

Mathematics of Finance – Notes on Options, Part 2 – 11/17/03

Robert V. Kohn, Courant Institute of Mathematical Sciences.

This material is extracted (with some editing) from my Derivative Securities lecture notes, (available in full at www.math.nyu.edu/faculty/kohn). Much of this material is in Luenberger, but I like my own way of presenting it better. Other good books providing more in-depth treatments include J.C. Hull, *Options, futures, and other derivatives*, R. Jarrow and S. Turnbull, *Derivative securities*, and N. Chriss, *Black Scholes and beyond*.

The one-period binomial model. We have learned that lognormal stock dynamics can be simulated by a “binomial tree”, representing the net effect of many coin flips. So let us learn to price options in such a marketplace.

The first, simple but essential step is to consider a one-period binomial model, i.e. an investment whose value at the end of a given time period can take only two values. In other words, we consider a market 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 $s_1, \dots, s_4$ are not arbitrary: the economy should




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

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.

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:

$$\begin{aligned} V(f) &\leq \text{the value of any portfolio whose payoff dominates } f; \\ V(f) &\geq \text{the value of any portfolio whose payoff is dominated by } f. \end{aligned}$$

In other words,

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

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

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

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

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.)

Multiperiod 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 examples here draw from a nice little book by Baxter and Rennie. Equivalent discussions can be found however in Luenberger (and elsewhere).

The multi-period binomial model generalizes the single-period binomial model we considered above. It has




Figure 3: States of a non-recombining binomial tree.

- just two securities: a risky asset (a “stock,” paying no dividend) and a riskless asset (“bond”);
- a series of times $0, \delta t, 2\delta t, \dots, N\delta t = T$ at which trades can occur;
- interest rate r_i during time interval i for the bond;
- a binomial tree of possible states for the stock prices.

The last statement means that for each stock price at time $j\delta t$, there are two possible values it can take at time $(j+1)\delta t$ (see Figure 3).

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 stock 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 4.




Figure 4: 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 stock 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, $s_{2j} < e^{r\delta t} s_j < s_{2j+1}$ for each j .

The tree in Figure 3 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 5 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 .)

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 6. 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 stock, when real markets have many stocks. But our goal is to price contingent claims based on considerations of arbitrage. If we succeed using just these two instruments (the stock and the riskless bond) then our conclusions necessarily apply to any larger market containing both instruments.




Figure 5: A simple recombining binomial tree.

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: $(S_T - K)_+$ and $(K - S_T)_+$ respectively). However our method is much more general. What really matters is that *the payoff of the claim depends entirely on the state of stock process at time T .*

Let's review what we found in the one-period binomial model. Our multiperiod 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

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

where

$$q = \frac{e^{r\delta t} s_{\text{now}} - s_{\text{down}}}{s_{\text{up}} - s_{\text{down}}}.$$

Here we're writing f_{now} for what we used to call $V(f)$, the (present) value of the contingent claim worth f_{up} or f_{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

$$\phi = \frac{f_{\text{up}} - f_{\text{down}}}{s_{\text{up}} - s_{\text{down}}}$$

units of stock, and a bond worth $f_{\text{now}} - \phi s_{\text{now}}$.

The valuation of a contingent claim in the multiperiod setting is an easy consequence of this formula. We need only “work backward through the tree,” applying the formula again and again.




Figure 6: A multiplicative recombinant binomial tree.

Consider, for example, the four-period recombinant tree shown in Figure 5. (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 $T = 3\delta t$. Working backward through the tree:

- The values at maturity are $(S_T - 100)_+ = 60, 20, 0, 0$ respectively.
- 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.

Let's show this in the example. Using our one-period rule, the replicating portfolio starts with $\phi = (25 - 5)/(120 - 80) = .5$ units of stock, worth $.5 \times 100 = 50$ dollars, and a bond worth $15 - 50 = -35$. (Viewed differently: the investment bank that sells the option collects 15 dollars; it should borrow another 35 dollars, and use these $15 + 35 = 50$ dollars to buy $1/2$ unit of stock.) 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 stock price does. Here is an example of a possible history, and how we would handle it:

Stock goes up to 120 The new ϕ is $(40 - 10)/(140 - 100) = .75$, so we need another .25 units of stock. We must buy this at the present price, 120 dollars per unit, and we do it by borrowing 30 dollars. Thus our debt becomes 65 dollars.

Stock goes up again to 140 The new ϕ is $(60 - 20)/(160 - 120) = 1$, so we buy another .25 unit at 140 dollars per unit. This costs another 35 dollars, bringing our debt to 100 dollars.




Figure 7: Value of the option as a function of stock price state.

Stock goes down to 120 At maturity we hold one share of stock and a debt of 100. So our portfolio is worth $120 - 100 = 20$, replicating the option. (Put differently: if the investment bank that sold the option followed our instructions, it could deliver the unit of stock, collect the 100, pay off its loan, and have neither a loss nor a gain.)

That wasn't a miracle. We'll explain it more formally pretty soon. But here's a second example – a different possible history – to convince you:

Stock goes down to 80 The revised ϕ is $(10 - 0)/(100 - 60) = .25$. So we should sell $1/4$ unit stock, receiving $80/4 = 20$. Our debt is reduced to 15.

Stock goes up to 100 The new ϕ is $(20 - 0)/(120 - 80) = .5$. So we must buy $1/4$ unit stock, spending $100/4 = 25$. Our debt goes up to 40.

Stock goes down again to 80 We hold 40 dollars worth of stock and we owe 40. Our position is worth $40 - 40 = 0$, replicating the option, which is worthless since it's out of the money. (Different viewpoint: the investment bank that sold the option can liquidate its position, selling the stock at market and using the proceeds to pay off the loan. This results in neither a loss nor a gain.)

Notice that the portfolio changes from one time to the next but the changes are *self-financing* – i.e. the total value of the portfolio before and after each trade are the same. (The investment bank neither receives or spends money except at the initial time, when it sells the option.)

Our example shows the importance of tracking ϕ , the number of units of stock 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:

$$\Delta_{\text{now}} = \text{our } \phi = \frac{f_{\text{up}} - f_{\text{down}}}{s_{\text{up}} - s_{\text{down}}}.$$

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

$$\text{value of claim at a node } j = \Delta_j s_j + b_j$$

by definition of the replicating portfolio (here b_j is the value of the bond holding at that moment). If the value of the stock changes by an amount ds while the bond holding stays fixed, 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 stock price.

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 multiperiod 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




Figure 8: General binomial tree with two time steps.

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:

$$q_1 = \frac{e^{r\delta t} s_1 - s_2}{s_3 - s_2}, \quad q_2 = \frac{e^{r\delta t} s_2 - s_4}{s_5 - s_4}, \quad q_3 = \frac{e^{r\delta t} s_3 - s_6}{s_7 - s_6}.$$

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 $f(j)$ for the value of the contingent claim f at node j ,

$$f(3) = e^{-r\delta t}[q_3f(7) + (1 - q_3)f(6)]$$

and

$$f(2) = e^{-r\delta t}[q_2f(5) + (1 - q_2)f(4)]$$

so

$$\begin{aligned} f(1) &= e^{-r\delta t}[q_1f(3) + (1 - q_1)f(2)] \\ &= e^{-2r\delta t}[q_1q_3f(7) + q_1(1 - q_3)f(6) + (1 - q_1)q_2f(5) + (1 - q_1)(1 - q_2)f(4)]. \end{aligned}$$

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

$$\text{initial value of the claim} = e^{-rN\delta t} \sum_{\text{final states}} [\text{probability of the associated path}] \times [\text{payoff of state}],$$

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 recombining 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 recombining tree and $s_{\text{up}} = us_{\text{now}}$, $s_{\text{down}} = ds_{\text{now}}$. In this case the formula becomes




Figure 9: Price states in a multiplicative recombining tree.

$$f(1) = e^{-2r\delta t}[(1 - q)^2f(4) + 2q(1 - q)f(5) + q^2f(6)]$$

with $q = (e^{r\delta t} - d)/(u - d)$, since there are two distinct paths leading to node 5.

The preceding calculation extends easily to recombining 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 need to specify the parameters u and d . We discussed that earlier, when we explained why trees mimic lognormal stock dynamics.) Consider an N -step recombining stock price model with $s_{\text{up}} = us_{\text{now}}$, $s_{\text{down}} = ds_{\text{now}}$, and s_0 =initial spot price. Then the present value of an option with payoff $f(S_T)$ is

$$e^{-rN\delta t} \sum_{j=0}^N \left[\binom{N}{j} q^j (1-q)^{N-j} f(s_0 u^j d^{N-j}) \right].$$

with $q = (e^{r\delta t} - d)/(u - d)$. 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

$$e^{-rN\delta t} \sum_{j=0}^N \left[\binom{N}{j} q^j (1-q)^{N-j} (s_0 u^j d^{N-j} - K)_+ \right];$$

a put with strike price K has present value

$$e^{-rN\delta t} \sum_{j=0}^N \left[\binom{N}{j} q^j (1-q)^{N-j} (K - s_0 u^j d^{N-j})_+ \right].$$