k \right) \right) - \sum_{j=2}^T \phi_{j-1} \cdot E_{j-1}^Q \left(\sum_{k=j}^T \tilde{\delta}_k \right) \right) \\
&= E^Q \left(\phi_0 \cdot \left(E_0^Q \sum_{k=1}^T \tilde{\delta}_k \right) \right) \\
&= E^Q \left(\phi_0 \cdot \tilde{S}_0 \right) \\
&= \phi_0 \cdot \tilde{S}_0 \blacksquare
\end{aligned}$$

Earlier, we established a one-to-one correspondence between consistent price functionals (normed to 1 at date 0) and equivalent martingale measures. Therefore we have also proved the following second fundamental theorem of asset pricing:

Corollary 1. *Assume the security model is arbitrage-free. Then the market is complete if and only if the equivalent martingale measure is unique.*

Another immediate consequence from the definition of consistent price functionals and equivalent martingale measures is the following

Corollary 2. *Let the security model defined by (S, δ) (including the money market account) on $(\Omega, P, \mathcal{F}, \mathbb{F})$ be arbitrage-free and complete. Then the augmented model obtained by adding a new pair (S^{N+1}, δ^{N+1}) of security price and dividend processes is arbitrage-free if and only if*

$$\tilde{S}_t^{N+1} = E_t^Q \left(\sum_{j=t+1}^T \tilde{\delta}_j^{N+1} \right) \tag{5.11}$$

i.e.

$$\frac{S_t^{N+1}}{R_{0,t}} = E_t^Q \left(\sum_{j=t+1}^T \frac{\delta_j^{N+1}}{R_{0,j}} \right)$$

where Q is the unique equivalent martingale measure for (S, δ) .

In the special case where the discount rate is deterministic the expression splitting index simplifies somewhat. For ease of notation assume that the spot interest rate is not only deterministic but also constant and let $R = 1 + \rho$. Then (5.11) becomes

$$\begin{aligned} S_t^{N+1} &= R^t E_t^Q \left(\sum_{j=t+1}^T \frac{\delta_j^{N+1}}{R_{0,j}} \right) \\ &= E_t^Q \left(\sum_{j=t+1}^T \frac{\delta_j^{N+1}}{R^{j-t}} \right) \end{aligned} \quad (5.12)$$

5.6 One-period submodels

Before we turn to applications we note a few results for which we do not give proofs. The results show that the one-period model which we analyzed earlier actually is very useful for analyzing multi-period models as well.

Given the market model with the N -dimensional security price process S and dividend process δ and assume that a money market account exists as well. Let $A_t \in \mathcal{P}_t$ and let

$$N(A_t) \equiv |\{B \in \mathcal{P}_{t+1} : B \subseteq A_t\}|.$$

This number is often referred to as the splitting index at A_t . In our graphical representation where the set A_t is represented as a node in a graph, the splitting index at A_t is simply the number of vertices leaving that node. At each such node we can define a one-period submodel as follows: Let

$$\pi(t, A_t) \equiv (1, S_t^1(A_t), \dots, S_t^N(A_t)).$$

Denote by $B_1, \dots, B_{N(A_t)}$ the members of \mathcal{P}_{t+1} which are subsets of A_t and define

$$D(t, A_t) \equiv \begin{pmatrix} 1 + \rho_t(A_t) & \cdots & 1 + \rho_t(A_t) \\ S_{t+1}^1(B_1) + \delta_{t+1}^1(B_1) & \vdots & S_{t+1}^1(B_{N(A_t)}) + \delta_{t+1}^1(B_{N(A_t)}) \\ \vdots & & \vdots \\ S_{t+1}^N(B_1) + \delta_{t+1}^N(B_1) & \cdots & S_{t+1}^N(B_{N(A_t)}) + \delta_{t+1}^N(B_{N(A_t)}) \end{pmatrix}.$$

Then the following results hold:

Proposition 11. *The security market model is arbitrage-free if and only if the one-period model $(\pi(t, A_t), D(t, A_t))$ is arbitrage-free for all (t, A_t) where $A_t \in \mathcal{P}_t$.*

Proposition 12. *The security market model is complete if and only if the one-period model $(\pi(t, A_t), D(t, A_t))$ is complete for all (t, A_t) where $A_t \in \mathcal{P}_t$.*

In the complete, arbitrage-free case we obtain from each one-period sub-model a unique state price vector $\psi(t, A_t)$ and by following the same procedure as outlined in chapter (4) we may decompose this into a discount factor, which will be $1 + \rho_t(A_t)$, and a probability measure $q_1, \dots, q_{N(A_t)}$. The probabilities thus obtained are then the conditional probabilities $q_i = Q(B_i | A_t)$ for $i = 1, \dots, N(A_t)$. From these conditional probabilities the martingale measure can be obtained.

The usefulness of these local results is that we often build multi-period models by repeating the same one-period structure. We may then check absence of arbitrage and completeness by looking at a one-period submodel instead of the whole tree.

5.7 The Standard Binomial Model

Theory. The binomial option pricing formula hitherto considered constitutes an elegant self-contained theory of no-arbitrage pricing, yet a number of relevant questions still remain unanswered. First, there is purely practical question of how we should go about calibrating our model to market data. Specifically, how do we most easily measure the up and down state coefficients u and d from a time series of the underlying asset price process? Secondly, from an inter-theoretic perspective, how do we bring the binomial model into congruence with the continuous time framework of the Black Scholes formula? As we shall see, an answer to the latter question immediately suggests an answer to the former.

A central assumption of the Black Scholes model is that the underlying stock process S obeys *Geometric Brownian Motion*. Specifically, over the time interval $[0, \Delta t]$ the stock price $S(\Delta t)$ is assumed related to the stock price $S(0) = S$ through the equation

$$S(\Delta t) = S e^{(\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}Z}, \quad (5.13)$$

where μ and σ are constant parameters, and Z is a standard normal random variable, i.e. $Z \sim N(0, 1)$. $S(\Delta t)$ is in other words assumed to be a log-normal random variable. Do not worry too much about the implications of this result (it will become much clearer in chapter 7): for now, simply take equation (5.13) at face value, whilst appreciating the following point: μ and

σ both have clear empirical meanings, which quickly become apparent upon scrutinising the moments of $S(\Delta t)$. In fact, using standard results² about the log-normal distribution we find that

$$E(S(\Delta t)) = Se^{\mu\Delta t}, \quad E(S(\Delta t)^2) = S^2 e^{(2\mu+\sigma^2)\Delta t}, \quad (5.14)$$

or, equivalently³,

$$E(\ln(S(\Delta t)/S)) = (\mu - \frac{1}{2}\sigma^2) \Delta t, \quad \text{Var}(\ln(S(\Delta t)/S)) = \sigma^2 \Delta t.$$

Hence, $\alpha := \mu - \frac{1}{2}\sigma^2$ codifies the expected log return per unit time, while σ^2 encodes the associated variance. Unbiased sample estimates of α and σ are given by

$$\hat{\alpha} = \frac{1}{N\Delta t} \sum_{t=1}^N \ln\left(\frac{S_{t+1}}{S_t}\right), \quad \hat{\sigma}^2 = \frac{1}{(N-1)\Delta t} \sum_{t=1}^N \left[\ln\left(\frac{S_{t+1}}{S_t}\right) - \hat{\alpha}\Delta t \right]^2, \quad (5.15)$$

where $\{S_t\}_1^{N+1}$ is a time series of stock prices sampled every Δt^{th} [unit of time]. Clearly, from these estimators we may also extract an expression for $\hat{\mu}$, although this turns out to be irrelevant for our present purposes.

Now, in order to align this machinery with binomial option pricing, it is quite clear that we must bring the moments into agreement i.e. the binomial moments

$$E[S(\Delta t)] = puS + (1-p)dS, \quad E[S(\Delta t)^2] = pu^2S^2 + (1-p)d^2S^2, \quad (5.16)$$

must be equated to the corresponding expression in (5.14). Defining $m := e^{\mu\Delta t}$ and $v := e^{(2\mu+\sigma^2)\Delta t}$ this means that

$$pu + (1-p)d = m, \quad pu^2 + (1-p)d^2 = v.$$

Clearly this system is underdetermined: to solve for (p, u, d) in terms of (m, v) an auxiliary constraint will need to be enforced. A particularly popular choice in this regard is to arbitrarily fix $p = \frac{1}{2}$ whence

²Here it is helpful to recall that if $X \sim N(\mu, \sigma^2)$ then the lognormal variable $Y = \exp(X)$ has the property that $E(Y) = \exp(\mu + \frac{1}{2}\sigma^2)$. Indeed, for any $n \in \mathbb{C}$ we have that $E(Y^n) = \exp(n\mu + \frac{1}{2}n^2\sigma^2)$, so we can readily compute any moment. Hence, you can also quickly compute the variance.

³Careful here: the operations \ln and E do not commute: $\ln(E(X)) \neq E(\ln(X))$.

$$u + d = 2m, \quad u^2 + d^2 = 2v.$$

Combining these results we readily find that

$$u^2 - 2mu + 2m^2 - v = 0$$

which is a quadratic equation in u . The only meaningful solution (why?) is

$$u = m + \sqrt{v - m^2} = e^{\mu\Delta t}(1 + \sqrt{e^{\sigma^2\Delta t} - 1}). \quad (5.17)$$

Analogously, one can solve for d to find

$$d = m - \sqrt{v - m^2} = e^{\mu\Delta t}(1 - \sqrt{e^{\sigma^2\Delta t} - 1}). \quad (5.18)$$

(5.17) and (5.18) may in turn be approximated⁴ by the slightly more appetising expressions

$$u = e^{(\mu - \frac{1}{2}\sigma^2)\Delta t + \sigma\sqrt{\Delta t}}, \quad d = e^{(\mu - \frac{1}{2}\sigma^2)\Delta t - \sigma\sqrt{\Delta t}},$$

at the cost of an $\mathcal{O}(\Delta t^2)$ error in the moment matching. We have thus succeeded in deriving a binomial model which is (I) consistent with the framework of the Black Scholes model, and which (II) is stated in terms of parameters which quickly can be estimated from market data. The gravity of this result is well worth restating in a theorem:

Theorem 5. *Let $p = \frac{1}{2}$ and let u and d be given by*

$$u = e^{\alpha\Delta t + \sigma\sqrt{\Delta t}}, \quad d = e^{\alpha\Delta t - \sigma\sqrt{\Delta t}}, \quad (5.19)$$

where $\alpha := \mu - \frac{1}{2}\sigma^2$. Then (I) up to $\mathcal{O}(\Delta t^2)$ we have a model which matches the first and second moment of Geometric Brownian Motion, and (II) can be calibrated to market data through the estimators $\hat{\alpha}$ and $\hat{\sigma}$ given (5.15). This is the so-called **standard binomial model**.

It is somewhat gratifying to observe that the specification (5.19) “preserves the moment matching” under the change of measure $P \rightarrow Q$. Explicitly, up to $\mathcal{O}(\Delta t^2)$ we find that the Q -expectation of $S(\Delta t)$ under Geometric Brownian Motion matches the Q -expectation of $S(\Delta t)$ under the binomial model (and similarly for the second moment). To see how this pans out, the reader should be aware (accept at face value) that we in continuous time have the results

⁴Taylor’s formula suggests that these expressions are $\mathcal{O}(\Delta t^{3/2})$ close to their exact counterparts.

$$E^Q(S(\Delta t)) = Se^{r\Delta t}, \quad E^Q(S(\Delta t)^2) = S^2 e^{(2r+\sigma^2)\Delta t}, \quad (5.20)$$

where r is the continuously compounded risk free rate. On the other hand, the binomial model clearly implies

$$E^Q(S(\Delta t)) = quS + (1-q)dS \quad E^Q(S(\Delta t)^2) = qu^2S^2 + (1-q)d^2S^2, \quad (5.21)$$

where $q = (R - d)/(u - d)$ and $R = e^{r\Delta t}$. Substituting in the expression for q we see that the first moment is trivially satisfied. The second moment requires a bit more work. To see that the error between the second moment in (5.20) and (5.21) is no greater than $\mathcal{O}(\Delta t^2)$ consider the Taylor expansion of the continuous case:

$$E^Q[S(\Delta t)^2/S^2] = 1 + (2r + \sigma^2)\Delta t + \mathcal{O}(\Delta t^2).$$

Next, observe that the binomial moment may be written as

$$\begin{aligned} E^Q[S(\Delta t)^2/S^2] &= q(u^2 - d^2) + d^2 = \frac{R-d}{u-d}(u-d)(u+d) + d^2, \\ &= (R-d)(u+d) + d^2 = R(u+d) - du, \end{aligned}$$

which on Taylor form reads

$$\begin{aligned} E^Q[S(\Delta t)^2/S^2] &= (1 + r\Delta t + \mathcal{O}(\Delta t^2))(2 + 2\mu\Delta t + \mathcal{O}(\Delta t^2)) \\ &\quad - (1 + 2\mu\Delta t - \sigma^2\Delta t + \mathcal{O}(\Delta t^2)) \\ &= 2(1 + \mu\Delta t + r\Delta t + \mathcal{O}(\Delta t^2)) - (1 + 2\mu\Delta t - \sigma^2\Delta t + \mathcal{O}(\Delta t^2)) \\ &= 1 + (2r + \sigma^2)\Delta t + \mathcal{O}(\Delta t^2). \end{aligned}$$

The result follows immediately.

Remark 5. *Whilst the standard binomial model emulates the moments of the underlying asset in the continuous time framework, we have yet to formally establish that binomial option prices converge to the prices of the Black Scholes formula as $\Delta t \rightarrow 0$. In rigorous terms this corresponds to showing that the binomial pricing formula*

$$V(t, S) = e^{-r\Delta t}(qV(t + \Delta t, uS) + (1-q)V(t + \Delta t, dS)), \quad (5.22)$$

converges to the Black Scholes PDE

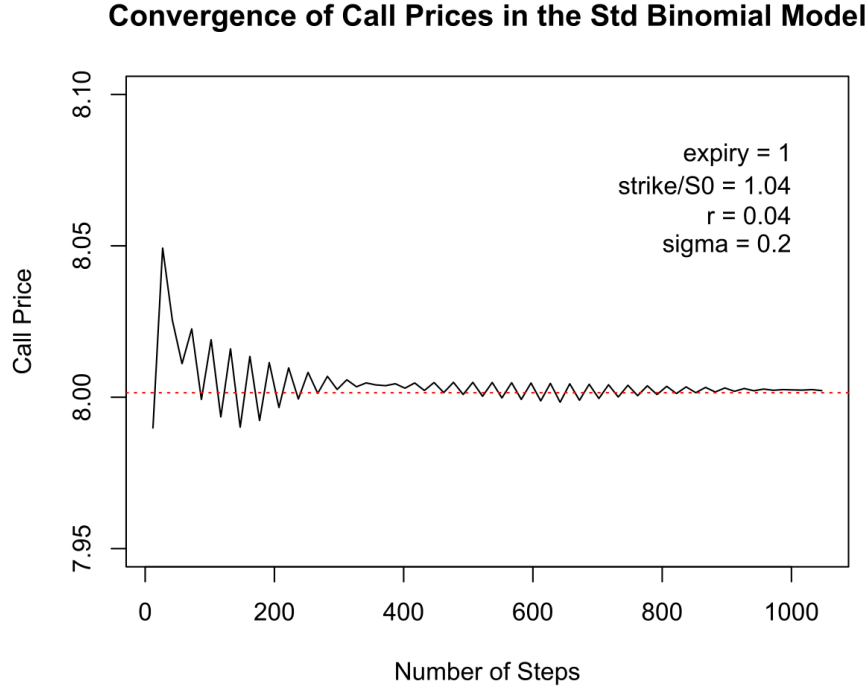


Figure 5.1: In the limit where $\Delta t \rightarrow 0$ the binomially computed call option price converges towards the Black Scholes price.

$$rV(t, S) = \frac{\partial V}{\partial t}(t, S) + rS \frac{\partial V}{\partial S}(t, S) + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 V}{\partial S^2}(t, S).$$

This can be done by performing a second order Taylor expansion of (5.22) and using the fact that the binomial moments are $\mathcal{O}(\Delta t^2)$ -close to their continuous counterparts. We leave this as an advanced exercise for the reader. To emphasise the validity of the purported convergence, we have in figure 5.1 plotted the binomial call option price for various step values $N = T/\Delta t$. Clearly, as $N \rightarrow \infty$ ($\Delta t \rightarrow 0$) the binomial price converges towards the Black Scholes price (illustrated by the dotted line).

Remark 6. Crucially, we notice that continuum limit does **not** depend on the parameter μ . This result is intimately connected to a deep result known as the First Fundamental Theorem of Asset Pricing which establishes that asset price processes under Q have a drift coefficient equal to the risk free rate, r . This gives us considerable freedom in choosing our α : Jarrow & Rudd e.g. set $\alpha = r - \frac{1}{2}\sigma^2$.

Empirics. We will end this section on a practical note, with some brief comments on calibration. First we might wonder how frequently we should sample our data in the estimation of $\hat{\alpha}$ and $\hat{\sigma}$. Is there any noticeable discrepancy between opting for daily, weekly or even monthly observations? The Δt parameter in (5.15) clearly suggests that the estimators are time scaled, and thus insensitive to our sampling frequency. To put this observation to the test, consider the table below in which we exhibit the raw returns and standard deviations (alongside their temporally scaled counterparts $\hat{\alpha}$, $\hat{\sigma}$) for the Danish C20 Index, computed on a daily, weekly and monthly basis over the horizon 1994-2014.

Horizon	Avg. raw return	$\hat{\alpha}$	St.d. raw return	$\hat{\sigma}$
1 day	0.00036	0.092	0.0122	0.195
5 days (~ 1 week)	0.00181	0.092	0.0280	0.198
21 days (~ 1 month)	0.00734	0.090	0.0578	0.200

Evidently, the variation in $\hat{\alpha}$ and $\hat{\sigma}$ is quite modest among the three sampling strategies. The presence of the temporal scaling factor Δt in $\hat{\alpha}$ and $\hat{\sigma}$ does indeed seem to calibrate the parameters “just right” (in a sense this corroborates the standard binomial model). Nonetheless, there is an important caveat to this, which should be illuminated. Due to the *telescoping* nature of the summation in $\hat{\alpha}$, this estimator will generally be unstable for different choices of Δt . Specifically, since

$$\begin{aligned}\hat{\alpha} &= \frac{1}{N\Delta t} \sum_{t=1}^N \ln \left(\frac{S_{t+1}}{S_t} \right) = \frac{1}{N\Delta t} \sum_{t=1}^N [\ln(S_{t+1}) - \ln(S_t)] \\ &= \frac{1}{N\Delta t} [\ln(S_2) - \ln(S_1) + \ln(S_3) - \ln(S_2) + \dots + \ln(S_{N+1}) - \ln(S_N)] \\ &= \frac{1}{N\Delta t} [\ln(S_{N+1}) - \ln(S_1)] = \frac{1}{N\Delta t} \ln \left(\frac{S_{N+1}}{S_1} \right),\end{aligned}$$

our daily, weekly and monthly estimators for $\hat{\alpha}$ will generally only be in agreement if they all employ the same initial and terminal stock values (which surely occurs if S_1 is chosen identically and $N \bmod 5 = N \bmod 21 = 0$). The upshot is that $\hat{\alpha}$ generally will fluctuate based on how you choose your sample. However, based on remark 6 you shouldn't be too perturbed by this. What matters is really that $\hat{\sigma}$ is stable (which is mostly the case despite the dependence of $\hat{\alpha}$ - instability in $\hat{\sigma}$ is mostly caused by lack of independence between the sample returns, which would give rise to non-zero covariance terms).

The second (interrelated) calibration issue we will discuss is the question how we should handle dates with no trading data such as weekends and public holidays? The short answer is that the totality of non-trading days are chopped up and distributed evenly across all trading days: e.g. the temporal distance between the two trading days Monday and Tuesday is not viewed as $(1/365.25)$ years, but rather $(1/252)$ years, cf. the calculation

$$\underbrace{365.25}_{\text{[days on avg. per year]}} \times \underbrace{\frac{5}{7}}_{\text{[proportion trading days per week]}} - \underbrace{9}_{\text{[trading holidays in the U.S.]}} \approx \underbrace{252}_{\text{[trading days per year]}} .$$

To put this idea to the test, let us consider what happens to the volatility of an Index (again, the Danish C20) as we move between consecutive trading days. If non-trading days are truly important, then we would expect the change in volatility to be about a factor of $\sqrt{3}$ higher over the weekend compared to the remaining days. From the data presented in the table this is clearly not the case! Two consecutive trading days nesting a weekend or a bank holiday can comfortably be approximated as though separated by a single day of trading.

day-to-day	$\hat{\sigma}$
Monday \rightarrow Tuesday	0.189
Tuesday \rightarrow Wednesday	0.200
Wednesday \rightarrow Thursday	0.192
Thursday \rightarrow Friday	0.178
Friday \rightarrow Monday	0.214

Chapter 6

Option pricing

The classical application of the arbitrage pricing machinery we have developed is to the pricing of options. The pricing models we obtain are used with minor modifications all over the world as the basis for trading billions of dollars worth of contracts every day. For students planning to become traders of financial derivatives this of course gives plenty of motivation for learning these models. But recent collapses of financial institutions have also reminded us that financial managers and executives must understand the way the derivatives markets work. A manager who understands the markets well may use them for effective risk management and will be able to implement effective control mechanisms within a firm to make sure that traders use the markets in accordance with the firm's overall objectives.

From a theoretical perspective, options are very important in several areas of finance. We will see later in the course how they are indispensable for our understanding of a firm's choice of capital structure. Also, a modern theory of capital budgeting relies critically on recognizing options involved in projects, so-called real options. And in actuarial science options appear when modelling reinsurance contracts.

6.1 Terminology

A European (American) call option on an underlying security S , with strike price K and expiration date T , gives the owner the right, but not the obligation, to buy S at a price of K at (up to and including) time T .

A European (American) put option on an underlying security S , with strike price K and expiration date T , gives the owner the right, but not the obligation, to sell S at a price of K at (up to and including) time T .

The strike price is also referred to as the exercise price, and using the

put-call parity
frictionless market

right to buy or sell is referred to as exercising the option.

There is no good reason for the American/European terminology - both types are traded in America and Europe.

In the definition above, we think of the person selling a call option (say), often referred to as the person writing an option, as actually delivering the underlying security to the option holder if the option holder decides to exercise. This is referred to as physical delivery. In reality, options are often cash settled. This means that instead of the option holder paying K to the writer of the call and the writer delivering the stock, the holder merely receives an amount $S_T - K$ from the option writer.

Some common examples of options are stock options in which the underlying security is a stock, currency options in which the underlying security is a foreign currency and where the strike price is to be thought of as an exchange rate, bond options which have bonds as underlying security and index options whose underlying security is not really a security but a stock market index (and where the contracts are then typically cash settled.) It will always be assumed that the underlying security has non-negative value.

6.2 Diagrams, strategies and put-call parity

Before we venture into constructing exact pricing models we develop some feel for how these instruments work. In this section we focus on what can be said about options if all we assume is that all securities (stocks, bonds, options) can be bought and sold in arbitrary quantities at the given prices with no transactions costs or taxes. This assumption we will refer to as an assumption of frictionless markets. We will also assume that at any time t and for any date $T > t$, there exists a zero coupon bond with maturity T in the market whose price at time t is $d(t, T)$.

An immediate consequence of our frictionless markets assumption is the following


Proposition 13. *The value of an American or European call option at the expiration date T is equal to $C_T = \max(S_T - K, 0)$, where S_T is the price of the underlying security at time T . The value of an American or European put option at the expiration date T is equal to $\max(K - S_T, 0)$.*

Proof. Consider the call option. If $S_T < K$, we must have $C_T = 0$, for if $C_T > 0$ you would sell the option, receive a positive cash flow, and there would be no exercise.¹ If $S_T \geq K$, we must have $C_T = S_T - K$. For if

¹Actually, here we need to distinguish between whether the person who bought the option is an idiot or a complete idiot. Both types are not very smart to pay something for

$C_T > S_T - K$ you would sell the option and buy the stock. After the option has been exercised, you are left with a total cash flow of $C_T - S_T + K > 0$, and you would have no future obligations arising from this trade. If $C_T < S_T - K$, buy the option, exercise it immediately, and sell the stock. The total cash flow is $-C_T + S_T - K > 0$, and again there would be no future obligations arising from this trade. The argument for the put option is similar.


We often represent payoffs of options at an exercise date using payoff diagrams, which show the value of the option as a function of the value of the underlying:




the option at time T . The idiot, however, would realize that there is no reason to pay K to receive the stock which can be bought for less in the market. The complete idiot would exercise the option. Then you as the person having sold the option would have to buy the stock in the market for S_T , but that would be more than financed by the K you received from the complete idiot.

moneyness

Of course, you can turn these hockey sticks around in which case you are looking at the value of a written option:



Note that we are only looking at the situation at an exercise date (i.e. date T for a European option). Sometimes we wish to take into account that the option had an initial cost at date 0, c_0 for a call, p_0 for a put, in which case we get the following profit diagrams:




Of course, we are slightly allergic to subtracting payments occurring at different dates without performing some sort of discounting. Therefore, one may also choose to represent the prices of options by their time T forward discounted values $\frac{c_0}{d(0,T)}$ and $\frac{p_0}{d(0,T)}$.

The world of derivative securities is filled with special terminology and here are a few additions to your vocabulary: A call option with strike price K is said to be (deep) in-the-money at time t if $S_t > K$ ($S_t \gg K$). The opposite situation $S_t < K$ ($S_t \ll K$) is referred to as the call option being (deep) out-of-the-money. If $S_t \approx K$, the option is said to be at-the-money. The same terminology applies to put options but with 'opposite signs': A

put option is in-the-money if $S_t < K$.

The diagrams we have seen so far considered positions consisting of just one option. We considered a long position, i.e. a position corresponding to holding the option, and we considered a short position, i.e. a position corresponding to having written an option. One of the attractive features of options is that they can be combined with positions in other options, the underlying security and bonds to produce more complicated payoffs than those illustrated in the profit diagrams above. We will see examples of this in the exercises. Note that you should think of the payoff diagram for holding the stock and the diagram for holding the bond as being represented by:



Until further notice we will assume that the stock does not pay any dividends in the time interval $[0, T]$. This means that if you own the stock you will not receive any cash unless you decide to sell the stock. With this assumption and the maintained assumption of frictionless markets we will give some restrictions on option prices which follow solely from arbitrage considerations.

The most important relation is the so-called put-call parity for European options. Consider the portfolio strategy depicted in the table below and the associated cash flows at time t and time T . Assume that both options are European, expire at date T and have strike price equal to K :

strategy/cashflow	date t	date $T, S_T \leq K$	date $T, S_T > K$
sell 1 call	c_t	0	$K - S_T$
buy 1 put	$-p_t$	$K - S_T$	0
buy stock	$-S_t$	S_T	S_T
sell K bonds	$Kd(t, T)$	$-K$	$-K$
total cash flow	must be 0	0	0

Note that we have constructed a portfolio which gives a payoff of 0 at time T no matter what the value of S_T . Since the options are European we need not consider any time points in (t, T) . This portfolio must have price 0, or else

put-call parity
forward price
forward contract

there would be an obvious arbitrage strategy. If, for example, the portfolio had positive value, we would sell the portfolio (corresponding to reversing the strategy in the table) and have no future obligations. In other words we have proved that in a frictionless market we have the following

Proposition 14. (*Put-call parity*) *The price c_t of a European call and the price p_t of a European put option with expiration date T and exercise price K must satisfy*

$$c_t - p_t = S_t - Kd(t, T).$$

Note one simple but powerful consequence of this result: When deciding which parameters may influence call and put prices the put-call parity gives a very useful way of testing intuitive arguments. If S_t, K and $d(t, T)$ are fixed, then a change in a parameter which produces a higher call price, must produce a higher put-price as well. One would easily for example be tricked into believing that in a model where S_T is stochastic, a higher mean value of S_T given S_t would result in a higher call price since the call option is more likely to finish in-the-money and that it would result in a lower put price since the put is more likely then to finish out-of-the money. But if we assume that S_t and the interest rate are held fixed, put-call parity tells us that this line of reasoning is wrong.

Also note that for $K = \frac{S_t}{d(t, T)}$ we have $c_t = p_t$. This expresses the fact that the exercise price for which $c_t = p_t$ is equal to the forward price of S at time t . A forward contract is an agreement to buy the underlying security at the expiration date T of the contract at a price of F_t . Note that F_t is specified at time t and that the contract unlike an option forces the holder to buy. In other words you can lose money at expiration on a forward contract. The forward price F_t is decided so that the value of the forward contract at date t is 0. Hence the forward price is not a price to be paid for the contract at date t . It is more like the exercise price of an option. Which value of F_t then gives the contract a value of 0 at date t ? Consider the following portfolio argument:

strategy / cashflow	date t	date T
buy 1 stock	$-S_t$	S_T
sell $\frac{S_t}{d(t, T)}$ bonds	S_t	$-\frac{S_t}{d(t, T)}$
sell 1 forward	0	$F_t - S_T$
total cash flow	0	$F_t - \frac{S_t}{d(t, T)}$

Note that the cash flow at time T is known at time t and since the cash flow by definition of the forward price is equal to 0 at date t , the cash flow at

date T must be 0 as well. Hence

$$F_t = \frac{S_t}{d(t, T)}.$$

Merton's tunnel
call option,
American

Note that buying a call and selling a put, both with exercise price K and expiration date T , is equivalent to buying forward at the price K . Therefore the convention that the forward contract has value 0 at date t is exactly equivalent to specifying K so that $c_t = p_t$.

6.3 Restrictions on option prices

In this section we derive some bounds on call prices which must be satisfied in frictionless markets. The line of reasoning used may of course be used on put options as well.

Consider a European call option with expiration date T and exercise price K . Assume that the underlying security does not pay any dividends during the life of the option. Then the value of the option c_t satisfies

$$S_t \geq c_t \geq \max(0, S_t - Kd(t, T)), \quad (6.1)$$

which is sometimes called Merton's tunnel.

Proof. Clearly, $c_t \geq 0$. Also, the corresponding put option satisfies $p_t \geq 0$. Hence

$$c_t \geq c_t - p_t = S_t - Kd(t, T) \quad (6.2)$$

where we have used put-call parity. To see that $S_t \geq c_t$, assume that $S_t < c_t$ and consider the strategy of buying the stock and selling the option. That gives a positive cash flow at time t . If at time T , $S_T > K$ and the option is exercised the stock is delivered to the option holder and K is received. If the option is not exercised, the stock can be sold at non-negative value. ■

It is clear that an American option is more valuable than the corresponding European option, hence we note that the price C_t of an American option also satisfies $C_t \geq S_t - Kd(t, T)$. If interest rates are positive, i.e. $d(t, T) < 1$, this produces the interesting result that the value of the American call is always strictly greater than the immediate exercise value $S_t - K$ when $t < T$. This shows the important result that an American option on a non-dividend paying stock should never be exercised early. Our inequalities above show that it will be better to sell the option. A corresponding result does not hold for put options. This is perhaps not so surprising considering that postponing the exercise of a put postpones the receipt of K , whereas delaying the exercise of a call delays the payment of K .

Typically, stocks pay dividends and it is important to take this into account when pricing options. It will often be the case that the option contract does not take into account whether the underlying stock pays dividends. A dividend payment will normally produce a drop in the stock price and an owner of a call option will be hurt by this drop without receiving the benefit of a dividend. A date t is denoted an ex-dividend date if purchasing the stock at time $s < t$ gives the new owner part in the next dividend payment whereas a purchase at time t does not. For simplicity, we assume in the following that the dividend payment takes place at the ex-dividend date. Furthermore, we will assume that the size of the dividend is known some time before the dividend date. In a world with no taxes it ought to be the case then that the drop in the stock price around the dividend date is equal to the size of the dividend. Assume, for example, that the drop in the stock price is less than D . Then buying the stock right before the dividend date for a price of S_{t-} and selling it for S_{t+} immediately after the dividend date will produce a cash flow of $S_{t+} + D - S_{t-} > 0$. This resembles an arbitrage opportunity and it is our explanation for assuming in the following that $S_{t-} = S_{t+} + D$.

Now let us consider the price at time 0 of a European call option on a stock which is known to pay one dividend D at time t . Then

$$c_0 \geq \max(0, S_0 - Kd(0, T) - Dd(0, t)).$$

Again, $c_0 \geq 0$ is trivial. Assume $c_0 < S_0 - Kd(0, T) - Dd(0, t)$. Then buy the left hand side and sell the right hand side. At time t , we must pay dividend D on the stock we have sold, but that dividend is exactly received from the D zero coupon bonds with maturity t . At time T the value of the option we have sold is equal to $\max(0, S_T - K)$. The value of the right hand side is equal to $S_T - K$. If $S_T \geq K$ the total position is 0. If $S_T < K$ the total position has value $K - S_T$. Hence we have constructed a positive cash flow while also receiving money initially. This is an arbitrage opportunity and hence we rule out $c_0 < S_0 - Kd(0, T) - Dd(0, t)$.

There are many possible variations on the dividend theme. If dividends are not known at time 0 we may assume that they fall within a certain interval and then use the endpoints of this interval to bound calls and puts. The reader may verify that the maximal dividend is important for bounding calls and the minimum dividend for bounding put prices.

However, we maintain the assumption of a known dividend and finish this section by another important observation on the early exercise of American calls on dividend paying stocks. Assume that the stock pays a dividend at time t and that we are at time $0 < t$. It is then not optimal to exercise the option at time 0 whereas it may be optimal right before time t . To see that it

is not optimal at time 0, note that the American option contains as a part of its rights an option with expiration date $s \in (0, t)$, and since this option is an option on a non-dividend paying stock we know that its value is larger than $S_0 - K$, which is the value of immediate exercise. Therefore, the American option is also more worth than $S_0 - K$ and there is no point in exercising before t . To see that it may be optimal to exercise right before t , consider a firm which pays a liquidating dividend to all its shareholders. The stock will be worthless after the liquidation and so will the call option. Certainly, the option holder is better off to exercise right before the dividend date to receive part of the liquidating dividend.

binomial model
binomial model
stock
option pricing

The picture is much more complicated for puts. In the next section we will see how to compute prices for American puts in binomial models and this will give us the optimal exercise strategy as well.

6.4 Binomial models for stock options

In this section we will go through the binomial model for pricing stock options. Our primary focus is the case where the underlying security is a non-dividend paying stock but it should be transparent that the binomial framework is highly flexible and will easily handle the pricing and hedging of derivative securities with more complicated underlying securities.

We consider a model with T periods and assume throughout that the following two securities trade:


1. A bank account with a constant short rate process ρ , so $1 + \rho_t = R$, and

$$R_{s,t} = R^{t-s} \text{ for } s < t.$$

2. A stock² S , which pays no dividends³, whose price at time 0 is S_0 and whose evolution under the measure P is described in the tree (where we have assumed that $u > R > d > 0$) shown below.

²Since there is only one stock we will write S instead of S^1 .

³To comply with the mathematical model of the previous chapter we should actually say that the stock pays a liquidating dividend of S_T at time T . We will however speak of S_T as the price at time T of the stock.



The mathematical description of the process is as follows: Let U_1, \dots, U_T be a sequence of i.i.d. Bernoulli variables, let $p = P(U_1 = 1)$ and define

$$N_t = \sum_{i=1}^t U_i.$$

Think of N_t as the number of up-jumps that the stock has had between time 0 and t . Clearly, this is a binomially distributed random variable. Let $u > R > d > 0$ be constants. Later, we will see how these parameters are chosen in practice. Then

$$S_t = S_0 u^{N_t} d^{t-N_t}. \tag{6.3}$$

Using the results on one-period submodels it is clear that the model is arbitrage free and complete and that the equivalent martingale measure is given in terms of conditional probabilities as

$$Q(S_t = uS_{t-1} | S_{t-1}) \equiv q = \frac{R - d}{u - d}$$

$$Q(S_t = dS_{t-1} | S_{t-1}) = 1 - q = \frac{u - R}{u - d}.$$

6.5 Pricing the European call

call option,
European

We now have the martingale measure Q in place and hence the value at time t of a European call with maturity T is given in an arbitrage-free model by

$$C_t = \frac{1}{R^{T-t}} E^Q (\max(0, S_T - K) | \mathcal{F}_t).$$

Using this fact we get the following

Proposition 15. *Let the stock and money market account be as described in section 6.4. Then the price of a European call option with exercise price K and maturity date T is given as*

$$C_t = \frac{1}{R^{T-t}} \sum_{i=0}^{T-t} \binom{T-t}{i} q^i (1-q)^{T-t-i} \max(0, S_t u^i d^{T-t-i} - K).$$

Proof. Since the money market account and S_0 are deterministic, we have that we get all information by observing just stock-prices, or equivalently the U 's, i.e. $\mathcal{F}_t = \sigma(S_1, \dots, S_t) = \sigma(U_1, \dots, U_t)$. By using (6.3) twice we can write

$$S_T = S_t u^{(N_T - N_t)} d^{(T-t) - (N_T - N_t)} = S_t u^Z d^{(T-t) - Z},$$

where $Z = N_T - N_t = \sum_{j=t+1}^T U_j \stackrel{Q}{\sim} \text{bi}(q; (T-t))$, and Z is independent of \mathcal{F}_t (because the U 's are independent). Therefore

$$R^{T-t} C_t = E^Q((S_T - K)^+ | \mathcal{F}_t) = E^Q((S_t u^Z d^{(T-t) - Z} - K)^+ | \mathcal{F}_t).$$

At this point in the narrative we need something called “the useful rule”. It states the following: Suppose we are given a function $f : \mathbb{R}^2 \mapsto \mathbb{R}$, a σ -algebra \mathcal{F} , an \mathcal{F} -measurable random variable X and a random variable Y that is independent of \mathcal{F} . Define the function $g : \mathbb{R} \mapsto \mathbb{R}$ by $g(x) = E(f(x, Y))$. Then $E(f(X, Y) | \mathcal{F}) = g(X)$. We then use this in the above expression with S_t playing the role of X , Z as Y , and $f(x, y) = (x u^y d^{(T-t) - y} - K)^+$. By using the general transformation rule for discrete random variables $E(h(Y)) = \sum_{y_i} h(y_i) P(Y = y_i)$, and the fact that Z is Q -binomially distributed we get in the notation of “the useful rule” that

$$g(x) = \sum_{i=0}^{T-t} \binom{T-t}{i} q^i (1-q)^{(T-t)-i} (x u^i d^{(T-t)-i} - K)^+,$$

and the desired result follows. ■

**hedging
call option,
European**

We rewrite the expression for C_0 using some handy notation. Let a be the smallest number of upward jumps needed for the option to finish in the money, i.e.

$$\begin{aligned} a &= \min_{j \in \mathbf{N}} \{j | S_0 u^j d^{T-j} > K\} \\ &= \min_{j \in \mathbf{N}} \{j | j \ln u + (T - j) \ln d > \ln(K/S_0)\} \\ &= \min_{j \in \mathbf{N}} \{j | j > \ln(K/(S_0 d^T)) / \ln(u/d)\} \\ &= \left\lceil \frac{\ln\left(\frac{K}{S_0 d^T}\right)}{\ln\left(\frac{u}{d}\right)} \right\rceil + 1. \end{aligned}$$

Letting

$$\Psi(a; T, q) = \sum_{i=a}^T \binom{T}{i} q^i (1-q)^{T-i},$$

we may write (you may want to check the first term on the RHS)

$$C_0 = S_0 \Psi(a; T, q') - \frac{K}{R^T} \Psi(a; T, q) \quad (6.4)$$

where

$$q' = \frac{u}{R} q.$$

Using put-call parity gives us the price of the European put:

Corollary 3. *The price of a European put option with T periods to maturity, exercise price K and the stocks as underlying security has a price at time 0 given by*

$$P_0 = \frac{K}{R^T} (1 - \Psi(a; T, q)) - S_0 (1 - \Psi(a; T, q'))$$

Note that our option pricing formulae use T to denote the number of periods until maturity. Later, we will be more explicit in relating this to actual calendar time.

6.6 Hedging the European call

We have already seen in a two period model how the trading strategy replicating a European call option may be constructed. In this section we simply state the result for the case with T periods and we then note an interesting

way of expressing the result. We consider the case with a bank account and one risky asset S and assume that the market is complete and arbitrage-free. The European call option has a payout at maturity of

*self-financing
trading strategy
replication
splitting index*

$$\delta_T^c = \max(S_T - K, 0).$$

Proposition 16. *A self-financing trading strategy replicating the dividend process of the option from time 1 to T is constructed recursively as follows: Find $\phi_{T-1} = (\phi_{T-1}^0, \phi_{T-1}^1)$ such that*

$$\phi_{T-1}^0 R + \phi_{T-1}^1 S_T = \delta_T^c.$$

For $t = T - 2, T - 3, \dots, 1$ find $\phi_t = (\phi_t^0, \phi_t^1)$ such that

$$\phi_t^0 R + \phi_t^1 S_{t+1} = \phi_{t+1}^0 + \phi_{t+1}^1 S_{t+1}.$$

The trading strategy is self-financing by definition, replicates the call and its initial price of $\phi_0^0 + \phi_0^1 S_0$ is equal to the arbitrage-free price of the option. We may easily extend to the case where both the underlying and the contingent claim have dividends other than the one dividend of the option considered above. In that case the procedure is the following: Find $\phi_{T-1} = (\phi_{T-1}^0, \phi_{T-1}^1)$ such that

$$\phi_{T-1}^0 R + \phi_{T-1}^1 (S_T + \delta_T) = \delta_T^c.$$

For $t = T - 2, T - 3, \dots, 1$ find $\phi_t = (\phi_t^0, \phi_t^1)$ such that

$$\phi_t^0 R + \phi_t^1 (S_{t+1} + \delta_{t+1}) = \phi_{t+1}^0 + \phi_{t+1}^1 S_{t+1} + \delta_{t+1}^c. \quad (6.5)$$

In this case the trading strategy is not self-financing in general but it matches the dividend process of the contingent claim, and the initial price of the contingent claim is still $\phi_0^0 + \phi_0^1 S_0$.

In general, Equation (6.5) is compact notation for a whole bunch of linear equations, namely one for each submodel; the number of unknowns is equal to $\dim(S)$, the number of equations is the splitting index, ie. the number of future states. And the equations must be solved recursively backwards. But in many applications things simple. If S is 1-dimensional, the splitting index is 2 (so we have simple binomial submodels; refer to their states as “up” and “down”), and neither S nor the derivative pay dividends, then

- The RHSs of (6.5) are just the option’s price in different states, say π^u and π^d .

delta hedging

- Subtract the up-equation from the down-equation to get the replication portfolio('s stock holdings)

$$\phi_t^1 = \frac{\pi^u - \pi^d}{S^u - S^d} = \text{“} \frac{\Delta\pi}{\Delta S} \text{”} := \Delta.$$

- This procedure/technique is called delta hedging. A very good mnemonic. The Δ in the equation above is called the option's (delta) hedge ratio.

Further insight into the hedging strategy is given by the proposition below. Recall the notation

$$\tilde{S}_t = \frac{S_t}{R_{0,t}}$$

for the discounted price process of the stock. Let C_t denote the price process of a contingent claim whose dividend process is δ^c and let

$$\begin{aligned} \tilde{C}_t &= \frac{C_t}{R_{0,t}} \\ \tilde{\delta}_t^c &= \frac{\delta_t^c}{R_{0,t}} \end{aligned}$$

denote the discounted price and dividend processes of the contingent claim. Define the conditional covariance under the martingale measure Q as follows:

$$\text{cov}^Q(X_{t+1}, Y_{t+1} | \mathcal{F}_t) = E^Q((X_{t+1} - X_t)(Y_{t+1} - Y_t) | \mathcal{F}_t)$$

The following can be shown (but we omit the proof):

Proposition 17. *Assume that the stock pays no dividends during the life of the option. The hedging strategy which replicates δ^c is computed as follows:*

$$\begin{aligned} \phi_t^1 &= \frac{\text{cov}^Q(\tilde{S}_{t+1}, \tilde{C}_{t+1} + \tilde{\delta}_{t+1}^c | \mathcal{F}_t)}{\text{var}^Q(S_{t+1} | \mathcal{F}_t)} & t = 0, 1, \dots, T-1 \\ \phi_t^0 &= \tilde{C}_t - \phi_t^1 \tilde{S}_t & t = 0, 1, \dots, T-1 \end{aligned}$$

Note the similarity with regression analysis. We will not go further into this at this stage. But this way of looking at hedging is important when defining so-called risk minimal trading strategies in incomplete markets.

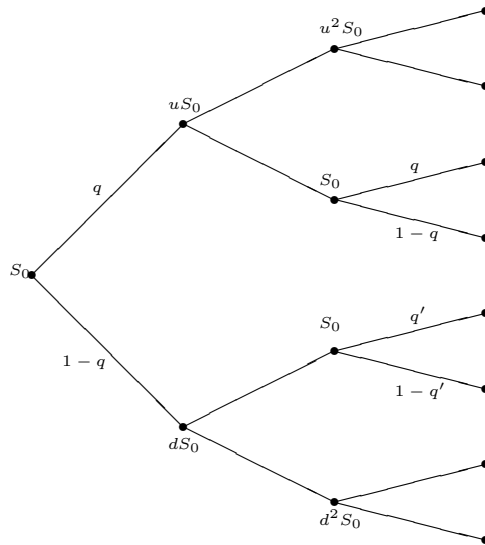


Figure 6.2: A tree that's not a lattice.

process at time $t + 1$. The binomial model of this chapter is clearly a Markov chain. An important consequence of this is that when $\mathcal{F}_t = \sigma(S_0, \dots, S_t)$ then for any (measurable) function f and time points $t < u$ there exists a function g such that

$$E^Q(f(S_u) | \mathcal{F}_t) = g(S_t). \quad (6.6)$$

In other words, conditional expectations of functions of future values given everything we know at time t can be expressed as a function of the value of S_t at time t . The way S arrived at S_t is not important. We used this fact in the formula for the price of the European call: There, the conditional expectation given time t information became a function of S_t . The past did not enter into the formula. We can therefore represent the behavior of the process S in a recombining tree, also known as a *lattice*, as shown in Figure 6.1 in which one node at time t represents exactly one value of S_t . Another way of stating this is to say that the tree keeps track of the number of up-jumps that have occurred, not the order in which they occurred. A full event tree would keep track of the exact timing of the up-jumps.

To see what can go wrong, Figure 6.2 shows a process that is not Markov. The problem is at time 2 when the value of the process is S_0 , we need to know the pre-history of S to decide whether the probability of going up to uS is equal to q or q' . In standard binomial models such behavior is normally precluded.

Note that now the number of nodes required at time t is only $t + 1$, and then using several hundred time periods is no problem for a spreadsheet.

A technical issue which we will not address here is the following: Normally we specify the process under the measure P , and it need not be the case that the Markov property is preserved under a change of measure. However, one may show that if the price process is Markov under P and the model is complete and arbitrage-free, then the price process is Markov under the equivalent martingale measure Q as well.

A second condition for using a recombining tree to price a contingent claim is a condition on the contingent claim itself:

Definition 37. *A contingent claim with dividend process δ^c is path independent if $\delta_t = f_t(S_t)$ for some (measurable) function f .*

Indeed if the claim is path independent and the underlying process is Markov, we have

$$\begin{aligned} C_t &= R_{0,t} E \left(\sum_{i=t+1}^T \tilde{\delta}_i^c \middle| \mathcal{F}_t \right) \\ &= R_{0,t} E \left(\sum_{i=t+1}^T f_i(S_i) \middle| \mathcal{F}_t \right) \\ &= R_{0,t} E \left(\sum_{i=t+1}^T f_i(S_i) \middle| S_t \right) \end{aligned}$$

and the last expression is a function of S_t by the Markov property. A European option with expiration date T is path-independent since its only dividend payment is at time T and is given as $\max(S_T - K, 0)$.

The Asian option is an example of a contingent claim which is not path-independent. An Asian option on the stock, initiated at time 0, expiration date T and exercise price K has a payoff at date T given by

$$C_T^{asian} = \max \left(0, \left(\frac{1}{T+1} \sum_{t=0}^T S_t \right) - K \right)$$

Hence the average of the stock price over the period determines the option price. Clearly, S_T is not sufficient to describe the value of the Asian option at maturity. To compute the average value one needs the whole path of S . As noted above, even in a binomial model keeping track of the whole path for, say, 50 periods becomes intractable.

contingent claim
Asian option
"see contingent claim, path dependent"
contingent claim, path dependent

put option,
American

6.8 The binomial model for American puts

We describe in this section a simple way of pricing the American put option in a binomial model. Strictly speaking, an American put is not a contingent claim in the sense we have thought of contingent claims earlier. Generally, we have thought of contingent claims as random variables or sometimes as processes but a put is actually not specified until an exercise policy is associated with the put. What we will do in the following is to simultaneously solve for the optimal exercise policy, i.e. the one that maximizes the expected, discounted value of the cash flows under the martingale measure, and the price of the option. The argument given is not a proof but should be enough to convince the reader that the right solution is obtained (it is fairly easy to show that another exercise policy will create arbitrage opportunities for the option writer).

The value of an American put at its maturity is easy enough:⁵

$$P_T = \max(0, K - S_T). \quad (6.7)$$

Now consider the situation one period before maturity. If the put has not been exercised at that date, the put option holder has two possibilities: Exercise the put at time $T - 1$ or hold the put to maturity. The value of holding the put to maturity is given as the discounted (back to time $T - 1$) value of (6.7), whereas the value at time $T - 1$ of exercising immediately is $K - S_{T-1}$ something only to be considered of course if $K > S_{T-1}$. Clearly, the put option holder has a contract whose value is given by the maximal value of these two strategies, i.e.⁶

$$P_{T-1} = \max \left(K - S_{T-1}, E^Q \left(\frac{P_T}{R} \middle| \mathcal{F}_{T-1} \right) \right).$$

Continue in this fashion by working backwards through the tree to obtain the price process of the American put option given by the recursion

$$P_{t-1} = \max \left(K - S_{t-1}, E^Q \left(\frac{P_t}{R} \middle| \mathcal{F}_{t-1} \right) \right) \quad t = 1, \dots, T.$$

Once this price process is given we see that the optimal exercise strategy is to exercise the put the first time t for which

$$K - S_t > E^Q \left(\frac{P_{t+1}}{R} \middle| \mathcal{F}_t \right).$$

⁵Or is it? As it stands P_t is really the value at time t given that the put has not been exercised at times $0, 1, t - 1$. But that will most often be exactly what we are interested in; if we exercised the put to years ago, we really don't care about it anymore.

⁶We do not need 0 in the list of arguments of max since positivity is assured by $P_T \geq 0$.

This way of thinking is easily translated to American call options on dividend paying stocks for which early exercise is something to consider.

**implied
volatility**
Black-Scholes
formula

6.9 Implied volatility

We assume in this section that the Black-Scholes formula is known to the reader: The price at time t of a European call option maturing at time T , when the exercise price is K and the underlying security is a non-dividend paying stock with a price of S_t , is given in the Black-Scholes framework by

$$C_t = S_t \Phi(d_1) - K e^{-r(T-t)} \Phi(d_2)$$

where

$$d_1 = \frac{\log\left(\frac{S_t}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)(T-t)}{\sigma\sqrt{T-t}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T-t}$$

where Φ is the cumulative distribution function of a standard normal distribution.

Consider the Black-Scholes formula for the price of a European call on an underlying security whose value at time 0 is S_0 : Recall that Φ is a distribution function, hence $\Phi(x) \rightarrow 1$ as $x \rightarrow \infty$ and $\Phi(x) \rightarrow 0$ as $x \rightarrow -\infty$. Assume throughout that $T > 0$. From this it is easy to see that $c_0 \rightarrow S_0$ as $\sigma \rightarrow \infty$. By considering the cases $S_0 < K \exp(-rT)$, $S_0 = K \exp(-rT)$ and $S_0 > K \exp(-rT)$ separately, it is easy to see that as $\sigma \rightarrow 0$, we have $c_0 \rightarrow \max(0, S_0 - K \exp(-rT))$. By differentiating c_0 with respect to σ , one may verify that c_0 is strictly increasing in σ . Therefore, the following definition makes sense:

Definition 38. *Given a security with price S_0 . Assume that the risk free rate (i.e. the rate of the money market account) is equal to r . Assume that the price of a call option on the security with exercise price K and time to maturity T is observed to have a price of c^{obs} with*

$$\max(0, S_0 - K \exp(-rT)) < c^{obs} < S_0.$$

Then the implied volatility of the option is the unique value of σ for which

$$c_0(S_0, K, T, \sigma, r) = c^{obs}. \quad (6.8)$$

In other words, the implied volatility is the unique value of the volatility which makes the Black-Scholes model 'fit' c^{obs} . Clearly, we may also associate an implied volatility to a put option whose observed price respects the appropriate arbitrage bounds.

There is no closed-form expression for implied volatility; Equation 6.8 must be solved numerically. Bisection works nicely (whereas a Newton-Raphson search without safety checks may diverge for deep out-of-the money options).

A very important reason for the popularity of implied volatility is the way in which it allows a transformation of option prices which are hard to compare into a common scale. Assume that the price of a stock is 100 and the riskfree rate is 0.1. If one observed a price of 9.58 on a call option on the stock with exercise price 100 and 6 months to maturity and a price of 2.81 on a put option on the stock with exercise price 95 and 3 months to maturity then it would require a very good knowledge of the Black-Scholes model to see if one price was in some way higher than the other. However, if we are told that the implied volatility of the call is 0.25 and the implied volatility of the put is 0.30, then at least we know that compared to the Black-Scholes model, the put is more expensive than the call. This way of comparing is in fact so popular that traders in option markets typically do not quote prices in (say) dollars, but use 'vols' instead.

If the Black-Scholes model were true the implied volatility of all options written on the same underlying security should be the same, namely equal to the volatility of the stock and this volatility would be a quantity we could estimate from historical data. In short, in a world where the Black-Scholes model holds, historical volatility (of the stock) is equal to implied volatility (of options written on the stock). In practice this is not the case - after all the Black-Scholes model is only a model. The expenses of hedging an option depend on the volatility of the stock during the life of the option. If, for example, it is known that, after a long and quiet period, important news about the underlying stock will arrive during the life of the option, the option price should reflect the fact that future fluctuations in the stock price might be bigger than the historical ones. In this case the implied volatility would be higher than the historical.

However, taking this knowledge of future volatility into account one could still imagine that all implied volatilities of options on the same underlying were the same (and equal to the 'anticipated' volatility). In practice this is not observed either. To get an idea of why, we consider the notion of portfolio insurance.

6.10 Portfolio insurance, implied volatility and crash fears

portfolio
insurance
implied
volatility
crash fear


Consider a portfolio manager who manages a portfolio which is diversified so that the value of her portfolio follows that of the market stock index. Assume that the value of her portfolio is 1000 times the value of the index which is assumed to be at 110. The portfolio manager is very worried about losing a large portion of the value of the portfolio over the next year - she thinks that there is a distinct possibility that the market will crash. On the other hand she is far from certain. If she were certain, she could just move the money to a bank at a lower but safer expected return than in the stock market. But she does not want to exclude herself from the gains that a surge in the index would bring. She therefore decides to buy portfolio insurance in such a way that the value of her portfolio will never fall below a level of (say) 90.000. More specifically, she decides to buy 1000 put options with one year to maturity and an exercise price of 90 on the underlying index. Now consider the value of the portfolio after a year as a function of the level of the index S_T :

value of index	$S_T \geq 90$	$S_T < 90$
value of stocks	$S_T \times 1000$	$S_T \times 1000$
value of puts	0	$1000 \times (90 - S_T)$
total value	$S_T \times 1000 > 90.000$	90.000

Although it has of course not been costless to buy put options, the portfolio manager has succeeded in preventing the value of her portfolio from falling below 90.000. Since the put options are far out-of-the-money (such contracts are often called “lottery tickets”) at the time of purchase they are probably not that expensive. And if the market booms she will still be a successful portfolio manager.

But what if she is not alone with her fear of crashes. We may then imagine a lot of portfolio managers interested in buying out-of-the-money put options hence pushing up the price of these contracts. This is equivalent to saying that the implied volatility goes up and we may experience the scenario shown in the graph below, in which the implied volatility of put options is higher for low exercise price puts:

volatility smile
firm
 debt
 equity



This phenomenon is called a “smirk”. If (as it is often seen from data) the implied volatility is increasing (the dotted part of the curve) for puts that are in the money, then we have what is known as a “smile”. Actually options that are deeply in-the-money are rarely traded, so the implied volatility figures used to draw “the other half” of the smile typically comes from out-of-the-money calls. (Why/how? Recall the put-call parity.)

A smirk has been observed before crashes and it is indicative of a situation where the Black-Scholes model is not a good model to use. The typical modification allows for stock prices to jump discontinuously but you will have to wait for future courses to learn about this.

6.11 Debt and equity as options on firm value

In this section we consider a very important application of option pricing. Our goal is to learn a somewhat simplified but extremely useful way of thinking about a firm which is financed by debt and equity (see below). A fundamental assumption in this section is that a firm has a market value given by a stochastic process V . In economies with Arrow-Debreu securities in which we know prices and production plans adopted by the firms, it is easy to define the value of a firm as the (net) value of its production. In reality things are of course a lot more complicated. It is hard to know, for example, what the value of NovoNordisk is - i.e. what is the market value of the firm’s assets (including know-how, goodwill etc.). Part of the problem is of course that it is extremely difficult to model future prices and production levels. But in a sense the actual value does not matter for this section in that the ‘sign’ of the results that we derive does not depend on what the value of the firm is - only the “magnitude” does.

The fundamental simplification concerns the capital structure of the firm. Assume that the firm has raised capital to finance its activities in two ways: It has issued stocks (also referred to as equity) and debt. The debt consists of zero coupon bonds with face value D maturing at time T . Legally what distinguishes the debt holders from the stock holders is the following: The stock holders control the firm and they decide at time T whether the firm should repay its debt to the bondholders. If the bondholders are not repaid in full they can force the firm into bankruptcy and take over the remaining assets of the firm (which means both controlling and owning it). The stocks will then be worthless. If the stockholders pay back D at maturity to the bondholders, they own the firm entirely. They may then of course decide to issue new debt to finance new projects but we will not worry about that now.

capital structure
equity
debt

It is clear that the stockholders will have an interest in repaying the bondholders precisely when $V_T > D$. Only then will the expense in paying back the debt be more than outweighed by the value of the firm. If $V_T < D$ (and there are no bankruptcy costs) the stockholders will default on their debt, the firm will go into bankruptcy and the bondholders will take over. In short, we may write the value of debt and equity at time T as

$$\begin{aligned} B_T &= \min(D, V_T) = D - \max(D - V_T, 0) \\ S_T &= \max(V_T - D, 0). \end{aligned}$$

In other words, we may think of equity as a call option on the value of the firm and debt as a zero coupon bond minus a put option on the value of the firm. Assuming then that V behaves like the underlying security in the Black-Scholes model and that there exists a money market account with interest rate r , we can use the Black-Scholes model to price debt and equity at time 0 :

$$\begin{aligned} B_0 &= D \exp(-rT) - p_0(V_0, D, T, \sigma, r) \\ S_0 &= c_0(V_0, D, T, \sigma, r) \end{aligned}$$

where p_0, c_0 are Black-Scholes put and call functions.

Let us illustrate a potential conflict between stockholders and bondholders in this model. Assume that at time 0 the firm has the possibility of adopting a project which will not alter the value of the firm at time 0, but which will have the effect of increasing the volatility of the process V . Since both the value of the call and the put increases when σ increases we see that the stockholders will like this project since it increases the value of the equity whereas the bondholders will not like the project since the put option which they have in a sense written will be a greater liability to them. This is a very

clear and very important illustration of so-called asset substitution, a source of conflict which exists between stock-and bondholders of a firm. This setup of analyzing the value of debt and equity is useful in a number of contexts and you should make sure that you understand it completely. We will return to this towards the end of the course when discussing corporate finance.

Chapter 7

The Black-Scholes formula

7.1 Black-Scholes as a limit of binomial models

So far we have not specified the parameters p, u, d and R which are of course critical for the option pricing model. Also, it seems reasonable that if we want the binomial model to be a realistic model for stock prices over a certain interval of time we should use a binomial model which divides the (calendar) time interval into many sub-periods. In this chapter we will first show that if one divides the interval into finer and finer periods and choose the parameters carefully, the value of the option converges to a limiting formula, the Black-Scholes formula, which was originally derived in a continuous time framework. We then describe that framework and show how to derive the formula in it.

Our starting point is an observed stock price whose logarithmic return satisfies

$$E^P \left[\ln \left(\frac{S_t}{S_{t-1}} \right) \right] = \mu$$

and

$$V^P \left(\ln \left(\frac{S_t}{S_{t-1}} \right) \right) = \sigma^2,$$

where S_t is the price of the stock t years after the starting date 0. Also, assume that the bank account has a continuously compounded interest rate of r , i.e. an amount of 1 placed in the bank grows to $\exp(r)$ in one year. Note that since $R^T = \exp(T \ln(R))$, a yearly rate of $R = 1.1$ (corresponding to a yearly rate of 10%) translates into the continuous compounding analogue $r = \ln(1.1)$ and this will be a number smaller than 0.1.

Suppose we construct a binomial model covering T years, and that we divide each year into n periods. This gives a binomial model with nT periods.

In each 1-period submodel choose

$$\begin{aligned} u_n &= \exp(\sigma/\sqrt{n}), \\ d_n &= \exp(-\sigma\sqrt{n}) = u_n^{-1}, \\ R_n &= \exp\left(\frac{r}{n}\right). \end{aligned}$$

Further, let us put

$$p_n = \frac{1}{2} + \frac{1}{2} \frac{\mu}{\sigma\sqrt{n}},$$

not that it matters much for our purposes. With the setup in the n th model specified above you may show by simple computation that the one-year logarithmic return satisfies

$$E^P \left[\ln \left(\frac{S_1}{S_0} \right) \right] = n \{ p_n \ln(u_n) + (1 - p_n) \ln(d_n) \} = \mu$$

and

$$V^P \left(\ln \left(\frac{S_1}{S_0} \right) \right) = \sigma^2 - \frac{1}{n} \mu^2,$$

so the log-return of the price process has the same mean and almost the same variance as the process we have observed. And since

$$V^P \left(\ln \left(\frac{S_1}{S_0} \right) \right) \rightarrow \sigma^2 \text{ for } n \rightarrow \infty,$$

it is presumably so that large values of n brings us closer to to “desired” model.

Let us now investigate precisely what happens to stock and call prices when n tends to infinity. For each n we may compute the price of an expiry- T call option the binomial model and we know that it is given as

$$C^n = S_0 \Psi(a_n; nT; q'_n) - \frac{K}{(R_n)^T} \Psi(a_n; nT; q_n) \quad (7.1)$$

where

$$q_n = \frac{R_n - d_n}{u_n - d_n}, \quad q'_n = \frac{u_n}{R_n} q_n$$

and a_n is the smallest integer larger than $\ln(K/(S_0 d_n^{Tn})) / \ln(u_n/d_n)$. Note that alternatively we may write(7.1) as

$$C^n = S_0 Q'(S_n(T) > K) - K e^{-rT} Q(S_n(T) > K) \quad (7.2)$$

where $S_n(T) = S_0 u_n^j d_n^{Tn-j}$ and $j \stackrel{Q}{\sim} \text{bi}(Tn, q_n)$ and $j \stackrel{Q'}{\sim} \text{bi}(Tn, q'_n)$. It is easy to see that

$$\begin{aligned} M_n^Q &:= E^Q(\ln S_n(T)) = \ln S_0 + Tn(q_n \ln u_n + (1 - q_n) \ln d_n) \\ V_n^Q &:= V^Q(\ln S_n(T)) = Tnq_n(1 - q_n)(\ln u_n - \ln d_n)^2, \end{aligned}$$

and that similar expressions (with q'_n instead of q_n) hold for Q' -moments. Now rewrite the expression for M_n^Q in the following way:

$$\begin{aligned} M_n^Q - \ln S_0 &= Tn \left(\frac{\sigma}{\sqrt{n}} \frac{e^{r/n} - e^{-\sigma/\sqrt{n}}}{e^{\sigma/\sqrt{n}} - e^{-\sigma/\sqrt{n}}} - \frac{\sigma}{\sqrt{n}} \frac{e^{\sigma/\sqrt{n}} - e^{r/n}}{e^{\sigma/\sqrt{n}} - e^{-\sigma/\sqrt{n}}} \right) \\ &= T\sqrt{n}\sigma \left(\frac{2e^{r/n} - e^{\sigma/\sqrt{n}} - e^{-\sigma/\sqrt{n}}}{e^{\sigma/\sqrt{n}} - e^{-\sigma/\sqrt{n}}} \right). \end{aligned}$$

Recall the Taylor-expansion to the second order for the exponential function: $\exp(\pm x) = 1 \pm x + x^2/2 + o(x^2)$. From this we get

$$\begin{aligned} e^{r/n} &= 1 + r/n + o(1/n) \\ e^{\pm\sigma/\sqrt{n}} &= 1 \pm \sigma/\sqrt{n} + \sigma^2/(2n) + o(1/n). \end{aligned}$$

Inserting this in the M_n^Q expression yields

$$\begin{aligned} M_n^Q - \ln S_0 &= T\sqrt{n}\sigma \left(\frac{2r/n - \sigma^2/n + o(1/n)}{2\sigma/\sqrt{n} + o(1/n)} \right) \\ &= T\sigma \left(\frac{2r - \sigma^2 + o(1)}{2\sigma + o(1/\sqrt{n})} \right) \\ &\rightarrow T \left(r - \frac{\sigma^2}{2} \right) \text{ for } n \rightarrow \infty. \end{aligned}$$

Similar Taylor expansions for V_n^Q , $M_n^{Q'}$ and $V_n^{Q'}$ show that

$$\begin{aligned} V_n^Q &\rightarrow \sigma^2 T, \\ M_n^{Q'} - \ln S_0 &\rightarrow T \left(r + \frac{\sigma^2}{2} \right) \text{ (note the change of sign on } \sigma^2), \\ V_n^{Q'} &\rightarrow \sigma^2 T. \end{aligned}$$

So now we know what the Q/Q' moments converge to. Yet another way to think of $\ln S_n(T)$ is as a sum of Tn independent Bernoulli-variables with possible outcomes $(\ln d_n, \ln u_n)$ and probability parameter q_n (or q'_n). This means that we have a sum of (well-behaved) independent random variables for which the first and second moments converge. Therefore we can use a

Central Limit
Theorem
Black-Scholes
formula

version of the Central Limit Theorem¹ to conclude that the limit of the sum is normally distributed, i.e.

$$\ln S_n(T) \xrightarrow{Q/Q'} N(\ln S_0 + (r \pm \sigma^2/2)T, \sigma^2 T).$$

This means (almost by definition of the form of convergence implied by CLT) that when determining the limit of the probabilities on the right hand side of (7.2) we can (or: have to) substitute $\ln S_n(T)$ by a random variable X such that

$$X \stackrel{Q/Q'}{\sim} N(\ln S_0 + (r \pm \sigma^2/2)T, \sigma^2 T) \Leftrightarrow \frac{X - \ln S_0 - (r \pm \sigma^2/2)T}{\sigma\sqrt{T}} \stackrel{Q/Q'}{\sim} N(0, 1).$$

The final analysis:

$$\begin{aligned} \lim_{n \rightarrow \infty} C^n &= \lim_{n \rightarrow \infty} (S_0 Q'(\ln S_n(T) > \ln K) - K e^{-rT} Q(\ln S_n(T) > \ln K)) \\ &= S_0 Q'(X > \ln K) - K e^{-rT} Q(X > \ln K) \\ &= S_0 Q' \left(\frac{X - \ln S_0 - (r + \sigma^2/2)T}{\sigma\sqrt{T}} > \frac{\ln K - \ln S_0 - (r + \sigma^2/2)T}{\sigma\sqrt{T}} \right) \\ &\quad - K e^{-rT} Q \left(\frac{X - \ln S_0 - (r - \sigma^2/2)T}{\sigma\sqrt{T}} > \frac{\ln K - \ln S_0 - (r - \sigma^2/2)T}{\sigma\sqrt{T}} \right) \end{aligned}$$

Now multiply by -1 inside the Q 's (hence reversing the inequalities), use that the $N(0, 1)$ -variables on the left hand sides are symmetric and continuous, and that $\ln(x/y) = \ln x - \ln y$. This shows that

$$\lim_{n \rightarrow \infty} C^n = S_0 \Phi(d_1) - K e^{-rT} \Phi(d_2),$$

where Φ is the standard normal distribution function and

$$\begin{aligned} d_1 &= \frac{\ln\left(\frac{S_0}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}, \\ d_2 &= \frac{\ln\left(\frac{S_0}{K}\right) + \left(r - \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}. \end{aligned}$$

This formula for the call price is called the Black-Scholes formula.

So far we can see it just as an artifact of going to the limit in a particular way in a binomial model. But the formula is so strikingly beautiful and simple

¹Actually, the most basic De Moivre-version will not quite do because we do not have a scaled sum of *identically* distributed random variables; the two possible outcomes depend on n . You need the notion of a triangular array and the Lindeberg-Feller-version of the Central Limit Theorem.

that there must be more to it than that. In particular, we are interested in the question: Does that exist a “limiting” model in which the above formula is the exact call option price? The answer is: Yes. In the next section we describe what this “limiting” model looks like, and show that the Black-Scholes formula gives the exact call price in the model. That does involve a number of concepts, objects and results that we cannot possibly make rigorous in this course, but the reader should still get a “net benefit” and hopefully an appetite for future courses in financial mathematics.

lognormal
distribution

7.2 The Black-Scholes model

The Black-Scholes formula for the price of a call option on a non-dividend paying stock is one of the most celebrated results in financial economics. In this section we will indicate how the formula is derived, or with the previous limiting argument in mind: A different way to derive the formula. A rigorous derivation requires some fairly advanced mathematics which is beyond the scope of this course. Fortunately, the formula is easy to interpret and to apply. Even if there are some technical details left over for a future course, the rigorous understanding we have from our discrete-time models of how arbitrage pricing works will allow us to apply the formula safely.

The formula is formulated in a continuous time framework with random variables that have continuous distribution. The continuous-time and infinite state space setup will not be used elsewhere in the course.² But let us mention that if one wants to develop a theory which allows random variables with continuous distribution and if one wants to obtain results similar to those of the previous chapters, then one has to allow continuous trading as well. By ‘continuous trading’ we mean that agents are allowed to readjust portfolios continuously through time.

If X is normally distributed $X \sim N(\alpha, \sigma^2)$, then we say that $Y := \exp(X)$ is *lognormally* distributed and write $Y \sim LN(\alpha, \sigma^2)$. There is one thing you must always remember about lognormal distributions:

$$\text{If } Y \sim LN(\alpha, \sigma^2) \text{ then } E(Y) = \exp\left(\alpha + \frac{\sigma^2}{2}\right).$$

If you have not seen this before, then you are strongly urged to check it. (With that result you should also be able to see why there is no need to

²A setup which combines *discrete time* and continuous distributions will be encountered later when discussing CAPM and APT, but the primary focus of these models will be to explain stock price behavior and not – as we are now doing – determining option prices for a given behavior of stock prices

Brownian motion

use “brain RAM” remembering the variance of a lognormally distributed variable.) Often the lognormal distribution is preferred as a model for stock price distributions since it conforms better with the institutional fact that prices of a stock are non-negative and the empirical observation that the logarithm of stock prices seem to show a better fit to a normal distribution than do prices themselves. However, specifying a distribution of the stock price at time t , say, is not enough. We need to specify the whole process of stock prices, i.e. we need to state what the joint distribution $(S_{t_1}, \dots, S_{t_N})$ is for any $0 \leq t_1 < \dots < t_N$. To do this the following object is central.

Definition 39. *A Brownian motion (BM) is a stochastic process $B = (B_t)_{t \in [0; \infty[}$ -i.e. a sequence of random variables indexed by t such that:*


1. $B_0 = 0$
2. $B_t - B_s \sim N(0, t - s) \forall s < t$
3. B has independent increments, i.e. for every N and a set of N time points $t_1 < \dots < t_N$, $B_{t_1}, B_{t_2} - B_{t_1}, B_{t_3} - B_{t_2}, \dots, B_{t_N} - B_{t_{N-1}}$ are independent random variables.

That these demands on a process can be satisfied simultaneously is not trivial. But don't worry, Brownian motion does exist. It is, however, a fairly “wild” object. The sample paths (formally the mapping $t \mapsto B_t$ and intuitively simply the graph you get by plotting “temperature/stock price/...” against time) of Brownian motion are continuous everywhere but differentiable nowhere. The figure shows a simulated sample path of a BM and should give an indication of this.

A useful fact following from the independent increment property is that for any measurable $f : \mathbb{R} \rightarrow \mathbb{R}$ for which $E[|f(B_t - B_s)|] < \infty$ we have

$$E[f(B_t - B_s) | \mathcal{F}_s] = E[f(B_t - B_s)] \quad (7.3)$$

where $\mathcal{F}_s = \sigma\{B_u : 0 \leq u \leq s\}$.



Black-Scholes
model
geometric
Brownian
motion

The fundamental assumption of the Black-Scholes model is that the stock price can be represented by

$$S_t = S_0 \exp(\alpha t + \sigma B_t) \quad (7.4)$$

where B_t is a Brownian motion. Such a process is called a geometric Brownian motion. Furthermore, it assumes that there exists a riskfree bank account that behaves like

$$\beta_t = \exp(rt) \quad (7.5)$$

where r is a constant (typically $r > 0$). Hence β_t is the continuous time analogue of $R_{0,t}$.

What does (7.4) mean? Note that since $B_t \sim N(0, t)$, S_t has a lognormal distribution and

$$\ln\left(\frac{S_{t_1}}{S_0}\right) = \alpha t_1 + \sigma B_{t_1},$$

$$\ln\left(\frac{S_{t_2}}{S_{t_1}}\right) = \alpha(t_2 - t_1) + \sigma(B_{t_2} - B_{t_1})$$

Since αt , $\alpha(t_2 - t_1)$, and σ are constant, we see that $\ln\left(\frac{S_{t_1}}{S_0}\right)$ and $\ln\left(\frac{S_{t_2}}{S_{t_1}}\right)$ are independent. The *return*, defined in this section as the logarithm of the price relative, that the stock earns between time t_1 and t_2 is independent

volatility

of the return earned between time 0 and time t_1 , and both are normally distributed. We refer to σ as the *volatility* of the stock - but note that it really describes a property of the logarithmic return of the stock. There are several reasons for modelling the stock price as geometric BM with drift or equivalently all logarithmic returns as independent and normal. First of all, unless it is blatantly unreasonable, modelling “random objects” as “*n iid*” is *the* way to start. Empirically it is often a good approximation to model the logarithmic returns as being normal with fixed mean and fixed variance through time.³ From a probabilistic point of view, it can be shown that if we want a stock price process with continuous sample paths and we want returns to be independent and stationary (but not necessarily normal from the outset), then geometric BM is the only possibility. And last but not least: It gives rise to beautiful financial theory.

If you invest one dollar in the money market account at time 0, it will grow as $\beta_t = \exp(rt)$. Holding one dollar in the stock will give an uncertain amount at time t of $\exp(\alpha t + \sigma B_t)$ and this amount has an expected value of

$$E \exp(\alpha t + \sigma B_t) = \exp(\alpha t + \frac{1}{2}\sigma^2 t).$$

The quantity $\mu = \alpha + \frac{1}{2}\sigma^2$ is often referred to as the *drift* of the stock. We have not yet discussed (even in our discrete models) how agents determine μ and σ^2 , but for now think of it this way: Risk averse agents will demand μ to be greater than r to compensate for the uncertainty in the stock’s return. The higher σ^2 is, the higher should μ be.

7.3 A derivation of the Black-Scholes formula

In this section we derive the Black-Scholes model taking as given some facts from continuous time finance theory. The main assertion is that the fundamental theorem of asset pricing holds in continuous time and, in particular, in the Black-Scholes setup:

$$\begin{aligned} S_t &= S_0 \exp(\alpha t + \sigma B_t) \\ \beta_t &= \exp(rt) \end{aligned}$$

What you are asked to believe in this section are the following facts:

- There is no arbitrage in the model and therefore there exists an *equivalent martingale measure* Q such that the discounted stock price $\frac{S_t}{\beta_t}$ is

³But skeptics would say many empirical analyses of financial data is a case of “believing is seeing” rather than the other way around.

a martingale under Q . (Recall that this means that $E^Q \left[\frac{S_t}{\beta_t} \mid \mathcal{F}_s \right] = \frac{S_s}{\beta_s}$). The probabilistic behavior of S_t under Q is given by

$$S_t = S_0 \exp \left(\left(r - \frac{1}{2} \sigma^2 \right) t + \sigma \widetilde{B}_t \right), \tag{7.6}$$

where \widetilde{B}_t is a SBM under the measure Q .

- To compute the price of a call option on S with expiration date T and exercise price K , we take the discounted expected value of $C_T = [S_T - K]^+$ assuming the behavior of S_t given by (7.6).

Recall that in the binomial model we also found that the expected return of the stock under the martingale measure was equal to that of the riskless asset. (7.6) is the equivalent of this fact in the continuous time setup. Before sketching how this expectation is computed note that we have not defined the notion of arbitrage in continuous time. Also we have not justified the form of S_t under Q . But let us check at least that the martingale behavior of $\frac{S_t}{\beta_t}$ seems to be OK (this may explain the “ $-\frac{1}{2}\sigma^2t$ ”-term which is in the expression for S_t). Note that

$$\begin{aligned} E^Q \left[\frac{S_t}{\beta_t} \right] &= E^Q \left[S_0 \exp \left(-\frac{1}{2} \sigma^2 t + \sigma \widetilde{B}_t \right) \right] \\ &= S_0 \exp \left(-\frac{1}{2} \sigma^2 t \right) E^Q \left[\exp \left(\sigma \widetilde{B}_t \right) \right]. \end{aligned}$$

But $\sigma \widetilde{B}_t \sim N(0, \sigma^2 t)$ and since we know how to compute the mean of the lognormal distribution we get that

$$E^Q \left[\frac{S_t}{\beta_t} \right] = S_0 = \frac{S_0}{\beta_0}, \text{ since } \beta_0 = 1.$$

By using the property (7.3) of the Brownian motion one can verify that

$$E^Q \left[\frac{S_t}{\beta_t} \mid \mathcal{F}_s \right] = \frac{S_s}{\beta_s}, \text{ (}\mathcal{F}_s \text{ = "information at time s")}$$

but we will not do that here.⁴

⁴If you want to try it yourself, use

$$\begin{aligned} E \left[\frac{S_t}{\beta_t} \mid \mathcal{F}_s \right] &= E \left[\frac{S_t \beta_s S_s}{S_s \beta_t \beta_s} \mid \mathcal{F}_s \right] \\ &= \frac{S_s}{\beta_s} E \left[\frac{S_t \beta_s}{S_s \beta_t} \mid \mathcal{F}_s \right] \end{aligned}$$

Accepting the fact that the call price at time 0 is

$$C_0 = \exp(-rT) E^Q \left[S_0 \exp \left(\left(r - \frac{1}{2} \sigma^2 \right) T + \sigma \tilde{B}_T \right) - K \right]^+$$

we can get the Black-Scholes formula: We know that $\sigma B_T \sim N(0, \sigma^2 T)$ and also “the rule of the unconscious statistician”, which tells us that to compute $E[f(X)]$ for some random variable X which has a density $p(x)$, we compute $\int f(x) p(x) dx$. This gives us

$$C_0 = e^{-rT} \int_{\mathbb{R}} \left[S_0 e^{(r-\sigma^2/2)T+x} - K \right]^+ \frac{1}{\sqrt{2\pi\sigma\sqrt{T}}} e^{-\frac{1}{2} \frac{x^2}{\sigma^2 T}} dx$$

The integrand is different from 0 when

$$S_0 e^{(r-\sigma^2/2)T+x} > K$$

i.e. when⁵

$$x > \ln(K/S_0) - (r - \sigma^2/2) T \equiv d$$

So

$$\begin{aligned} C_0 &= e^{-rT} \int_d^\infty \left(S_0 e^{(r-\frac{1}{2}\sigma^2)T+x} - K \right) \frac{1}{\sqrt{2\pi\sigma\sqrt{T}}} e^{-\frac{1}{2} \frac{x^2}{\sigma^2 T}} dx \\ &= e^{-rT} S_0 \underbrace{\int_d^\infty \frac{1}{\sqrt{2\pi\sigma\sqrt{T}}} e^{(r-\frac{1}{2}\sigma^2)T+x} e^{-\frac{1}{2} \frac{x^2}{\sigma^2 T}} dx}_{:=A} - K e^{-rT} \underbrace{\int_d^\infty \frac{1}{\sqrt{2\pi\sigma\sqrt{T}}} e^{-\frac{1}{2} \frac{x^2}{\sigma^2 T}} dx}_{:=B}. \end{aligned}$$

It is easy to see that $B = K e^{-rT} \text{Prob}(Z > d)$, where $Z \sim N(0, \sigma^2 T)$. So by using symmetry and scaling with $\sigma\sqrt{T}$ we get that

$$B = K e^{-rT} \Phi(d_2),$$

where (as before)

$$d_2 = -\frac{d}{\sigma\sqrt{T}} = \frac{\ln\left(\frac{S_0}{K}\right) + \left(r - \frac{1}{2}\sigma^2\right) T}{\sigma\sqrt{T}}.$$

and then see if you can bring (7.3) into play and use

$$E[\exp(\sigma(B_t - B_s))] = \exp\left(\frac{1}{2}\sigma^2(t-s)\right).$$

⁵This should bring up memories of the quantity a which we defined in the binomial model.

So “we have half the Black-Scholes formula”. The A -term requires a little more work. First we use the change of variable $y = x/(\sigma\sqrt{T})$ to get (with a few rearrangements, a completion of the square, and a further change of variable ($z = y - \sigma\sqrt{T}$))

$$\begin{aligned} A &= S_0 e^{-T\sigma^2/2} \int_{-d_2}^{\infty} \frac{1}{\sqrt{2\pi}} e^{\sigma\sqrt{T}y - y^2/2} dy \\ &= S_0 e^{-T\sigma^2/2} \int_{-d_2}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-(y - \sigma\sqrt{T})^2/2 + T\sigma^2/2} dy \\ &= S_0 \int_{-d_1}^{\infty} \frac{1}{\sqrt{2\pi}} e^{-z^2/2} dz, \end{aligned}$$

where as per usual $d_1 = d_2 + \sigma\sqrt{T}$. But the last integral we can write as $\text{Prob}(Z > d_1)$ for a random variable $Z \sim N(0, 1)$, and by symmetry we get

$$A = S_0 \Phi(d_1),$$

which yields the “promised” result.

Theorem 6. *The unique arbitrage-free price of a European call option on a non-dividend paying stock in the Black-Scholes model is given by*

$$C_0 = S_0 \Phi(d_1) - Ke^{-rT} \Phi(d_2)$$

where

$$d_1 = \frac{\ln\left(\frac{S_0}{K}\right) + \left(r + \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}$$

and

$$d_2 = d_1 - \sigma\sqrt{T},$$

where Φ is the distribution function of a standard normal distribution.

As stated, the Black-Scholes formula says only what the call price is at time 0. But it is not hard to guess what happens if we want the price at some time $t \in [0; T]$: The same formula applies with S_0 substituted by S_t and T substituted by $T - t$. You may want to “try your hand” with conditional expectations and properties of Brownian motion by proving this.

7.3.1 Hedging the call

There is one last thing about the Black-Scholes model/formula you should know. Just as in the binomial model the call option can be hedged in the Black-Scholes model. This means that there exists a self-financing trading

delta-hedging

strategy involving the stock and the bond such that the value of the strategy at time T is exactly equal to the payoff of the call, $(S_T - K)^+$. (This is in fact the very reason we can talk about a *unique* arbitrage-free price for the call.) It is a general fact that if we have a contract whose price at time t can be written as

$$\pi(t) = F(t, S_t)$$

for some deterministic function F , then the contract is hedged by a continuously adjusted strategy consisting of (the two last equation is just notation; deliberately deceptive as all good notation)

$$\phi^1(t) = \left. \frac{\partial F}{\partial x}(t, x) \right|_{x=S_t} = \frac{\Delta\pi}{\Delta S} := \Delta(t)$$

units of the stock and $\phi^0(t) = \pi(t) - \phi^1(t)S_t$ \$ in the bank account. This is called delta hedging.

For the Black-Scholes model this applies to the call with

$$\begin{aligned} F^{BScall}(t, x) &= x\Phi\left(\frac{\ln\left(\frac{x}{K}\right) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right) \\ &\quad - Ke^{-r(T-t)}\Phi\left(\frac{\ln\left(\frac{x}{K}\right) + (r - \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right). \end{aligned}$$

The remarkable result (and what you must forever remember) is that the partial derivative (wrt. x) of this lengthy expression is simple:⁶

$$\frac{\partial F^{BScall}}{\partial x}(t, x) = \Phi\left(\frac{\ln\left(\frac{x}{K}\right) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right) = \Phi(d_1),$$

where the last part is standard and understandable but slightly sloppy notation. So to hedge the call option in a Black-Scholes economy you have to hold (at any time t) $\Phi(d_1)$ units of the stock. This quantity is called *the delta* (or: Δ) *hedge ratio* for the call option. The “lingo” comes about because of the intimate relation to partial derivatives; Δ is approximately the amount that the call price changes, when the stock price changes by 1. In this course we will use computer simulations to illustrate, justify, and hopefully to some degree understand the result.

⁶At one time or another you are bound to be asked to verify this, so you may as well do it right away. Note that if you just look at the B-S formula, forget that S_0 (or x) also appears inside the Φ 's, and differentiate, then you get the right result with a wrong proof.

Chapter 8

Stochastic Interest Rates

8.1 Introduction

After the brief encounter with continuous time modelling in Chapter 7 we now return to the discrete time, finite state space models of Chapter 5. They still have a great deal to offer.

One of the most widespread applications of arbitrage pricing in the multi-period finite state space model is in the area of term structure modelling. We saw in Chapter 3 how the term structure could be defined in several equivalent ways through the discount function, the yields of zero coupon bonds and by looking at forward rates. In this chapter we will think of the term structure as the yield of zero coupon bonds as a function of time to maturity. In Chapter 3 we considered the term structure at a fixed point in time. In this chapter our goal is to look at dynamic modelling of the evolution of the term structure. This topic could easily occupy a whole course in itself so here we focus merely on explaining a fundamental method of constructing arbitrage-free systems of bond prices. Once this method is understood the reader will be able to build models for the evolution of the term structure and price interest rate related contingent claims.

We also consider a few topics which are related to term structure modelling and which we can discuss rigorously with our arbitrage pricing technology. These topics are the difference between forwards and futures and the role of 'convexity effects' - or Jensen's inequality - can rule out various properties of term structure evolutions. We also look briefly at so-called swap contracts which are quite important in bond markets.

zero coupon bond,
ZCB
short rate process

8.2 Constructing an arbitrage free model

Our goal is to model prices of zero coupon bonds of different maturities and through time. Let $P(t, T_i)$, $0 \leq t \leq T_i \leq T$, denote the price at time t of a zero coupon bond with maturity T_i . To follow the notation which is most commonly used in the literature we will deviate slightly from the notation of Chapter 5. To be consistent with Chapter 5 we should write $P(t, T_i)$ for the price of the bond prior to maturity. i.e. when $t < T_i$ and then have a dividend payment $\delta(T_i) = 1$ at maturity and a price process satisfying $P(t, T_i) = 0$ for $t \geq T_i$. We will instead write the dividend into the price and let

$$P(t, t) = 1$$

for all t . (You should have gotten used to this deceptive notation in Chapters 6 and 7.)

We will consider models of bond prices which use the short rate process $\rho = (\rho_t)_{t=0, \dots, T-1}$ as the fundamental modelling variable. Recall that the money market account is a process with value 1 and dividend at date $t < T$ given by ρ_{t-1} and a dividend of $1 + \rho_T$ at time T . We will need our simple notation for returns obtained by holding money over several periods in the money market account:

Definition 40. *The return of the bank account from period t to u is*

$$R_{t,u} = (1 + \rho_t)(1 + \rho_{t+1}) \cdots (1 + \rho_{u-1}), \quad \text{for } t < u$$

Make sure you understand that $R_{t,t+1}$ is known at time t , whereas $R_{t,t+2}$ is not.

From the fundamental theorem of asset pricing (Theorem 4) we know that the system consisting of the money market account and zero coupon bonds will be arbitrage free if and only if

$$\left(\frac{P(t, T_i)}{R_{0,t}} \right)_{0 \leq t \leq T_i}$$

is a martingale for every T_i under some measure Q . Here, we use the fact that the zero coupon bonds only pay one dividend at maturity and we have denoted this dividend $P(T_i, T_i)$ for the bond maturing at date T_i . It is not easy, however, to specify a family of sensible and consistent bond prices. If T is large there are many maturities of zero coupon bonds to keep track of. They all should end up having price 1 at maturity, but that is about all we know. How do we ensure that the large system of prices admits no arbitrage opportunities?

What is often done is the following: We simply construct zero coupon short rate lattice bond prices as expected discounted values of their terminal price 1 under a measure Q which we specify in advance (as opposed to derived from bond prices). More precisely:

Proposition 18. *Given a short rate process $\rho = (\rho_t)_{t=0, \dots, T-1}$. Let*

$$\mathcal{F}_t = \sigma(\rho_0, \rho_1, \dots, \rho_T).$$

For a given Q define

$$P(t, T_i) = E_t^Q \left[\frac{1}{R_{t, T_i}} \right] \text{ for } 0 \leq t \leq T_i \leq T,$$

where $E_t^Q[\cdot]$ is short hand for $E^Q[\cdot | \mathcal{F}_t]$. Then the system consisting of the money market account and the bond price processes $(P(t, T_i))_{t=0, \dots, T}$ is arbitrage free.


Proof. The proof is an immediate consequence of the definition of prices, since

$$\frac{P(t, T_i)}{R_{0,t}} = \frac{1}{R_{0,t}} E_t^Q \left[\frac{1}{R_{t, T_i}} \right] = E_t^Q \left[\frac{1}{R_{0, T_i}} \right]$$

and this we know defines a martingale for each T_i by Lemma 2. ■

It is important to note that we take Q as given. Another way of putting this is that a P -specification of the short rate (however well it may fit the data) is not enough to determine Q , bond prices and the Q -dynamics of the short rate. If you only have a short rate process, the only traded asset is the bank account and you cannot replicate bonds with that. Later courses will explain this in more detail.

Example 20. Here is a simple illustration of the procedure in a model where the short rate follows a binomial process.



The short rate at time 0 is 0.10. At time 1 it becomes 0.11 with probability $\frac{1}{2}$ and 0.09 with probability $\frac{1}{2}$ (both probabilities under Q) Given that it is 0.09 at time 1, it becomes either 0.10 or 0.08 at time 2, both with probability $\frac{1}{2}$. The bond prices have been computed using Proposition 18. Note that a consequence of Proposition 18 is that (check it!)

$$P(t, T_i) = \frac{1}{1 + \rho_t} E_t^Q [P(t + 1, T_i)]$$

and therefore the way to use the proposition is to construct bond prices working backwards through the tree. For a certain maturity T_i we know $P(T_i, T_i) = 1$ regardless of the state. Now the price of this bond at time $T_i - 1$ can be computed as a function of ρ_{T_i-1} , and so forth. The term structure at time 0 is now computed as follows

$$r(0, 1) = \frac{1}{P(0, 1)} - 1 = 0.1$$

$$r(0, 2) = \left(\frac{1}{P(0, 2)} \right)^{\frac{1}{2}} - 1 = 0.09995$$

$$r(0, 3) = \left(\frac{1}{P(0, 3)} \right)^{\frac{1}{3}} - 1 = 0.0998$$

using definitions in Chapter 3. So the term structure in this example is decreasing in t - which is not what is normally seen in the market (but it does happen, for instance in Denmark in 1993 and in the U.S. in 2000). In fact, one calls the term structure "inverted" in this case. Note that when the Q -behavior of r has been specified we can determine not only the current term structure, we can find the term structure in any node of the tree. (Since the model only contains two non-trivial zero-coupon bonds at time 1, the term structure only has two points at time 1.)

So Example 20 shows how the term structure is calculated from a Q -tree of the short rate. But what we (or: practitioners) are really interested in is the reverse question: Given today's (observed) term structure, how do we construct a Q -tree of the short rate that is consistent with the term structure? (By consistent we mean that if we use the tree for ρ in Example 20-fashion we match the observed term structure at the first node.) Such a tree is needed for pricing more complicated contracts (options, for instance).

**short rate
lattice,
calibration**

First, it is easy to see that generally such an “inversion” is in no way unique; a wide variety of ρ -trees give the same term structure. But that is not bad; it means that we impose a convenient structure on the ρ -process and still fit observed term structures. Two such conveniences are that the development of ρ can be represented in a recombining tree (a lattice), or in other words that ρ is Markovian, and that the Q -probability $1/2$ is attached to all branches. (It may not be totally clear that we can do that, but it is easily seen from the next example/subsection.)

8.2.1 Constructing a Q -tree/lattice for the short rate that fits the initial term structure

Imagine a situation where two things have been thrust upon us.

1. The almighty (“God “or “The Market”) has determined today's term structure,

$$(P(0, 1), P(0, 2), \dots, P(0, T)).$$

2. Our not-so-almighty boss has difficulties understanding probability beyond the tossing of a fair coin and wants answers fast, so he(s secretary) has drawn the ρ -lattice in Figure 8.1.

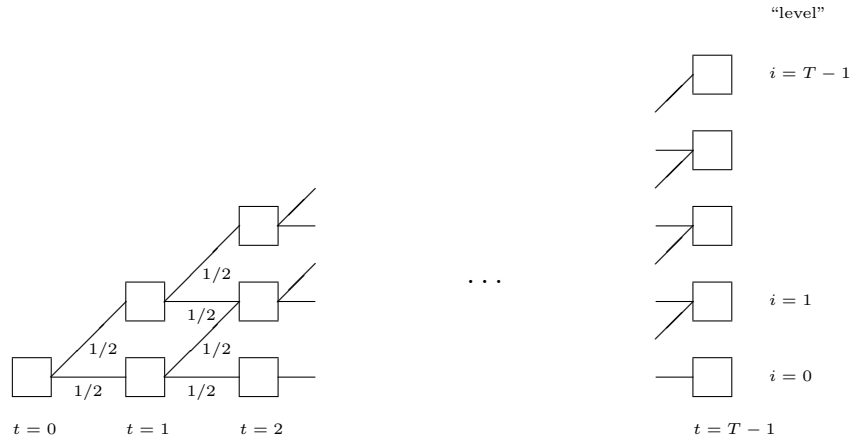
All we have to do is “fill in the blanks”. Optimistically we start, and in the box corresponding to $(t = 0, i = 0)$ we have no choice but to put

$$\rho_0(0) = \frac{1}{P(0, 1)} - 1.$$

To fill out boxes corresponding to $(t = 1, i = 0)$ and $(t = 1, i = 1)$ we have the equation

$$P(0, 2) = \frac{1}{1 + \rho_0(0)} \left(\frac{1}{2} \times \frac{1}{1 + \rho_1(0)} + \frac{1}{2} \times \frac{1}{1 + \rho_1(1)} \right), \quad (8.1)$$

which of course has many solutions. (Even many sensible ones.) So we can/have to put more structure on the problem. Two very popular ways of

Figure 8.1: The ρ -lattice we must complete.

doing this are these functional forms: ¹

$$\begin{aligned} \text{Ho/Lee-specification:} \quad & \rho_t(i) = a_{imp}(t) + b_{hist}i \\ \text{Black/Derman/Toy-specification:} \quad & \rho_t(i) = a_{imp}(t) \exp(b_{hist}i) \end{aligned}$$

For each t we fit by choosing an appropriate a_{imp} , while b_{hist} is considered a known constant. b_{hist} is called a volatility parameter and is closely related (as you should be able to see) to the conditional variance of the short rate (or its logarithm). This means that it is fairly easy to estimate from historical time series data of the short rate. With b_{hist} fixed, (8.1) can be solved hence determining what goes in the two “ $t = 1$ ”-boxes. We may have to solve the equation determining $a_{imp}(1)$ numerically, but monotonicity makes this an easy task (by bisection or Newton-Raphson, for instance).

And now can can do the same for $t = 2, \dots, T - 1$ and we can put our computer to work and go to lunch. Well, yes and no. Even though we take a long lunch there is a good chance that the computer is not finished when we get back. Why? Note that as it stands, every time we make a guess at $a_{imp}(t)$ (and since a numerical solution is involved we are likely to be making a number of these) we have to work our way backward through the lattice all the way down to 0. And this we have to do for each t . While not a computational catastrophe (a small calculation shows that the computation time grows as T^3), it does not seem totally efficient. We would like to go through the lattice only once (as it was the case when the initial term structure was determined

¹Of course there is a reason for the names attached. As so often before, this is for later courses to explain.

from a known ρ -lattice). Fortunately there is a way of doing this. We need the following lemma.

Lemma 3. *Consider the binomial ρ -lattice in Figure 8.1. Let $\psi(t, i)$ be the price at time 0 of a security that pays 1 at time t if state/level i occurs at that time. Then $\psi(0, 0) = 1$, $\psi(0, i) = 0$ for $i > 0$ and the following forward equation holds:*

$$\psi(t+1, i) = \begin{cases} \frac{\psi(t, i)}{2(1+\rho_t(i))} + \frac{\psi(t, i-1)}{2(1+\rho_t(i-1))} & 0 < i < t+1, \\ \frac{\psi(t, i-1)}{2(1+\rho_t(i-1))} & i = t+1, \\ \frac{\psi(t, i)}{2(1+\rho_t(i))} & i = 0. \end{cases}$$

Proof. We do the proof only for the “ $0 < i < t+1$ ”-case, the others are similar. Recall that we can think of \mathcal{F}_t -measurable random variables (of the type considered here) as vectors in \mathbb{R}^{t+1} . Since conditional expectation is linear, we can (for $s \leq t$) think of the \mathcal{F}_s -conditional expectation of an \mathcal{F}_t -measurable random variable as a linear mapping from \mathbb{R}^{t+1} to \mathbb{R}^{s+1} . In other words it can be represented by a $(s+1) \times (t+1)$ -matrix. In particular the time $t-1$ price of a contract with time t price X can be represented as

$$E_t^Q \left(\frac{X}{1 + \rho_{t-1}} \right) = \Pi_{t-1} X$$

Now note that in the binomial model there are only two places to go from a given point, so the Π_{t-1} -matrices have the form

$$\Pi_{t-1} = \underbrace{\left[\begin{array}{cccccc} \frac{1-q}{1+\rho_{t-1}(0)} & \frac{q}{1+\rho_{t-1}(0)} & & & & 0 \\ & \frac{1-q}{1+\rho_{t-1}(1)} & \frac{q}{1+\rho_{t-1}(1)} & & & \\ & & \ddots & \ddots & & \\ 0 & & & \frac{1-q}{1+\rho_{t-1}(t-1)} & \frac{q}{1+\rho_{t-1}(t-1)} & \end{array} \right]}_{t+1 \text{ columns}} \left. \vphantom{\left[\begin{array}{cccccc} \end{array} \right]} \right\} t \text{ rows}$$

Let $e_i(t)$ be the i 'th vector of the standard base in \mathbb{R}^t . The claim that pays 1 in state i at time $t+1$ can be represented in the lattice by $e_{i+1}(t+2)$ and by iterated expectations we have

$$\psi(t+1, i) = \Pi_0 \Pi_1 \cdots \Pi_{t-1} \Pi_t e_{i+1}(t+2).$$

But we know that multiplying a matrix by $e_i(t)$ from the right picks out the i 'th column. For $0 < i < t+1$ we may write the $i+1$ 'st column of Π_t as (look at $i=1$)

$$\frac{1-q}{1+\rho_t(i-1)} e_i(t+1) + \frac{q}{1+\rho_t(i)} e_{i+1}(t+1).$$

term structure
of interest
rates, parallel
shifts

Hence we get

$$\begin{aligned}\psi(t+1, i) &= \Pi_0 \Pi_1 \cdots \Pi_{t-1} \left(\frac{1-q}{1+\rho_t(i-1)} e_i(t+1) + \frac{q}{1+\rho_t(i)} e_{i+1}(t+1) \right) \\ &= \frac{1-q}{1+\rho_t(i-1)} \underbrace{\Pi_0 \Pi_1 \cdots \Pi_{t-1} e_i(t+1)}_{\psi(t, i-1)} \\ &\quad + \frac{q}{1+\rho_t(i)} \underbrace{\Pi_0 \Pi_1 \cdots \Pi_{t-1} e_{i+1}(t+1)}_{\psi(t, i)},\end{aligned}$$

and since $q = 1/2$, this ends the proof. ■

Since $P(0, t) = \sum_{i=0}^t \psi(t, i)$, we can use the following algorithm to fit the initial term structure.

1. Let $\psi(0, 0) = 1$ and put $t = 1$.
2. Let $\lambda_t(a_{imp}(t-1)) = \sum_{i=0}^t \psi(t, i)$ where $\psi(t, i)$ is calculated from the $\psi(t-1, \cdot)$'s using the specified $a_{imp}(t-1)$ -value in the forward equation from Lemma 3.
Solve $\lambda_t(a_{imp}(t-1)) = P(0, t)$ numerically for $a_{imp}(t-1)$.
3. Increase t by one. If $t \leq T$ then go to 2., otherwise stop.

An inspection reveals that the computation time of this procedure only grows as T^2 , so we have “gained an order”, which can be quite significant when T is large. And don't worry: There will be exercises to help you understand and implement this algorithm.

8.3 Flat shifts of flat term structures

Now let us demonstrate that in our term structure modelling framework it is impossible to have only parallel shifts of a flat term structure. In other words, in a model with no arbitrage we cannot have bond prices at time 0 given as

$$P(0, t) = \frac{1}{(1+r)^t}$$

for some $r \geq 0$, $t = 1, \dots, T$ and

$$P(1, t) = \frac{1}{(1+\tilde{r})^{t-1}}, \quad t = 2, \dots, T,$$

where \tilde{r} is a random variable (which takes on at least two different values with positive probability). To assign meaning to a "flat term structure" at time 1 we should have $T \geq 3$.

Now consider the zero-coupon bonds with maturity dates 2 and 3. If the term structure is flat at time 0 we have for some $r \geq 0$

$$P(0, 2) = \frac{1}{(1+r)^2} \text{ and } P(0, 3) = \frac{1}{(1+r)^3}$$

and if it remains flat at time 1, there exist a random variable \tilde{r} such that

$$P(1, 2) = \frac{1}{1+\tilde{r}} \text{ and } P(1, 3) = \frac{1}{(1+\tilde{r})^2}.$$

Furthermore, in an arbitrage-free model it will be the case that

$$\begin{aligned} P(0, 2) &= \frac{1}{1+r} E^Q [P(1, 2)] \\ &= \frac{1}{1+r} E^Q \left[\frac{1}{1+\tilde{r}} \right] \end{aligned}$$

and

$$\begin{aligned} P(0, 3) &= \frac{1}{1+r} E^Q [P(1, 3)] \\ &= \frac{1}{1+r} E^Q \left[\frac{1}{(1+\tilde{r})^2} \right] \end{aligned}$$

Combining these results, we have

$$\frac{1}{1+r} = E^Q \left[\frac{1}{(1+\tilde{r})} \right]$$

and

$$\frac{1}{(1+r)^2} = E^Q \left[\frac{1}{(1+\tilde{r})^2} \right]$$

which contradicts Jensen's inequality, for if

$$\frac{1}{1+r} = E^Q \left[\frac{1}{(1+\tilde{r})} \right]$$

then since $u \mapsto u^2$ is strictly convex and \tilde{r} not constant we must have

$$\frac{1}{(1+r)^2} < E^Q \left[\frac{1}{(1+\tilde{r})^2} \right].$$

forward
contract
futures contract

Note that the result does not say that it is impossible for the term structure to be flat. But it is inconsistent with no arbitrage to have a flat term structure and only have the possibility of moves to other flat term structures.

This explains what goes “wrong” in the example in Section 3.6.1. There the term structure was flat. We then created a position that had a value of 0 at that level of interest rates, but a strictly positive value with at flat term structure at any other level. But if interest rates are really stochastic then an arbitrage-free model cannot have only flat shifts of flat structure.

8.4 Forwards and futures

A forward and a futures contract are very similar contracts: The buyer (seller) of either type of contract is obligated to buy (sell) a certain asset at some specified date in the future for a price - the delivery price - agreed upon today. The forward/futures price of a certain asset is the delivery price which makes the forward/futures contract have zero value initially. It is very important to see that a forward/futures price is closer in spirit to the exercise price of an option than to the price of an option contract. Whereas an option always has positive value (and usually strictly positive) initially, both futures and forwards have zero value initially because the delivery price is used as a balancing tool.

The following example might clarify this: If a stock trades at \$100 today and we were to consider buying a futures contract on the stock with delivery in three months and if we had an idea that this stock would not move a lot over the next three months, then we would be happy to pay something for a contract which obligated us to buy the stock in three months for, say, \$50. Even though things could go wrong and the stock fall below \$50 in three months we consider that a much smaller risk of loss than the chance of gaining a lot from the contract. Similarly, we would not obligate ourselves to buying the stock in three months for, say, \$150 without receiving some money now. Somewhere in between \$50 and \$150 is a delivery price at which we would neither pay nor insist on receiving money to enter into the contract.

In a market with many potential buyers and sellers there is an equilibrium price at which supply meets demand: The number of contracts with that delivery price offered at zero initial cost equals the number of contracts demanded. This equilibrium price is the forward/futures price (depending on which contract we consider). In the following we will look at this definition in a more mathematical way and we will explain in what sense futures and forwards are different. Although they produce different cash flows (see below) that only results in a price difference when interest rates are stochas-

tic. Therefore, we will illustrate this difference with an example involving futures price forward price futures/forwards on bonds. We will ignore margin payments (i.e. payments that one or both sides of the contract have to make initially to guarantee future payments) in this presentation.

First, let us look at the key difference between forwards and futures by illustrating the cash flows involved in both types of contracts: Let F_t denote the forward price at time t for delivery of an underlying asset at time T and let Φ_t denote the futures price of the same asset for delivery at T , where $t \leq T$. Strictly speaking, we should write $F_{t,T}$ and $\Phi_{t,T}$ instead of F_t and Φ_t respectively, since it is important to keep track of both the date at which the contract is entered into and the delivery date. But we have chosen to consider the particular delivery date T and then keep track of how the futures and forward prices change as a function of t . The cash flows produced by the two types of contracts, if bought at time t , are as follows:

	t	t+1	t+2	...	T-1	T
Forward	0	0	0	...	0	$S_T - F_t$
Futures	0	$\Phi_{t+1} - \Phi_t$	$\Phi_{t+2} - \Phi_{t+1}$...	$\Phi_{T-1} - \Phi_{T-2}$	$S_T - \Phi_{T-1}$

where S_T is the price of the underlying asset at time T . The forward cash flow is self-explanatory. The futures cash flow can be explained as follows: If you buy a futures contract at date t you agree to buy the underlying asset at time T for Φ_t . At time $t + 1$ markets may have changed and the price at which futures trade changed to Φ_{t+1} . What happens is now a resettlement of the futures contract. If Φ_{t+1} is bigger than Φ_t you (the buyer of the futures at time t) receive the amount $\Phi_{t+1} - \Phi_t$ from the seller at time $t+1$ whereas you pay the difference between Φ_{t+1} and Φ_t to the seller if $\Phi_{t+1} < \Phi_t$. The story continues as shown in the figure.

We have already seen that if the underlying asset trades at time t and a zero coupon bond with maturity T also trades then the forward price is given as

$$F_t = \frac{S_t}{P(t, T)}$$

i.e.

$$F_t = S_t (1 + r(t, T))^{T-t} \quad (8.2)$$

where $r(t, T)$ is the internal rate of return on the zero coupon bond.

To see what Φ_t is requires a little more work: First of all to avoid arbitrage we must have $\Phi_T = S_T$. Now consider Φ_{T-1} . In an arbitrage free system there exists an equivalent martingale measure Q . The futures price Φ_{T-1} is such that the cash flow promised by the contract (bought at $T - 1$) has value 0.

We must therefore have

$$0 = E_{T-1}^Q \left[\frac{S_T - \Phi_{T-1}}{R_{T-1,T}} \right]$$

but since $R_{T-1,T}$ is \mathcal{F}_{T-1} -measurable this implies

$$0 = \frac{1}{R_{T-1,T}} E_{T-1}^Q [S_T - \Phi_{T-1}]$$

i.e.

$$\Phi_{T-1} = E_{T-1}^Q [S_T] \quad (8.3)$$

Since Q is a martingale measure recall that

$$\frac{S_{T-1}}{R_{0,T-1}} = E_{T-1}^Q \left[\frac{S_T}{R_{0,T}} \right]$$

i.e.

$$S_{T-1} = \frac{1}{1 + \rho_{T-1}} E_{T-1}^Q [S_T]$$

hence we can write (8.3) as

$$\Phi_{T-1} = (1 + \rho_{T-1}) S_{T-1}$$

and that is the same as (8.2) since the yield on a one period zero coupon bond is precisely the short rate. So we note that with one time period remaining we have $\Phi_{T-1} = F_{T-1}$. But that also follows trivially since with one period remaining the difference in cash flows between forwards and futures does not have time to materialize.

Now consider Φ_{T-2} . By definition Φ_{T-2} should be set such that the cash flow of the futures contract signed at $T - 2$ has zero value:

$$0 = E_{T-2}^Q \left[\frac{\Phi_{T-1} - \Phi_{T-2}}{R_{T-2,T-1}} + \frac{S_T - \Phi_{T-1}}{R_{T-2,T}} \right] \quad (8.4)$$

Now note that using the rule of iterated expectations and the expression for Φ_{T-1} we find

$$\begin{aligned} & E_{T-2}^Q \left[\frac{S_T - \Phi_{T-1}}{R_{T-2,T}} \right] \\ &= \frac{1}{R_{T-2,T-1}} E_{T-2}^Q \left[E_{T-1}^Q \left[\frac{S_T - \Phi_{T-1}}{R_{T-1,T}} \right] \right] \\ &= 0 \end{aligned}$$

so (8.4) holds precisely when

$$\begin{aligned} 0 &= E_{T-2}^Q \left[\frac{\Phi_{T-1} - \Phi_{T-2}}{R_{T-2,T-1}} \right] \\ &= \frac{1}{R_{T-2,T-1}} E_{T-2}^Q [\Phi_{T-1} - \Phi_{T-2}] \end{aligned}$$

i.e. we have

$$\Phi_{T-2} = E_{T-2}^Q [\Phi_{T-1}] = E_{T-2}^Q [S_T].$$

This argument can be continued backwards and we arrive at the expression

$$\Phi_t = E_t^Q [S_T] \quad (8.5)$$

Note that (8.5) is not in general equal to (8.2):

Under Q , we have $S_t = E_t^Q \left[\frac{S_T}{R_{t,T}} \right]$ so if $\frac{1}{R_{t,T}}$ and S_T are uncorrelated under Q we may write

$$S_t = E_t^Q \left[\frac{1}{R_{t,T}} \right] E_t^Q [S_T] = P(t, T) \Phi_t$$

which would imply that

$$\Phi_t = \frac{S_t}{P(t, T)} = F_t$$

Hence, if $\frac{1}{R_{t,T}}$ and S_T are uncorrelated under Q , the forward price F_t and the futures price Φ_t are equal. A special case of this is when interest rates are deterministic, i.e. all future spot rates and hence $R_{t,T}$ are known at time t .

Note that in general,

$$\begin{aligned} \Phi_t - F_t &= \frac{1}{P(t, T)} \left(P(t, T) E_t^Q [S_T] - S_t \right) \\ &= \frac{1}{P(t, T)} \left(E_t^Q \left[\frac{1}{R_{t,T}} \right] E_t^Q [S_T] - S_t \right) \\ &= \frac{1}{P(t, T)} \left(E_t^Q \left(\frac{S_T}{R_{t,T}} \right) - Cov_t^Q \left(\frac{1}{R_{t,T}}, S_T \right) - S_t \right) \\ &= \frac{-1}{P(t, T)} \left(Cov_t^Q \left(\frac{1}{R_{t,T}}, S_T \right) \right). \end{aligned}$$

Note that margin payments go to the holder of a futures contract when spot prices rise, i.e. in states where S_T is high. If $\frac{1}{R_{t,T}}$ is negatively correlated with S_T , then interest rates tend to be high when the spot price is high and hence the holder of a futures contract will receive cash when interest rates are high. Hence a futures contract is more valuable in that case and the futures price should therefore be set higher to keep the contract value at 0.

swap contract

8.5 Swap contracts

A swap contract is an agreement to exchange one stream of payments for another. A wide variety of swaps exists in financial markets; they are often tailor-made to the specific need of a company/an investor and can be highly complex. However, we consider only the valuation of the simplest² interest rate swap where fixed interest payments are exchanged for floating rate interest payments.

This swap you may see referred to as anything from “basis” to “forward starting ???monthly payer swap settled in arrears”. Fortunately the payments are easier to describe. For a set of equidistant dates $(T_i)_{i=0}^n$, say δ apart, it is a contract with cash flow (per unit of notational principal)

$$\left(\underbrace{\frac{1}{P(T_{i-1}, T_i)} - 1}_{\text{floating leg}} - \underbrace{\delta\kappa}_{\text{fixed leg}} \right) \text{ at date } T_i \text{ for } i = 1, \dots, n,$$

where κ is a constant (an interest rate with δ -compounding quoted on yearly basis.) You should convince yourself why the so-called floating leg does in fact correspond to receiving floating interest rate payments. The term $(1/P(T_{i-1}, T_i) - 1)/\delta$ is often called the $(12*\delta)$ -month LIBOR (which an acronym for London Interbank Offer Rate, and does not really mean anything nowadays, it is just easy to pronounce). Note that the payment made at T_i is known at T_{i-1} .

It is clear that since the payments in the fixed leg are deterministic, they have a value of

$$\delta\kappa \sum_{i=1}^n P(t, T_i).$$

The payments in the floating leg are not deterministic. But despite this, we can find their value without a stochastic model for bond prices/interest rates. Consider the following simple portfolio strategy:

Time	Action	Net cash flow
t	Sell 1 T_i -ZCB Buy 1 T_{i-1} -ZCB	$P(t, T_i) - P(t, T_{i-1})$
T_{i-1}	Use principal received from T_{i-1} -ZCB to buy $1/P(T_{i-1}, T_i)$ T_i -ZCBs	0
T_i	Close position	$1/P(T_{i-1}, T_i) - 1$

²Simple *objects* are often referred to as plain vanilla *objects*. But what is seen as simple depends very much on who is looking.

This means that the T_i -payment in the floating leg has a value of $P(t, T_{i-1}) - \text{swap rate} P(t, T_i)$, so when summing over i see that the value of the floating leg is

$$P(t, T_0) - P(t, T_n).$$

In the case where $t = T_0$ this is easy to remember/interpret. A bullet-like bond that has a principal of 1 pays a coupon that is the short rate must have a price of 1 (lingo: “it is trading at par”). The only difference between this contract and the floating leg is the payment of the principal at time T_n ; the time t value of this is $P(t, T_n)$ hence the value of the floating leg is $1 - P(t, T_n)$.

All in all the swap has a value of

$$V = P(t, T_0) - P(t, T_n) - \delta \kappa \sum_{i=1}^n P(t, T_i).$$

But there is a further twist; these basis swaps are only traded with one κ (for each length; each n), namely the one that makes the value 0. This rate is called the swap rate (at a given date for a given maturity)

$$\kappa_n(t) = \frac{P(t, T_0) - P(t, T_n)}{\delta \sum_{i=1}^n P(t, T_i)}. \quad (8.6)$$

In practice (8.6) is often used “backwards”, meaning that swap rates for swaps of different lengths (called the “swap curve”) are used to infer discount factors/the term structure. Note that this is easy to do recursively if we can “get started”, which is clearly the case if $t = T_0$.³

The main point is that the basis swap can be priced without using a full dynamic model, we only need today’s term structure. But it takes only minor changes in the contract specification for this conclusion to break down. For instance different dynamic models with same current term structure give different swap values if the i th payment in the basis swap is transferred to date T_{i-1} (where it is first known; this is called settlement in advance) or if we swap every 3 months against the 6-month LIBOR.

The need for a swap-market can also be motivated by the following example showing swaps can offer comparative advantages. In its swap-formulation it is very inspired by Hull’s book, but you should recognize the idea from introductory economics courses (or David Ricardo’s work of 1817, whichever came first). Consider two firms, A and B, each of which wants to borrow

³There should be a “don’t try this at work” disclaimer here. In the market different day count conventions are often used on the two swap legs, so things may not be quite what they seem.

\$10M for 5 years. Firm A prefers to pay a floating rate, say one that is adjusted every year. It could be that the cash-flows generated by the investment (that it presumably needs the \$10M for) depend (positively) on the interest rate market conditions. So from their point of view a floating rate loan removes risk. Firm B prefers to borrow at a fixed rate. In this way it knows in advance exactly how much it has to pay over the 5 years, which it is quite conceivable that someone would want. The firms contact their banks and receive the following loan offers: (Lingo: “bp” means basispoints (pronounced “beeps” if you’re really cool) and is one hundredth of a percentage point, i.e. “100bp = 1%”)

Firm	Fixed	Floating
A	5Y-ZCB-rate + 50bp	1Y-ZCB-rate + 30bp
B	5Y-ZCB-rate + 170bp	1Y-ZCB-rate + 100bp

So B gets a systematically “worse deal” than A, which could be because of lower credit quality than A. But “less worse” for a floating rate loan, where they only have to pay 70bp more than A compared to 120bp for a fixed rate loan. So A could take the floating rate offer and B the fixed rate offer, and everybody is mildly happy. But consider the following arrangement: A takes the fixed rate offer from the bank and B the floating rate. A then offers to lend B the 10M as a fixed rate loan “at the 5Y-ZCB-rate + 45bp”, whereas B offers to lend A its 10M floating rate loan “at the 1Y-ZCB-rate” (and would maybe add “flat” to indicate that there *is* no spread). In other words A and B are exchanging, or swapping, their bank loans. The result:

A: Pays (5Y-ZCB-rate + 50bp) (to bank), Pays 1Y-ZCB-rate (to B) and receives (5Y-ZCB-rate + 45bp) (from B). In net-terms: Pays 1Y-ZCB-rate+5bp

B: Pays (1Y-ZCB-rate + 100bp) (to bank), Pays (5Y-ZCB-rate + 45bp) (to A) and receives (1Y-ZCB-rate) (from A). In net-terms: Pays 5Y-ZCB-rate+145bp

So this swap-arrangement has put both A and B in a better position (by 25bp) than they would have been had they only used the bank.

But when used in the finance/interest rate context, there is somewhat of a snag in this story. We argued that the loans offered reflected differences in credit quality. If that is so, then it must mean that default (“going broke”) is a possibility that cannot be ignored. It is this risk that the bank is “charging extra” for. With this point of view the reason why the firms get better deals after swapping is that each chooses to take on the credit risk from the other party. If firm B defaults, firm A can forget about (at least part of) what’s in the “receives from B”-column, but will (certainly with this construction)

only be able to get out of its obligations to B to a much lesser extent. So the firms are getting lower rates by taking on default risk, which a risk of the type “a large loss with a small probability”. One can quite sensibly ask if that is the kind of risks that individual firms want to take.

One could try to remedy the problem by saying that we set up a financial institution through which the swapping takes place. This institution should ensure payments to the non-defaulting party (hence taking “credit risk” \times 2), in return for a share of the possible “lower rate”-gain from the swap, and hope for some “law of large numbers”-diversification effect. But that story is questionable; isn’t that what the bank is doing in the first place?

So the morale is two-fold: *i*) If something seems to be too good to be true it usually is. Also in credit risk models. *ii*) The only way to see if the spreads offered to firms A and B are set such that there is no gain without extra risk, i.e. consistent with no arbitrage, is to set up a real dynamic stochastic model of the defaults (something that subsequent courses will do), just as stochastic term structure models help us realize that non-flat yield curves do not imply arbitrage.

**term structure
of interest
rates, parallel
shifts**

8.6 Flat shifts of flat term structures

Now let us demonstrate that in our term structure modelling framework it is impossible to have only parallel shifts of a flat term structure. In other words, in a model with no arbitrage we cannot have bond prices at time 0 given as

$$P(0, t) = \frac{1}{(1 + r)^t}$$

for some $r \geq 0$, $t = 1, \dots, T$ and

$$P(1, t) = \frac{1}{(1 + \tilde{r})^{t-1}}, t = 2, \dots, T,$$

where \tilde{r} is a random variable (which takes on at least two different values with positive probability). To assign meaning to a “flat term structure” at time 1 we should have $T \geq 3$.

Now consider the zero-coupon bonds with maturity dates 2 and 3. If the term structure is flat at time 0 we have for some $r \geq 0$

$$P(0, 2) = \frac{1}{(1 + r)^2} \text{ and } P(0, 3) = \frac{1}{(1 + r)^3}$$

and if it remains flat at time 1, there exist a random variable \tilde{r} such that

$$P(1, 2) = \frac{1}{1 + \tilde{r}} \text{ and } P(1, 3) = \frac{1}{(1 + \tilde{r})^2}.$$

Furthermore, in an arbitrage-free model it will be the case that

$$\begin{aligned} P(0, 2) &= \frac{1}{1+r} E^Q [P(1, 2)] \\ &= \frac{1}{1+r} E^Q \left[\frac{1}{1+\tilde{r}} \right] \end{aligned}$$

and

$$\begin{aligned} P(0, 3) &= \frac{1}{1+r} E^Q [P(1, 3)] \\ &= \frac{1}{1+r} E^Q \left[\frac{1}{(1+\tilde{r})^2} \right] \end{aligned}$$

Combining these results, we have

$$\frac{1}{1+r} = E^Q \left[\frac{1}{(1+\tilde{r})} \right]$$

and

$$\frac{1}{(1+r)^2} = E^Q \left[\frac{1}{(1+\tilde{r})^2} \right]$$

which contradicts Jensen's inequality, for if

$$\frac{1}{1+r} = E^Q \left[\frac{1}{(1+\tilde{r})} \right]$$

then since $u \mapsto u^2$ is strictly convex and \tilde{r} not constant we must have

$$\frac{1}{(1+r)^2} < E^Q \left[\frac{1}{(1+\tilde{r})^2} \right].$$

Note that the result does not say that it is impossible for the term structure to be flat. But it is inconsistent with no arbitrage to have a flat term structure and only have the possibility of moves to other flat term structures.

This explains what goes "wrong" in the example in Section 3.6.1. There the term structure was flat. We then created a position that had a value of 0 at that level of interest rates, but a strictly positive value with a flat term structure at any other level. But if interest rates are really stochastic then an arbitrage-free model cannot have only flat shifts of flat structure.

Chapter 9

Portfolio Theory

Matrix Algebra

First we need a few things about matrices. (A very useful reference for mathematical results in the large class imprecisely defined as “well-known” is Berck & Sydsæter (1992), “Economists’ Mathematical Manual”, Springer.)

- When $\mathbf{x} \in \mathbb{R}^n$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ then

$$\frac{\partial}{\partial \mathbf{x}} (\mathbf{x}^\top \mathbf{V} \mathbf{x}) = (\mathbf{V} + \mathbf{V}^\top) \mathbf{x}$$

- A matrix $\mathbf{V} \in \mathbb{R}^{n \times n}$ is said to be *positive definite* if $\mathbf{z}^\top \mathbf{V} \mathbf{z} > 0$ for all $\mathbf{z} \neq \mathbf{0}$. If \mathbf{V} is positive definite then \mathbf{V}^{-1} exists and is also positive definite.
- Multiplying (appropriately) partitioned matrices is just like multiplying 2×2 -matrices.
- Covariance is bilinear. Or more specifically: When X is an n -dimensional random variable with covariance matrix Σ then

$$\text{Cov}(\mathbf{A}X + \mathbf{B}, \mathbf{C}X + \mathbf{D}) = \mathbf{A}\Sigma\mathbf{C}^\top,$$

where \mathbf{A} , \mathbf{B} , \mathbf{C} , and \mathbf{D} are deterministic matrices such that the multiplications involved are well-defined.

Basic Definitions and Justification of Mean/Variance Analysis

We will consider an agent who wants to invest in the financial markets. We look at a simple model with only two time-points, 0 and 1. The agent has an initial wealth of W_0 to invest. We are not interested in how the agent

rate of return
relative portfolio
weights

determined this amount, it's just there. There are n financial assets to choose from and these have prices

$$S_{i,t} \text{ for } i = 1, \dots, n \text{ and } t = 0, 1,$$

where $S_{i,1}$ is stochastic and not known until time 1. The rate of return on asset i is defined as

$$r_i = \frac{S_{i,1} - S_{i,0}}{S_{i,0}},$$

and $r = (r_1, \dots, r_n)^\top$ is the vector of rates of return. Note that r is stochastic.

At time 0 the agent chooses a portfolio, that is he buys a_i units of asset i and since all in all W_0 is invested we have

$$W_0 = \sum_{i=1}^n a_i S_{i,0}.$$

(If $a_i < 0$ the agent is selling some of asset i ; in most of our analysis short-selling will be allowed.)

Rather than working with the absolute number of assets held, it is more convenient to work with relative portfolio weights. This means that for the i th asset we measure the value of the investment in that asset relative to total investment and call this w_i , i.e.

$$w_i = \frac{a_i S_{i,0}}{\sum_{i=1}^n a_i S_{i,0}} = \frac{a_i S_{i,0}}{W_0}.$$

We put $\mathbf{w} = (w_1, \dots, w_n)^\top$, and have that $\mathbf{w}^\top \mathbf{1} = 1$. In fact, *any* vector satisfying this condition identifies an investment strategy. Hence in the following a portfolio is a vector whose coordinate sum to 1. Note that in this one period model a portfolio \mathbf{w} is not a stochastic variable (in the sense of being unknown at time 0).

The terminal wealth is

$$\begin{aligned} W_1 &= \sum_{i=1}^n a_i S_{i,1} = \sum_{i=1}^n a_i (S_{i,1} - S_{i,0}) + \sum_{i=1}^n a_i S_{i,0} \\ &= W_0 \left(1 + \sum_{i=1}^n \frac{S_{i,0} a_i}{W_0} \frac{S_{i,1} - S_{i,0}}{S_{i,0}} \right) \\ &= W_0 (1 + \mathbf{w}^\top r), \end{aligned} \tag{9.1}$$

so if we know the relative portfolio weights and the realized rates of return, we know terminal wealth. We also see that

$$E(W_1) = W_0 (1 + \mathbf{w}^\top E(r))$$

and

utility function

$$\text{Var}(W_1) = W_0^2 \text{Cov}(\mathbf{w}^\top r, \mathbf{w}^\top r) = W_0^2 \mathbf{w}^\top \underbrace{\text{Cov}(r)}_{n \times n} \mathbf{w}.$$

In this chapter we will look at how agents should choose \mathbf{w} . We will focus on how to choose \mathbf{w} such that for a given expected rate of return, the variance on the rate of return is minimized. This is called mean-variance analysis. Intuitively, it sounds reasonable enough, but can it be justified?

An agent has a utility function, u , and let us for simplicity say that he derives utility from directly from terminal wealth. (So in fact we are saying that we can eat money.) We can expand u in a Taylor series around the expected terminal wealth,

$$\begin{aligned} u(W_1) &= u(E(W_1)) + u'(E(W_1))(W_1 - E(W_1)) \\ &\quad + \frac{1}{2}u''(E(W_1))(W_1 - E(W_1))^2 + R_3, \end{aligned}$$

where the remainder term R_3 is

$$R_3 = \sum_{i=3}^{\infty} \frac{1}{i!} u^{(i)}(E(W_1))(W_1 - E(W_1))^i,$$

“and hopefully small”. With appropriate (weak) regularity condition this means that expected terminal wealth can be written as

$$E(u(W_1)) = u(E(W_1)) + \frac{1}{2}u''(E(W_1))\text{Var}(W_1) + E(R_3),$$

where the remainder term involves higher order central moments. As usual we consider agents with increasing, concave (i.e. $u'' < 0$) utility functions who maximize expected wealth. This then shows that to a second order approximation there is a preference for expected wealth (and thus, by (9.1), to expected rate of return), and an aversion towards variance of wealth (and thus to variance of rates of return).

But we also see that mean/variance analysis cannot be a completely general model of portfolio choice. A sensible question to ask is: What restrictions can we impose (on u and/or on r) to ensure that mean-variance analysis is fully consistent with maximization of expected utility?

An obvious way to do this is to assume that utility is quadratic. Then the remainder term is identically 0. But quadratic utility does not go too well with the assumption that utility is increasing and concave. If u is concave (which it has to be for mean-variance analysis to hold ; otherwise our interest

mean/variance
analysis
Markowitz
analysis

would be in maximizing variance) there will be a point of satiation beyond which utility decreases. Despite this, quadratic utility is often used with a “happy-go-lucky” assumption that when maximizing, we do not end up in an area where it is decreasing.

We can also justify mean-variance analysis by putting distributional restrictions on rates of return. If rates of return on individual assets are normally distributed then the rate of return on a portfolio is also normal, and the higher order moments in the remainder can be expressed in terms of the variance. In general we are still not sure of the signs and magnitudes of the higher order derivatives of u , but for large classes of reasonable utility functions, mean-variance analysis can be formally justified.

9.1 Mathematics of Minimum Variance Portfolios

9.1.1 The case with no riskfree asset

First we consider a market with no riskfree asset and n risky assets. Later we will include a riskfree asset, and it will become apparent that we have done things in the right order.

The risky assets have a vector of rates of return of r , and we assume that

$$E(r) = \mu, \quad (9.2)$$

$$\text{Cov}(r) = \Sigma, \quad (9.3)$$

where Σ is positive definite (hence invertible) and not all coordinates of μ are equal. As a covariance matrix Σ is always positive semidefinite, the definiteness means that there does not exist an asset whose rate of return can be written as an affine function of the other $n - 1$ assets' rates of return. Note that the existence of a riskfree asset would violate this.

Consider the following problem:

$$\min_{\mathbf{w}} \frac{1}{2} \underbrace{\mathbf{w}^\top \Sigma \mathbf{w}}_{\sigma_P^2} \quad \text{subject to} \quad \mathbf{w}^\top \mu = \mu_P \\ \mathbf{w}^\top \mathbf{1} = 1$$

Analysis of such a problem is called mean/variance analysis, or Markowitz analysis after Harry Markowitz who studied the problem in the 40'ies and 50'ies. (He won the Nobel prize in 1990 together with William Sharpe and Merton Miller both of whom we'll meet later.)

Our assumptions on μ and Σ ensure that a unique finite solution exists for any value of μ_P . The problem can be interpreted as choosing portfolio weights (the second constraint ensures that \mathbf{w} is a vector of portfolio weights) such that the variance portfolio's rate return ($\mathbf{w}^\top \Sigma \mathbf{w}$; the "1/2" is just there for convenience) is minimized given that we want a specific expected rate of return (μ_P ; "P is for portfolio").

To solve the problem we set up the Lagrange-function with multipliers

$$\mathcal{L}(\mathbf{w}, \lambda_1, \lambda_2) = \frac{1}{2} \mathbf{w}^\top \Sigma \mathbf{w} - \lambda_1 (\mathbf{w}^\top \mu - \mu_P) - \lambda_2 (\mathbf{w}^\top \mathbf{1} - 1).$$

The first-order conditions for optimality are

$$\frac{\partial \mathcal{L}}{\partial \mathbf{w}} = \Sigma \mathbf{w} - \lambda_1 \mu - \lambda_2 \mathbf{1} = 0, \tag{9.4}$$

$$\mathbf{w}^\top \mu - \mu_P = 0, \tag{9.5}$$

$$\mathbf{w}^\top \mathbf{1} - 1 = 0. \tag{9.6}$$

Usually we might say "and these are linear equations that can easily be solved", but working on them algebraically leads to a deeper understanding and intuition about the model. Invertibility of Σ gives that we can write (9.4) as (check for yourself)

$$\mathbf{w} = \Sigma^{-1} [\mu \ \mathbf{1}] \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix}, \tag{9.7}$$

and (9.5)-(9.6) as

$$[\mu \ \mathbf{1}]^\top \mathbf{w} = \begin{bmatrix} \mu_P \\ 1 \end{bmatrix}. \tag{9.8}$$

Multiplying both sides of (9.7) by $[\mu \ \mathbf{1}]^\top$ and using (9.8) gives

$$\begin{bmatrix} \mu_P \\ 1 \end{bmatrix} = [\mu \ \mathbf{1}]^\top \mathbf{w} = \underbrace{[\mu \ \mathbf{1}]^\top \Sigma^{-1} [\mu \ \mathbf{1}]}_{=: \mathbf{A}} \begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix}. \tag{9.9}$$

Using the multiplication rules for partitioned matrices we see that

$$\mathbf{A} = \begin{bmatrix} \mu^\top \Sigma^{-1} \mu & \mu^\top \Sigma^{-1} \mathbf{1} \\ \mu^\top \Sigma^{-1} \mathbf{1} & \mathbf{1}^\top \Sigma^{-1} \mathbf{1} \end{bmatrix} =: \begin{bmatrix} a & b \\ b & c \end{bmatrix}$$

We now show that \mathbf{A} is positive definite, in particular it is invertible. To this end let $\mathbf{z}^\top = (z_1, z_2) \neq \mathbf{0}$ be an arbitrary non-zero vector in \mathbb{R}^2 . Then

$$\mathbf{y} = [\mu \ \mathbf{1}] \begin{bmatrix} z_1 \\ z_2 \end{bmatrix} = z_1 \mu + z_2 \mathbf{1} \neq \mathbf{0},$$

minimum variance
portfolio
minimum variance
portfolio

because the coordinates of μ are not all equal. From the definition of \mathbf{A} we get

$$\forall \mathbf{z} \neq \mathbf{0} : \mathbf{z}^\top \mathbf{A} \mathbf{z} = \mathbf{y}^\top \Sigma^{-1} \mathbf{y} > 0,$$

because Σ^{-1} is positive definite (because Σ is). In other words, \mathbf{A} is positive definite. Hence we can solve (9.9) for the λ 's,

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \end{bmatrix} = \mathbf{A}^{-1} \begin{bmatrix} \mu_P \\ 1 \end{bmatrix},$$

and insert this into (9.7) in order to determine the optimal portfolio weights

$$\widehat{\mathbf{w}} = \Sigma^{-1} [\mu \ \mathbf{1}] \mathbf{A}^{-1} \begin{bmatrix} \mu_P \\ 1 \end{bmatrix}. \quad (9.10)$$

The portfolio $\widehat{\mathbf{w}}$ is called the minimum variance portfolio for a given mean μ_P . (We usually can't be bothered to say the correct full phrase: "minimum variance on rate of return for a given mean rate on return μ_P ".) The minimal portfolio return variance is

$$\begin{aligned} \widehat{\sigma}_P^2 &= \widehat{\mathbf{w}}^\top \Sigma \widehat{\mathbf{w}} \\ &= [\mu_P \ 1] \mathbf{A}^{-1} [\mu \ \mathbf{1}]^\top \Sigma^{-1} \Sigma \Sigma^{-1} [\mu \ \mathbf{1}] \mathbf{A}^{-1} [\mu_P \ 1]^\top \\ &= [\mu_P \ 1] \mathbf{A}^{-1} \underbrace{([\mu \ \mathbf{1}]^\top \Sigma^{-1} [\mu \ \mathbf{1}])}_{=\mathbf{A} \text{ by def.}} \mathbf{A}^{-1} [\mu_P \ 1]^\top \\ &= [\mu_P \ 1] \mathbf{A}^{-1} \begin{bmatrix} \mu_P \\ 1 \end{bmatrix}, \end{aligned}$$

where symmetry (of Σ and \mathbf{A} and their inverses) was used to obtain the second line. But since

$$\mathbf{A}^{-1} = \frac{1}{ac - b^2} \begin{bmatrix} c & -b \\ -b & a \end{bmatrix},$$

we have

$$\widehat{\sigma}_P^2 = \frac{a - 2b\mu_P + c\mu_P^2}{ac - b^2}. \quad (9.11)$$

In (9.11) the relation between the variance of the minimum variance portfolio for a given r_p , $\widehat{\sigma}_P^2$, is expressed as a parabola and is called the *minimum variance portfolio frontier* or *locus*.

Note that we have not just solved one "minimize variance" problem, but a whole bunch of them, namely one for each conceivable expected rate of return.

In mean-standard deviation-space the relation is expressed as a hyperbola. Figure 9.1 illustrates what things look like in mean-variance-space. (When using graphical arguments you should be quite careful to use “the right space”; for instance lines that are straight in one space, are not straight in the other.) The upper half of the curve in Figure 9.1 (the solid line) identifies the set of portfolios that have the highest mean return for a given variance; these are called mean-variance *efficient portfolios*.

efficient portfolio
two-fund
separation

Figure 9.1 also shows the *global minimum variance portfolio*, the portfolio with the smallest possible variance for any given mean return. Its mean, μ_G , is found by minimizing (9.11) with respect to μ_P , and is $\mu_{gmv} = \frac{b}{c}$. By substituting this in the general $\hat{\sigma}^2$ -expression we obtain

$$\hat{\sigma}_{gmv}^2 = \frac{a - 2b\mu_{gmv} + c\mu_{gmv}^2}{ac - b^2} = \frac{a - 2b(b/c) + c(b/c)^2}{ac - b^2} = \frac{1}{c},$$

while the general formula for portfolio weights gives us

$$\hat{\mathbf{w}}_{gmv} = \frac{1}{c}\boldsymbol{\Sigma}^{-1}\mathbf{1}.$$

Example 21. Consider the case with 3 assets (referred to as A , B , and C) and

$$\boldsymbol{\mu} = \begin{bmatrix} 0.1 \\ 0.12 \\ 0.15 \end{bmatrix}, \quad \boldsymbol{\Sigma} = \begin{bmatrix} 0.25 & 0.10 & -0.10 \\ 0.10 & 0.36 & -0.30 \\ -0.10 & -0.30 & 0.49 \end{bmatrix}.$$

The all-important \mathbf{A} -matrix is then

$$\mathbf{A} = \begin{bmatrix} 0.33236 & 2.56596 \\ 2.565960 & 20.04712 \end{bmatrix},$$

which means that the locus of mean-variance portfolios is given by

$$\hat{\sigma}_P^2 = 4.22918 - 65.3031\mu_P + 255.097\mu_P^2.$$

The locus is illustrated in Figure 9.2 in both in (variance, expected return)-space and (standard deviation, expected return)-space.

An important property of the set of minimum variance portfolios is is so-called two-fund separation. This means that the minimum variance portfolio frontier can be generated by any two distinct minimum variance portfolios.

Proposition 19. Let \mathbf{x}_a and \mathbf{x}_b be two minimum variance portfolios with mean returns μ_a and μ_b , $\mu_a \neq \mu_b$. Then every minimum variance portfolio, \mathbf{x}_c is a linear combination of \mathbf{x}_a and \mathbf{x}_b . Conversely, every portfolio that is




Figure 9.1: The minimum variance portfolio frontier.

a linear combination of \mathbf{x}_a and \mathbf{x}_b (i.e. can be written as $\alpha\mathbf{x}_a + (1 - \alpha)\mathbf{x}_b$) is a minimum variance portfolio. In particular, if \mathbf{x}_a and \mathbf{x}_b are efficient portfolios, then $\alpha\mathbf{x}_a + (1 - \alpha)\mathbf{x}_b$ is an efficient portfolio for $\alpha \in [0; 1]$.

Proof. To prove the first part let μ_c denote the mean return on a given minimum variance portfolio \mathbf{x}_c . Now choose α such that $\mu_c = \alpha\mu_a + (1 - \alpha)\mu_b$, that is $\alpha = (\mu_c - \mu_b)/(\mu_a - \mu_b)$ (which is well-defined because $\mu_a \neq \mu_b$). But since \mathbf{x}_c is a minimum variance portfolio we know that (9.10) holds, so

$$\begin{aligned} \mathbf{x}_c &= \Sigma^{-1}[\mu \ \mathbf{1}]\mathbf{A}^{-1} \begin{bmatrix} \mu_c \\ 1 \end{bmatrix} \\ &= \Sigma^{-1}[\mu \ \mathbf{1}]\mathbf{A}^{-1} \begin{bmatrix} \alpha\mu_a + (1 - \alpha)\mu_b \\ \alpha + (1 - \alpha) \end{bmatrix} \\ &= \alpha\mathbf{x}_a + (1 - \alpha)\mathbf{x}_b, \end{aligned}$$

where the third line is obtained because \mathbf{x}_a and \mathbf{x}_b also fulfill (9.10). This proves the first statement. The second statement is proved by “reading from right to left” in the above equations. This shows that $\mathbf{x}_c = \alpha\mathbf{x}_a + (1 - \alpha)\mathbf{x}_b$ is




Figure 9.2: The minimum variance frontiers and individual assets

the minimum variance portfolio with expected return $\alpha\mu_a + (1 - \alpha)\mu_b$. From this, the validity of the third statement is clear. \square

Another important notion is *orthogonality* of portfolios. We say that two portfolios \mathbf{x}_P and \mathbf{x}_{zP} (“z is for zero”) are orthogonal if the covariance of their rates of return is 0, i.e.

$$\mathbf{x}_{zP}^\top \Sigma \mathbf{x}_P = 0. \tag{9.12}$$

Often \mathbf{x}_{zP} is called \mathbf{x}_P ’s 0- β portfolio (we’ll see why later).

Proposition 20. *For every minimum variance portfolio, except the global minimum variance portfolio, there exists a unique orthogonal minimum variance portfolio. Furthermore, if the first portfolio has mean rate of return μ_P , its orthogonal one has mean*

$$\mu_{zP} = \frac{a - b\mu_P}{b - c\mu_P}.$$

Proof. First note that μ_{zP} is well-defined for any portfolio except the global minimum variance portfolio. By (9.10) we know how to find the minimum variance portfolios with means μ_P and $\mu_{zP} = (a - b\mu_P)/(b - c\mu_P)$. This leads to

$$\begin{aligned}
\mathbf{x}_{zP}^\top \Sigma \mathbf{x}_P &= [\mu_{zP} \ 1] \mathbf{A}^{-1} [\mu \ \mathbf{1}]^\top \Sigma^{-1} \Sigma \Sigma^{-1} [\mu \ \mathbf{1}] \mathbf{A}^{-1} [\mu_P \ 1]^\top \\
&= [\mu_{zP} \ 1] \mathbf{A}^{-1} \underbrace{([\mu \ \mathbf{1}]^\top \Sigma^{-1} [\mu \ \mathbf{1}])}_{=\mathbf{A} \text{ by def.}} \mathbf{A}^{-1} [\mu_P \ 1]^\top \\
&= [\mu_{zP} \ 1] \mathbf{A}^{-1} \begin{bmatrix} \mu_P \\ 1 \end{bmatrix} \\
&= \begin{bmatrix} a - b\mu_P & 1 \\ b - c\mu_P & 1 \end{bmatrix} \frac{1}{ac - b^2} \begin{bmatrix} c & -b \\ -b & a \end{bmatrix} \begin{bmatrix} \mu_P \\ 1 \end{bmatrix} \\
&= \frac{1}{ac - b^2} \begin{bmatrix} a - b\mu_P & 1 \\ b - c\mu_P & 1 \end{bmatrix} \begin{bmatrix} c\mu_P - b \\ a - b\mu_P \end{bmatrix} \\
&= 0,
\end{aligned} \tag{9.13}$$

which was the desired result. \square

Proposition 21. Let \mathbf{x}_{mv} ($\neq \mathbf{x}_{gmv}$, the global minimum variance portfolio) be a portfolio on the mean-variance frontier with rate of return r_{mv} , expected rate of return μ_{mv} and variance σ_{mv}^2 . Let \mathbf{x}_{zmv} be the corresponding orthogonal portfolio, \mathbf{x}_P be an arbitrary portfolio, and use similar notation for rates of return on these portfolios. Then the following holds:

$$\mu_P - \mu_{zmv} = \beta_{P,mv}(\mu_{mv} - \mu_{zmv}),$$

where

$$\beta_{P,mv} = \frac{\text{Cov}(r_P, r_{mv})}{\sigma_{mv}^2}.$$

Proof. Consider first the covariance between return on asset i and \mathbf{x}_{mv} . By using (9.10) we get

$$\begin{aligned}
\text{Cov}(r_i, r_{mv}) &= \mathbf{e}_i^\top \Sigma \mathbf{x}_{mv} \\
&= \mathbf{e}_i^\top [\mu \ \mathbf{1}] \mathbf{A}^{-1} \begin{bmatrix} \mu_{mv} \\ 1 \end{bmatrix} \\
&= [\mu_i \ 1] \mathbf{A}^{-1} \begin{bmatrix} \mu_{mv} \\ 1 \end{bmatrix}.
\end{aligned}$$

From calculations in the proof of Proposition 20 we know that the covariance between \mathbf{x}_{mv} and \mathbf{x}_{zvp} is given by (9.13). We also know that it is 0.

Subtracting this 0 from the above equation gives

$$\begin{aligned} \text{Cov}(r_i, r_{mv}) &= [\mu_i - \mu_{zmv} \ 0] \mathbf{A}^{-1} \begin{bmatrix} \mu_{mv} \\ 1 \end{bmatrix} \\ &= (\mu_i - \mu_{zmv}) \underbrace{\frac{c\mu_{mv} - b}{ac - b^2}}_{:=\gamma}, \end{aligned} \tag{9.14}$$

where we have used the formula for \mathbf{A}^{-1} . Since this holds for all individual assets and covariance is bilinear, it also holds for portfolios. In particular for \mathbf{x}_{mv} ,

$$\sigma_{mv}^2 = \gamma(\mu_{mv} - \mu_{zmv}),$$

so $\gamma = \sigma_{mv}^2 / (\mu_{mv} - \mu_{zmv})$. By substituting this into (9.14) we get the desired result for individual assets. But then linearity ensures that it holds for all portfolios. ■

Proposition 21 says that the expected excess return on any portfolio (over the expected return on a certain portfolio) *is a linear function* of the expected excess return on a minimum variance portfolio. It also says that the expected excess return is proportional to covariance.

The converse of Proposition 21 holds in the following sense: If there is a candidate portfolio x_C and a number μ_{zC} such that for any individual asset i we have

$$\mu_i - \mu_{zC} = \beta_{i,C}(\mu_C - \mu_{zC}), \tag{9.15}$$

with $\beta_{i,C} = \text{Cov}(r_i, r_C) / \sigma_C^2$, then x_C is a minimum-variance portfolio. To see why, put $\gamma_i = \sigma_C^2(\mu_i - \mu_{zC}) / (\mu_C - \mu_{zC})$, note that we have $\gamma = \Sigma x_C$, and that that uniquely determines the candidate portfolio. But by Proposition 32 we know that the minimum variance portfolio with expected rate of return μ_C is (the then) one (and only) portfolio for which (9.15) holds.

9.1.2 The case with a riskfree asset

We now consider a portfolio selection problem with $n + 1$ assets. These are indexed by $0, 1, \dots, n$, and 0 corresponds to the riskfree asset with (deterministic) rate of return μ_0 . For the risky assets we let μ_i^e denote the *excess* rate of return over the riskfree asset, i.e. the actual rate of return less μ_0 . We let μ^e denote the mean excess rate of return, and Σ the variance (which is of course unaffected). A portfolio is now a $n + 1$ -dimensional vector whose coordinate sum to unity. But in the calculations we let \mathbf{w} denote the vector of weights w_1, \dots, w_n corresponding to the risky assets and write $w_0 = 1 - \mathbf{w}^\top \mathbf{1}$.

capital market
line, CML
tangent portfolio

With these conventions the mean excess rate of return on a portfolio P is

$$\mu_P^e = \mathbf{w}^\top \boldsymbol{\mu}^e$$

and the variance is

$$\sigma_P^2 = \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w}.$$

Therefore the mean-variance portfolio selection problem with a riskless asset can be stated as

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^\top \boldsymbol{\Sigma} \mathbf{w} \quad \text{subject to} \quad \mathbf{w}^\top \boldsymbol{\mu}^e = \mu_P^e.$$

Note that $\mathbf{w}^\top \mathbf{1} = 1$ is not a constraint; some wealth may be held in the riskless asset.

As in the previous section we can set up the Lagrange-function, differentiate it, and solve to first order conditions. This gives the optimal weights

$$\hat{\mathbf{w}} = \frac{\mu_P^e}{(\boldsymbol{\mu}^e)^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e} \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e, \quad (9.16)$$

and the following expression for the variance of the minimum variance portfolio with mean excess return μ_P :

$$\hat{\sigma}_P^2 = \frac{(\mu_P^e)^2}{(\boldsymbol{\mu}^e)^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e}. \quad (9.17)$$

So we have determined the efficient frontier. For required returns above the riskfree rate, the efficient frontier in standard deviation-mean space is a straight line passing through $(0, \mu_0)$ with a slope of $\sqrt{(\boldsymbol{\mu}^e)^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e}$. This line is called the capital market line (CML).

The tangent portfolio, \mathbf{x} , is the minimum variance portfolio with all wealth invested in the risky assets, i.e. $\mathbf{x}_{tan}^\top \mathbf{1} = 1$. The mean excess return on the tangent portfolio is

$$\mu_{tan}^e = \frac{(\boldsymbol{\mu}^e)^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e}{\mathbf{1}^\top \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}^e},$$

which may be positive or negative. It is economically plausible to assert that the riskless return is lower than the mean return of the global minimum variance portfolio of the risky assets. In this case the situation is as illustrated in Figure 9.3, and that explains why we use the term ‘‘tangency’’. When $\mu_{tan}^e > 0$, the tangent portfolio is on the capital market line. But the tangent portfolio must also be on the ‘‘risky assets only’’ efficient frontier. So the straight line (the CML) and the hyperbola intersect at a point corresponding

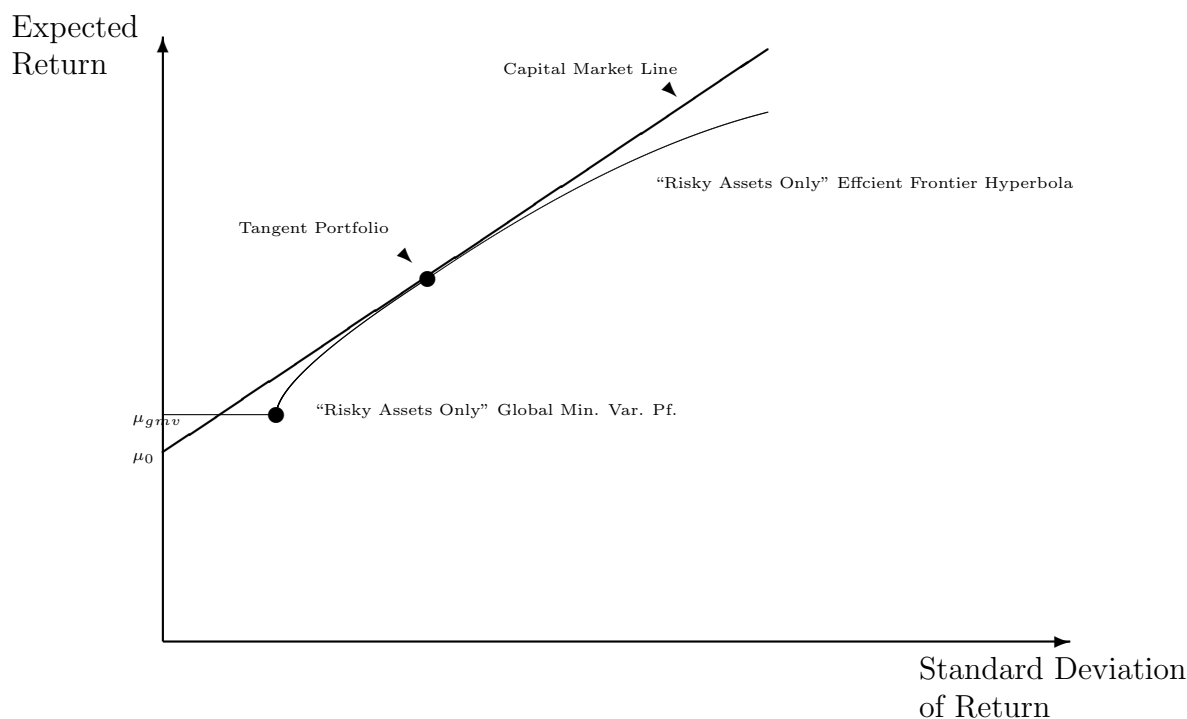


Figure 9.3: The capital market line.

to the tangency portfolio. But clearly the CML must be above the efficient frontier hyperbola (we are minimizing variance with an extra asset). So the CML is a tangent to the hyperbola.

For any portfolio, P we define the *Sharpe-ratio* (after William Sharpe) as excess return relative to standard deviation,

$$\text{Sharpe-ratio}_P = \frac{\mu_P - \mu_0}{\sigma_P}.$$

In the case where $\mu_{tan}^e > 0$, we see note from Figure 9.3 that the tangency portfolio is the “risky assets only”-portfolio with the highest Sharpe-ratio since the slope of the CML is the Sharpe-ratio of tangency portfolio. (Generally/”strictly algebraically” we should say that \mathbf{x}_{tan} has maximal squared Sharpe-ratio.) The observation that “Higher Sharpe-ratio is better. End of story.” makes it a frequently used tool for evaluating/comparing the performance for investment funds.

Note that a portfolio with full investment in the riskfree asset is orthogonal to any other portfolio; this means that we can prove the following result in exactly the manner as Proposition 21 (and its converse).

Proposition 22. Let \mathbf{x}_{mv} be a portfolio on the mean-variance frontier with rate of return r_{mv} , expected rate of return μ_{mv} and variance σ_{mv}^2 . Let \mathbf{x}_P be an arbitrary portfolio, and use similar notation for rates of return on these portfolios. Then the following holds:

$$\mu_P - \mu_0 = \beta_{P,mv}(\mu_{mv} - \mu_0),$$

where

$$\beta_{P,mv} = \frac{\text{Cov}(r_P, r_{mv})}{\sigma_{mv}^2}.$$

Conversely, a portfolio for which these equations hold for all individual assets is on the mean-variance frontier.

9.2 The Capital Asset Pricing Model (CAPM)

With the machinery of portfolio optimization in place, we are ready to formulate one of the key results of modern finance theory, the CAPM-relation. Despite the clearly unrealistic assumptions on which the result is built it still provides invaluable intuition on what factors determine the price of assets in equilibrium. Note that until now, we have mainly been concerned with pricing (derivative) securities when taking prices of some basic securities as given. Here we try to get more insight into what determines prices of securities to begin with.

We consider an economy with n risky assets and one riskless asset. Here, we let μ_i denote the rate of return on the i 'th risky asset and we let $\mu_0 = r_0$ denote the riskless rate of return. We assume that μ_0 is strictly smaller than the return of the global minimum variance portfolio.

Just as in the case of only risky assets one can show that with a riskless asset the expected return on any asset or portfolio can be expressed as a function of its beta with respect to an efficient portfolio. In particular, since the tangency portfolio is efficient we have

$$Er_i - \mu_0 = \beta_{i,tan}(E(r_{tan}) - \mu_0) \quad (9.18)$$

where

$$\beta_{i,tan} = \frac{\text{Cov}(r_i, r_{tan})}{\sigma_{tan}^2} \quad (9.19)$$

The critical component in deriving the CAPM is the identification of the tangency portfolio as the *market portfolio*. The market portfolio is defined as follows: Assume that the initial supply of risky asset j at time 0 has a value of P_0^j . (So P_0^j is the number of shares outstanding times the price per

share.) The market portfolio of risky assets then has portfolio weights given as

$$w^j = \frac{P_0^j}{\sum_{i=1}^n P_0^i} \quad (9.20)$$

equilibrium
capital asset
pricing model,
CAPM

Note that it is quite reasonable to think of a portfolio with these weights as reflecting “the average of the stock market”.

Now if all (say K) agents are mean-variance optimizers (given wealths of $W_i(0)$ to invest), we know that since there is a riskless asset they will hold a combination of the tangency portfolio and the riskless asset since two fund separation applies. Hence all agents must hold the same mix of risky assets as that of the tangency portfolio. This in turn means that in equilibrium where market clearing requires all the risky assets to be held, the market portfolio (which is a convex combination of the individual agents’ portfolios) has the same mixture of assets as the tangency portfolio. Or in symbols: Let ϕ_i denote the fraction of his wealth that agent i has invested in the tangency portfolio. By summing over all agents we get

$$\begin{aligned} \text{Total value of asset } j &= \sum_{i=1}^K \phi_i W_i(0) \mathbf{x}_{tan}(j) \\ &= \mathbf{x}_{tan}(j) \times \text{Total value of all risky assets,} \end{aligned}$$

where we have used that market clearing condition that all risky assets must be held by the agents. This is a very weak consequence of equilibrium; some would just call it an accounting identity. The main *economic assumption* is that agents are mean-variance optimizers so that two fund separation applies. Hence we may as well write the market portfolio in equation (9.18). This is the CAPM:

$$E(r_i) - \mu_0 = \beta_{i,m}(E(r_m) - \mu_0) \quad (9.21)$$

where $\beta_{i,m}$ is defined using the market portfolio instead of the tangency portfolio. Note that the type of risk for which agents receive excess returns are those that are correlated with the market. The intuition is as follows: If an asset pays off a lot when the economy is wealthy (i.e. when the return of the market is high) that asset contributes wealth in states where the marginal utility of receiving extra wealth is small. Hence agents are not willing to pay very much for such an asset at time 0. Therefore, the asset has a high return. The opposite situation is also natural at least if one ever considered buying insurance: An asset which moves opposite the market has a high pay off in states where marginal utility of receiving extra wealth is high. Agents are willing to pay a lot for that at time 0 and therefore the asset has a low return. Indeed it is probably the case that agents are willing to accept a return on

an insurance contract which is below zero. This gives the right intuition but the analogy with insurance is actually not completely accurate in that the risk one is trying to avoid by buying an insurance contract is not linked to market wide fluctuations.

Note that one could still view the result as a sort of relative pricing result in that we are pricing everything in relation to the given market portfolio. To make it more clear that there is an equilibrium type argument underlying it all, let us see how characteristics of agents help in determining the risk premium on the market portfolio. Consider the problem of agent i in the one period model. We assume that returns are multivariate normal and that the utility function is twice differentiable and concave¹:

$$\begin{aligned} \max_{\mathbf{w}} E(u_i(W_1^i)) \\ \text{s.t. } W_1^i = W_0(\mathbf{w}^\top \mathbf{r} + (1 - \mathbf{w}^\top \mathbf{1})r_0) \end{aligned}$$

When forming the Lagrangian of this problem, we see that the first order condition for optimality is that for each asset j and each agent i we have

$$E(u_i'(W_1^i)(r_j - r_0)) = 0 \quad (9.22)$$

Remembering that $\text{Cov}(X, Y) = EXY - EXEY$ we rewrite this as

$$E(u_i'(W_1^i)) E(r_j - r_0) = -\text{Cov}(u_i'(W_1^i), r_j)$$

A result known as Stein's lemma says that for bivariate normal distribution (X, Y) we have

$$\text{Cov}(g(X), Y) = Eg'(X)\text{Cov}(X, Y)$$

and using this we have the following first order condition:

$$E(u_i'(W_1^i)) E(r_j - r_0) = -Eu_i''(W_1^i)\text{Cov}(W_1^i, r_j)$$

i.e.

$$\frac{-E(u_i'(W_1^i)) E(r_j - r_0)}{Eu_i''(W_1^i)} = \text{Cov}(W_1^i, r_j)$$

Now define the following measure of agent i 's absolute risk aversion:

$$\theta_i := \frac{-Eu_i''(W_1^i)}{Eu_i'(W_1^i)}.$$

¹This derivation follows Huang and Litzenberger: *Foundations for Financial Economics*. If prices are positive, then returns are bigger than -1 , so normality must be an approximation.

Summation over all agents gives us

$$\begin{aligned} E(r_j - r_0) &= \frac{1}{\sum_{i=1}^K \frac{1}{\theta_i}} \text{Cov}(W_1, r_j) \\ &= \frac{1}{\sum_{i=1}^K \frac{1}{\theta_i}} W_0 \text{Cov}(r_m, r_j) \end{aligned}$$

where the total wealth at time 1 held in risky assets is $W_1 = \sum_{i=1}^K W_1^i$, W_0 is the total wealth in risky assets at time 0, and

$$r_m = \frac{W_1}{W_0} - 1$$

therefore is the return on the market portfolio. Note that this alternative representation tells us more about the risk premium as a function of the aggregate risk aversion across agents in the economy. By linearity we also get that

$$Er_m - \mu_0 = W_0^M \text{Var}(r_m) \frac{1}{\sum_{i=1}^K \frac{1}{\theta_i}},$$

which gives a statement as to the actual magnitude expected excess return on the market portfolio. A high θ_i corresponds to a high risk aversion and this contributes to making the risk premium larger, as expected. Note that if one agent is very close to being risk neutral then the risk premium (holding that person's initial wealth constant) becomes close to zero. Can you explain why that makes sense?

The derivation of the CAPM when using returns is not completely clear in the sense that finding an equilibrium return does not separate out what is found exogenously and what is found endogenously. One should think of the equilibrium argument as determining the initial price of assets given assumptions on the distribution of the price of the assets at the end of the period. A sketch of how the equilibrium argument would run is as follows:

1. Let the expected value and the covariance of end-of-period asset prices for all assets be given.
2. Suppose further that we are given a utility function for each investor which depends only on mean and variance of end-of-period wealth. Assume that utility decreases as a function of variance and increases as a function of mean. Assume also sufficient differentiability
3. Let investor i have an initial fraction of the total endowment of risky asset j .

Roll's critique

4. Assume that there is riskfree lending and borrowing at a fixed rate r . Hence the interest rate is exogenous.
5. Given initial prices of all assets, agent i chooses portfolio weights on risky assets to maximize end of period utility. The difference in price between the initial endowment of risky assets and the chosen portfolio of risky assets is borrowed n/placed in the money market at the riskless rate. (In equilibrium where all assets are being held this implies zero net lending/borrowing.)
6. Compute the solution as a function of the initial prices.
7. Find a set of initial prices such that markets clear, i.e such that the sum of the agents positions in the risky assets sum up to the initial endowment of assets.
8. The prices will reflect characteristics of the agents' utility functions, just as we saw above.
9. Now it is possible to derive the CAPM relation by computing expected returns etc. using the endogenously determined initial prices. This is a purely mathematical exercise translating the formula for prices into formulas involving returns.

Hence CAPM is to be thought of as an equilibrium argument explaining asset prices.

There are of course many unrealistic assumptions underlying the CAPM. The distributional assumptions are clearly problematic. Even if basic securities like stocks were well approximated by normal distributions there is no hope that options would be well approximated due to their truncated payoffs. An answer to this problem is to go to continuous time modelling where 'local normality' holds for very broad classes of distributions but that is outside the scope of this course. Note also that a conclusion of CAPM is that all agents hold the same mixture of risky assets which casual inspection show is not the case.

A final problem, originally raised by Roll, and thus referred to as Roll's critique², concerns the observability of the market portfolio and the logical equivalence between the statement that the market portfolio is efficient and the statement that the CAPM relation holds. To see that observability is a problem think for example of human capital. Economic agents face many

²R. Roll (1977): A critique of the asset pricing theory's test; Part I, Journal of Financial Economics, 4:pp 129 - 76

decisions over a life time related to human capital - for example whether it is worth taking a loan to complete an education, weighing off leisure against additional work which may increase human capital etc. Many empirical studies use all traded stocks (and perhaps bonds) on an exchange as a proxy for the market portfolio but clearly this is at best an approximation. And what if the test of the CAPM relation is rejected using that portfolio? At the intuitive level, the relation (9.18) tells us that this is equivalent to the inefficiency of the chosen portfolio. Hence one can always argue that the reason for rejection was not that the model is wrong but that the market portfolio is not chosen correctly (i.e. is not on the portfolio frontier). Therefore, it becomes very hard to truly test the model. While we are not going to elaborate on the enormous literature on testing the CAPM, note also that even at first glance it is not easy to test what is essentially a one period model. To get estimates of the fundamental parameters (variances, covariances, expected returns) one will have to assume that the model repeats itself over time, but when firms change the composition of their balance sheets they also change their betas.

Hence one needs somehow to accommodate betas which change over time and this inevitably requires some statistical compromises.

9.3 Relevant but unstructured remarks on CAPM

9.3.1 Systematic and non-systematic risk

This section follows Huang and Litzenberger's Chapters 3 and 4. We have two versions of the capital asset pricing model. The most "popular" version, where we assumed the existence of a riskless asset whose return is r_0 , states that the expected return on any asset satisfies

$$Er_i - r_0 = \beta_{i,m}(Er_m - r_0). \quad (9.23)$$

This version we derived in the previous section. The other version is the so-called zero-beta CAPM, which replaces the return on the riskless asset by the expected return on m 's zero-covariance portfolio:

$$Er_i - Er_{zm} = \beta_{i,m}(Er_m - Er_{zm}).$$

This version is proved by assuming mean-variance optimizing agents, using that two-fund separation then applies, which means that the market portfolio is on the mean-variance locus (note that we cannot talk about a tangent portfolio in the model with no riskfree asset) and using Proposition 21. Note

that both relations state that excess returns (i.e. returns in addition to the riskless returns) are linear functions of β_{im} .

From now on we will work with the case in which a riskless asset exists, but it is easy to translate to the zero-beta version also. Dropping the expectations (and writing “error terms” instead) we have also seen that if the market portfolio m is efficient, the return on any portfolio (or asset) q satisfies

$$r_q = (1 - \beta_{q,m})r_f + \beta_{q,m}r_m + \epsilon_{q,m}$$

where

$$E\epsilon_{q,m} = E\epsilon_{q,m}r_m = 0.$$

Hence

$$\text{Var}(r_q) = \beta_{q,m}^2 \text{Var}(r_m) + \text{Var}(\epsilon_{q,m}).$$

This decomposes the variance of the return on the portfolio q into its *systematic risk* $\beta_{qm}^2 \text{Var}(r_m)$ and its *non-systematic* or *idiosyncratic risk* $\text{Var}(\epsilon_{q,m})$. The reason behind this terminology is the following: We know that there exists a portfolio which has the same expected return as q but whose variance is $\beta_{qm}^2 \text{Var}(r_m)$ - simply consider the portfolio which invests $1 - \beta_{qm}$ in the riskless asset and $\beta_{q,m}$ in the market portfolio. On the other hand, since this portfolio is efficient, it is clear that we cannot obtain a lower variance if we want an expected return of $E r_q$. Hence this variance is a risk which is correlated with movements in the market portfolio and which is non-diversifiable, i.e. cannot be avoided if we want an expected return of $E r_q$. On the other hand as we have just seen the risk represented by the term $\text{VAR}(\epsilon_{q,m})$ can be avoided simply by choosing a different portfolio which does a better job of diversification without changing expected return.

9.3.2 Problems in testing the CAPM

Like any model CAPM builds on simplifying assumptions. The model is popular nonetheless because of its strong conclusions. And it is interesting to try and figure out whether the simplifying assumptions on the behavior of individuals (homogeneous expectations) and on the institutional setup (no taxation, transactions costs) of trading are too unrealistic to give the model empirical relevance. What are some of the obvious problems in testing the model?

First, the model is a one period model. To produce estimates of mean returns and standard deviations, we need to observe years of price data. Can we make sure that the distribution of returns over several years remain the same³?

³Multiperiod versions exist, but they also face problems with time varying parameters.

Second (and this a very important problem) what is the 'market portfolio'? Since investments decisions of firms and individuals in real life are not restricted to stocks and bonds but include such things as real estate, education, insurance, paintings and stamp collections, we should include these assets as well, but prices on these assets are hard to get and some are not traded at all.

A person rejecting the CAPM could always be accused of not having chosen the market portfolio properly. However, note that if 'proper choice' of the market portfolio means choosing an efficient portfolio then this is mathematically equivalent to having the CAPM hold.

This point is the important element in what is sometimes referred to as Roll's critique of the CAPM. When discussing the CAPM it is important to remember which facts are mathematical properties of the portfolio frontier and which are behavioral assumptions. The key behavioral assumption of the CAPM is that the market portfolio is efficient. This assumption gives the CAPM-relation mathematically. Hence it is impossible to separate the claim 'the portfolio m is efficient' from the claim that 'CAPM holds with m acting as market portfolio'.

9.3.3 Testing the efficiency of a given portfolio

Since the question of whether CAPM holds is intimately linked with the question of the efficiency of a certain portfolio it is natural to ask whether it is possible to devise a statistical test of the efficiency of a portfolio with respect to a collection of assets. If we knew expected returns and variances exactly, this would be a purely mathematical exercise. However, in practice parameters need to be estimated and the question then takes a more statistical twist: Given the properties of estimators of means and variances, can we reject at (say) a 5% level that a certain portfolio is efficient? Gibbons, Ross and Shanken (Econometrica 1989, 1121-1152) answer this question - and what follows here is a sketch of their test.

Given a portfolio m and N assets whose excess returns are recorded in T time periods. It is assumed that a sufficiently clear concept of riskless return can be defined so that we can really determine excess returns for each period. NOTE: We will change our notation in this section slightly and assume that r_p , Er_p and μ_p refer to *excess* returns, mean excess returns and estimated mean excess returns of an asset or portfolio p . Hence using this notation the CAPM with a riskless asset will read

$$Er_p = \beta_{p,m} Er_m.$$

We want to test this relation or equivalently whether m is an efficient portfolio

in a market consisting of N assets. Consider the following statistical model for the excess returns of the assets given the excess return on the portfolio m :

$$r_{it} = \alpha_i + \gamma_i r_{mt} + \epsilon_{it}$$

$$i = 1, \dots, N \text{ and } t = 1, \dots, T$$

where r_{it} is the (random) *excess* return⁴ of asset i in the t 'th period, r_{mt} is the observed *excess* return on the portfolio in the t 'th period, α_i, γ_i are constants and the ϵ_{it} 's are normally distributed with $Cov(\epsilon_{it}, \epsilon_{jt}) = \sigma_{ij}$ and $Cov(\epsilon_{it}, \epsilon_{is}) = 0$ for $t \neq s$. Given these data a natural statistical representation of the question of whether the portfolio m is efficient is the hypothesis that $\alpha_1 = \dots = \alpha_N = 0$. This condition must hold for (9.23) to hold.

To test this is not difficult in principle (but there are some computational tricks involved which we will not discuss here): First compute the MLE's of the parameters. It turns out that in this model this is done merely by computing Ordinary Least Squares estimators for α, γ and the covariance matrix for each period Σ . A so-called Wald test of the hypothesis $\alpha = 0$ can then be performed by considering the test statistic

$$W_0 = \hat{\alpha} Var(\hat{\alpha}) \hat{\alpha}^{-1}$$

which you will learn more about in a course on econometrics. Here we simply note that the test statistic measures a distance of the estimated value of α from the origin. Normally, this type of statistics leads to an asymptotic chi squared test, but in this special model the distribution can be found explicitly and even more interesting from a finance perspective, it is shown in GRS that W_0 has the following form

$$W_0 = \frac{(T - N - 1)}{N} \frac{\left(\frac{\hat{\mu}_q^2}{\hat{\sigma}_q^2} - \frac{\hat{\mu}_m^2}{\hat{\sigma}_m^2} \right)}{\left(1 + \frac{\hat{\mu}_m^2}{\hat{\sigma}_m^2} \right)}$$

where the symbols require a little explanation: In the minimum variance problem with a riskless asset we found that the excess return of any portfolio satisfies

$$Er_p = \beta_{pm} Er_m.$$

We refer to the quantity

$$\frac{Er_p}{\sigma(r_p)}$$

⁴Note this change to excess returns.

as the Sharpe ratio for portfolio p . The Sharpe ratio in words compares excess return to standard deviation. Note that using the CAPM relation we can write

$$\frac{Er_p}{\sigma(r_p)} = \frac{\sigma(r_m)\rho_{mp}}{\sigma^2(r_m)}(Er_m)$$

where ρ_{mp} is the correlation coefficient between the return of portfolios p and m . From this expression we see that the portfolio which maximizes the Sharpe ratio is (proportional) to m . Only portfolios with this Sharpe ratio are efficient. Now the test statistic W_0 compares two quantities: On one side, the maximal Sharpe ratio that can be obtained when using for parameters in the minimum variance problem the estimated covariance matrix and the estimated mean returns for the economy consisting of the N assets and the portfolio m . On the other side, the Sharpe ratio for the particular portfolio m (based on its estimated mean return and standard deviation).

Large values of W_0 will reject the hypothesis of efficiency and this corresponds to a case where the portfolio m has a very poor expected return per unit of standard deviation compared to what is obtained by using all assets.

Index

- geometric Brownian motion, 127

- abstract Bayes formula, 83
- adapted process, 74
- alfahage; $\alpha_n]R$, 24
- annuity, 22, 23
- arbitrage opportunity, 16, 54, 76
- arbitrage, odds, 17
- Arrow-Debreu prices, 54
- Asian option *see* contingent claim, path dependent 113

- bank account, 48
- barbell strategy, 42
- binomial model, 47, 105, 121
- binomial model, one-period, 48
- Black-Scholes formula, 115, 121, 124
- Black-Scholes model, 127
- Brownian motion, 126
- bullet bond, 22, 24

- call option, American, 103
- call option, European, 48, 70, 107, 108
- capital asset pricing model, CAPM, 165
- capital budgeting, 32
- capital market line, CML, 162
- capital structure, 119
- Central Limit Theorem, 124
- clean price, 26
- complete market, 18, 55, 80
- compounding periods, 21
- conditional expectation, 81

- contingent claim, 56, 113
- contingent claim, path dependent, 113
- continuously compounded interest rate, 22, 121
- convexity, 37
- crash fear, 117

- debt, 118, 119
- delta hedging, 110
- delta-hedging, 132
- derivative security, 56, 72
- dirty price, 26
- discount factors, 18, 19
- dividend process, 74, 83
- dividend stream; δ_0^ϕ , 76
- duration, 41
- duration, Fisher-Weil, 41
- duration, Macaulay, 36
- duration, modified, 38

- edge, 12
- efficient portfolio, 157
- equilibrium, 9, 165
- equity, 118, 119
- equivalent martingale measure, 83
- equivalent measure, 51
- equivalent probability measures , 82

- filtration, 72
- financial market, 16, 54
- firm, 118
- Fisher Separation, 11
- forward contract, 102, 142
- forward price, 102, 143

- forward rates, 20
- frictionless market, 98
- fundamental theorem of asset pricing,
 - first, 84
- fundamental theorem of asset pricing,
 - second, 87
- futures contract, 142
- futures price, 143
- Gordon's growth formula, 32
- hedging, 72, 108
- immunization, 40
- implied volatility, 115, 117
- information, 72
- internal rate of return, 29
- Iterated expectations, 81
- Kelly criterion, 13
- local characterization of martingale measure, 83
- locally riskfree asset, 75
- logarithmic return, 121
- lognormal distribution, 125
- long position, 16
- marginal utility, 63
- Markov property, 111
- Markowitz analysis, 154
- martingale, 82
- mean/variance analysis, 154
- measurable, 74
- Merton's tunnel, 103
- minimum variance portfolio, 156
- moneyness, 100
- multi-period model, 69
- net present value, NPV, 15, 30
- NPV criterion, 30
- numeraire asset, 51
- odds, 12
- option pricing, 47, 97, 105
- partition, 72
- pay-off diagram, 99
- payment stream, 16
- portfolio, 16, 54
- portfolio insurance, 117
- positivity of vectors, 15
- price process, 75, 83
- pricing functional, 78
- probability measure, 53
- put option, American, 97, 114
- put option, European, 97
- put-call parity, 98, 102
- rate of return, 152
- redundant, 55
- relative portfolio weights, 152
- replication, 49, 72, 109
- Riding the yield curve, 44
- Riesz' representation theorem, 79
- risk neutral measure, 51
- risk-neutral pricing, 56
- riskfree asset, 54
- Roll's critique, 168
- security price system, 16, 54
- self-financing trading strategy, 76, 109
- separating hyperplane, 79
- serial loan, 22, 24
- Sharpe-ratio, 163
- short position, 16
- short rate, 20
- short rate lattice, 135
- short rate lattice, calibration, 137
- short rate process, 75, 134
- splitting index, 88, 109
- state-price vector, 54, 62
- Stiemke's lemma, 18
- stochastic interest rate, 133
- stock, 48, 105, 121
- swap contract, 146

- swap rate, 147
- tangent portfolio, 162
- term structure of interest rates, 21, 133
- term structure of interest rates, parallel shifts, 140, 149
- trading strategy, ϕ , 76
- two-fund separation, 157
- utility function, 9, 60, 153
- volatility, 128
- volatility smile, 118
- yield to maturity, 21
- zero coupon bond, ZCB, 19, 133, 134
- zero-beta portfolio, 159