STOCHASTIC FINANCIAL MODELS

1 Utility and mean-variance analysis.

Why should financial models be stochastic? Randomness is an inescapable feature of financial markets; although some agents may be better at predicting the future behaviour of markets, even the most successful make losses from time to time. Our modelling therefore must involve probabilistic elements.

A market is the interaction of agents trading goods and services, and the actions and choices of individual agents are shaped by *preferences* over different *contingent claims.* A contingent claim is simply a well-specified random payment, mathematically, a random variable. We shall suppose that agents' preferences are expressed through an *expected utility representation*, that is,

$$X \preceq Y \quad (Y \text{ is preferred to } X) \qquad \iff \qquad EU(X) \leq EU(Y), \tag{1.1}$$

where the function $U : \mathbb{R} \to [-\infty, \infty)$ is *increasing* (that is, non-decreasing), and may differ from one agent to another.

Remarks. (i) We insist that the utility function U is non-decreasing, because it is reasonable to assume that more is preferred to less. One can imagine situations where this may not be true (is a ton of ice-cream *really* better that a litre?), but they tend to be somewhat artificial, and so we exclude them from consideration[1]. (ii) The utility may be allowed to take the value $-\infty$, both for mathematical convenience, and to represent the possibility of an unacceptable outcome.

It is generally agreed that if an agent is offered the choice of a contingent claim X, or a certain payment of the mean value EX, then he will prefer the second; this property is called *risk-aversion*, for obvious reasons. In terms of the expected utility representation of the agent's preferences, this amounts to saying that U is *concave.*

Definition 1.1. *A function $U : \mathbb{R} \to [-\infty, \infty)$ is said to be* concave *if for all x, $y \in \mathbb{R}$, for all $p \in [0, 1]$,*

$$pU(x) + (1-p)U(y) \leq U(px + (1-p)y).$$

We shall use the notation

$$\mathcal{D}(U) = \{x : U(x) > -\infty\}$$

and

$$\mathcal{D}^o(U) = \text{int}(\mathcal{D}(U))$$

[1]... though not entirely; quadratic U is sometimes considered because of its tractability, even though it is not increasing.

REMARKS. (i) By Jensen's inequality,

$$EU(X) \leq U(EX),$$

so an agent will always prefer the certain mean of a contingent claim to the claim itself.

(ii) In the special case where the agent's utility function is linear, we say that the agent is *risk neutral.*

Examples.

(i) The function

$$U(x) = -\exp(-\gamma x) \quad (x \in \mathbb{R}),$$

where $\gamma > 0$ is a constant, is a concave increasing function, and commonly used as a utility, called the *constant absolute risk aversion (CARA)* utility.

(ii) The function

$$\begin{aligned} U(x) &= \frac{x^{1-R}}{1-R} \qquad (x \geq 0) \\ &= -\infty \qquad (x < 0) \end{aligned}$$

where $R > 0$, $R \neq 1$, is concave and increasing, and is the *constant relative risk aversion (CRRA)* utility.

(iii) The function

$$\begin{aligned} U(x) &= \log x \qquad (x \geq 0) \\ &= -\infty \qquad (x < 0) \end{aligned}$$

is the *logarithmic* utility. It can be thought of as the CRRA utility[2] with $R = 1$.

(iv) For fixed $\alpha \in [0, 1)$, the function

$$U(x) = \min(x, \alpha x)$$

is concave and increasing.

(v) The function

$$U(x) = -\frac{1}{2}x^2 + ax$$

for $a \geq 0$ is concave, but not increasing.

(vi) If U_1 and U_2 are utilities, and α_1 and α_2 are positive, then $U \equiv \alpha_1 U_1 + \alpha_2 U_2$ is again a utility.

[2]For $R \geq 1$, we understand $U(0) = -\infty$.

(vii) If $\{U_\lambda : \lambda \in \Lambda\}$ is a family of utilities, then

$$U(x) \equiv \inf_{\lambda \in \Lambda} U_\lambda(x)$$

defines another utility U.

(viii) In fact, a converse of (vii) holds: if U is concave and *upper semicontinuous* (USC)[3] then U can be expressed as the infimum of a family of (not just concave but) *linear* functions:

$$U(x) = \inf_\lambda \{\tilde{U}(\lambda) + \lambda x\}.$$

Here, $\tilde{U}$ is the *dual* function of U, defined by

$$\tilde{U}(\lambda) \equiv \sup_x \{U(x) - \lambda x\}. \tag{1.2}$$

How do we know that all the functions given above are actually concave? Calculus gives us a way to check; on the way to this, we need the following useful characterisation of concavity.

Proposition 1.2. *A function $U : \mathbb{R} \to [-\infty, \infty)$ is concave if and only if for all points $x_1 < y_1 \leq x_2 < y_2$ at which U is finite, the inequality*

$$\frac{U(y_1) - U(x_1)}{y_1 - x_1} \geq \frac{U(y_2) - U(x_2)}{y_2 - x_2} \tag{1.3}$$

holds.

Proof. Firstly, suppose that U is concave. Clearly it is sufficient to suppose that $y_1 = x_2$, so for brevity write $x_1 = x < y_1 = x_2 = z < y_2 = y$. The inequality to be proved is

$$\frac{U(z) - U(x)}{z - x} \geq \frac{U(y) - U(z)}{y - z},$$

or equivalently

$$(y - x)U(z) \geq (z - x)U(y) + (y - z)U(x), \tag{1.4}$$

which is easily seen to be implied by concavity. Conversely, if (1.3) holds, then (1.4) also holds, and implies concavity.

Corollary 1.3. *For concave U and $x \in \mathcal{D}^o(U)$, the limits*

$$\lim_{h \downarrow 0} \frac{U(x+h) - U(x)}{h} = \uparrow \lim_{h \downarrow 0} \frac{U(x+h) - U(x)}{h} \equiv DU(x+)$$

[3] A function f is upper semicontinuous if for all $a \in \mathbb{R}$ the set $\{x : f(x) < a\}$ is open. Clearly if f is the infimum of a family of continuous functions, then it must be USC.

and

$$\lim_{h\downarrow 0}\frac{U(x)-U(x-h)}{h} = \downarrow\lim_{h\downarrow 0}\frac{U(x)-U(x-h)}{h} \equiv DU(x-)$$

exist, and satisfy

$$DU(x-) \geq DU(x+).$$

The right and left derivatives of U exist at every point of $\mathcal{D}^o(U)$, and are equal except on a set which is at most countable.

Corollary 1.4. *If U is C^2, then $U''(x) \leq 0$ for all x if and only if U is concave.*

Preferences lead nowhere without the possibility of choice. From now on, unless explicitly mentioned to the contrary, we shall make the (small) assumption that

all utilities are strictly increasing.

If we consider an agent with wealth w and C^2 utility U who is contemplating whether or not to accept a contingent claim X, then he will do so provided

$$EU(w+X) > U(w).$$

If we suppose that X is 'small' so that we may perform a Taylor expansion, this condition is approximately the same as the condition

$$U(w) + U'(w)E(X) + \frac{1}{2}U''(w)EX^2 > U(w).$$

Since $U'(w) > 0$ - utility is increasing - and $U''(w) < 0$ - utility is concave - the benefits of a positive mean EX are offset by the disadvantage of positive variance; the balance is just right (to this order of approximation) when

$$\frac{2EX}{EX^2} = -\frac{U''(w)}{U'(w)}, \tag{1.5}$$

where the right-hand side is the so-called *Arrow-Pratt coefficient of absolute risk aversion.* If we consider instead the effect of the proposed gamble to be multiplicative rather than additive, the decision for the agent will be to accept if

$$EU(w(1+X)) > U(w).$$

Assuming that $w > 0$, a similar argument shows that to this order of approximation the agent should accept when

$$\frac{2EX}{EX^2} \geq -\frac{wU''(w)}{U'(w)}, \tag{1.6}$$

where the right-hand side is the so-called *Arrow-Pratt coefficient of relative risk aversion.* The names of the first two examples of utilities are thus explained.

1.1 Reservation and marginal prices.

Although the derivations above are not rigorous, they do build our intuition. Developing this intuitive theme a bit further, let us consider an agent with utility U who is able to choose any contingent claim X from an admissible set $\mathcal{A}$; he will naturally choose X to achieve

$$\sup_{X\in\mathcal{A}} EU(X); \tag{1.7}$$

we shall suppose that the supremum is achieved at $X^* \in \mathcal{A}$. In the special case where $\mathcal{A}$ is an affine space[4] $\mathcal{A} = X^* + \mathcal{V}$, we have therefore that for all $\xi \in \mathcal{V}$ and all $t \in \mathbb{R}$

$$EU(X^*) \geq EU(X^* + t\xi),$$

formally differentiating with respect to t gives us the conclusion

$$\boxed{E[\, U'(X^*)\xi\,] = 0 \quad \text{for all } \xi \in \mathcal{V}.} \tag{1.8}$$

Definition 1.5. *The* utility-indifference price $\pi(Y)$ *for a contingent claim Y is defined by*

$$\sup_{X\in\mathcal{A}} E\,U(X + Y - \pi(Y)) = \sup_{X\in\mathcal{A}} E\,U(X). \tag{1.9}$$

The interpretation of this definition is clear; if you were to agree to pay some price p in order to receive the contingent claim Y, then $\pi(Y)$ is the largest price you would be willing to pay for it.

The first thing to know about the utility-indifference price is that it is concave.

Proposition 1.6. *The map $Y \mapsto \pi(Y)$ is concave.*

Proof. Suppose that Y_1, Y_2 are two random variables, and p_1, p_2 are two non-negative reals, $p_1 + p_2 = 1$. To simplify, suppose that the supremum in (1.7) is achieved at X^*, and that

$$E\,U(X_1^* + Y_1 - \pi(Y_1)) = E\,U(X_2^* + Y_2 - \pi(Y_2)) = E\,U(X^*). \tag{1.10}$$

Now we argue as follows (with $\bar{X} = p_1 X_1^* + p_2 X_2^*$ and $\bar{Y} = p_1 Y_1 + p_2 Y_2$):

$$\begin{aligned}
\sup_{X\in\mathcal{A}} E\,U(X + \bar{Y} - \pi(\bar{Y})) &= E\,U(X^*) \\
&= p_1 E\,U(X_1^* + Y_1 - \pi(Y_1)) + p_2 E\,U(X_2^* + Y_2 - \pi(Y_2)) \\
&\leq E\,U(\bar{X} + \bar{Y} - p_1\pi(Y_1) - p_2\pi(Y_2)) \\
&\leq \sup_{X\in\mathcal{A}} E\,U(X + \bar{Y} - p_1\pi(Y_1) - p_2\pi(Y_2)).
\end{aligned}$$

[4]That is, for any $X_1, X_2 \in \mathcal{A}$ and $t \in \mathbb{R}$, $tX_1 + (1-t)X_2 \in \mathcal{A}$. Equivalently, there exists a vector space $\mathcal{V}$ such that for any $X \in \mathcal{A}$ we have $\mathcal{A} = X + \mathcal{V}$.

Monotonicity of U forces the conclusion

$$p_1\pi(Y_1) + p_2\pi(Y_2) \le \pi(\bar{Y}) \tag{1.11}$$

as required. □

If we now consider some non-negative contingent claim Y, then the map $t \mapsto \pi(tY)$ is concave, and obviously non-decreasing. Equally obviously, $\pi(0) = 0$, so by concavity we learn that the map

$$t \mapsto f_Y(t) \equiv \frac{\pi(tY)}{t}$$

defined on $\mathbb{R}\backslash\{0\}$ is monontone decreasing, and so in particular limits at zero from either side exist. If we now suppose that $\mathcal{A}$ is an affine space, and that for each $t \neq 0$

$$\sup_{X\in\mathcal{A}} EU(X - tY + \pi(tY)) = EU(X_t^* - tY + \pi(tY)),$$

then

$$\begin{aligned}
EU(X^*) &= EU(X_t^* - tY + \pi(tY)) \\
&= EU(X^* + (X_t^* - X^*) - tY + tf_Y(t)) \\
&= E[\,U(X^*) + U'(X^*)\big\{(X_t^* - X^*) - tY + tf_Y(t)\big\}\,] + o(t) \\
&= E[\,U(X^*) + U'(X^*)\big\{-tY + tf_Y(t)\big\}\,] + o(t)
\end{aligned}$$

where to pass from the third to the fourth line we have used (1.8), and hence

$$\boxed{\frac{E[\,U'(X^*)Y\,]}{E[\,U'(X^*)\,]} = \lim_{t\to 0}\frac{\pi(tY)}{t}} \tag{1.12}$$

This expression is the agent's *marginal price* for Y, that is, the price per unit at which he would be prepared to buy or sell an infinitesimal amount of Y. Notice that the marginal price is *linear* in the contingent claim, in contrast to the bid and ask prices. If prices had been derived from some economic equilibrium, and the contingent claim Y was one which was marketed, then the market price of Y would have to equal the marginal price of Y given by (1.12), and this would have to hold for every agent. This is not to say that for every agent the marginal utility of optimal wealth would have to be the same; in general they are not. But the prices obtained by each agent from their marginal utility of optimal wealth via (1.12) would have to agree on all marketed contingent claims.

This heuristic discussion provides us with firm guidance for our intuition, and the form of the prices frequently fits (1.12). Although there are many steps where the analysis could fail, where we assume that suprema are attained, or that we can differentiate under the expectation, the most common reason for the above analysis to fail is that $\mathcal{A}$ is not an affine space!

1.2 Mean-variance analysis and the efficient frontier.

Looking at (1.5) and (1.6), it is natural[5] to think that if we are given a choice of contingent claims X, all with the same mean, then we should take the one with smallest variance. To take an explicit situation, consider a single-period model with d assets in which an agent may invest; at time 0 the price of the j^{th} asset is S_0^j, non-random, and at time 1 this asset delivers a random quantity S_1^j of the only consumption good in the economy; at time 1, agents receive their due goods and consume them. Introducing the notation

$$\mu = ES_1, \quad V = \text{cov}(S_1) \equiv E\Big[(S_1 - ES_1)(S_1 - ES_1)^T\Big],$$

if at time 0 the agent chooses to hold θ^j units of asset j ($j = 1, \dots, d$), then at time 1 his portfolio is worth

$$w_1 = \theta \cdot S_1 \equiv \sum_{j=1}^{d} \theta^j S_1^j.$$

Thus

$$Ew_1 = \theta \cdot \mu, \quad \text{var}(w_1) = \theta \cdot V\theta.$$

If the agent now requires to choose θ to give a predetermined mean value $Ew_1 = m$ and to have minimal variance, then his optimisation problem is to find

$$\min \frac{1}{2}\,\theta \cdot V\theta \quad \text{subject to } \theta \cdot \mu = m, \quad \theta \cdot S_0 = w_0. \tag{1.13}$$

The second constraint is the *budget constraint*, that the cost at time 0 of the chosen portfolio must equal the agent's wealth at time 0. The problem can be expressed more compactly as

$$\min \frac{1}{2}\,\theta \cdot V\theta \quad \text{subject to } A^T\theta = b, \tag{1.14}$$

where

$$A = \begin{pmatrix} \mu & S_0 \end{pmatrix}, \qquad b = \begin{pmatrix} m \\ w_0 \end{pmatrix} \tag{1.15}$$

To solve this, we introduce the Lagrange multiplier $\lambda = (\lambda_1, \lambda_2)^T$ and the Lagrangian

$$L = \frac{1}{2}\,\theta \cdot V\theta + \lambda \cdot (b - A^T\theta)$$

Assuming V is non-singular, this is minimised by choosing

$$\theta = V^{-1}A\lambda; \tag{1.16}$$

The undetermined multiplier λ is fixed by the constraint values in (1.14); assuming that μ is not a multiple of S_0, we obtain

$$\lambda = (A^T V^{-1} A)^{-1} b \tag{1.17}$$

[5]... but in general not correct ...

solved by ($\Delta \equiv (\mu \cdot V^{-1}\mu)(S_0 \cdot V^{-1}S_0) - (S_0 \cdot V^{-1}\mu)^2$)

$$\lambda = \Delta^{-1} \begin{pmatrix} S_0 \cdot V^{-1}S_0 & -S_0 \cdot V^{-1}\mu \\ -\mu \cdot V^{-1}S_0 & \mu \cdot V^{-1}\mu \end{pmatrix} \begin{pmatrix} m \\ w_0 \end{pmatrix}$$

The variance is thus

$$\theta \cdot V\theta = \lambda^T A^T V^{-1} A\lambda = \lambda \cdot b$$

which is more simply

$$\Delta^{-1}(m^2 S_0 \cdot V^{-1}S_0 - 2mw_0 S_0 \cdot V^{-1}\mu + w_0^2 \mu \cdot V^{-1}\mu), \tag{1.18}$$

which is *quadratic* in the required mean m. The variance is minimised to value

$$\frac{w_0^2}{S_0 \cdot V^{-1}S_0}$$

when we take $m = w_0(S_0 \cdot V^{-1}\mu)/(S_0 \cdot V^{-1}S_0)$. We can display the conclusions of this analysis graphically, as in Figure 1.2. For any chosen value of the mean m, corresponding to a given level in the plot Figure 1.2, values of the portfolio variance corresponding to points to the left of the parabola are not achievable, wheras points on and to the right of the parabola are. The parabola is called the *mean-variance efficient frontier.*

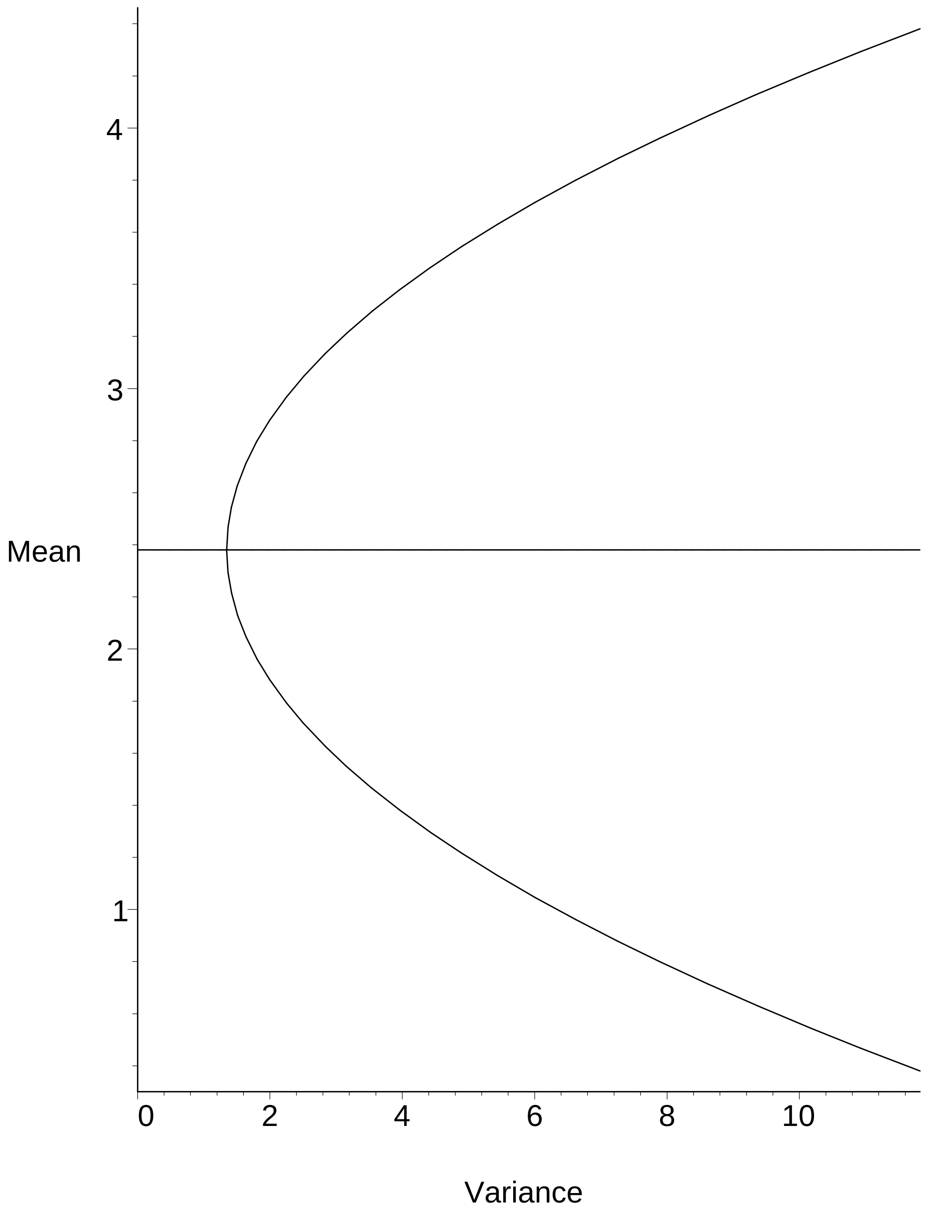

Figure 1: Mean-variance efficient frontier; achievable combinations of the mean and variance are to the right of the parabola.

Remarks. We noted in an earlier footnote that while it is natural to think that from among all available contingent claims with given mean we should choose the one with the smallest variance, this is not in general correct. The reason is not hard to see; assuming our agent has expected-utility preferences, if two contingent claims are considered equally desirable if they have the same mean and the same variance, then the utility must be a function *only* of the mean and the variance - in effect, the 'utility' is quadratic, which is disqualified because it is not increasing.

Despite this, the kind of mean-variance analysis set forth above, and graphs such as Figure 1.2 of the efficient frontier, are ubiquitous in the practice of portfolio management. Why should this be? There are two reasons:

(i) This analysis is just about as sophisticated as you can expect to put across to the mathematically untrained;

(ii) In one very special situation, when S_1 is multivariate normal, the mean-variance analysis in effect amounts to the correct expected-utility maximisation. We study this situation right now, assuming that agents have CARA utilities.

Example: CARA/Gaussian problem, no riskless asset. Recall the setup: this is a single-period model, in which an agent invests in d assets. At time 0, the j^{th} asset's price is S_0^j, and at time 1 it delivers S_1^j of the consumption good; we will speak of $S_t = (S_t^1, \ldots, S_t^d)^T$ as the vector of asset prices at time t, though there is no reason why prices should be comparable across the two time periods, and indeed at period 0 there is a scaling indeterminacy, as it is only the *ratio* of asset prices which is significant.

If at time 0 the agent chooses to hold θ^j units of asset j ($j = 1, \ldots, d$), then at time 1 his portfolio is worth

$$w_1 = \theta \cdot S_1 \equiv \sum_{j=1}^{d} \theta^j S_1^j.$$

Suppose that the agent has CARA utility, and so aims to maximise

$$E - \exp(-\gamma w_1). \tag{1.19}$$

Concerning the returns on the assets, we shall assume that the vector S_1 has a multivariate normal distribution, with mean vector μ and covariance matrix V. *We shall assume V is non-singular:* presently we look at what happens when one of the assets is riskless. It is well known that for $x \in \mathbb{R}^n$,

$$E \exp(x \cdot S_1) = \exp(x \cdot \mu + \frac{1}{2} x \cdot Vx),$$

where $x \cdot y \equiv x^T y$ denotes the scalar product of x and y, so the agent's objective is to minimise

$$E \exp(-\gamma\theta \cdot S_1) = \exp(-\gamma\theta \cdot \mu + \frac{1}{2}\gamma^2\theta \cdot V\theta).$$

Of course, the portfolio θ cannot be chosen unrestrictedly; the time-0 value must equal the time-0 wealth of the agent:

$$\theta \cdot S_0 = w_0.$$

The agent therefore faces the constrained optimisation problem

$$\min\{-\gamma\theta \cdot \mu + \frac{1}{2}\gamma^2\theta \cdot V\theta\} \quad \text{subject to} \quad \theta \cdot S_0 = w_0.$$

Using the Lagrangian method, we convert this problem into the unconstrained minimisation of

$$-\gamma\theta \cdot \mu + \frac{1}{2}\gamma^2\theta \cdot V\theta + \gamma\lambda(w_0 - \theta \cdot S_0).$$

Differentiating with respect to θ gives us the equation

$$\gamma V\theta = \mu + \lambda S_0, \tag{1.20}$$

which is solved by taking

$$\theta = \gamma^{-1}V^{-1}(\mu + \lambda S_0).$$

To match the constraint, we take

$$\lambda = \frac{\gamma w_0 - S_0 \cdot V^{-1}\mu}{S_0 \cdot V^{-1}S_0},$$

and hence the optimal θ has the explicit form

$$\boxed{\theta = \gamma^{-1}V^{-1}\mu + \frac{\gamma w_0 - S_0 \cdot V^{-1}\mu}{\gamma S_0 \cdot V^{-1}S_0}V^{-1}S_0.} \tag{1.21}$$

REMARKS. (i) Notice that the optimal portfolio (1.2) is a weighted average of just two portfolios, the *minimum-variance* portfolio $V^{-1}S_0$, which minimises the variance of S_1 subject to the initial budget constraint $\theta \cdot S_0 = w_0$, and the *diversified portfolio* $V^{-1}\mu$. This is an example of a *mutual fund theorem*.

(ii) Why did we choose to maximise objective (1.19), instead of maximising the utility of the expected *gain*

$$E - \exp(-\gamma\theta \cdot (S_1 - S_0)), \tag{1.22}$$

say? It is arguable that this is the criterion that an investor should be most concerned with; by taking objective (1.22), we now free the admissible portfolios from the initial budget constraint $\theta \cdot S_0 = w_0$, and in effect allow borrowing at time 0 to fund the portfolio choice. But this is exactly the problem; so far in the model as described, *there is no mechanism for borrowing!* There is no way to transfer obligations from one time period to another other than through holding the shares. The value $\theta \cdot S_0$ is the value of a portfolio at time 0, and cannot be compared with $\theta \cdot S_1$, the value one period later; they are denominated in completely different units, time-0 consumption and time-1 consumption.

Example: CARA/Gaussian problem with a riskless asset. Let us take exactly the situation of the previous example, but add one more asset, sometimes referred to as the *bank account* and denoted S^0, whose return is riskless[6]. We shall suppose that $S_0^0 = 1$, $S_1^0 = 1 + r$, where r is the

[6] That is, the variance of S_1^0 is zero. We shall therefore use the equivalent notations μ^0 and S_1^0 interchangeably.

rate of interest on the bank account. Asset zero now permits us to borrow risklessly; we may borrow amount x at time 0, provided we pay back amount $(1+r)x$ at time 1.

If the agent chooses at time 0 a portfolio $\theta = (\theta^1, \dots \theta^d)$ of the risky assets, the cost of this will be $\theta \cdot S_0$, so the optimization problem for the agent will be to

$$\min_\theta E \exp(-\gamma\{\, \theta \cdot S_1 + (1+r)(w_0 - \theta \cdot S_0)\}), \tag{1.23}$$

where we think that the bank account, initially at wealth w_0 gets changed to $w_0 - \theta \cdot S_0$ when the agent buys his desired portfolio, and this quantity of money gets scaled up by $(1+r)$ by time 1. Simple calculation turns this problem into the problem

$$\min_\theta \big[\, \tfrac{1}{2}\gamma\theta \cdot V\theta + \theta \cdot (\mu - (1+r)S_0)\, \big], \tag{1.24}$$

which is solved by

$$\boxed{\theta = \gamma^{-1}\theta_M \equiv \gamma^{-1}V^{-1}(\mu - (1+r)S_0). \qquad (1.25)}$$

Remarks. (i) Once again, the optimal portfolio is a weighted average of the minimum-variance portfolio and the diversified portfolio, though this time the weights are in fixed proportions. The portfolio θ_M is referred to as the *market portfolio,* for reasons we shall explain shortly.

(ii) Notice that in contrast to the solution (1.2) to the previous example, the best value of θ *does not depend on* w_0, the initial wealth of the agent. How can this be reconciled with the initial budget constraint? Very simply: the agent takes up the portfolio (1.25) in the risky assets, and his holding θ^0 of the riskless asset adjusts to pay for it!

(iii) Looking at (1.25), we see that the more risk-averse the agent is (that is, the larger γ), the less he invests in the risky assets - evidently sensible. If we took the simple special case where V were diagonal, we see that the position in asset j is

$$\frac{\mu^j - (1+r)S_0^j}{\gamma V_{jj}},$$

proportional to the *excess mean return* $\mu^j - (1+r)S_0^j$ of asset j, that is, the average amount by which investing in asset j improves upon investing the same initial amount S_0^j in the riskless asset. We also see that the higher the variance of asset j, the less we are prepared to invest in it, again evidently sensible.

1.3 Capital Asset Pricing Model (CAPM).

Simple algebra takes us from the result of the previous example all the way to a Nobel prize-winning discovery! For each asset i, we define the *beta* of that asset

$$\begin{aligned}\beta_i &\equiv \frac{\mathrm{cov}(S_1^i, \theta_M \cdot S_1)}{\mathrm{var}(\theta_M \cdot S_1)} \\ &= \frac{(V\theta_M)^i}{\theta_M \cdot V\theta_M} \\ &= \frac{(\mu - (1+r)S_0)^i}{\theta_M \cdot V\theta_M}\end{aligned}$$

Now consider an agent who at time 0 has wealth $\theta_M \cdot S_0$; if he chooses to invest that wealth in the risky assets according to the market portfolio, he will have $\theta_M \cdot S_1$ at time 1, a random wealth with mean

$$\mu_M \equiv \theta_M \cdot \mu.$$

On the other hand, he could invest his initial wealth in the riskless asset, in which case at time 1 he would have a certain

$$(1+r)\theta_M \cdot S_0.$$

On average, then, by investing in the risky assets he is better off by

$$\mu_M - (1+r)\theta_M \cdot S_0 = \theta_M \cdot V\theta_M.$$

Simply rearranging the definition of β_i then,

$$\begin{aligned}\mu^i - (1+r)S_0^i &= \text{excess return of asset } i \\ &= \beta_i(\mu_M - (1+r)\theta_M \cdot S_0) \qquad (1.26) \\ &= \beta_i \times (\text{excess return of the market portfolio})\end{aligned}$$

Is this a profound result, or merely a tautologous reworking of the definition of β_i? It is both; the profundity lies in the fact that (1.26) expresses a relation between on the one hand the mean rates of return of individual assets and of the market portfolio, and on the other, the variances and covariances of asset returns, *which could all be estimated very easily from market data*[7], thereby providing a test of the CAPM analysis. It is rare to find a verifiable prediction from economic theory; sadly, it turns out in practice to be very hard to make reliable estimates of rates of return.

[7] The market portfolio would be taken to be a major share index.

ASIDE. Let us look at some typical figures to substantiate the preceding comment. Suppose that we model an asset as a log Brownian motion[8], so that $\log(S_t) = \sigma W_t + \alpha t$ for some constants α and σ. If one unit of time corresponds to one year, then typical values for α would be of the order of $10\% - 30\%$ and for σ of the order of $20\% - 80\%$. To fix ideas, let us suppose that $\sigma = 0.2 = \alpha$. Suppose we observe the asset at intervals of $\delta = 1/n$, so that we see a sequence of IID random normal variables $\log(S_{j\delta}) - \log(S_{(j-1)\delta}) \equiv X_j$ for $j = 1, 2, \ldots, N$. The common mean is $\alpha\delta$ and the common variance is $\sigma^2\delta$. For concreteness, we shall suppose that $n = 250$ as there are approximately 250 trading days in any year. Then the MLE of α based on these observations will be

$$\hat{\alpha} = n\bar{X} \equiv \frac{n}{N}\sum_{j=1}^{N} X_j \sim N(\alpha, \frac{n\sigma^2}{N}).$$

If we want N to be so large that the 95% confidence interval is $[\hat{\alpha} - 0.01, \hat{\alpha} + 0.01]$ (roughly, we are 95% certain that the mean is between 0.19 and 0.21), then we need

$$1.96\sigma\sqrt{\frac{n}{N}} = 0.01,$$

or again

$$\frac{N}{n} = (\frac{196}{5})^2 = 1536.64;$$

so for this degree of certainty concerning the rate of return of the asset, *we need to observe the process for over 1500 years!!* We can similarly estimate the accuracy of the MLE of the variance, and this is typically much better; more to the point, by increasing the frequency of observation we can improve the accuracy of our estimate of σ arbitrarily, but *increasing the frequency of observation does not affect the accuracy of our estimate of* α, since the sum of an IID Gaussian sample is sufficient for the mean.

Again, let us consider how long we would have to observe in order for 0 to be outside the one-sided 95% confidence interval with probability at least 0.95. The 95-percentile of the standard Gaussian distribution is at $\theta = 1.6449$, and we shall take the MLE

$$\hat{\alpha} = \frac{\log(S_T) - \log(S_0)}{T} \sim N(\alpha, \sigma^2/T)$$

of α. We are therefore asking how big T must be in order that

$$\mathbb{P}(\hat{\alpha} > \theta\sigma/\sqrt{T}) = 0.95.$$

In terms of a standard normal variable Z, we want

$$0.95 = \mathbb{P}(\alpha + Z\sigma/\sqrt{T} > \theta\sigma/\sqrt{T}) = \mathbb{P}(Z > \theta - \alpha\sqrt{T}/\sigma),$$

implying that $\alpha\sqrt{T}/\sigma = 2\theta$; when $\alpha = \sigma = 0.2$, this means that *we have to wait 10.823 years even to be 95% certain that* $\alpha > 0$ *with probability 0.95!*

[8] We will have more to say on Brownian motion later; the remarks here will be entirely self-contained

The moral of these little calculations is that we know the mean with very little precision; so while the optimal portfolios calculated above may be correct if we happen to *know* the true mean, we cannot expect that plugging a (very erroneous) point estimate into a formula calculated on the assumption that the mean was known will necessarily be much good in practice. It is not; and we really need to redo the entire analysis on the assumption that the mean μ of S_1 is itself normally distributed. Not surprisingly, it will turn out that under this more realistic assumption the agent will invest less in the risky asset. But the size of the effect on the maximised utility is not as bad as we might expect from the error in the estimate of μ; the reason for this is that the expected utility depends smoothly on the portfolio weights chosen, so at the maximum a small change of $O(h)$ in the portfolio weights produces a (much smaller) change of $O(h^2)$ in the expected utility!

1.4 Equilibrium pricing.

We cannot pass by this place without pausing to see how central ideas of economic equilibrium work in the particular example we are studying. We have been looking at a vector of n assets whose values at time 1 are multivariate normal random variables, and whose values at time 0 are given constants; but where did those constants come from? How were they determined? Why are they not just the means of the contingent claims at time 1? An economist would answer these questions by saying that the time-0 prices are *equilibrium* prices, determined by the agents in the market and their interaction. If the time-0 prices are given to us, we have just seen how to compute for each agent the optimal holding of the various assets; the central idea of equilibrium analysis (due to Arrow and Debreu) is that *we now adjust the prices until the markets clear;* that is, the supply and demand are matched.

So let us suppose in the context of the previous example that each of the risky assets, thought of as shares in various enterprises, are in unit net supply; there is one unit of enterprise 1, one of enterprise 2, ... Concerning the riskless asset, let us suppose it is in *zero* net supply; riskless borrowing requires one agent to give another a promise to pay a named sum at time 1, and the total of all such promises held exactly equals the total of all such promises made. Since there is an indeterminacy of scale in the time-0 prices (prices could be in USD, EUR, JPY, ...), let us suppose that $S_0^0 = 1$. Concerning the agents who make up the market, we shall suppose that there are K of them, and that each has a CARA utility, the k^{th} having coefficient of absolute risk aversion γ_k.

From (1.25), agent k is going to hold the portfolio

$$\theta_k = \gamma_k^{-1}\theta_M$$

in the risky assets, so that the total holdings of all the agents will be

$$\sum_{k=1}^{K} \theta_k = \Gamma^{-1}\theta_M = \Gamma^{-1}V^{-1}(\mu - (1+r)S_0),$$

where $\Gamma^{-1} = \sum_k \gamma_k^{-1}$. Market clearing therefore requires that

$$\mathbf{1} = \sum_{k=1}^{K} \theta_k = \Gamma^{-1} V^{-1} (\mu - (1+r) S_0),$$

and hence we deduce that

$$S_0 = (\mu - \Gamma V \mathbf{1}) / \mu^0. \tag{1.27}$$

2 Arbitrage pricing theory in discrete time.

Orientation

In the examples studied in Chapter 1, we worked with a single period model and Gaussian returns; in this Chapter, we shall drop these assumptions and investigate very simple discrete-time models, beginning with single-period models, and later moving on to multi-period models. We shall continue with the notation already introduced: there are d risky assets, the price of the k^{th} at (integer) time t being denoted S_t^k, and we shall also suppose that there is a *strictly positive* zero'th asset, referred to as a *numeraire* asset. We use the notations

$$S_t \equiv (S_t^1, \ldots, S_t^d)^T, \quad \bar{S}_t \equiv (S_t^0, S_t^1, \ldots, S_t^d)^T.$$

We are going to consider portfolio processes

$$\bar{\theta}_t = (\theta_t^0, \theta_t^1, \ldots, \theta_t^d)^T$$

where we interpret θ_t^j as the number of units of asset j chosen on day $t-1$ and held through to day t. In this notation then, the gain realised on day t for the investment from day $t-1$ through to day t will simply be

$$\bar{\theta}_t \cdot \Delta \bar{S}_t \equiv \bar{\theta}_t \cdot (\bar{S}_t - \bar{S}_{t-1}). \tag{2.1}$$

Naturally, we would expect that when we choose the portfolio $\bar{\theta}_t$, on day $t-1$, we should only be able to use information available to us on day $t-1$; we shall call such a process *previsible*, though the formal definition will have to wait til later when we have some more notation at our disposal. We shall also restrict attention to *self-financing* portfolios.

Definition 2.1. *A portfolio* $(\bar{\theta}_t)_{t\geq 0}$ *is called* self-financing *if*

$$(\bar{\theta}_{t+1} - \bar{\theta}_t) \cdot \bar{S}_t = 0. \tag{2.2}$$

REMARK. The interpretation is clear; when we change our portfolio at time t from $\bar{\theta}_t$ to $\bar{\theta}_{t+1}$ there should be no change in our wealth. Using this, it is easy to verify that for a self-financing portfolio process with associated wealth process $w_t \equiv \bar{\theta}_t \cdot \bar{S}_t$ we shall have

$$w_T - w_0 = \sum_{t=1}^{T} \bar{\theta}_t \cdot \Delta \bar{S}_t \equiv (\bar{\theta} \cdot \bar{S})_T \, ; \tag{2.3}$$

that is, *the change in wealth equals the total gains from trade.* The process $(\bar{\theta} \cdot \bar{S})$ is called the *gains-from-trade process.* Notice the notational difference $\bar{\theta}_t \cdot \bar{S}_t \neq (\bar{\theta} \cdot \bar{S})_t$, though as we see from (2.3) for a self-financing portfolio the two sides differ only by w_0.

We introduce the notation

$$\tilde{S}_t \equiv \bar{S}_t / S_t^0, \qquad \tilde{w}_t \equiv w_t / S_t^0 \tag{2.4}$$

for the assets and wealth denominated in units of the numeraire asset S^0. We shall speak of the *discounted* wealth process $\tilde{w}$, because often the numeraire asset will be a bank account growing at a constant rate. We have the following simple observation.

Proposition 2.2. *The discounted gains from trade of a self-financing portfolio process satisfies*

$$\tilde{w}_t - \tilde{w}_{t-1} = \bar{\theta}_t \cdot (\tilde{S}_t - \tilde{S}_{t-1}) = \sum_{j=1}^{d} \theta_t^j (\tilde{S}_t^j - \tilde{S}_{t-1}^j) \tag{2.5}$$

PROOF. We have

$$\begin{aligned} \tilde{w}_t - \tilde{w}_{t-1} &= \frac{\bar{\theta}_t \cdot \bar{S}_t}{S_t^0} - \frac{\bar{\theta}_{t-1} \cdot \bar{S}_{t-1}}{S_{t-1}^0} \\ &= \frac{\bar{\theta}_t \cdot \bar{S}_t}{S_t^0} - \frac{\bar{\theta}_t \cdot \bar{S}_{t-1}}{S_{t-1}^0} \\ &= \bar{\theta}_t \cdot \Delta \tilde{S}_t. \end{aligned}$$

The second equality of (2.5) is clear because $\tilde{S}_t^0 = 1$ for all t. □

It is now clear that any statement about wealth processes and self-financing portfolios can be translated into equivalent statements in terms of discounted wealth and discounted assets. This has a number of advantages:

(i) The numeraire asset in discounted terms is identically 1;

(ii) *Any* portfolio process $\theta_t = (\theta_t^1, \dots, \theta_t^d)$ can be turned into a self-financing portfolio by defining
$$\theta_t^0 = \theta_{t-1}^0 + (\theta_{t-1} - \theta_t) \cdot (\tilde{S}_t^1, \dots, \tilde{S}_t^d)^T. \tag{2.6}$$

(iii) The key concept of an *arbitrage* is stated in terms of the discounted assets.

Definition 2.3. *A (self-financing previsible) portfolio process $(\bar{\theta}_t)_{t\geq 0}$ is an* arbitrage *for the asset price process $(\bar{S}_t)_{t\geq 0}$ over $[0,T]$ if*

$$P[\, \tilde{w}_T - \tilde{w}_0 \geq 0 \,] = 1, \tag{2.7}$$
$$P[\, \tilde{w}_T - \tilde{w}_0 > 0 \,] > 0. \tag{2.8}$$

REMARK. Why did we insist on working with discounted wealth, and not with wealth? If we simply replaced w with $\tilde{w}$, there would always be an arbitrage as soon as there was an interest-bearing bank account! Indeed, it we put 1 unit of money in the bank account at time 0, then by time T it will be worth $(1+r)^T > 1$ and this would be an arbitrage if we restated (2.3) in terms of w, not $\tilde{w}$. The point of an arbitrage is that it is an opportunity to do better than the riskless bank account with certainty, and if you could do that, then you would borrow a vast sum of money, invest it in the arbitrage, and then when it came to pay back your borrowings you could be sure to do this, with probably some additional wealth also.

So what now? For the rest of this discussion, we will just assume that we have rebased everything in terms of the numeraire asset zero, so that $S^0_t = 1$ for all t, so that the gains from trade of a self-financing portfolio will be just

$$(\theta \cdot S)_T = \sum_{t=1}^{T} \theta_t \cdot (S_t - S_{t-1}). \tag{2.9}$$

An agent who starts from wealth w_0 and plays the market until time T may therefore generate any wealth in

$$\mathcal{A} \equiv \{w_0 + (\theta \cdot S)_T : \theta \text{ previsible, self-financing}\} \tag{2.10}$$

which is manifestly an affine space. Now we are going to invoke the results on reservation and utility-indifference prices from Section 1.1; if there is some $X^* \in \mathcal{A}$ which maximizes $E[U(X)]$, then as at (1.8), we have

$$E[U'(X^*)(\theta \cdot S)_T] = 0 \tag{2.11}$$

for any θ. This then leads us to define a *pricing measure* Q equivalent to P by the recipe

$$\frac{dQ}{dP} \propto U'(X^*). \tag{2.12}$$

Straightforward results from martingale theory allow us to conclude from (2.11) that *S is a Q-martingale*. If there is no maximizer of $E[U(X)]$, then as we shall see[9], there is an arbitrage. These statements constitute the very important Fundamental Theorem of Asset Pricing (FTAP).

2.1 Single-period FTAP.

We are going to establish the FTAP in the simplest setting, where there are just two times, time 0 and time 1. When returns were jointly Gaussian, we already studied what happens if we have an agent with CARA utility, but now we are not making any joint Gaussian assumption, but are working with a completely general joint distribution of returns.

Definition 2.4. *A probability $\mathbb{Q}$ is* absolutely continuous *with respect to $\mathbb{P}$ if there is an integrable non-negative function f such that for all events A*

$$\mathbb{Q}(A) = \int_A f \, d\mathbb{P}.$$

The function f is referred to as the density *of $\mathbb{Q}$ with respect to $\mathbb{P}$, and the notation*

$$f = \frac{d\mathbb{Q}}{d\mathbb{P}}$$

is used. When $f > 0$ $\mathbb{P}$-almost-surely, we say that $\mathbb{Q}$ and $\mathbb{P}$ are equivalent.

[9].. for the special case of CARA utility ...

REMARKS. The Radon-Nikodym Theorem states that $\mathbb{Q}$ is absolutely continuous with respect to $\mathbb{P}$ if and only if every $\mathbb{P}$-null event (that is, an event A for which $\mathbb{P}(A) = 0$) is also $\mathbb{Q}$-null.

Theorem 2.5. *Assume that $S_t^0 = 1$ for all $t \geq 0$. Then the following are equivalent:*

(i) There is no arbitrage;

(ii) There exists some probability $\mathbb{Q}$ equivalent to $\mathbb{P}$ such that

$$E^{\mathbb{Q}}\big[S_1\big] = S_0 \tag{2.13}$$

When this condition holds, we may take

$$\frac{d\mathbb{Q}}{d\mathbb{P}} \propto \exp(-\theta \cdot S_1 - \tfrac{1}{2}|S_1|^2) \tag{2.14}$$

for some $\theta \in \mathbb{R}^n$.

PROOF. Write $X \equiv S_1 - S_0$ for brevity. We shall without loss of generality make the non-degeneracy assumption that $\mathbb{P}(\theta \cdot X = 0) < 1$ for all non-zero $\theta \in \mathbb{R}^n$, for otherwise X lies in some proper subspace of R^d, and there are linear dependencies among the assets. We could then discard redundant assets and reduce to a set for which the non-degeneracy assumption holds.

$(ii) \Rightarrow (i)$: If θ were an arbitrage, we shall have from (2.13) that

$$E^{\mathbb{Q}}\big[\theta \cdot S_1\big] = \theta \cdot S_0, \tag{2.15}$$

and the left-hand side is non-negative, the right-hand side is non-positive, so both must be zero. But this means that $\mathbb{P}(\theta \cdot S_1 = 0) = 1 = \mathbb{P}(\theta \cdot S_0 = 0)$, so θ is not an arbitrage.

$(i) \Rightarrow (ii)$: We must prove the existence of some equivalent martingale measure assuming that there is no arbitrage. This is somewhat more involved, but we will actually construct such a measure, using the principle of maximisation of expected utility; the result (1.8) is in effect what we need. For this, define the function

$$\theta \mapsto \varphi(\theta) \equiv \frac{E \exp(-\theta \cdot X - \frac{1}{2}|X|^2)}{E \exp(-\frac{1}{2}|X|^2)}.$$

This function is finite-valued[10], non-negative, continuous, convex and differentiable. If $\inf_\theta \varphi(\theta)$ is attained at some θ^*, then by differentiating we learn that

$$E\big[\exp(-\theta^* \cdot X - \tfrac{1}{2}|X|^2)\, X\big] = 0,$$

[10]The interpretation of the proof in terms of a utility-maximisation argument is far more direct if we had simply used $\varphi_0(\theta) = E \exp(-\theta \cdot X)$, for then we are *literally* maximising the CARA utility of $\theta \cdot X$. The snag is that this expectation need not be finite for all θ, whereas the definition given for φ is certain to be finite-valued. If φ_0 were finite-valued, we could have used this instead of φ.

so defining $\mathbb{Q}$ by

$$\begin{aligned} \frac{d\mathbb{Q}}{d\mathbb{P}} &= c^{-1} \exp(-\theta^* \cdot X - \tfrac{1}{2}|X|^2) \\ c &= E \exp(-\theta^* \cdot X - \tfrac{1}{2}|X|^2) \end{aligned}$$

gives (2.13).

So the only thing that could go wrong is that the infimum $\inf_\theta \varphi(\theta)$ is *not* attained. We shall now prove that this can only happen if there is arbitrage, in contradiction of our hypothesis; it follows then that the infimum *is* attained, and we *do* have an equivalent martingale measure.

If we consider the sets

$$F_\alpha \equiv \{\theta \in \mathbb{R}^d : |\theta| = 1, \varphi(\alpha\theta) \leq 1\}, \quad (\alpha \geq 0)$$

we see that these are closed subsets of a compact subset of $\mathbb{R}^d$, and it is not hard to see[11] that $F_\beta \subseteq F_\alpha$ for all $0 \leq \alpha \leq \beta$. By the Finite Intersection Property, *either* the intersection $\cap_\alpha F_\alpha$ is non-empty, *or* for some α, $F_\alpha = \emptyset$. If the infimum is not attained, then it is less than $1 = \varphi(0)$ and there exist a_k such that $\varphi(a_k)$ decrease to the infimum; these a_k cannot be bounded, else some subsequence would converge to a point where the infimum is attained, so we must have a sequence of points tending to infinity where φ is less than 1, and so $\cap_\alpha F_\alpha$ is non-empty. Thus there is some unit vector a such that

$$\varphi(ta) = \frac{E \exp(-ta \cdot X - \frac{1}{2}|X|^2)}{E \exp(-\frac{1}{2}|X|^2)} \leq 1$$

for all $t \geq 0$, and this can only happen if

$$\mathbb{P}[a \cdot X < 0] = 0.$$

Thus

$$a \cdot X = a \cdot (S_1 - S_0) \geq 0, \tag{2.16}$$

and with positive $\mathbb{P}$-probability (non-degeneracy!) this inequality is strict. We therefore take a portfolio consisting of a in the risky assets, and $-a \cdot S_0$ in the riskless asset; at time 0 this is worth nothing, and at time 1 it is worth $a \cdot X$. Because of (2.16), this portfolio is an arbitrage.

□

REMARKS. The assumption that S^0 is identically 1 is restrictive, asymmetric and unnecessary; the notion of arbitrage for any $(d+1)$-vector $\bar{S}$ of assets does not require this, and in fact we can deduce a far more flexible form of the above result.

Corollary 2.6. (Fundamental Theorem of Asset Pricing, 0). *Let $(\bar{S}_t)_{t \in \mathbb{Z}^+}$ be a $(d+1)$-vector of asset prices, and assume:*

[11]... by convexity of φ and the fact that $\varphi(0) = 1$ ...

ASSUMPTION (N): Among the assets $S^0, \ldots, S^d$, there is one which is strictly positive.

Select a strictly positive asset N from the $d+1$ assets. Then the following are equivalent:

(i) There is no arbitrage;

(ii) There exists some probability $\mathbb{Q}$ equivalent to $\mathbb{P}$ such that

$$E^Q\left[\, S_1/N_1 \,\right] = S_0/N_0. \tag{2.17}$$

The probability $\mathbb{Q}$ is referred to as an equivalent martingale measure *(or sometimes an equivalent martingale probability.)*

PROOF. It is evident that $\bar{\theta}$ is an arbitrage for $\bar{S}$ if and only if it is an arbitrage for $\tilde{S}$, where we define

$$\tilde{S}^i_t \equiv S^i_t/N_t.$$

The result follows by applying Theorem 2.5 to $\tilde{S}$ (assume without loss of generality that $N = S^0$.)

□

REMARKS. (i) The strictly positive asset N used above is referred to as a *numeraire.* We have often considered a situation where there is a single riskless asset (referred to variously as the money-market account, the bond, the bank account, ..) in the market, and it is very common to use this asset as numeraire. It turns out that this will serve for our present applications, but there are occasions when it is advantageous to use other numeraires. Note that the Fundamental Theorem of Asset Pricing does *not* require the existence of a riskless asset.

(ii) Note that the Fundamental Theorem of Asset Pricing does not make any claim about uniqueness of $\mathbb{Q}$ when there is no arbitrage. This is because situations where there is a unique $\mathbb{Q}$ are rare and special; when $\mathbb{Q}$ is unique, the market is called *complete.* We shall have more to say about this presently.

(iii) Theorem 2.6 tells us in a single-period setting that when there is no arbitrage, there exists an equivalent martingale measure. The meaning of the term 'equivalent' has been defined, but we need to explain what a martingale is.

Definition 2.7. *A stochastic process $(X_n)_{n\geq 0}$ is called*[12] *a* supermartingale *if for each $n \geq 0$*

$$X_n \geq E[X_{n+1}|X_0, X_1, \ldots, X_n]. \tag{2.18}$$

If $(-X_n)_{n\geq 0}$ is a supermartingale, the process $(X_n)_{n\geq 0}$ is called a submartingale*, and a process which is both a supermartingale and a submartingale is called a* martingale*.*

[12]We give a definition of a martingale which is slightly less general than the correct one; see, for example, the book of Williams () for the full story, which requires concepts from measure-theoretic probability. The current account tells no lies, however; anything which is a martingale in the sense of the definition we have given here will indeed be a martingale according to the proper definition.

Of course, X is a martingale if the inequality (2.18) is an equality for all n. Equipped with this terminology, we can now state the full discrete-time Fundamental Theorem of Asset Pricing.

Theorem 2.8. (Fundamental Theorem of Asset Pricing). *Let $(\bar{S}_t)_{t\in\mathbb{Z}^+}$ be a $(d+1)$-vector of asset prices, and assume:*

ASSUMPTION (N): Among the assets $S^0, \ldots, S^d$, there is one which is strictly positive.

Select a strictly positive asset N from the $d+1$ assets. Then the following are equivalent:

(i) There is no arbitrage;

(ii) There exists some probability $\mathbb{Q}$ locally[13] *equivalent to $\mathbb{P}$ such that*

$$\left(\frac{\bar{S}_t}{N_t}\right)_{t\in\mathbb{Z}^+} \quad \text{is a } \mathbb{Q}\text{-martingale.} \tag{2.19}$$

The probability $\mathbb{Q}$ is referred to as an equivalent martingale measure.

(iv) We have just proved a very general form of the Fundamental Theorem of Asset Pricing in discrete time, though only in the single-period situation. Its extension to the multi-period situation is not essentially difficult, though there are some technical points to be handled[14] to give the result in its simplest and strongest form. There is an analogous result in continuous time, but this is quite deep and subtle (see Delbaen & Schachermayer ()); the first subtlety is in framing the definition of arbitrage correctly! We shall not dwell on the details of extending to the multi-period case in a general context (see Rogers () for these), but shall for the rest of this chapter consider only very simple and explicit models where we can characterise the equivalent martingale measure completely, and perform calculations.

(v) To link the statement of Theorem 2.8 with the discussion at the beginning of this Chapter, we need to observe that if we define

$$Z_t \equiv \left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_t},$$

then Z is a $\mathbb{P}$-martingale, and for any $\mathbb{Q}$ martingale M the product ZM is a $\mathbb{P}$-martingale. None of these facts is hard to prove, but they do require a basic familiarity with definition and properties of conditional expectation which we are not here assuming. But given these facts, we then define

$$\zeta_t = \frac{Z_t}{N_t},$$

and the pricing expression (**??**) is seen to amount to the same as (2.19).

[13] If $\mathcal{F}_T$ denotes the set (σ-field) of all events which are known at time T, then $\mathbb{Q}$ will be equivalent to $\mathbb{P}$ on $(\Omega, \mathcal{F}_T)$ for every T. It can happen that $\mathbb{Q}$ is not equivalent to $\mathbb{P}$ on the set of *all* events, which is why we have to qualify the statement with the adjective 'locally'.

[14] .. relating to measurable selection of maximising portfolios in the case of non-uniqueness ...

Aside: axiomatic derivation of the pricing equation. For this little aside, we use the language of measure-theoretic probability, but this is not essential. We shall show how the pricing expression (**??**) can be derived very quickly from four simple axioms which a family of pricing operators should naturally obey.

Suppose that we have pricing operators $(\pi_{tT})_{0 \le t \le T}$ for contingent claims; if Y is some $\mathcal{F}_T$ measurable contingent claim to be paid at time T, the time-t 'market' price will be

$$\pi_{tT}(Y),$$

which may be random, but must be $\mathcal{F}_t$-measurable.

We shall assume that the pricing operators $(\pi_{tT})_{0 \le t \le T}$ satisfy certain axioms:

(A1) Each π_{tT} is a bounded positive linear operator from $L^\infty(\mathcal{F}_T)$ to $L^\infty(\mathcal{F}_t)$;

(A2) If $Y \in L^\infty(\mathcal{F}_T)$ is almost surely 0, then $\pi_{0T}(Y)$ is 0, and if $Y \in L^\infty(\mathcal{F}_T)$ is non-negative and not almost surely 0, then $\pi_{0T}(Y) > 0$;

(A3) For $0 \le s \le t \le T$ and each $X \in L^\infty(\mathcal{F}_t)$ we have

$$\pi_{st}(X\pi_{tT}(Y)) = \pi_{sT}(XY);$$

(A4) For each $t \ge 0$ the operator π_{0t} is bounded monotone-continuous - which is to say that if $Y_n \in L^\infty(\mathcal{F}_t)$, $|Y_n| \le 1$ for all n, and $Y_n \uparrow Y$ as $n \to \infty$, then $\pi_{0t}(Y_n) \uparrow \pi_{0t}(Y)$ as $n \to \infty$.

Axiom (A1) says that the price of a non-negative contingent claim will be non-negative, and the price of a linear combination of contingent claims will be the linear combination of their prices - which are reasonable properties for a market price. Axiom (A2) says that a contingent claim that is almost surely worthless when paid, will be almost surely worthless at all earlier times (and conversely) - again reasonable. The third axiom, (A3), is a 'consistency' statement; the market prices at time s for XY at time T, or for X times the time-t market price for Y at time t, should be the same, for any X which is known at time t. The final axiom is a natural 'continuity' condition which is needed for technical reasons.

Let's see where these axioms lead us. Firstly, for any $T > 0$ we have that the map

$$A \mapsto \pi_{0T}(I_A)$$

defines a non-negative measure on the σ-field $\mathcal{F}_T$, from the linearity and positivity (A1) and the continuity property (A4). Moreover, this measure is absolutely continuous with respect to $\mathbb{P}$, in view of (A2). Hence there is a non-negative $\mathcal{F}_T$-measurable random variable ζ_T such that

$$\pi_{0T}(Y) = E[\zeta_T Y]$$

for all $Y \in L^\infty(\mathcal{F}_T)$. Moreover, $\mathbb{P}[\zeta_T > 0] > 0$, because of (A2) again. Now we exploit the consistency condition (A3); we have

$$\pi_{0t}(X\pi_{tT}(Y)) = E[X\zeta_t\pi_{tT}(Y)] = \pi_{0T}(XY) = E[XY\zeta_T].$$

Since $X \in L^\infty(\mathcal{F}_t)$ is arbitrary, we deduce that

$$\pi_{tT}(Y) = E_t[Y\zeta_T]/\zeta_t,$$

which shows that the pricing operators π_{st} are actually given by a risk-neutral pricing recipe, with the state-price density process ζ.

3 The (Cox-Ross-Rubinstein) binomial model.

The model we are about to study is the most important example to understand in the whole subject. It is technically very simple, and can be analysed using only arithmetic; it displays the main qualitative features of the more sophisticated Brownian model which we discuss later; and it serves as a computational workhorse for computing (approximately) the prices of derivatives written on (log-Brownian) shares.

The story is very simple. We start with a two-period model, time 0 and time 1. There are two assets in the model, S^0 and S^1, and both are worth 1 at time 0;

$$S_0^0 = 1 = S_0^1.$$

At time 1, asset 0 (the riskless asset, or *bond* for short) is worth the sure amount $1+r$, whereas asset 1 (the share) is worth u if the time period was good, and $d < u$ if the time period was bad. We shall assume that[15]

$$d < 1 + r < u$$

and we assume that the period was good with probability $p_0 \equiv 1 - q_0 \in (0, 1)$.

[15]If $1+r$ were not between u and d, then one of the assets would be dominated, and no-one would hold it.

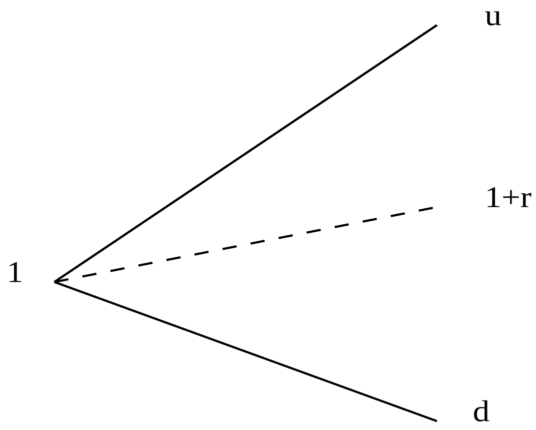

Figure 2: Single-period binomial model.

The first question is '*How do we price contingent claims in this model?*' So suppose a contingent claim Y delivers amount a in a good year, and amount b in a bad year; what would be the price of this contingent claim at time 0? Because of a special feature of this model, we can answer this explicitly. The point is that we can build a *replicating portfolio* for Y, that is, a portfolio of x bonds and y shares which exactly replicates the payoff of Y at time 1. How would this be done? We need to choose x and y so that

$$\begin{aligned} x(1+r) + yu &= a, \\ x(1+r) + yd &= b; \end{aligned}$$

Solving the simultaneous linear equations gives the replicating portfolio to be

$$\begin{aligned} x &= \frac{ub - ad}{(u-d)(1+r)} \\ y &= \frac{a-b}{u-d} \end{aligned}$$

and the time-0 cost of this replicating portfolio is

$$x + y = \frac{(u-(1+r))b + (1+r-d)a}{(u-d)(1+r)} = \frac{1}{1+r}(pa+qb), \tag{3.1}$$

where

$$p = 1 - q = \frac{1+r-d}{u-d}\ . \tag{3.2}$$

The cost of this replicating portfolio is the *only* possible time-0 price for the contingent claim Y, because if it were to trade for any other price there would be riskless profit to be made. The market is said to be *complete* because any contingent claim can be perfectly replicated by taking positions in the traded assets[16]. It is important to notice the interpretation of (3.1);

$$x + y = E^{\mathbb{Q}}[Y/(1+r)] \tag{3.3}$$

where under the probability $\mathbb{Q}$ the probability of a good year is p; $\mathbb{Q}$ is the equivalent martingale measure whose existence was guaranteed by Theorem 2.8.

Remarks. Previously in Section 1.1 we arrived at the notions of bid and ask prices, and of marginal prices, by considering how an individual agent would optimally choose his portfolio, yet here we have arrived at a unique price for a contingent claim Y by a completely different replication argument. Let us see how the utility-maximisation approach leads to exactly the same conclusions.

An agent who begins with initial wealth w_0 can create any wealth at time 1 which has the representation

$$w_1 = w_0 S_1^0 + t(S_1^1 - S_1^0) = w_0(1+r) + t\{S_1^1 - (1+r)\}$$

for some $t \in \mathbb{R}$, so the set of all achievable wealths is an affine set. The optimisation problem (1.7) is therefore to

$$\max_t \Big[p_0 U(w_0(1+r) + t(u-1-r)) + q_0 U(w_0(1+r) + t(d-1-r)) \Big];$$

differentiating and setting to zero determines the optimal t^* via the equation[17]

$$p_0(u-1-r)U'(w_g) = q_0(1+r-d)U'(w_b)$$

where we abbreviate the wealths resulting at the ends of good and bad periods:

$$\begin{aligned} w_g &= w_0(1+r) + t^*(u-1-r), \\ w_b &= w_0(1+r) + t^*(d-1-r). \end{aligned}$$

The expression (1.12) for the price of the contingent claim Y which delivers a in a good period and b in a bad period therefore gives us

$$\begin{aligned} \frac{E[\,U'(w_1)Y\,]}{(1+r)E[\,U'(w_1)\,]} &= \frac{p_0 U'(w_g)a + q_0 U'(w_b)b}{\{p_0 U'(w_g) + q_0 U'(w_b)\}(1+r)} \\ &= \frac{(1+r-d)a + (u-1-r)b}{(u-d)(1+r)}, \end{aligned}$$

exactly as at (3.1). Notice the extra factor $(1+r)$ in the denominator - this is because if the agent pays out π at time 0 to buy Y, then his time-1 wealth is reduced by $\pi(1+r)$.

[16] We do not expect completeness to hold in any practical situation, but it is a simple enough model to allow us to get started.

[17] There will be a unique optimal t^* provided that U is strictly concave, C^1, and satisfies $\sup U'(x) = \infty$, $\inf U'(x) = 0$.

The bid and ask prices for Y for any agent will agree with the price found by the above argument, essentially because Y can be replicated.

So we have seen two arguments to arrive at the price of a general contingent claim Y, the first one proved in a few lines by solving two simultaneous linear equations, the other by more involved expected-utility maximisation. The simplicity of the first is an attractive feature, but we have to realise that the argument *only* works because of the completeness of this very simple market; if we consider the slightly more complicated situation where a period can be good, bad or indifferent, then the entire first argument collapses, and *only* the expected-utility maximisation approach is open to us. It has to be admitted straight away that there are now almost no examples where anything can be computed in closed form, and we are forced to resort to numerical methods. In such a situation, the price of a contingent claim, the *equilibrium* price, will depend on the preferences of the agents, and may in principle fail to be unique.

3.1 The multi-period binomial model.

The single-period binomial model just described extends easily to a multiperiod model. To see how, let us just examine the special case of the two-period model, and consider how we may replicate (and hence price) a contingent claim that will be paid at time 2.

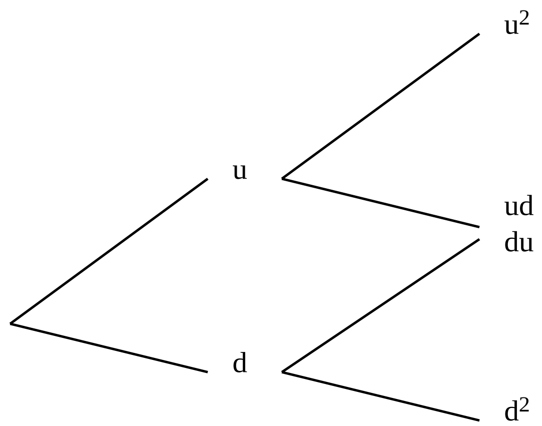

Figure 3: Two-period binomial model.

At time 2, there are 4 nodes of the tree, and to replicate a contingent claim which takes values y_{uu}, y_{ud}, y_{du} and y_{dd} at the four nodes (using an obvious notation), we can simply apply the preceding argument three times. Indeed, at the end of period 1, assuming that the share moved to u, we are faced with the problem of replicating the outcome of the next period, and this is the problem we have just solved; the random payment of y_{ud} if the share goes up in period 2, or y_{ud}

if it goes down, can be replicated by taking a position in the share and the bond which costs

$$y_u = \frac{py_{uu} + qy_{ud}}{1+r},$$

and similarly if the first period resulted in a downward movement of the share, we can replicate the coming period by a position which costs

$$y_u = \frac{py_{du} + qy_{dd}}{1+r};$$

see Figure 4

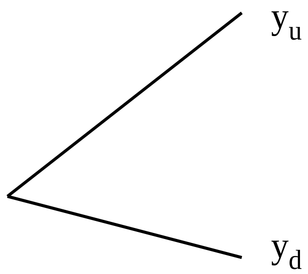

Figure 4: Two-period binomial model reduced to a single period.

Thus at time zero, we can replicate the time-2 claim by taking a position which costs

$$\frac{py_u + qy_d}{1+r}.$$

It is clear now how a multi-period version of the binomial model will be handled.

Worked examples. Let us now work out the prices of a few derivatives in the context of the two-period binomial model. We will take $u = 6/5$, $d = 9/10$, and $r = 1/10$, so that the risk-neutral 'up'-probability p is computed from (3.2) to be

$$p = \frac{1 + \frac{1}{10} - \frac{9}{10}}{\frac{6}{5} - \frac{9}{10}} = \frac{2}{3}.$$

European call option. We price a European call option with strike $K \in (81/100, 27/25)$. The holder of this option has the right but not the obligation to buy one share for the strike price K at expiry, time 2. The option will be exercised if and only if the prevailing price of the share exceeds the strike; in all other cases, the option is worthless. The payoffs at the four terminal nodes are therefore

$$\begin{aligned} y_{uu} &= u^2 - K = \frac{36}{25} - K, \\ y_{ud} &= ud - K = \frac{27}{25} - K, \\ y_{du} &= ud - K = \frac{27}{25} - K, \\ y_{dd} &= 0. \end{aligned}$$

The time-0 price of this is therefore

$$C = \frac{p^2 y_{uu} + pq y_{ud} + qp y_{du} + q^2 y_{dd}}{(1+r)^2} = \frac{100}{121}\left(\frac{28}{25} - \frac{8}{9}K\right). \tag{3.4}$$

When the strike is 1, the call is see to be worth

$$\frac{208}{1089} = 0.191001.$$

European put option. We price a European put option with strike $K \in (81/100, 27/25)$. The holder of this option has the right but not the obligation to sell one share for the strike price K at expiry, time 2. We may of course just repeat the calculation of the European call price *mutatis mutandis*; this calculation is particularly easy, as only one of the terminal nodes has positive value, and the time-0 price is easily seen to be

$$P = q^2\left(K - \frac{81}{100}\right)/(1+r)^2 = \frac{1}{9}.\frac{100}{121}\left(K - \frac{81}{100}\right) \tag{3.5}$$

However, it is more elegant to deduce the price of the put by *put-call parity*. This is based on the simple observation that if you are long one call and short one put (that is, you hold 1 call and hold -1 puts), then at time 2 the aggregate value of your position will just be one share minus the strike; the call will be worth $(S_2 - K)^+$ and the put will be worth $(K - S_2)^+$, and the difference is just $S_2 - K$. But this contingent claim could be financed by buying one share at time 0, and borrowing $K/(1+r)^2$ at time 0; thus the difference between the call and put prices at time 0 must be the cost of this alternative replicating portfolio:

$$C - P = S_0 - \frac{K}{(1+r)^2} = 1 - \frac{100}{121}K,$$

which is easily confirmed by looking at (3.4) and (3.5).

Lookback option. The lookback option pays at expiry the largest value achieved by the underlying up to that time. This contingent claim therefore has values

$$\begin{aligned} y_{uu} &= u^2 = \frac{36}{25}, \\ y_{ud} &= u = \frac{6}{5}, \\ y_{du} &= ud = \frac{27}{25}, \\ y_{dd} &= 1 \end{aligned}$$

at time 1, so its time-0 price is obtained as

$$\frac{p^2 y_{uu} + pq y_{ud} + qp y_{du} + q^2 y_{dd}}{(1+r)^2} = \frac{100}{121}\frac{283}{225} = \frac{1132}{1089} = 1.039486$$

which is only a little bigger than the value of the share.

Down-and-out call option. A down-and-out call is a European call option which is 'knocked out' (that is, becomes worthless) if at any time the price of the underlying asset falls below the *barrier*. For concreteness, let us suppose that we take the strike price to be $K = 1$, with barrier at 0.95. This time, the values at expiry are

$$\begin{aligned} y_{uu} &= u^2 - 1 = \frac{11}{25}, \\ y_{ud} &= ud - 1 = \frac{2}{27}, \\ y_{du} &= 0, \\ y_{dd} &= 0, \end{aligned}$$

so we compute the price as

$$\frac{100}{121}\left(\frac{4}{9}\frac{11}{25} + \frac{2}{9}\frac{2}{27}\right) = \frac{5152}{29403} = 0.1752202.$$

As expected, this is less than the price of the standard European call with the same strike, and is in fact about 8.26% less. There are of course a range of different types of barrier option; for example, a down-and-in call is a European call which only becomes worth a positive amount if the minimum price is below the barrier, and there is an obvious parity relation:

$$\text{Down-and-out call} + \text{Down-and-in call} = \text{call}$$

The binomial model is often used as an approximation to a continuous-time continuous-state model, and prices for that model are often approximated by the corresponding prices for the binomial model. It is worth remarking that this approximation can be *very* poor for barrier options.

****

Something about trinomial models??

****

4 Dynamic programming.

As we have already seen, solving optimisation problems is an important element of pricing, and of deriving the behaviour of agents in a market. When that market is evolving in time, we naturally find ourselves looking at various *dynamic* optimisation problems, and one of the best methods for attacking such problems is *dynamic programming.* The method is very simple conceptually, and also easy to implement; let's immediately look at an example which displays the key idea.

The trees in an orchard are arranged in a regular rectangular grid pattern, as shown in the array below; there are 10 columns each containing 11 trees, and the numbers of apples on each tree are given in the array below:

2	5	6	1	9	4	3	3	2	9
5	3	8	2	1	4	7	7	1	1
4	9	2	1	4	5	5	7	4	3
1	5	3	3	3	2	4	5	3	7
8	3	4	5	1	2	1	4	1	1
0	2	5	7	8	1	3	1	9	2
3	1	5	6	2	9	4	1	1	1
7	2	3	2	4	5	1	6	5	9
4	3	5	6	1	1	1	2	2	3
8	8	4	5	2	5	7	7	4	2
3	4	2	4	1	9	9	7	1	1

You start at the tree in the leftmost column of the array, at the tree with no apples. You may now move through the orchard from one column to the next, but you may choose at each stage whether to go straight to the tree in front of you, or the one to the right, or the one to the left. Thus at your first step, you can choose to go to a tree with 1 apple, 2 apples or 3 apples; if you choose the tree with 3 apples, then next step you get to choose a tree with 3, 4 or 5 apples. Supposing that you keep all the apples from every tree that you visit in this way, how many apples would you collect if you used the best route? If you were allowed to select the tree that you started at, which one would it be? In the next paragraph I shall explain how to solve these problems, but if you read on now before you have solved it for yourself, you should consider this a personal defeat!

The method is very easy to describe; at each tree, we compute the *value* at that tree, that is, the maximum number of apples that could be collected starting from that tree. But how? By working back from the end! So if we started at one of the trees in the tenth column, the value is just the number of apples on that tree. If we now go back to the previous column, we can work out the value at each tree; for example, at the tree with 9 apples, the value is the largest of 9+1, 9+2, 9+1, corresponding to the three possible moves which could be made from that tree. So at that tree, the value is 11.

Proceeding in this way, we can compute the value at every tree, and the answers are displayed in this array:

55	53	46	39	38	29	21	14	11	9
63	51	48	40	31	29	25	18	10	1
62	58	42	35	34	30	23	18	11	3
59	57	49	37	33	25	22	16	10	7
65	56	52	46	26	24	17	15	8	1
56	55	53	48	41	25	18	12	11	2
58	54	53	47	35	33	24	15	10	1
62	55	50	39	37	29	21	20	14	9
61	53	49	44	33	26	21	16	11	3
65	57	48	43	38	32	25	18	7	2
60	52	45	42	37	36	27	14	3	1

Now we can easily read off the answers to our questions; we would collect 56 apples in total if we used the best route through the orchard, and if we were allowed to choose where to start, we would start at either of the trees in the first column which had 8 apples on them. As we computed the values, we would also determine the best route through the orchard: starting from the tree with no apples, we would have $+ - 00 - 0 - 00$ as the optimal route, where $+$ means we move up, $-$ means that we move down, and 0 means that we go straight across.

It would of course be very easy to solve similarly a stochastic version of the above orchard problem, where with probability p you actually move to your chosen tree, otherwise you move to one of the other possible trees, choosing uniformly from the (one or two) possibilities.

Problem formulation. In general, we consider a discrete-time finite-horizon Markov decision process X taking values in a statespace $\mathcal{X}$. There is a set $\mathcal{A}$ of allowable actions; if the process is in state $x \in \mathcal{X}$ and we use action $a \in \mathcal{A}$, then we receive an immediate reward $r(x,a)$. The process then jumps to another state, according to the distribution $P(x,a;\cdot)$. The objective is to choose actions in such a way as to obtain

$$\sup E[\sum_{n=0}^{N-1} \beta^n r(X_n, a_n) + \beta^N R(X_N)\,], \tag{4.1}$$

where $\beta \in (0,1]$ is the *discount factor*, and R is the terminal reward function. The action a_n chosen at time n must of course be a function only of what has been observed by time n, and not of observations yet to come. The discount factor β reflects the common realisation that money now is better than money tomorrow, and also turns out to be mathematically convenient in the so-called *infinite horizon* case where $N = \infty$.

Problem solution. If we define the *value function* by

$$\begin{aligned} V_n(x) &\equiv \sup E[\sum_{i=n}^{N-1} \beta^{i-n} r(X_i, a_i) + \beta^{N-n} R(X_N)\,], \quad (n < N); \\ V_N(x) &= R(x), \end{aligned}$$

then the value function at successive times is related by

$$V_n(x) = \sup_{a \in \mathcal{A}} \left\{ r(x,a) + \beta \int V_{n+1}(y) P(x,a;dy) \right\} \tag{4.2}$$

This is the famous *Bellman equation* of dynamic programming, and is the most widely-used approach to solving a given dynamic programming problem. It has to be said that although text-books present diverse examples which can be solved in closed form, in virtually every example encountered outside a text-book some numerical solution will be required!

The Bellman equation says informally that the best you can do starting at time n is to to what is best if you were to take a single step and then behave optimally from time $n+1$. This is really not surprising, so we shall not try to present a proof here; you can find everything done carefully in ??????. If you take a look at that proof, the only issues are in how to set up the mathematical notation so as to establish the result. More interesting for us is how this equation can be applied to solving problems from finance.

An optimal investment/consumption problem. An agent has wealth w_n at the start of day n. He first chooses an amount c_n to consume, then divides his remaining wealth between a riskless asset (which earns simple interest at rate r per day), and a risky asset, with a random return $X_n > 0$. The X_i are independent with common distribution. Thus his wealth evolves according to

$$w_{n+1} = (w_n - c_n)\big\{(1+r)(1-\theta_n) + \theta_n X_n\big\}, \tag{4.3}$$

where θ_n is the proportion of his wealth which gets invested in the risky asset. His objective is to obtain

$$\sup E\left[\sum_{n=0}^{N-1} \beta^n U(c_n) + A\beta^N U(w_N)\right] \tag{4.4}$$

The Bellman equations for this problem will take the form

$$\begin{aligned} V_N(w) &= Au(w) \\ V_n(w) &= \sup_{c,\theta}\Big\{U(c) + \beta E(V_{n+1}((w-c)[(1+r)(1-\theta)+\theta X]))\Big\}. \end{aligned} \tag{4.5}$$

In general, it will be impossible to find any closed-form solution to the Bellman equations, but if we suppose that

$$U(x) = x^{1-R}/(1-R)$$

for some positive $R \neq 1$, then we can make some progress, in view of the scaling properties of U. Indeed, we can prove inductively that

$$V_n(w) = a_N U(w)$$

for all n. This is clearly true for $n = N$, with $a_N = A$. The inductive step requires us to solve the optimisation problem

$$\begin{aligned} V_n(w) &= \sup_{c,\theta}\Big\{U(c) + \beta a_{n+1} E[U((w-c)[(1+r)(1-\theta)+\theta X])]\Big\} \\ &= \sup_{c,\theta}\Big\{U(c) + \beta a_{n+1} (w-c)^{1-R} E[U((1+r)(1-\theta)+\theta X)]\Big\} \\ &= \sup_{c}\Big\{U(c) + \beta a_{n+1} \Gamma\, \frac{(w-c)^{1-R}}{1-R}\Big\} \\ &= \sup_{c}\Big\{U(c) + \beta a_{n+1} \Gamma U(w-c)\Big\} \end{aligned}$$

where

$$\Gamma \equiv (1-R)\sup_\theta EU((1+r)(1-\theta)+\theta X),$$

some positive constant. The optimal choice of θ therefore remains the same for all time. Thus we find that

$$\begin{aligned} V_n(w) &= w^{1-R}\sup_x\Big\{U(x) + \beta a_{n+1}\Gamma U(1-x)\Big\}, \\ &\equiv a_n U(w) \end{aligned}$$

where $x = c/w$ is optimally chosen to be

$$x = \frac{1}{1+(\beta a_{n+1}\Gamma)^{1/R}}.$$

A few calculations lead to the form of a_n:

$$a_n = (1+(\beta a_{n+1}\Gamma)^{1/R})^R.$$

More simply, this tells us that

$$a_n^{1/R} = 1 + (\beta\Gamma)^{1/R} a_{n+1}^{1/R},$$

which is a simple linear recursion for $a_n^{1/R}$. If $\beta\Gamma < 1$, then there is a finite limiting value, and otherwise the a_n tend monotonically to infinity.

Pricing an American option. We have seen in the context of the binomial model for an asset how we may price a European option (that is, one whose value depends only on the price of the underlying at expiry), and we have also shown how to compute the prices of certain path-dependent options, such as barrier options and lookbacks. But there are other very common types of option, known as *American* options, whose defining feature is that they may be exercised at *any* time prior to expiry. So the holder of an American put option with strike K and expiry T on an asset S may exercise at any time $\tau \le T$, and will at that time receive $(K - S_\tau)^+$. We insist (of course) that the time τ is chosen without knowledge of how the asset performs in the future, that is, τ is a *stopping time*. Computing the price of such an option requires us to understand the nature of the optimal exercise policy of the holder, which can be cast as a dynamic programme.

In more detail, we take the set $\mathcal{X}$ of states to be

$$\mathcal{X} = \{(n,j) : 0 \leq n \leq T, 0 \leq j \leq n\} \cup \{\Delta\},$$

where the state (n,j) corresponds to time-period n, when there have been j good periods, and exercise has yet to happen. Thus $(2,1) \in \mathcal{X}$ corresponds to the middle vertices $\{ud, du\}$ in Figure 3. The state Δ is the state we pass to once the option is exercised. Notice that all we need to keep track of is the current value of the asset, not the path by which we arrived at that value, so we can think of a *recombining* lattice instead of a non-recombining tree:

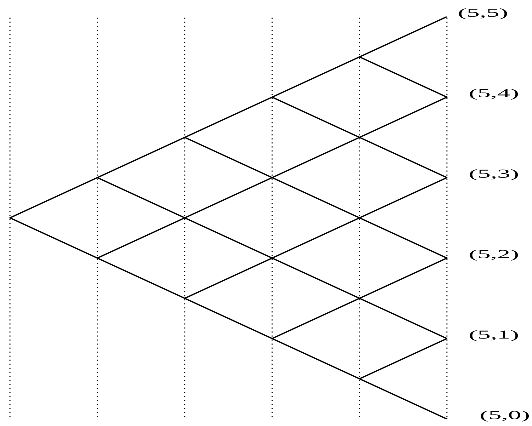

Figure 5: Recombining binomial lattice

if we had for example been trying to price an up-and-out American put, this statespace would not have been sufficient. The action space $\mathcal{A}$ consists of just two actions, s (stop) and c (continue). To cast the problem in the form (4.1), we specify

$$\begin{aligned}
\beta &= \frac{1}{1+r} \\
r((n,j),c) &= 0 \\
r((n,j),s) &= (K - S_0 u^j d^{n-j})^+ \\
r(\Delta,s) &= 0 = r(\Delta,c) \\
R(T,j) &= (K - S_0 u^j d^{T-j})^+ \\
P((n,j),c;(n+1,j+1) &= p \\
P((n,j),c;(n+1,j) &= 1-p \equiv q \\
P((n,j),s;\Delta) &= 1
\end{aligned}$$

Doing this, the Bellman equations (4.2) take the form

$$V(n,j) = \max\{\ (K - S_0 u^j d^{n-j})^+,\ \beta[pV(n+1,j+1) + qV(n+1,j)]\ \}$$

for $n < T$, with the boundary condition $V(T, j) = R(T, j)$. We can now in principle solve this system of equations numerically backwards from the final time, and this is indeed possible and often done. When we come to study the log-Brownian model for a share, we shall see that the American put option is the first example of an option whose price does not have a closed-form expression; the price has to be computed numerically, and the dynamic-programming method just described is one way of computing it. Not surprisingly, much better problem-specific numerical methods are known, and many approximations, and expansions are also known; nevertheless, the simple binomial method just described deals reasonably well with this celebrated example, and is robust enough to cope with a wide range of different examples too.

FOR LATER: Value improvement? Policy improvement?

5 Brownian motion.

MOTIVATION. The binomial model for a share is a discrete-time model, and as such it is a poor approximation to the reality of a market, where trading happens in a continuous fashion. We might try to make the binomial model describe such a market better by thinking of the time period as being very short, such as one second, or even one microsecond; if we did this, there would be a very large number of moves of the share in an hour. The log share price in the binomial model is a random walk (its steps are independent identically-distributed random variables), and in view of the Central Limit Theorem, it would not be surprising if there existed some (distributional) limit of the binomial random walk as the time periods became ever shorter. It would also be expected that the Gaussian distribution should feature largely in that limit process, and indeed it does. This Chapter introduces the basic ideas about a continuous-time process called *Brownian motion* (in American, *Wiener process*), in terms of which the most common continuous-time model of a share is defined. Using this model, various derivative prices can be computed in closed form; the celebrated Black-Scholes formula for the price of a European call option is the prime example. More generally, prices of more complicated options on the share can be computed by solving a partial differential equation (PDE).

It might be thought that we can now operate entirely with our sophisticated Brownian motion model, and forget about the much simpler binomial model, but this is far from being the case. Usually, the pricing PDE which arises for a given derivative must be solved numerically, and this requires us to discretise the PDE in some way. To do this, we can *either* discretise the derivatives in the PDE using some finite-difference approximation and then solve the resulting system of linear equations, *or* we can approximate the Brownian motion process by some random walk, and then solve the pricing problem for that process by the usual dynamic-programming methodology. The first computes an approximation to the solution to the problem we wanted to solve; the second computes the exact solution to an approximating problem. These slightly different points of view are both valuable; for basic pricing problems, the PDE technology is faster and more accurate, but if the problem is more complicated, the second approach is more robust.

End of Motivation.

———————————————————————————-

Defn; existence statement; covariance structure; scaling; time-inversion; sup infinite a.s., and consequences;

Stopping times; reflection principle; law of sup and last value

change of measure; joint law of sup and final value for drifting BM