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.

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.
- A *money market 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:

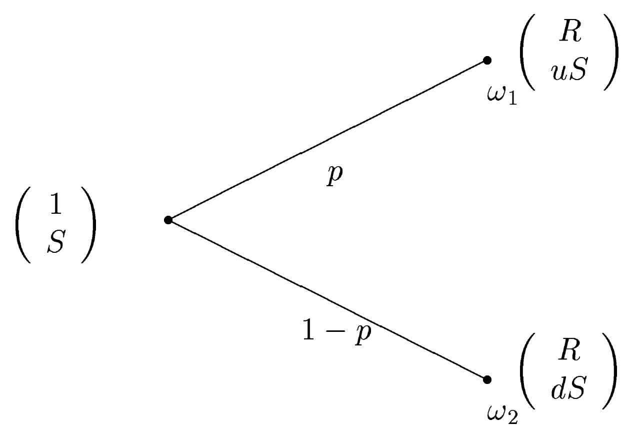

Now assume that we introduce into the economy a *European call option on the stock with exercise price K and maturity 1.* At time 1 the value of this call is equal to (where the notation $[y]^+$ (or sometimes $(y)^+$) means $\max(y, 0)$)

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

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.

To simplify notation, let $C_u = C_1(\omega_1)$ and $C_d = C_1(\omega_2)$. The question is: What should the price of this call option be at time 0? A simple portfolio 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.

Some interesting features of this example will be much clearer as we go along:

- The probability p plays no role in the expression for C_0.
- A new set of probabilities

$$q = \frac{R-d}{u-d} \qquad \text{and} \qquad 1-q = \frac{u-R}{u-d}$$

  emerges (this time we also use that $d < R < u$) and with this set of probabilities we may write the value of the call as

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

  i.e. an expected value *using* q of the discounted time 1 value of the call.

- If we compute the expected value *using* q of the discounted time 1 stock price we find

$$E^q\left[\frac{S(\omega)}{R}\right] = \left(\frac{R-d}{u-d}\right)\frac{1}{R}(uS) + \left(\frac{u-R}{u-d}\right)\frac{1}{R}(dS) = S$$

The method of pricing the call really did not use the fact that C_u and C_d were call-values. Any security with a time 1 value depending on ω_1 and ω_2 could have been priced.

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

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

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 *security price system* is a vector $\pi \in \mathbb{R}^N$ and an $N \times S$ matrix D where we interpret the i'th row $(d_{i1}, \ldots, d_{iS})$ of D as the payoff at time 1 of the i'th security in states $1, \ldots, S$, 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 1 *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*

$$\pi = D\psi$$

Clearly, we have already proved the following in Chapter 3:

Proposition 6 *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 riskless asset if there exists a linear combination of the rows of* D *which gives us* $(1, \ldots, 1) \in \mathbb{R}^S$.

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 riskless asset.* Note that when a riskless 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 riskless asset.

Now define

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

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:

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

If we read this coordinate by coordinate it says that

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

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 riskless asset exists.

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

Proposition 7 *Let an arbitrage-free system* (π, D) *and a redundant security* c *by given. The augmented system* $\left(\widehat{\pi}, \widehat{D}\right)$ *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*

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

Proof. Assume $\pi_c < \psi_1 c_1 + \ldots + \psi_S c_S$. 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 $\pi_c > \psi_1 c_1 + \ldots + \psi_S c_S$ 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*

$$D^\top \theta = y$$

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

Proposition 8 *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 do contingent claims pricing! 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" $c = (c_1, \ldots, c_S)$. Now compute the price of the contingent claim as

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

where $q_i = \frac{\psi_i}{d_0} \equiv R_0 \psi_i$. The portfolio generating the claim is the solution to $D^\top \theta_c = c$, and since we can always in a complete model reduce the matrix to an $S \times S$ invertible matrix without changing the model this can be done by matrix inversion.

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

$$\left\{ \begin{pmatrix} 1 \\ S \end{pmatrix}, \begin{pmatrix} R & R \\ uS & dS \end{pmatrix} \right\}.$$

For this to be arbitrage-free, proposition (6) tells us that there must be a solution (ψ_1, ψ_2) with $\psi_1 > 0$ and $\psi_2 > 0$ to the equation

$$\begin{pmatrix} 1 \\ S \end{pmatrix} = \begin{pmatrix} R & R \\ uS & dS \end{pmatrix} \begin{pmatrix} \psi_1 \\ \psi_2 \end{pmatrix}.$$

$u \neq d$ ensures that the matrix $\begin{pmatrix} R & R \\ uS & dS \end{pmatrix}$ has full rank. $u > d$ can be assumed without loss of generality. We find

$$\psi_1 = \frac{R - d}{R(u - d)}$$

$$\psi_2 = \frac{u - R}{R\,(u - d)}$$

and note that the solution is strictly positive precisely when $u > R > d$ (given our assumption that $u > d > 0$).

Clearly the riskfree asset has a return of R, and

$$q_1 = R\psi_1 = \frac{R - d}{u - d}$$

$$q_2 = R\psi_2 = \frac{u - R}{u - d}$$

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.

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.

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.

Consider[1] an agent who has an endowment $e = (e_1, \ldots, e_S) \in \mathbb{R}^S_+$. This vector represents the random wealth that the agent will have at time 1. The agent has a utility function $U : \mathbb{R}^S_+ \to \mathbb{R}$ 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

$$\begin{aligned} \max_{\theta} U(e &+ D^\top \theta) \\ \text{s.t. } \pi^\top \theta &\leq 0. \end{aligned} \tag{4.1}$$

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. 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.

Now make sure that you can prove the following

Proposition 9 *If there exists a portfolio θ with $D^\top \theta > 0$ then the agent can find a solution to the maximization problem if and only if (π, D) is arbitrage-free.*

[1] This closely follows Darrell Duffie: Dynamic Asset Pricing Theory. Princeton University Press. 1996

The 'only if' part of this statement shows how no arbitrage is a necessary condition for existence of a solution to the agent's problem and hence for the existence of equilibrium for economies where agents have increasing utility (no continuity assumptions are needed here). The 'if' part uses continuity and compactness (why?) to ensure existence of a maximum, but of course to discuss equilibrium would require more agents and then we need some more of our general equilibrium apparatus to prove existence.

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

Proposition 10 *Assume that there exists a portfolio θ with $D^\top\theta > 0$. If there exists a solution θ^* to (4.1) and the associated optimal consumption is given by $c^* := e + D^\top\theta^* >> 0$, 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

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

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 $U(c^*)$ is proportional to a state-price vector. Choosing a θ^+ with $D^\top\theta^+ > 0$ we know from no arbitrage that $\pi^\top\theta^+ > 0$ and from the assumption that the utility function is strictly increasing, we have $\nabla U(c^*)D^\top\theta^+ > 0$. Hence μ must be positive. ■

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 $(p_1, \ldots, p_S)$ 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

$$\psi_i = \lambda p_i u'(c_i^*), \;\; i = 1, \ldots, S. \tag{4.2}$$

where λ is some constant of proportionality. 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 (4.2) 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.