Derivative Securities – Section 2
Notes by Robert V. Kohn, Courant Institute of Mathematical Sciences.

Binomial, trinomial, and more general one-period models. These notes explore the implications of arbitrage for the pricing of contingent claims. Our goal is to understand the power and limitations of risk-neutral pricing. The essence of the matter, as we shall see, is the duality theory of linear programming. Our focus is on one-period market models. This discussion lays the foundation for our treatment of multiperiod models.

The analysis of the binomial market is standard textbook material; I like Baxter-Rennie but you'll also find it in Jarrow-Turnbull and Hull. The analysis of trinomial and more general models is standard but not commonly considered textbook material; for a treatment rather different from mine, emphasizing portfolio selection and the maximization of utility, see John Cochrane's forthcoming book *Asset Pricing*, available on the web at http://gsbwww.uchicago.edu/fac/john.cochrane/research/Papers/. Another treatment, closer to mine but more difficult, is in Darrell Duffie's book *Dynamic Asset Pricing Theory* (Princeton University Press, 1996).

****

The binomial model. We consider an economy which has

- just two securities: a stock (paying no dividend, initial unit price per share s_1 dollars) and a bond (interest rate r, one bond pays one dollar at maturity).
- just one maturity date T
- just two possible states for the stock price at time T: s_2 and s_3, with $s_2 < s_3$

(see the figure).

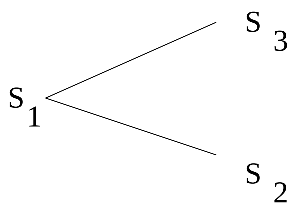

Figure 1: Prices in the one-period binomial market model.

We could suppose we know the probability p that the stock will be worth s_3 at time T. This would allow us to calculate the expected value of any contingent claim. However we will make no use of such knowledge. Pricing by arbitrage considerations makes no use of information about probabilities – it uses just the list of possible events.

The reasonable values of s_1, s_2, s_3 are not arbitrary: the economy should permit no arbitrage. This requires that

$$s_2 < s_1 e^{rT} < s_3.$$

It's easy to see that if this condition is violated then an arbitrage is possible. The converse is extremely plausible; a simple proof will be easy to give a little later.

In this simple setting a contingent claim can be specified by giving its payoff when $S_T = s_2$ and when $S_T = s_3$. For example, a long call with strike price K has payoff $f_2 = (s_2 - K)_+$ in the first case and $f_3 = (s_3 - K)_+$ in the second case. The most general contingent claim is specified by a vector $f = (f_2, f_3)$ giving its payoffs in the two cases.

Claim 1: In this model every contingent claim has a replicating portfolio. Thus arbitrage considerations determine the value of every contingent claim. (A market with this property is said to be "complete".)

In fact, consider the portfolio consisting of ϕ shares of stock and ψ bonds. Its initial value is

$$\phi s_1 + \psi e^{-rT}.$$

Its value at maturity replicates the contingent claim $f = (f_2, f_3)$ if

$$\begin{aligned} \phi s_2 + \psi &= f_2 \\ \phi s_3 + \psi &= f_3. \end{aligned}$$

This is a system of two linear equations for the two unknowns ϕ, ψ. The solution is

$$\phi = \frac{f_3 - f_2}{s_3 - s_2}, \quad \psi = \frac{s_3 f_2 - s_2 f_3}{s_3 - s_2}.$$

The initial value of the contingent claim f is necessarily the initial value of the replicating portfolio:

$$V(f) = \phi s_1 + \psi e^{-rT} = \frac{f_3 - f_2}{s_3 - s_2} s_1 + \frac{s_3 f_2 - s_2 f_3}{s_3 - s_2} e^{-rT}.$$

Claim 2: The value can conveniently be expressed as

$$V(f) = e^{-rT}[(1-q)f_2 + qf_3] \quad \text{where} \quad q = \frac{s_1 e^{rT} - s_2}{s_3 - s_2}.$$

Moreover, the condition that the market admit no arbitrage is $0 < q < 1$, which is equivalent to $s_2 < s_1 e^{rT} < s_3$.

The formula for $V(f)$ in terms of q is a matter of algebraic rearrangement. This simplification seems mysterious right now, but we'll see a natural reason for it later.

The condition that the market supports no arbitrage has two parts:

(i) a portfolio with nonnegative payoff must have a nonnegative value; and

(ii) a portfolio with nonnegative and sometimes positive payoff must have positive value.

In the binomial setting every payoff (f_2, f_3) is replicated by a portfolio, so we may replace "portfolio" by "contingent claim" in the preceding statement without changing its impact. Part (i) says $f_2, f_3 \geq 0 \Rightarrow (1-q)f_2 + qf_3 \geq 0$. This is true precisely if $0 \leq q \leq 1$. Part (ii) forces the sharper inequalities $q > 0$ and $q < 1$.

Notice the form of Claim 2. It says the present value of a contingent claim is obtained by taking its "expected final value" $(1-q)f_2 + qf_3$ then discounting (multiplying by e^{-rT}). However the "expected final value" has nothing to do with the probability of the stock going up or down. Instead it must be taken with respect to a special probability measure, assigning weight $1-q$ to state s_2 and q to state s_3, where q is determined by s_1, s_2, s_3 and r as above. This special probability measure is known as the "risk-neutral probability" associated with the market.

I like to view q as nothing more than a convenient way of representing $V(f)$. However the term "risk-neutral probability" can be understood as follows. Real-life investors prefer a guaranteed return at rate r to an uncertain one with expected rate r. We may nevertheless imagine the existence of "risk-neutral" investors, who are indifferent to risk. Such investors would consider these two alternatives to be equivalent. In a world where all investors were risk-neutral, investments would have to be valued so that their expected return agrees with the risk-free interest rate r. The formula for $V(f)$ has this form, if we assume in addition that the "expectations" of the risk-neutral investors are expressed by the probabilities $1-q$ and q:

$$\text{expected payoff} = (1-q)s_2 + qs_3 = e^{rT}V(f) = e^{rT} \cdot \text{initial investment}.$$

****

The trinomial model. This is the simplest example of an *incomplete* economy. It resembles the binomial model in having

- just two securities: a stock (paying no dividend, initial unit price per share s_1 dollars) and a bond (interest rate r, one bond pays one dollar at maturity).
- just one maturity date T.

However it differs by having three final states rather than two:

- The stock price at time T can take values s_2, s_3, or s_4, with $s_2 < s_3 < s_4$

(see the figure). The reasonable values of $s_1, \ldots, s_4$ are not arbitrary: the economy should permit no arbitrage. This requires that

$$s_2 < s_1 e^{rT} \quad \text{and} \quad s_1 e^{rT} < s_4.$$

In other words, the stock must be able to do better than or worse than the risk-free return on an initial investment of s_1 dollars. It's easy to see that if this condition is violated then an arbitrage is possible.

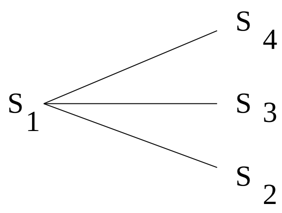

Figure 2: Prices in the one-period trinomial market model.

In this case a contingent claim is specified by a 3-vector $f = (f_2, f_3, f_4)$; here f_j is the payoff at maturity if the stock price is s_j. Question: which contingent claims are replicatable? Answer: those for which the system

$$\begin{aligned} \phi s_2 + \psi &= f_2 \\ \phi s_3 + \psi &= f_3 \\ \phi s_4 + \psi &= f_4 \end{aligned}$$

has a solution. This specifies a two-dimensional space of f's. So the market is not complete, and "most" contingent claims are not replicatable.

If a contingent claim f is not replicatable then arbitrage does not specify its price $V(f)$. However arbitrage considerations still *restrict* its price:

$V(f) \leq$ the value of any portfolio whose payoff dominates f;
$V(f) \geq$ the value of any portfolio whose payoff is dominated by f.

In other words,

$$\left.\begin{aligned} \phi s_2 + \psi &\geq f_2 \\ \phi s_3 + \psi &\geq f_3 \\ \phi s_4 + \psi &\geq f_4 \end{aligned}\right. \Longrightarrow V(f) \leq \phi s_1 + e^{-rT}\psi$$

$$\left.\begin{aligned} \phi s_2 + \psi &\leq f_2 \\ \phi s_3 + \psi &\leq f_3 \\ \phi s_4 + \psi &\leq f_4 \end{aligned}\right. \Longrightarrow V(f) \geq \phi s_1 + e^{-rT}\psi.$$

We obtain the strongest possible consequences for $V(f)$ by solving a pair of linear programming problems:

$$\max_{\substack{\phi s_j + \psi \leq f_j \\ j=2,3,4}} \phi s_1 + e^{-rT}\psi \leq V(f) \leq \min_{\substack{\phi s_j + \psi \geq f_j \\ j=2,3,4}} \phi s_1 + e^{-rT}\psi.$$

These bounds capture *all* the information available from arbitrage concerning the price of the contingent claim f. (The actual price observed in the market must be determined by additional considerations besides arbitrage. The standard "equilibrium" approach to understanding prices uses utility maximization – to be discussed in the course Capital Markets and Portfolio Theory.)

Linear programming duality. In the binomial model $V(f)$ had a convenient expression in terms of a special "risk-neutral probability." To derive the analogous result here, we use the duality theory of linear programming. The following discussion should be accessible even to those who know nothing about linear programming – but it should make you want to learn something about this important topic. My favorite text is V. Chvatal, *Linear Programming*. A more recent text, more up-to-date on current developments such as interior point methods, is R. Vanderbei, *Linear programming: foundations and extensions*. Also of note is Peter Lax's book *Linear Algebra*, whose chapter on linear programming takes exactly the min-max viewpoint I use below.

Let us concentrate on the upper bound for $V(f)$. It is

$$\begin{aligned}
\min_{\phi s_j+\psi\geq f_j} \phi s_1 + e^{-rT}\psi &= \min_{\phi,\psi}\max_{\pi_j\geq 0} \phi s_1 + e^{-rT}\psi + \sum_{j=2}^{4}\pi_j(f_j - \phi s_j - \psi) \\
&= \max_{\pi_j\geq 0}\min_{\phi,\psi} \phi s_1 + e^{-rT}\psi + \sum \pi_j(f_j - \phi s_j - \psi) \\
&= \max_{\pi_j\geq 0}\min_{\phi,\psi} \phi(s_1 - \sum \pi_j s_j) + \psi(e^{-rT} - \sum \pi_j) + \sum \pi_j f_j \\
&= \max_{\substack{\sum \pi_j s_j = s_1 \\ \sum \pi_j = e^{-rT}}} \sum \pi_j f_j .
\end{aligned}$$

The first line holds because

$$\max_{\pi_j\geq 0} \pi_j(f_j - \phi s_j - \psi) = \begin{cases} 0 & \text{if } \phi s_j + \psi \geq f_j \\ +\infty & \text{otherwise.} \end{cases}$$

The second line holds by the duality theorem of linear programming, which says in this setting that $\min\max = \max\min$. The third line is obtained by rearrangement, and the fourth line by an argument similar to the first. Making the minor change of variables $\hat{\pi}_j = e^{rT}\pi_j$, we conclude that

$$\begin{aligned}
V(f) \leq \max\{e^{-rT}[\hat{\pi}_2 f_2 + \hat{\pi}_3 f_3 + \hat{\pi}_4 f_4] \ &: \ \hat{\pi}_2 s_2 + \hat{\pi}_3 s_3 + \hat{\pi}_4 s_4 = e^{rT} s_1 \\
&\quad\ \hat{\pi}_2 + \hat{\pi}_3 + \hat{\pi}_4 = 1, \quad \hat{\pi}_j \geq 0\}.
\end{aligned}$$

The lower bound is handled similarly. The only difference is that we are maximizing in ϕ, ψ and minimizing in π_j. An argument parallel to the one given above shows

$$\begin{aligned}
V(f) \geq \min\{e^{-rT}[\hat{\pi}_2 f_2 + \hat{\pi}_3 f_3 + \hat{\pi}_4 f_4] \ &: \ \hat{\pi}_2 s_2 + \hat{\pi}_3 s_3 + \hat{\pi}_4 s_4 = e^{rT} s_1 \\
&\quad\ \hat{\pi}_2 + \hat{\pi}_3 + \hat{\pi}_4 = 1, \quad \hat{\pi}_j \geq 0\}.
\end{aligned}$$

Thus the upper and lower bounds on $V(f)$ are obtained by maximizing and minimizing the "discounted expected return" $e^{-rT}[\hat{\pi}_2 f_2 + \hat{\pi}_3 f_3 + \hat{\pi}_4 f_4]$ over an appropriate class of "risk-neutral probabilities" $(\hat{\pi}_2, \hat{\pi}_3, \hat{\pi}_4)$. The incompleteness of the market is reflected in the fact that there is more than one risk-neutral probability: in the present trinomial setting the 3-vector $(\hat{\pi}_2, \hat{\pi}_3, \hat{\pi}_4)$ is constrained by two inequalities, so the class of risk-neutral probabilities is one-dimensional (a line segment).

We noted before the condition $s_2 < s_1 e^{rT} < s_4$, which is required for the economy to be "reasonable" – i.e. not to admit an arbitrage. This is precisely the condition that there be at least one "risk-neutral probability", i.e. a vector $(\hat{\pi}_2, \hat{\pi}_3, \hat{\pi}_4)$ such that $\sum \hat{\pi}_j = 1$, $\sum \hat{\pi}_j s_j = e^{rT} s_1$, and $\hat{\pi}_j > 0$ for each j.

****

The general one-period market model. The binomial and trinomial models are special cases of a more general theory, which we now present. Main purposes of this discussion:

- deeper understanding of risk-neutral probabilities; and
- more careful treatment of the "principle of no arbitrage."

In considering one-period models with few assets and many states, we are close to the question of portfolio analysis: which of the many possible portfolios should an investor hold? We'll resist addressing this, however, concentrating instead on the narrower goal of understanding risk-neutral pricing and the role of arbitrage.

The general one-period market has

- N securities, $i = 1, \ldots, N$
- M final states, $\alpha = 1, \ldots, M$
- fixed initial values: one unit of security i is worth p_i dollars
- state-dependent final values: if the final state is α then one unit of security i is worth $D_{i\,\alpha}$.

An investor can hold any *portfolio*: θ_i units of security i. It has initial value $\langle p, \theta \rangle = \sum_i p_i \theta_i$. If the final state is α then its final value is $\langle \theta, D_{\cdot\,\alpha} \rangle = \sum_i \theta_i D_{i\,\alpha}$.

Examples:

Binomial model: $p = (e^{-rT}, s_1)$, $D = \begin{pmatrix} 1 & 1 \\ s_2 & s_3 \end{pmatrix}$

Trinomial model: $p = (e^{-rT}, s_1)$, $D = \begin{pmatrix} 1 & 1 & 1 \\ s_2 & s_3 & s_4 \end{pmatrix}$

In general, if security 1 is a riskless bond then

$$p = (e^{-rT}, p_2, \ldots, p_N), \quad D = \begin{pmatrix} 1 & \cdots & 1 \\ D_{21} & \cdots & D_{2M} \\ \cdot & & \cdot \\ \cdot & & \cdot \\ D_{N1} & \cdots & D_{NM} \end{pmatrix}$$

Here's a careful statement of the **Principle of no arbitrage**:

(a) $\sum_i \theta_i D_{i\,\alpha} \geq 0$ for all $\alpha \implies \sum_i \theta_i p_i \geq 0$

(b) if we have both $\sum_i \theta_i D_{i\,\alpha} \geq 0$ for all α and $\sum_i \theta_i p_i = 0$ then we must have $\sum_i \theta_i D_{i\,\alpha} = 0$ for every α.

These capture with precision the informal statements that (a) a portfolio with nonnegative payoff has nonnegative value; and (b) a portfolio with nonnegative and sometimes positive payoff has strictly positive value.

The key result relating risk-neutral probabilities to lack of arbitrage is this:

Arbitrage Theorem: The economy satisfies (a) iff there exist $\pi_\alpha \geq 0$ such that

$$\sum_\alpha D_{i\,\alpha} \pi_\alpha = p_i, \quad i = 1, \ldots, N.$$

It satisfies both (a) and (b) if in addition the π_α can be chosen to be all strictly positive.

The theorem is trivial in one direction: assuming the existence of π_α we can easily prove the absence of arbitrage. In fact, for any portfolio θ_i we have

$$\begin{aligned} \sum_i \theta_i D_{i\,\alpha} \geq 0 \text{ for all } \alpha \quad &\implies \quad \sum_{i,\alpha} \theta_i D_{i\,\alpha} \pi_\alpha \geq 0 \\ &\implies \quad \sum_i \theta_i p_i \geq 0 \end{aligned}$$

since $\pi_\alpha \geq 0$. If $\pi_\alpha > 0$ for each α then the conclusion can hold with $=$ only if each hypothesis holds with $=$ rather than $\geq$. Thus existence of $\pi_\alpha \geq 0$ implies part (a) of the no-arbitrage principle; and if each π_α is strictly positive then we also get part (b) of the principle.

The other half of the theorem (no arbitrage implies existence of π_α) is decidedly nontrivial. We shall sketch the proof presently; but first let us make contact with what we did earlier in the trinomial setting. Suppose security 1 is a risk-less bond. Then $p_1 = e^{-rT}$ and the first row of $D_{i\,\alpha}$ is filled with 1's. The statement of the theorem becomes: the market permits no arbitrage iff there exist positive π_α such that

$$\begin{aligned} \pi_1 + \cdots + \pi_M &= e^{-rT} \\ \sum_\alpha \pi_\alpha D_{i\,\alpha} &= p_i, \quad i = 2, \ldots, N. \end{aligned}$$

Writing $\hat\pi_\alpha = e^{rT}\pi_\alpha$ we see that the this is equivalent to the existence of positive $\hat\pi_\alpha$ such that

$$\begin{aligned}
\hat\pi_1 + \cdots + \hat\pi_M &= 1 \\
\sum_\alpha \hat\pi_\alpha D_{i\,\alpha} &= e^{rT} p_i, \quad i = 2, \dots, N.
\end{aligned}$$

These $\hat\pi_\alpha$ are the *risk-neutral probabilities.*

For the trinomial market, we showed how arbitrage considerations restrict the initial value of any contingent claim f. The same max/min argument works in general, for any market in which Security 1 is a riskless bond. The conclusion is

$$\min_{\substack{\text{risk-neutral}\\ \text{probs } \hat\pi}} e^{-rT} \sum_\alpha \hat\pi_\alpha f_\alpha \le V(f) \le \max_{\substack{\text{risk-neutral}\\ \text{probs } \hat\pi}} e^{-rT} \sum_\alpha \hat\pi_\alpha f_\alpha.$$

We immediately see that

$$\begin{aligned}
\text{market completeness} \;&\Leftrightarrow\; \text{arbitrage determines the value of every contingent claim} \\
&\Leftrightarrow\; \text{there is a unique risk-neutral probability.}
\end{aligned}$$

****

Sketch of a proof of the Arbitrage Theorem. The rest of this section sketches a proof of (the nontrivial part of) the Arbitrage Theorem, based on the following

Fundamental Lemma of Linear Programming: If a collection of linear inequalities implies another linear inequality then it does so "trivially," i.e. the conclusion is a (nonnegative) linear combination of the hypotheses.

The name "fundamental lemma of linear programming" is my own; the proper name of this result is Farkas' Lemma. See e.g. V. Chvatal, *Linear Programming*, pg. 248, for this and related results.

Our first task is to show that part (a) of the no-arbitrage principle implies the existence of $\pi_\alpha \ge 0$. Now, part (a) says that the collection of linear inequalities $\sum_i \theta_i D_{i\,\alpha} \ge 0$ for $\alpha = 1, \dots, M$ implies another linear inequality $\sum_i \theta_i p_i \ge 0$. By the Fundamental Lemma of Linear Programming this occurs only if there is a "trivial" proof, i.e. if there exists $\pi_\alpha \ge 0$ such that $\sum_i \theta_i p_i = \sum_{i,\alpha} \theta_i D_{i\,\alpha}\pi_\alpha$ for all θ_i. But that means $\sum D_{i\,\alpha}\pi_\alpha = p_i$.

Our second task is to show that if the economy satisfies both parts (a) and (b) of the no-arbitrage principle then we can take $\pi_\alpha > 0$ for all α. If the π_α already identified are all positive then we're done. If not, then renumbering states if necessary we may suppose $\pi_1, \dots, \pi_{M'} > 0$ and $\pi_{M'+1} = \ldots = \pi_M = 0$.

Let's concentrate for a moment on index $M' + 1$. If $D_{\cdot\, M'+1} = (D_{1\, M'+1}, \ldots, D_{N\, M'+1})$ is a linear combination of $D_{\cdot\, 1}, \ldots, D_{\cdot\, M'}$ then we can easily modify π_α to make $\pi_{M'+1} > 0$. In fact, suppose $D_{\cdot\, M'+1} = b_1 D_{\cdot\, 1} + \ldots + b_{M'} D_{\cdot\, M'}$. Then

$$\begin{aligned} p_i &= \sum_{\alpha=1}^{M'} D_{i\,\alpha} \pi_\alpha \\ &= \epsilon D_{i\, M'+1} + \sum_{\alpha=1}^{M'} D_{i\,\alpha}(\pi_\alpha - \epsilon b_\alpha), \end{aligned}$$

so replacing $\pi = (\pi_1, \ldots, \pi_{M'}, 0, \ldots, 0)$ with $(\pi_1 - \epsilon b_1, \ldots, \pi_{M'} - \epsilon b_{M'}, \epsilon, 0, \ldots, 0)$ does the trick when ϵ is sufficiently small.

Essentially the same argument shows that if *any* positive combination of $D_{\cdot\, M'+1}, \ldots, D_{\cdot\, M}$ lies in the span of $D_{\cdot\, 1}, \ldots, D_{\cdot\, M'}$ then we can modify π_α to make additional components positive.

Applying the preceding argument finitely many times, we either arrive at a new π with strictly positive components, or we find ourselves in a situation (with a new value of M') where no positive combination of $D_{\cdot\, M'+1}, \ldots, D_{\cdot\, M}$ lies in the span of $D_{\cdot\, 1}, \ldots, D_{\cdot\, M'}$. We claim the second alternative cannot happen when the economy has property (b).

This is another application of the Fundamental Lemma of Linear Programming. Our "second alternative" is that

$$\sum_{\alpha=M'+1}^{M} a_\alpha D_{\cdot\,\alpha} = \sum_{\alpha=1}^{M'} b_\alpha D_{\cdot\,\alpha}, \quad a_\alpha \geq 0 \implies a_\alpha = 0, \alpha = M'+1, \ldots, M.$$

The "trivial consequences" of the hypotheses are obtained by taking linear combinations. This amounts to taking the inner product with a vector $\theta \in R^N$. Thus the trivial consequences of the hypotheses are

$$\sum_{\alpha=M'+1}^{M} a_\alpha \langle D_{\cdot\,\alpha}, \theta \rangle = \sum_{\alpha=1}^{M'} b_\alpha \langle D_{\cdot\,\alpha}, \theta \rangle.$$

For this (coupled with $a_\alpha \geq 0$) to give a trivial proof that $a_\alpha = 0$ we must have

$$\begin{aligned} \langle D_{\cdot\,\alpha}, \theta \rangle = \sum_i \theta_i D_{i\,\alpha} &= 0 \quad \alpha = 1, \ldots, M' \\ \langle D_{\cdot\,\alpha}, \theta \rangle = \sum_i \theta_i D_{i\,\alpha} &> 0 \quad \alpha = M'+1, \ldots, M. \end{aligned}$$

But then θ represents a portfolio with no downside, some upside, and value 0 since $\sum_i \theta_i p_i = \sum_i \sum_{\alpha=1}^{M'} \theta_i D_{i\,\alpha} \pi_\alpha = 0$. This contradicts our assumption that the economy admits no arbitrage.