ARBITRAGE PRICING THEORY : THE ONE-PERIOD MODEL

This lecture describes the basic principles of derivative security valuation. The ideas are presented here can be applied to most valuation problems – from the simplest ones, involving straightforward compound interest-rate calculations, to the most complicated, such as the valuation of exotic options. For simplicity, we shall discuss a model for a securities market with finitely many final states and with a single trading period. In this model, the main definitions and results can be formulated in terms of elementary mathematics. The key idea behind asset valuation in markets with uncertainty is the notion of *absence of arbitrage opportunities.*

Suppose that an investor takes a "position" in the marketplace, through buying and selling securities, which has zero net cost and which guarantees i) no losses in the future and ii) some chance of making a profit. In this hypothetical situation, the investor has a positive probability of realizing a profit without taking a risk. This situation is known as an *arbitrage opportunity* or simply an *arbitrage.* Although arbitrage opportunities may arise sporadically in financial markets, they cannot last long. In fact, an arbitrage can be viewed as a relative mispricing between correlated assets. If this mispricing becomes known to sufficiently many investors, the prices will be affected as they move to take advantage of such opportunity. As a consequence, prices will change and the arbitrage will disappear. This principle can be stated as follows: *in an efficient market there are no (permanent) arbitrage opportunities.*

Example 1. Suppose that the current (spot) price of an ounce gold is \$398 and that the three-month forward price is \$390. Furthermore, suppose that the annualized three-month interest rate for borrowing gold (known as the "convenience yield") is 10% and that the interest rate on 3-month deposits is 4% (annualized). This situation gives rise to an arbitrage opportunity. In fact, an arbitrageur can borrow one ounce of gold, sell it at its current price of \$398 (go short 1 ounce), lend this money for three-months and enter into a three-month forward contract to buy one ounce of gold at \$390. Since the cost of borrowing the ounce of gold is $\$398 \times 10\%/4 = 398 \times 2.5\% = \$9.95$ and the interest on the 3-month deposit amounts to $\$398 \times (.01) = \$3.98$, the total financing cost for this operation is \$5.97. He will therefore have 398- 5.97 = 392.03 dollars in his bank account after three months. By purchasing the ounce of gold in three months at the forward price of \$390 and returning it, he will make a profit of \$2.03. (This argument neglects transaction costs and assumes that interests are paid after the lending period.)

Typeset by $\mathcal{A}\mathcal{M}\mathcal{S}$-TEX

1. The Arrow-Debreu Model

We shall consider the simplest possible model of a securities market with uncertainty, the *Arrow Debreu model.* We assume that there are N securities, $s_1, s_2 \ldots, s_N$, that can be held long or short by any investor. Initially, an investor takes a position in the "market" by acquiring a portfolio of securities. He holds this position during a "trading period" (a specified period of time), which gives him the right to claim/owe the dividends paid by the securities (capital gains/losses). He liquidates the position after the trading period and incurs in a net profit/loss from the change in the market value of his holdings (market gains/losses). The model is completely specified in terms of the *price vector* for the N securities

$$\mathbf{p} \;=\; (p_1\,, p_2 \ldots p_N\,) \tag{1}$$

and the *cash-flows matrix*

$$\mathbf{D} = (D_{ij}),\; 1 \leq \; i \; \leq \; N\,, 1 \leq \; j \; \leq \; M \tag{2}$$

where M is the number of all possible "states" of the market at the end of the trading period. The j^{th} row of the matrix $\mathbf{D}$ represents the M possible cash-flows associated with holding one unit of the j^{th} security, including dividend payments and market profit/losses.[1] We shall assume that the matrix D is known to all investors but that the final state of the market, represented by the M different columns, is not known in advance.

A *portfolio* of securities is represented by a vector

$$\theta \;=\; (\theta_1\,, \theta_2 \ldots \theta_N\,)\,. \tag{3}$$

Here, θ_i represents the number of units of the i^{th} security held in the portfolio. If θ_i is positive, the investor is *long* the security and hence has acquired the right to receive the corresponding cash-flow $\theta_i\, D_{ij}$ at the end of the period. If θ_i is negative, the investor is *short* the security and thus will have a liability at the end of the trading period. (Short positions are taken by borrowing securities and selling them at the market price). It is assumed that all investors can take short and long positions in arbitrary amounts of securities. Transaction costs, commissions and tax implications associated with trading are neglected. For simplicity, we assume that the amounts θ_i held long or short are not necessarily integers, but instead arbitrary real numbers.

The price of a portfolio θ is

$$\theta \cdot \mathbf{p} \;=\sum_1^N \theta_i\, p_i$$

[1] Prices and cashflows are expressed in the same unit of account. For simplicity, we take this unit to be the dollar.

and the cash-flow for this portfolio in the j^{th} "state" of the market will be is

$$\theta \cdot \mathbf{D}_{.j} = \sum_1^N \theta_i D_{ij} . \tag{4}$$

We can express mathematically the concept of arbitrage opportunity within this simple model.

Definition. *An arbitrage portfolio is a portfolio θ such that either (i)*

$$\theta \cdot \mathbf{p} = 0 ,$$

$$\theta \cdot \mathbf{D}_{.j} \geq 0 \quad \textit{for all} \quad 1 \leq j \leq M ,$$

and

$$\theta \cdot \mathbf{D}_{.j} > 0 \quad \textit{for some} \quad 1 \leq j \leq M$$

or

$$\theta \cdot \mathbf{p} < 0 ,$$

and

$$\theta \cdot \mathbf{D}_{.j} \geq 0 \quad \textit{for all} \quad 1 \leq j \leq M .$$

In plain words, an arbitrage portfolio is a position in the market that either (i) has zero initial cost, has no "down side" regardless of the market outcome, and offers a possibility of realizing a profit or (ii) realizes an immediate profit for the investor and which has no down side. We remark here that the distinction between the two cases is not really important: it is a consequence of the general form of the model, in which the nature of the "securities" is not specified. If it is possible to lend money (buy bonds), then the second case reduces to the first, because the investor can lend out the initial profit and then realize a positive cash-flow at the end of the trading period.

Theorem 1. *If there exists a vector of positive numbers*

$$\pi = (\pi_1 \, \pi_2 \ldots \pi_M)$$

such that

$$\mathbf{p} = \mathbf{D} \cdot \pi , \tag{5a}$$

i.e., if

$$p_i = \sum_1^M D_{ij} \, \pi_j \quad \textit{for all } 1 \leq i \leq N \tag{5b}$$

there exist no arbitrage portfolios. Conversely, if there are no arbitrage portfolios, there exists a vector π with positive entries satisfying (5).

Proof[2] **:** The first statement is easy to verify. If (5) holds then, for any portfolio θ,

$$\begin{aligned} \theta \cdot \mathbf{p} &= \theta \cdot \big(\mathbf{D} \cdot \pi\big) \\ &= \big(\theta \cdot \mathbf{D}\big) \cdot \pi . \end{aligned} \tag{6}$$

Suppose that θ is an arbitrage portfolio. Then, its initial value is non-positive and its cash-flows are non-negative for all final states. Furthermore, either (i) at least one cash flow is positive, or (ii) the initial cost is negative. Clearly, equation (6) tells us that neither case can ocurr. In fact, since the π_j's are all positive, the initial cost will be positive if at least one of the cash-flows is positive and it will be zero if all the cash-flows are zero.

We pass to the proof of the converse statement: no-arbitrage implies the existence of a vector with positive entries satisfying (5). Let $\mathbf{R}^{M+1}$ denote the vector space of $M+1$-tuples $\mathbf{x} = (x_0 \dots, x_M)$ and let $\mathbf{R}_+^{M+1}$ represent the *convex cone*

$$\mathbf{R}_+^{M+1} \equiv \{\mathbf{x} : x_j \geq 0 \text{ , for all } 0 \leq j \leq M\}.$$

Let L be the linear subspace of $\mathbf{R}^{M+1}$ defined by

$$L \equiv \{(-\theta \cdot \mathbf{p}, \theta \cdot \mathbf{D}_{.1}, \dots, \theta \cdot \mathbf{D}_{.M}), \theta \in \mathbf{R}^N\}.$$

The non-existence of arbitrage portfolios implies that the subspace L and the cone $\mathbf{R}_+^{M+1}$ intersect only at the origin, $(0, \dots 0)$. From Convex Analysis, (cf. Rockafellar, Princeton University Press, 1990), it is known that there must exist a *hyper-plane*, i.e. a linear subspace of $\mathbf{R}^{M+1}$ of dimension M, which contains L and meets $\mathbf{R}_+^{M+1}$ only at the origin. The general equation for a hyper-plane in $\mathbf{R}^{M+1}$ is

$$H \equiv \{\mathbf{x} : \sum_0^M \lambda_j x_j = 0\},$$

where $\lambda = (\lambda_0 \dots \lambda_M)$ is a vector in $\mathbf{R}^{M+1}$. It is easy to verify that the condition that such hyper-plane meets $\mathbf{R}_+^{M+1}$ only at the origin is equivalent to having $\lambda_j > 0$ for all j or $\lambda_j < 0$ for all j. (Note that λ represents the normal direction to H). But, since L is contained in H, we conclude that for all $\theta = (\theta_1 \dots, \theta_N)$

$$-\lambda_0 \theta \cdot \mathbf{p} + \sum_1^M \lambda_j \theta \cdot \mathbf{D}_{.j} = 0 .$$

This implies that

$$-\lambda_0 \mathbf{p} + \sum_1^M \lambda_j \mathbf{D}_{.j} = \mathbf{0},$$

[2] from D. Duffie: *Dynamical Asset Pricing Theory*, Princeton University Press, 1992

or

$$\mathbf{p} = \sum_{1}^{M} \frac{\lambda_j}{\lambda_0} \mathbf{D}_{.j}$$
$$= \sum_{1}^{M} \pi_j \mathbf{D}_{.j} ,$$

with $\pi_j = \frac{\lambda_j}{\lambda_0}$. This is precisely what we wanted to show.[3]

This theorem says prices and cash-flows must satisfy certain restrictions in a no-arbitrage market. The positive coefficients π_j $1 \leq j \leq M$ are usually called *state-prices*. To give a financial interpretation to this theorem, we define the *risk-neutral probabilities* or *risk-adjusted probabilities* (the terminology will become clear later)

$$\hat{\pi}_j = \frac{\pi_j}{(\sum_{1}^{M} \pi_k)} \qquad 1 \leq j \leq M .$$

These coefficients are all positive and have sum 1 so, mathematically, they can be viewed as probabilities. Also, set

$$1 + R = 1/(\sum_{1}^{M} \pi_j).$$

Now, suppose that there exists an investment opportunity which guarantees a risk-less payoff of \$1 at the end of the period – a bond or money-market deposit. In terms of the model, the bond payoff can be represented as the vector (1, 1, ..., 1) in $\mathbf{R}^M$. According to (5), the value of such bond must be

$$p_{bond} = \sum_{1}^{M} \pi_j = 1/(1+R) ,$$

so we have

$$R = \text{ bond yield.}$$

We can rewrite relation (5a) as

$$p_i = \frac{1}{1+R} \sum_{1}^{M} D_{ij} \hat{\pi}_j \qquad (7)$$

or,

$$p_i = \frac{1}{1+R} \mathbf{E}\{ \mathbf{D}_i \}$$

[3] A more intuitive geometric interpretation of this theorem is given in paragraph §2 below.

where $\mathbf{E}$ is the expectation-value operator associated with the probabilities $\hat{\pi}_j$, $1 \leq j \leq M$. We have established the following corollary of Theorem 1:

Theorem 2. *Assume that the market admits no arbitrage portfolios and that there exists riskless lending/borrowing at rate R %. Then, there exists a probability measure defined on the set of possible market outcomes, {1,2,...,M}, such that the value of any security is equal to the expected value of its future cash flows discounted at the riskless lending rate.*

This is an important general principle satisfied by markets in equilibrium. It has several remarkable features. First of all, notice that in our original model, we did not make any assumptions about the frequency at which each of the M "states" occurred. These frequencies could, in principle, be determined statistically, by observing the market over many time periods. One could then write

$$\text{Prob.}\{\text{ state j ocurrs }\} \;=\; f_j$$

for $1 \leq j \leq M$. This raises the following question: what is the relation between the risk-neutral probabilities $\hat{\pi}_j$ of the no-arbitrage theorem and the "statistical" probabilities which arise by observing the frequency of the different states? Interestingly enough, the two probabilities can be different. The market value of a given security will *not* be equal, in general, to its discounted expected cash-flows under the *statistical* probabilities. This has to do with investors' perception of the *risk* of holding different securities given the present information. Thus, the market may attach *economic values* to future states which are not proportional to their observed frequency in the past. If we have

$$p_i \;=\; \frac{1}{1+R}\sum_{1}^{M} D_{ij}\, f_j\,, \tag{8}$$

we say that the market is *risk-neutral*: the importance attached by the investors to the cash flows in the different future states is proportional to their frequency. If (8) holds, the different states are equally "important" after adjusting for frequency. On the other hand, writing the pricing equation (7) in the form

$$p_i \;=\; \frac{1}{1+R}\sum_{1}^{M} D_{ij}\,\Big(\frac{\hat{\pi}_j}{f_j}\Big)\, f_j\,,$$

we see that the prices of securities are *weighted* statistical averages of future cash-flows discounted at the riskless rate. The "weights" $\frac{\hat{\pi}_j}{f_j}$ reflect investor's preferences towards the different states; they are usually called *state-price deflators.*

To make this interpretation of state-prices more specific, suppose that an additional set of M state-contingent "elementary securities" $s_{N+1}\,, s_{N+2}\,, \ldots\, s_{N+M}$ is introduced in the market. For each j, the security s_j has cash-flow \$1 in state j and \$0 otherwise. Notice that a portfolio containing D_{i1} units of s_{N+1}, D_{i2} units of s_{N+2}, etc. has cash-flows $(D_{i1}\,, D_{i2},\, \ldots\, D_{iM}\,)$ according to the M possible final states. Thus, it provides the *same return* as the i^{th} "standard" security s_i. If there are no arbitrage opportunities, then the value of such portfolio should be equal to the value of the security. This is readily seen as follows: if the price of the portfolio is less than p_i, then an investor can short the portfolio and purchase the security, making an immediate profit. After the trading period, the cash-flows from the security exactly compensate the short position in the portfolio, and hence the investor will be able to make a profit without taking risk. A similar arbitrage can be constructed if the portfolio is traded at a lower price than the security. Therefore, if we set

$$\hat{\pi}_j = \text{market value of } s_{N+j}$$

for each j, we must have

$$p_i = \sum_1^M D_{ij}\,\hat{\pi}_j$$

for all i, which is precisely equation (5b). We conclude that *state-prices can be interpreted as market prices for "state-contingent claims" which pay* \$1 in state j and zero otherwise, for $1 \le j \le M$, supporting the notion that state-prices correspond to the prices of wealth in the different states. Notice then that the risk-neutral probabilities are those that "make the investor risk-neutral", given the prices of wealth of the different states.

Example 2. Imagine a hypothetical inflationary economy, with a monetary unit which we call the "eagle". At some future date, the central bank is expected to devalue the eagle (with respect to gold) by printing more currency. For each j, let state j correspond to a devaluation of $d_j\%$ of the eagle. The market does not know beforehand which level of devaluation will be chosen by the central bank, and all states are equally likely (thus $f_j = 1/M$). Assume that, for each j, there exists a security (s_j) which pays 1 eagle if state j occurs and nothing otherwise. Assume also that there is a bond with yield R %. In terms of "real" wealth (measured in terms of gold, say) the present value of s_j is proportional to $1/(1+d_j)$ for each j. Also, the present value of s_j is $\hat{\pi}_j\,/(1+R)$, according to our theory. It is therefore reasonable to assume, under these circumstances, the risk-neutral probabilities are given by

$$\hat{\pi}_j = \frac{1/(1+d_j)}{\sum_k 1/(1+d_k)} \qquad \text{for } 1 \le j \le M\,.$$

In general, the f_j's are probabilities in a statistical sense, whereas the risk-neutral probabilities π_j's are "mathematical probabilities" used to calculate the market values of all securities, including "state-contingent" claims (i.e. derivatives), from their cash-flows. These two notions of probability should not be confused. The correct market values of securities are determined from the risk-neutral probabilities.

2. Security-space diagram: a geometric interpretation of Theorem 1

The Arrow-Debreu model with N traded securities and M final states can be visualized geometrically. This interpretation is slightly different that the one presented in the proof of Theorem 1: let $\mathbf{R}^N$ represent N-dimensional Euclidean space. Since there are M final states, we can represent the final outcomes as M vectors $\mathbf{D}_{\cdot 1}, \ldots \mathbf{D}_{\cdot M}$, where the entries of each vector represent the final cash-flows of each traded security. (See Figure 1a).

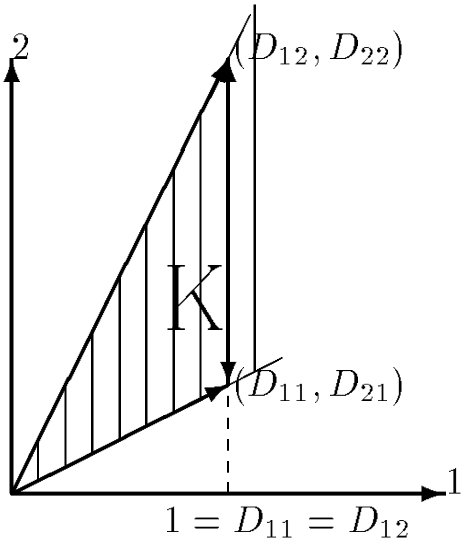

Figure 1a

In addition, we can draw the price vector $\mathbf{p}$ as an element of this space. The absence of arbitrage corresponds to the price vector $\mathbf{p}$ lying in the interior of the convex cone $\mathbf{K}$ generated by the M cash-flow vectors, as shown in Figure 1b.

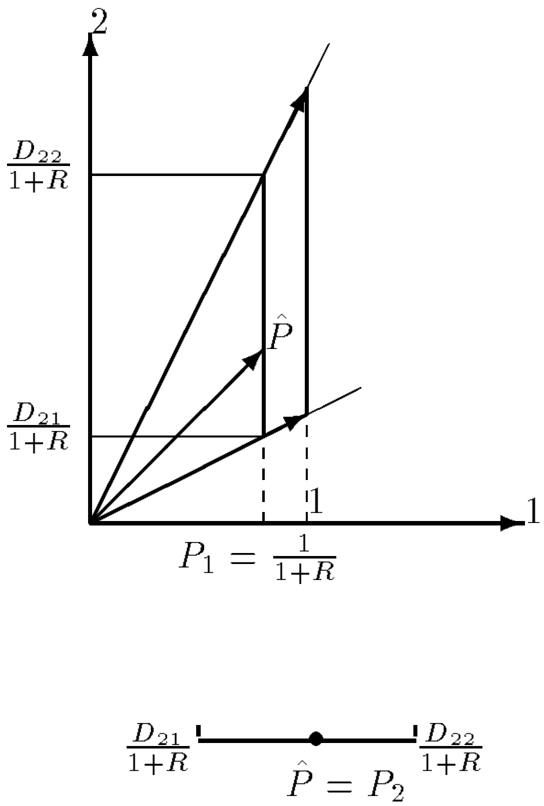

Figure 1b

In fact, if $\mathbf{p}$ would lie in the exterior or on the boundary of $\mathbf{K}$, the separating hyperplane theorem ensures that there exists linear subspace described by the equation

$$\theta \in \mathbf{R}^N \quad : \qquad \theta \cdot \mathbf{n} \;=\; 0\;,$$

where $\mathbf{n}$ is an N-vector, which *separates* $\mathbf{p}$ from the interior of the convex cone generated by the $\mathbf{D}_{\cdot\, j}$ s. The separation property would imply (possibly after a change of the sign of $\mathbf{n}$) the inequalities

$$\mathbf{p}\cdot\mathbf{n} \;\leq\; 0\;, \qquad \mathbf{D}_{\cdot\, j}\cdot\mathbf{n} \geq\; 0$$

with at least one inner product positive for some j – i.e. an arbitrage.

As in the discussion preceding Theorem 2, the existence of a bond or money-market deposit can be used to pass from a convex cone to a bounded convex set in dimension $N-1$. In fact, assume without loss of generality that the bond corresponds to the security with $i \;=\; 1$, so that all the $\mathbf{D}$-vectors have the first entry equal to one ($D_{1\,j} \;=\; 1$). Denoting a generic vector in $\mathbf{R}^N$ by

$$\theta \;=\; (\theta_1, \hat{\theta}\,) \quad , \quad \hat{\theta} \;=\; (\theta_2\,, ...\theta_N)\;,$$

we can consider the intersection of the convex cone $\mathbf{K}$ with the hyper-plane in $\mathbf{R}^N$ of vectors with first coordinate equal to p_1, the price of the bond. This intersection is the convex set $\widehat{\mathbf{K}}$ generated by the vectors

$$p_1 \, \widehat{\mathbf{D}}_{\cdot\, 1}\,,\ p_1 \, \widehat{\mathbf{D}}_{\cdot\, 2}\,,\ ...p_1 \, \widehat{\mathbf{D}}_{\cdot\, M}\,,\ .$$

The no-arbitrage condition is then equivalent to stating that $\hat{\mathbf{p}}$ lies in the interior of the convex set generated by these vectors, i.e.

$$\hat{\mathbf{p}} \;=\; \sum_{j=0}^{M} \hat{\pi}_j \, p_1 \, \hat{\mathbf{D}}_{\cdot\, j}$$

$$=\; p_1 \, \sum_{j=0}^{M} \hat{\pi}_j \, \hat{\mathbf{D}}_{\cdot\, j}$$

where the weights $\hat{\pi}_j$ are positive and sum to 1. (In the special case of Figure 1b the convex set $\widehat{\mathbf{K}}$ is a line segment.) Since, by definition, we have $p_1 \;=\; \frac{1}{1+R}$, the last formula is equivalent to equation (7).

3. Replication

The interpretation of state-prices as the values of elementary state-contingent claims is the basis for the valuation of derivative securities in a no-arbitrage market.

Given a security s, and a set of securities $s_1\,, s_2\,, ..., s_K$, we say that the portfolio $(\,\theta_1\,, \theta_2\,, ..., \theta_K)$, (representing holdings in each of the K securities) *replicates* s if the security and the portfolio have identical cash-flows. Under no-arbitrage conditions, the value of the security and of the replicating portfolio must be the same. Otherwise, an arbitrage could be realized either by shorting the portfolio and buying the security or, alternatively, by shorting the security and buying the portfolio. If the value of the portfolio is less than the value of the security, the first strategy is an arbitrage. If the portfolio is worth more than the security, the second strategy is an arbitrage. This argument gives rise to a simple but important valuation principle.

Proposition. *In a no-arbitrage market, if a security admits a replicating portfolio of traded securities, its value is equal to the value of the replicating portfolio.*

Here are some elementary applications of this principle.

Example 3: Forward prices. Suppose that a security has (spot) price P and that the yield for riskless lending over the trading period is R. Consider a *forward contract*, which consists of an agreement to purchase the security at the end of the trading period at price K. Assume that the security pays no dividends over the trading period. The no-arbitrage price of the forward contract is

$$Q = P - K/(1+R).$$

To see this, consider a portfolio consisting of being long one unit of the security and short $K/(1+R)$ worth of riskless bonds. After the trading period, the holder of this portfolio will own the security and will owe $\$K$. Therefore, if he receives K at the end of the final period he will be able to meet his cash obligation and delivers the security, he will have no profit/loss. The portfolio is equivalent to having long position in the forward contract. In practice, forward contracts are designed so that they have zero initial cost ($Q = 0$). The *forward price*, which is price for delivery of the security *after* the trading period, is

$$F \equiv (1+R)\,P\,,$$

because $K = F$ makes the initial cost zero.

Example 4. Put-Call Parity. This important example of replication involves options. Suppose that a security with price S is traded as well as a call option and a put option with exercise price (strike price) K. Recall, from I.1, that a *call option* is a contract which gives the holder the right, but not the obligation, to purchase (one unit of) the security at price K, at a stipulated date (expiration date). A *put option* gives the holder the right to sell the underlying security at price K at the maturity date. In this example, we assume that the call and the put have same strike prices and expiration dates. Denote the market prices of the call and the put by C and P, respectively. Let R be the interest paid for riskless lending over the duration of the option (compounded simply). Suppose that an investor has the following position: long one call, short one unit of the underlying security and long $K/(1+R)$ worth of riskless bonds. Let us examine the cash-flows arising from this position at the maturity date of the call. If, on the one hand, the price of the security at the expiration date, S_T, is greater than K, the investor can exercise the call, purchasing the security at the strike price K and then return the security held short. This leaves him with zero profit/loss. On the other hand, if the price of the underlying security is less than K, he will not exercise the call. At the option's expiration date his new position is short one unit of underlying security and he has K in cash, for a total value of $K - S_T$. If, instead, the investor would have held a put struck at K initially, his position after the trading period would have been neutral if $S_T \geq K$ (the put goes unexercised) and worth $K - S_T$ if $S_T < K$, since he can exercise the put. Therefore, the portfolio "long call, short underlying asset, long $K/(1+R)$ in bonds " replicates the put. We conclude that

$$P = C - S + K/(1+R).$$

This is the *put-call parity* relation. It shows how to construct a "synthetic put " via a portfolio. Since there are four variables in this equation, we can produce similar portfolios to replicate other assets as well. Accordingly, the position " long underlying security, long put, short $K/(1+R)$ in bonds " replicates a call. The position "long one call, short one put, long $K/(1+R)$ in bonds "replicates the cash-flows of the underlying security. Finally, the position " long one unit of security, short one call, long one put " replicates a riskless bond paying $\$K$ at the expiration date of the option.

4. The binomial model

We present the simplest case of the Arrow-Debreu model. Consider an ideal situation in which there are only two states ($M = 2$) and two securities: a bond with yield R % and a security (s) with price P. We assume that the cash-flows of this security in states 1 and 2 are $P\,U$ and and $P\,D$, respectively, where U and D are given numbers and $D < U$. If there are no arbitrage opportunities, we must have

$$P = \frac{1}{1+R}\left(\hat{\pi}_1\,P\,U + \hat{\pi}_2\,P\,D\right)$$

where $\hat{\pi}_1$ and $\hat{\pi}_2$ are positive and $\hat{\pi}_1 + \hat{\pi}_2 = 1$. Thus, the risk-neutral probabilities satisfy

$$\begin{cases} \hat{\pi}_1 + \hat{\pi}_2 = 1 \\ \\ \hat{\pi}_1\,U + \hat{\pi}_2\,D = 1+R\,. \end{cases}$$

It is easy to see that this system will have positive solutions if and only if

$$D < 1+R < U\,. \tag{9}$$

This intuitively clear: if $1+R \geq U$, the return on riskless lending is greater than of equal to the return for investing in the risky security, regardless of the final state. An arbitrage could then be achieved by shorting the security and lending out the proceeds at the riskless rate. Similarly, if $1+R \leq D$, the investor can make a riskless profit by borrowing at the riskless rate and purchasing the security. Suppose that (9) holds. The solution of the linear system is

$$\hat{\pi}_1 = \frac{1+R-D}{U-D} \quad , \quad \hat{\pi}_2 = \frac{U-1-R}{U-D}\,. \tag{10}$$

Thus, the risk-neutral probabilities are entirely determined from the parameters of the binomial model. The actual probabilities of occurrence of each state are irrelevant in the pricing process (as long as neither one is zero — we must assume that both states can occur).

This has interesting consequences. Suppose now that we augment the number of traded securities (the number of final states is still $M = 2$.) Then, the price of any security is completely determined by future cash-flows because the risk-neutral probabilities are still given by (10). A "state-contingent" security which has cash-flows D_1 in state 1 and D_2 in state 2 must have value V, where

$$V = \frac{1}{1+R}\left(\hat{\pi}_1\,D_1 + \hat{\pi}_2\,D_2\right) \tag{11}$$

This idea can be applied to price any security contingent on the value of s. For instance, a "call option" on s with exercise price K (with $P\,D < K < P\,U$) has cash-flows

$$D_1 = P\,U - K \quad , \quad D_2 = 0$$

since the holder will gain the difference between the market value and the exercise price if this difference is positive. We conclude that the arbitrage-free value of this call option on s is

$$V_{call} = \frac{1}{1+R} \cdot \frac{1+R-D}{U-D} \cdot (P\,U - K).$$

We can also illustrate the idea of replicating portfolios in the binomial model. Consider a portfolio (θ_1, θ_2) representing an investor's holdings in the security s (with price P) and the \$1 discount bond (with price $1/R$), respectively. This portfolio yields exactly the same returns as a security with cash-flows D_1 and D_2 provided that

$$\begin{cases} \theta_1\,P\,U + \theta_2 = D_1 \\ \\ \theta_1\,P\,D + \theta_2 = D_2\,. \end{cases}$$

Solving for θ_1 and θ_2, we find that

$$\theta_1 = \frac{D_1 - D_2}{P\,U - P\,D} \quad \text{and} \quad \theta_2 = \frac{D}{U-D}\,(D_2 - D_1)\,. \tag{12}$$

Since holding the portfolio (θ_1, θ_2) gives the investor the same returns as holding the security with cash-flows D_1 and D_2, the two positions should have equal value. We conclude that $V = \theta_1\,P + \theta_2/(1+R)$. Substituting the values of the θs from (12), we recover the price (11). Notice that this calculation, based on replicating the payoff of the contingent claim, did not require knowing the risk-neutral probabilities. This is because the binomial model admits a *unique* set of state-prices/ risk-neutral probabilities.

5. Complete and Incomplete Markets

In the binomial model, *any* vector of future cash flows $(D_1\,,\,D_2\,)$ (the index labels the final state) can be replicated in terms of a portfolio of the basic security and a riskless bond. This property can be generalized to the setting of the N-securities/M-states model.

Definition: *A securities market with M states is said to be complete if, for any cash-flow vector* $(D_1, D_2, ..., D_M)$*, there exists a portfolio of traded securities* $(\theta_1, \theta_2, ..., \theta_N)$ *which has cash-flow* D_j *in state* j*, for all* $1 \le j \le M$.

Market completeness is therefore equivalent to having a cash-flows matrix $\mathbf{D}=(D_{ij})$ with the property that the system of linear equations

$$\theta \cdot \mathbf{D} = D\,,$$

or

$$\sum_1^N \theta_i \, D_{ij} \; = \; D_j \, , \qquad 1 \, \leq \, j \, \leq \, M \, ,$$

has a solution $\theta \in \mathbf{R}^N$ for any $D \in \mathbf{R}^M$. From Linear Algebra, we know that this property is satisfied if and only if

$$rank \; \mathbf{D} \; = \; M \, ,$$

which is equivalent to saying that the column vectors of the matrix D span the entire space $\mathbf{R}^M$. Market completeness is a very strong assumption, which greatly simplifies the valuation of derivative securities. Derivative securities can be represented by general cash-flow vectors $(D_1 \, , D_2 \, , ..., \, D_M)$, as opposed to the N "standard securities" which have cash-flow vectors $(D_{i1} \, , D_{i2} \, , ..., \, D_{iM})$, for $1 \, \leq \, i \, \leq \, N$. Since any derivative security is equivalent to a portfolio of standard traded assets, its price is fully determined from $\mathbf{D}$ in the absence of arbitrage. More formally, we have

Proposition. *Suppose that the market is complete and that there are no arbitrage opportunities. Then there is a unique set of state-prices* $(\pi_1 \, , \pi_2 \, , \ldots \, , \pi_M)$ *satisfying (5) and hence a unique set of risk-neutral probabilities* $(\hat{\pi}_1 \, , \hat{\pi}_2 \, , \ldots \, , \hat{\pi}_M)$*. Conversely, if there is a unique set of state-prices, then the market is complete.*

Proof: Market completeness implies that the price of a state-contingent claim which pays \$1 in state j and 0 otherwise is determined for all j. Therefore, there can be at most one set of state-prices. If they exist, state-prices are unique.

To prove the converse statement, suppose that there exists a unique vector of state prices $\pi \; = \; (\pi_1 \, , \pi_2 \, , ..., \, \pi_M \,)$ (with strictly positive entries) such that (5) is satisfied. We will argue that the market complete by contradiction. In fact, if the market is not complete, then $rank \; \mathbf{D} \; < \; M$. From Linear Algebra, we know that the matrix $\mathbf{D}$ must have a non-empty *right-nullspace*, i.e., there exists $\lambda \; = \; (\lambda_1 \, , \lambda_2 \, , ..., \lambda_M \,)$ such that

$$\mathbf{D} \cdot \lambda \; = \mathbf{0} \, , \tag{13}$$

or, equivalently,

$$\sum_1^M D_{ij} \, \lambda_j \; = \; 0 \, , \qquad \text{for all} \, 1 \, \leq \, i \, \leq \, N \, .$$

Using the no-arbitrage relation (5a) and (13), we conclude that

$$\mathbf{D} \cdot (\pi \; + \rho \, \lambda) \; = \mathbf{p} \, ,$$

for all real numbers ρ. Since the entries of π are strictly positive, we can choose ρ sufficiently small so that $\pi_j \, + \rho \, \lambda_j$ is positive for all j. Therefore, we have constructed a new state-price vector, contradicting our hypothesis. We conclude that, in a no-arbitrage market, uniqueness of state-prices implies that the market is complete.

The above argument can also be used to characterize the set of state-price vectors corresponding to an incomplete market with given price vector $\mathbf{p}$ and cash-flow matrix $\mathbf{D}$. Let $k = rank\ \mathbf{D}$. In an incomplete market, we have $M - k > 0$. Notice that k is the dimension of the right-nullspace of $\mathbf{D}$. Given two state-price vectors $\pi^{(1)}$ and $\pi^{(2)}$ and ρ such that $0 \leq \rho \leq 1$, the the convex combination $\rho\, \pi^{(1)} + (1 - \rho)\, \pi^{(2)}$ is also a state-price vector. Hence, the set of all state-price vectors is convex. Moreover, in the proof of the Proposition, we showed that each state-price vector is contained in a $(M - k)$-dimensional neighborhood inside this cone. Consequently, *the set of admissible state-price vectors is an open cone in an* $(M - k)$*-dimensional affine subspace of* $\mathbf{R}^M$.

What is the financial meaning of a complete/incomplete market? In a complete market there is a unique set of state-prices. However, in real markets it is usually impossible to completely identify the set of final states or quantify the investors' preferences towards different states. Market completeness is a convenient idealization of the behavior of securities markets. Incomplete markets — with many possible price structures satisfying the no-arbitrage condition — are the rule rather than the exception.

6. The trinomial model

We describe a simple example of an incomplete market. Assume that there are two securities — a riskless bond with with yield R paying \$1 at the end of the trading period and a security s. There are three states, which correspond to different cash-flows for s. We assume that the price of s is P and that the cash-flows of s are $P\,U$ in state 1, $P\,M$ in state 2 and $P\,D$ in state 3, with

$$D < M < U\,.$$

Clearly, this market is incomplete, because the dimension of the cash-flows matrix is 3×2 and there are three states. (If there are more states than traded securities, the market is always incomplete). We can investigate the conditions for the existence of state-prices. Since we have assumed that there is riskless lending, we have

$$\frac{1}{1+R} = \pi_1 + \pi_2 + \pi_3\,, \tag{14}$$

for any admissible set of state-prices. Since there is no arbitrage, we must have

$$P = P\,U\,\pi_1 + P\,M\,\pi_2 + P\,D\,\pi_3\,,$$

or

$$1 = U\,\pi_1 + M\,\pi_2 + D\,\pi_3\,. \tag{15}$$

Admissible state-price vectors π must satisfy equations (14) and (15) and have positive entries. These equations show that there is no arbitrage if and only if

$$D < 1 + R < U\,, \tag{16}$$

the interpretation of which was given in the discussion of the binomial model. Assuming that condition (16) holds, the set of state-prices can be visualized as a line segment corresponding to the intersection of the planes described by (14) and (15) and the positive quadrant in $\mathbf{R}^3$. The two extreme points of this segment are

$$\pi_1 = \frac{(1+R)-D}{(1+R)(U-D)}, \quad \pi_2 = 0, \quad \pi_3 = \frac{U-(1+R)}{(1+R)(U-D)} \tag{17}$$

and

$$\pi_1 = 0, \quad \pi_2 = \frac{M-(1+R)}{(1+R)(M-D)}, \quad \pi_3 = , \frac{(1+R)-D}{(1+R)(M-D)}, \tag{18a}$$

if $M \geq (1+R)$, or

$$\pi_1 = \frac{(1+r)-M}{(1+R)(U-M)}, \quad \pi_2 = \frac{U-(1+R)}{(1+R)(U-M)}, \quad \pi_3 = 0, \tag{18b}$$

if $M < (1+R)$.

Because the prices of traded assets are linear functions of the state-prices, according to eq. (5b), this calculation can be used to derive *bounds* on the values of contingent claims based on the available information. In fact, we know that a linear function defined on a closed convex set attains its maximum and minimum values at the extreme points of this set. (Actually, since the set of state-prices is an *open* convex set, the maxima and minima are attained at " degenerate " state-price vectors which have entries equal to zero and hence are not, strictly speaking, state prices.) To illustrate this, consider the case of a call option on the basic security s with strike price K. To fix ideas, assume that $PM < K < PU$. Then, the cash-flows for this option are $PU - K$ in state 1 and 0 in states 2 and 3. Its no-arbitrage value is $C = \pi_1 (PU - K)$. Therefore, using (17) and (18), we find that

$$C^+ = \frac{(1+R)-D}{(1+R)(U-D)} (PU-K)$$

is an upper bound for the price of the option. If $M \geq (1+R)$ the lower bound on the price is $C^- = 0$, and if $M < (1+R)$, the lower bound is

$$C^- = \frac{(1+R)-M}{(1+R)(U-M)} (PU-K).$$

(Of course, the upper and the lower bounds coincide when $M = D$, and we recover the result of the binomial model.)

This example shows an important application of state-prices as a tool for contingent claim valuation in incomplete markets. State-prices are not unique but can nevertheless be used to obtain partial information about fair prices. This result can be interpreted financially in terms of *risk-aversion* of different agents. A market participant that does not want to incur in *any* risk will bid (be willing to buy) the security at the price corresponding to the lower bound C^- and will offer (be willing to sell) the security at the price corresponding to the upper bound C^+. Since there is incomplete information about Arrow-Debreu prices, transactions made between the bounds imply a risk for the buyer as well as for the seller.

7. Exercises

1. Consider a hypothetical country where the government has declared a "currency band" policy, in which the exchange rate between the domestic currency, denoted by XYZ, and the U.S. dollar is guaranteed to fluctuate in a prescribed band, namely

$$USD\ 0.95 \ \leq \ XYZ \ \ \leq \ USD\ 1.05\ ,$$

for at least one year. Suppose also that the government has issued 1-year notes denominated in the domestic currency which pay a simply compounded annualized rate of 30 %. Assuming that the corresponding interest rate for U.S. deposits is 6%, show that this market is not arbitrage-free in the "pure" sense. Describe the situation in terms of the Arrow-Debreu model. Propose some realistic scenarios that could make this pure arbitrage disappear in practice.

2. (i) Show that the set of all probability measures on a finite state-space of M elements can be represented as a convex subset $\mathbf{P}$ of the Euclidean space $\mathbf{R}^M$. Given a security s defined by its price and its cash-flows, verify that the set of measures which are risk-neutral for this security corresponds to the intersection of the set $\mathbf{P}$ with a hyper-plane in $\mathbf{R}^M$. Similarly, show that the set of admissible risk-neutral measures for a securities market with N securities corresponds to the intersection of $\mathbf{P}$ with N hyper-planes.

(ii) Apply this analysis to the trinomial model of §6 – assuming that $S = \$100$, $U = 1.10$, $M = 1.00$, $D = .80$, $R = .05$ and that a call option with strike \$105 is trading at a premium of $C = $ \$3.80. Show that if, instead, $C = $ \$1.00, there is an arbitrage opportunity.

3. On the week of Sept. 7, 1996, Ladbroke, a London betmaker, gave the following odds regarding the upcoming U.S. presidential election: Clinton 1-6, Dole 7-2, Perot 1-50. [For instance, Ladbroke pays one pound for every 6 pounds bet on Clinton if he wins.]. Calculate the corresponding risk-neutral probabilities for the victory of each candidate assuming that one of them will necessarily win.