Derivative Securities -- Spring 2007 — Section 3

Notes by Robert V. Kohn, modified by Steve Allen
Courant Institute of Mathematical Sciences

Multi-period Binomial Trees. We turn to the valuation of derivative securities in a time-dependent setting. We focus for now on multi-period binomial models, i.e. binomial trees. This setting is simple enough to let us do everything explicitly, yet rich enough to approximate many realistic problems.

The material covered in this section is very standard (and very important). The treatment here is essentially that of Baxter and Rennie (Chapter 2). The same material is also in Hull (Chapter 11 in the 6th edition).. In the next section of notes we'll discuss how the parameters should be chosen to mimic the conventional (Black-Scholes) hypothesis of lognormal stock prices, and we'll pass to the continuous-time limit.

Binomial trees are widely used in practice, in part because they are easy to implement numerically (also because the scheme can easily be adjusted to price American options). For a nice discussion of alternative numerical implementations, seethe article ``Nine ways to implement the binomial method for option valuation in Matlab," by D.J. Higham, SIAM Review 44, no. 4, 661-677.

****

The multi-period binomial model generalizes the single-period binomial model we considered in Section 2. It has

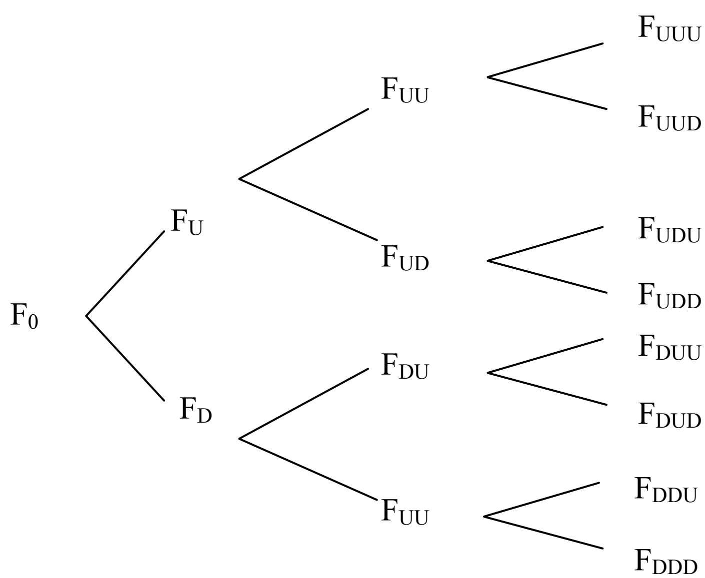

Figure 1: States of a non-recombining binomial tree.

- just two securities: a risky asset (a forward) and a riskless asset ("bond")
- a series of times 0, δt, $2\delta t$,…,$N\delta t = T$ at which trades can occur;
- interest rate r_i during time interval i for the bond;

- a binomial tree of possible states for the forward prices

The last statement means that for each forward price at time $j\delta t$, there are two possible values it can take at time $(j+1)\ \delta t$ (see Figure 1).

The interest rate environment is described by specifying the interest rates r_i. We restrict our attention for now to the case of a constant interest rate: $r_i = r$ for all i.

The forward price dynamics is described by assigning a price s_j to each state in the tree. Strictly speaking we should also assign (subjective) probabilities p_j to the branches (the two branches emerging from a given node should have probabilities summing to 1): see Figure 2.

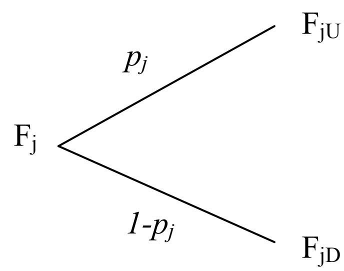

Figure 2: One branch of a binomial tree

Actually, we will make no use of the subjective probabilities p_j; our arguments are based on arbitrage, so they depend only on the list of possible states not on their probabilities. However our pricing formula will make use of *risk-neutral probabilities* q_j. These ``look like'' subjective probabilities, except that they are determined by the stock prices and the interest rate.

Our forward prices must be ``reasonable'' in the sense that the market support no arbitrage. Motivated by the one-period model, we (correctly) expect this condition to take the form:

- starting from any node, the stock price may do better than or worse than the risk-free rate during the next period.

In other words, $F_{jD} < Fj < F_{jU}$ for each j.

The tree in Figure 1 is the most general possible. At the n^{th} time step it has 2^n possible states. That's a lot of states, especially when n is large. It's often convenient to let selected states have the same prices in such a way that the list of distinct prices forms a *recombinant tree*. Figure 3 gives an example of a 4-stage recombinant tree, with stock prices marked for each state: (A recombinant tree has just $n+1$ possible states at time step n.)

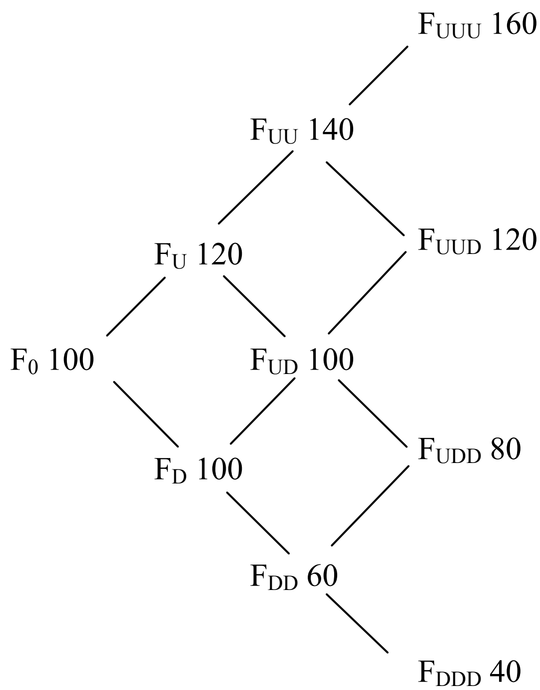

Figure 3: A simple recombinant binomial tree

A special class of recombinant trees is obtained by assuming the stock price goes up or down by fixed multipliers *u* or *d* at each stage: see Figure 4.

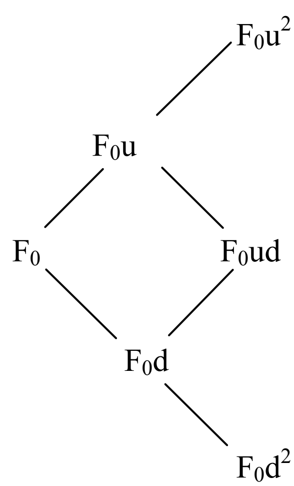

Figure 4: A multiplicative recombinant binomial tree

This last class may seem terribly special relative to the general binomial tree. But we shall see it is general enough for many practical purposes -- just as a random walk (consisting of many steps, each of fixed magnitude but different in direction) can approximate Brownian motion. And it has the advantage of being easy to specify -- one has only to give the values of u and d.

It may seem odd that we consider a market with just one forward, when real markets have many instruments. But our goal is to price contingent claims based on considerations of arbitrage. If we succeed using just these two instruments (the forward and the riskless bond), then our conclusions necessarily apply to any larger market containing both instruments.

****

Our goal is to determine the value (at time 0) of a contingent claim. We will consider American options later; for the moment we consider only European ones, i.e. early redemption is prohibited. The most basic examples are European calls and puts (payoffs: $(F_T - K)_+$ and $(K - F_T)_+$ respectively). However our method is much more general. What really matters is that *the payoff of the claim depends entirely on the state of forward process at time T.*

Let's review what we found in the one-period binomial model. Our multi-period model consists of many one-period models, so it is convenient to introduce a flexible labeling scheme. Writing "now" for what used to be the initial state, and "up, down" for what used to be the two final states, our risk-neutral valuation formula was

$$f_{now} = e^{-r\delta t} [q f_{up} + (1-q) f_{down}]$$

where

$$q = \frac{F_{now} - F_{down}}{F_{up} - F_{down}}$$

Here we're writing V_{now} for the (present) value of the contingent claim worth V_{up} or V_{down} at the next time step if the stock price goes up or down respectively. This formula was obtained by replicating the payoff with a combination of stock and bond; the replicating portfolio used

$$\phi = \frac{V_{up} - V_{down}}{F_{up} - F_{down}}$$

units of forward.

The valuation of a contingent claim in the multi-period setting is an easy consequence of this formula. We need only "work backward through the tree," applying the formula again and again.

Consider, for example, the four-period recombinant tree shown in Figure 3. (This example, taken straight from Baxter and Rennie, has the nice feature of very simple arithmetic.) Suppose the interest rate is r=0, for simplicity. Then q=1/2 at each node (we chose the prices to keep this calculation simple). Let's find the value of a European call with strike price 100 and maturity T=3δt. Working backward through the tree:

- The values at maturity are $(S_T - 100)_+$ = 60, 20, 0, 0respectively.
- The values one time step earlier are 40, 10, and 0 respectively, each value being obtained by an application of the one-period formula.
- The values one time step earlier are 25 and 5.
- The value at the initial time is 15.

Easy. But is it right? Yes, because these values can be replicated. However the replication strategies are more complicated than in the one-period case: the replicating portfolio must be adjusted at each trading time, taking into account the new stock price.

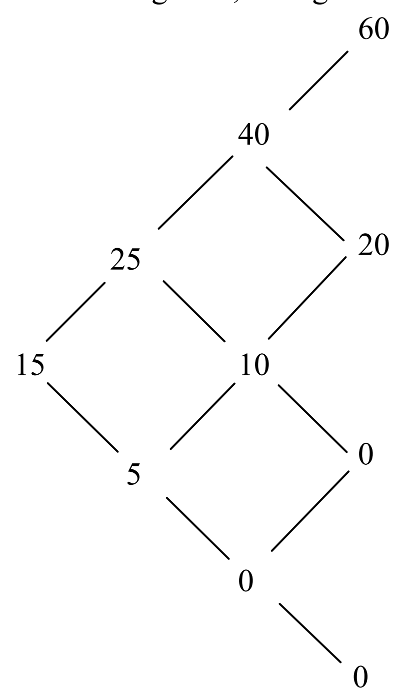

Figure 5: Value of the option as a function of forward price state

Let’s show this in the example. Using our one-period rule, the replicating portfolio starts with ø = (25 − 5) / (120 − 80) = .5 units of the forward, worth.5 x 100 = 50 dollars. The claim is that by trading intelligently at each time-step we can adjust this portfolio so it replicates the payoff of the option no matter what the forward price does. Here is an example of a possible history, and how we would handle it:

Forward goes up to 120. The new ø is (40 − 10) / (140 − 100) = .75, so we need another .25 units of the forward. We must buy this at the present price, 120 dollars per unit, so it costs 30 dollars.

Forward goes up again to 140. The new ø is (60 - 20) / (160 − 120) = 1, so we buy another .25 unit at 140 dollars per unit. This costs another 35 dollars.

Forward goes down to 120. At maturity we hold one share of the forward, replicating the option. Our total costs have been 50 + 30 + 35 = 115 and we deliver the forward for 100. We have a loss of 15, exactly offset by the 15 we charged for the option.

That wasn't a miracle. It follows from the fact that we have determined all of our hedge ratios so as to replicate the value of the option. But here's a second example -- a different possible history – to convince you:

Forward goes down to 80. The revised ø is (10 - 0) / (100 - 60) = .25. So we should sell .25 units of the forward, receiving 80 /4 = 20.

Forward goes up to 100. The new ø is (20 - 0) / (120 - 80) = .5. So we must buy .25 units of the forward, costing 100 / 4 = 25.

Forward goes down again to 80. Our position is worth 40 – 40 = 0, replicating the option, which is worthless since it's out of the money. We bought the forward at 50 + 25 = 75 and sold at 20 + 40 = 60 for a loss of 75 – 60 = 15, exactly offsetting the 15 we charged for the option.

Our example shows the importance of tracking ø, the number of units of the forward to be held as you leave a given node. It characterizes the replicating portfolio (the "hedge"). Its value is known as the *Delta* of the claim. Thus:

$$\Delta_{now} = our\phi = \frac{V_{up} - V_{down}}{F_{up} - F_{down}}$$

To understand the meaning Δ observe that as you leave node *j*

value of claim at a node j = $\Delta_j \, s_j$

by definition of the replicating portfolio If the value of the forward changes by an amount *ds*, then the value of the replicating portfolio changes by Δds. Thus Δ is a sort of derivative:

Δ= rate of change of replicating portfolio value, with respect to change of the forward price.

****

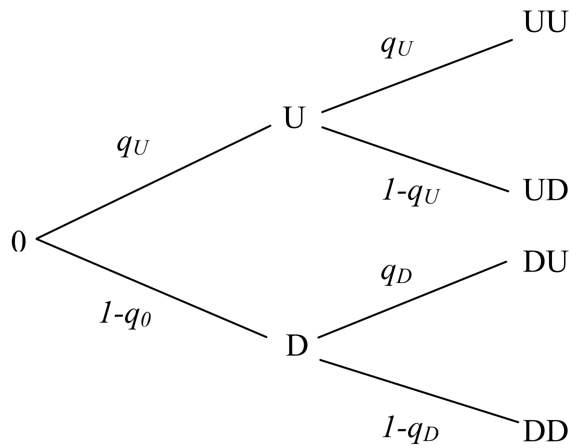

Figure 6. General binomial tree with two time steps

Our valuation algorithm is easy to implement. But in the one-period setting we had more than an algorithm: we also had a *formula* for the value of the option, as the discounted expected value using a risk-neutral probability. A similar formula exists in the multi-period setting. To see this, it is most convenient to work with a general binomial tree. Consider, for example, a tree with two time steps. The risk-neutral probabilities q_j, $1 - q_j$ are determined by the embedded one-period models. (Remember, the risk-neutral probabilities are characteristic of the market; they don't depend on the contingent claim under consideration.) In this case:

$$q_D = \frac{F_0 - F_D}{F_U - F_D}, \quad q_D = \frac{F_D - F_{DD}}{F_{DU} - F_{DD}}, \quad q_U = \frac{F_U - F_{UD}}{F_{UU} - F_{UD}}$$

As we work backward through the tree, we get a formula for the value of the contingent claim at each node, as a discounted weighted average of its values at maturity. In fact, writing V_j for the value of the contingent claim V at node j,

$$V_U = e^{-r\delta t}\,[q_U V_{UU} + (1 - q_U)V_{UD}]$$

and

$$V_D = e^{-r\delta t}\,[q_D V_{DU} + (1 - q_D)V_{DD}]$$

so

$$\begin{aligned} V_0 &= e^{-r\delta t}\,[q_0 V_U + (1 - q_0)V_D] \\ &= e^{-2r\delta t}\,[q_0 q_U V_{UU} + q_0(1 - q_U)V_{UD} + (1 - q_0)q_D V_{DU} + (1 - q_0)(1 - q_D)V_{DD}] \end{aligned}$$

It should be clear now what happens, for a binomial tree with any number of time periods:

initial value of the claim = $e^{-rN\delta t} \sum_{final\ states}$ [probability of the associated path] x [payoff of state], where the probability of any path is the product of the probabilities of the individual risk-neutral probabilities along it. (Thus: the different risk-neutral probabilities must be treated as if they described independent random variables.)

A similar rule applies to recombinant trees, since they are just special binomial trees in disguise. We must simply be careful to count the paths with proper multiplicities. For example, consider a two-period model with a recombinant tree and

$$F_{up} = uF_{now},\ F_{down} = dF_{now}$$

In this case the formula becomes

$$V_0 = e^{-2r\delta t}\left[(1-q)^2 V_{DD} + 2q(1-q)V_{UD} + q^2 V_{UU}\right]$$

with q = (1 − d)/(u − d), since there are two distinct paths leading to node UD.

The preceding calculation extends easily to recombinant trees with any number of time steps. The result is one of the most famous and important results of the theory: an explicit formula for the value of a European option. This is in a sense the binomial tree version of the Black-Scholes formula. (To really use it, of course, we'll need to know how to specify the parameters u and d; we'll come to that soon.) Consider an N-step recombinant stock price model with $F_{up} = uF_{now}$, $F_{down} = dF_{now}$, and F_0 =initial forward price. Then the present value of an option with payoff $V(F_T)$ is

$$e^{-rN\delta t}\sum_{j=0}^{N}\binom{N}{j} q^j (1-q)^{N-j} V\left(F_0 u^j d^{N-j}\right)$$

with $q = (1 - d) / (u - d)$. This holds because there are $\binom{N}{j}$ different ways of accumulating j ups and $N - j$ downs in N time-steps (just as there are $\binom{N}{j}$ different ways of getting heads exactly j times out of N coin flips.) Making this specific to European puts and calls: a call with strike price K has present value

$$e^{-rN\delta t}\sum_{j=0}^{N}\left[\binom{N}{j} q^j (1-q)^{N-j} \left(F_0 u^j d^{N-j} - K\right)_+\right]$$

a put with strike price K has present value

$$e^{-rN\delta t}\sum_{j=0}^{N}\left[\binom{N}{j} q^j (1-q)^{N-j} \left(K - F_0 u^j d^{N-j}\right)_+\right]$$

Since $N\delta t = T$, we may restate this as:

$$V_0 = e^{-rT} E_{RN} [V(F_T)]$$

where $E_{RN} [V(F_T)]$ is the expected final payoff, computed with respect to the risk-neutral probability:

$$E_{RN}\left[V\left(F_T\right)\right] = \sum_{j=0}^{N} \binom{N}{j} q^j \left(1-q\right)^{N-j} V\left(F_0 u^j d^{N-j}\right)$$