Derivative Securities – Fall 2012– Section 7
Notes by Robert V. Kohn, extended and improved by Steve Allen.
Courant Institute of Mathematical Sciences.

Reprise and some consequences of the continuous time viewpoint, then a return to pricing options on trees. This section discusses:

(a) The link between risk-neutral expectation and the Black-Scholes PDE.

(b) An alternative (more general) viewpoint, based on martingales.

(c) What happens if you rebalance only at discrete times?

(d) Path-dependent options.

(e) American options revisited.

****

(a) The link between risk-neutral expectations and PDE's. We have discussed two apparently different approaches to the valuation of a European option: (i) take the discounted risk-neutral expected payoff, or (ii) solve the Black-Scholes PDE. Let's show now that these two approaches are equivalent.

First, consider options on a forward price. We saw long ago that in the discrete time setting, the forward process satisfies

$$\mathcal{F}_t = E_{\mathrm{RN}}[\mathcal{F}_T]$$

for any $t < T$. (In the terminology we'll introduce soon, $\mathcal{F}_t$ is a *martingale* when we calculate expectations using the risk-neutral measure.) In the continuous-time setting, this means the SDE for $\mathcal{F}$ under the risk-neutral measure has no dt term. If in addition $\mathcal{F}$ is lognormal then its SDE under the risk-neutral measure must be

$$d\mathcal{F} = \sigma \mathcal{F} dw. \tag{1}$$

I claim that if $\mathcal{F}$ satisfies this SDE, and $V(F, t)$ satisfies

$$V_t + \tfrac{1}{2}\sigma^2 F^2 V_{FF} - rV = 0 \quad \text{for } t < T\text{, with } V(F, T) = \phi(F), \tag{2}$$

then

$$V(\mathcal{F}(0), 0) = e^{-rT} E[\phi(\mathcal{F}_T)]. \tag{3}$$

To see why, let's apply the Ito calculus to the function $H(F, t) = e^{r(T-t)} V(F, t)$. We get

$$dH(\mathcal{F}(t), t) = e^{r(T-t)} \left[V_t + \tfrac{1}{2}\sigma^2 F^2 V_{FF} - rV \right] dt + e^{r(T-t)} \sigma \mathcal{F} V_F \, dw$$

with the understanding that each term on the right is evaluated at $(F, t) = (\mathcal{F}(t), t)$. *According to the PDE (1) the dt term has coefficient zero.* Therefore

$$H(\mathcal{F}(T), T) - H(\mathcal{F}_0, 0) = \int_0^T (\text{stuff}) \, dw.$$

Now take the expected value of both sides. On the right we get 0. The second term on the left is known at time 0. As for the first term: remembering the definition of H and the final-time condition in the definition of V it is is

$$H(\mathcal{F}(T), T) = e^0 V(\mathcal{F}(T), T) = \phi(\mathcal{F}_T).$$

We thus conclude that

$$E[\phi(\mathcal{F}_T)] = H(\mathcal{F}_0, 0) = e^{rT} V(\mathcal{F}_0, 0)$$

which is equivalent to (3).

A similar calculation applies to options on a non-dividend-paying stock. We learned in Section 4 that if $s(t)$ is lognormal, then under the risk-neutral measure $s(t) = s_0 e^X$ where X is Gaussian with mean $(r - \frac{1}{2}\sigma^2)t$ and variance $\sigma^2 t$. Put differently: $s_t = s_0 \exp\left[(r - \frac{1}{2}\sigma^2)t + \sigma w(t)\right]$. We now know (from Section 6) an equivalent statement using SDE's: under the risk-neutral measure, s solves the SDE

$$ds = rs\,dt + \sigma s\,dw. \tag{4}$$

I claim that if s satisfies this SDE, and $V(F, t)$ satisfies

$$V_t + rsV_s + \tfrac{1}{2}\sigma^2 s^2 V_{ss} - rV = 0 \quad \text{for } t < T, \text{ with } V(s, T) = \phi(s), \tag{5}$$

then

$$V(s_0, 0) = e^{-rT} E[\phi(s_T)]. \tag{6}$$

The argument is just as before: we apply the Ito calculus to $H(F, t) = e^{r(T-t)} V(s, t)$. We get

$$dH(s(t), t) = e^{r(T-t)} \left[V_t + rsV_s + \tfrac{1}{2}\sigma^2 s^2 V_{ss} - rV\right] dt + e^{r(T-t)} \sigma s V_s\,dw$$

with the understanding that each term on the right is evaluated at $(s, t) = (s(t), t)$. The PDE assures us that the dt term has coefficient zero, so that

$$H(s(T), T) - H(s_0, 0) = \int_0^T (\text{stuff})\,dw.$$

Taking the expectation of both sides, we find

$$E[\phi(s_T)] = H(s_0, 0) = e^{rT} V(s_0, 0)$$

which is equivalent to (6).

****

(b) An alternative viewpoint, using martingales. The SDE's (1) and (4) are clearly fundamental. The argument we gave above for the former was hopefully convincing. The argument we gave for the latter was perhaps less so. We'll do better now, as we reformulate what we've been doing in terms of *martingales*. This viewpoint has many advantages; in particular, it explains the origin of our SDE's, and it extends easily to stochastic interest rates (which we'll begin addressing very soon).

The basic prescription for working backward in a binomial tree was this: if V is the value of a tradeable non-dividend-paying security (such as an option) then

$$V_{\text{now}} = e^{-r\delta t}[qV_{\text{up}} + (1-q)V_{\text{down}}] = e^{-r\delta t}E_{\text{RN}}[V_{\text{next}}]$$

and if $\mathcal{F}$ is the futures price of a tradeable security then

$$\mathcal{F}_{\text{now}} = [q\mathcal{F}_{\text{up}} + (1-q)\mathcal{F}_{\text{down}}] = E_{\text{RN}}[\mathcal{F}_{\text{next}}],$$

where q is the risk-neutral probability of the "up" state. (I wrote "futures" rather than "forward" on purpose. When the interest rate is deterministic, futures prices and forward prices are the same. But when the interest rate is random, it is the futures price not the forward price that satisfies the preceding equation.)

When the risk-free rate is constant the factors of $e^{-r\delta t}$ don't bother us – we just bring them out front. When the risk-free rate is stochastic, however, we must handle them differently. To this end it is convenient to introduce a *money market account* which earns interest at the risk-free rate. Let $A(t)$ be its balance, with $A(0) = 1$. In the constant interest rate setting obviously $A(t) = e^{rt}$; in the variable interest rate setting we still have $A(t + \delta t) = e^{r\delta t}A(t)$, however r might vary from time to time, and even (if interest rates are stochastic) from one binomial subtree to another. With this this convention, the prescription for determining the *price of a tradeable security* becomes

$$V_{\text{now}}/A_{\text{now}} = E_{\text{RN}}[V_{\text{next}}/A_{\text{next}}]$$

since $A_{\text{now}}/A_{\text{next}} = e^{-r\delta t}$ where r is the risk-free rate. (This relation is valid even if the risk-free rate varies from one subtree to the next). Working backward in the tree, this relation generalizes to one relating the option value at any pair of times $0 \le t < t' \le T$:

$$V(t)/A(t) = E_{\text{RN}}[V(t')/A(t')].$$

Here, as usual, the risk-neutral expectation weights each state at time t' by the probability of reaching it via a coin-flipping process starting from time t – with independent, biased coins at each node of the tree, corresponding to the risk-neutral probabilities of the associated subtrees.

The preceding results say, in essence, that certain processes are *martingales*. Concentrating on binomial trees, a "process" is just a function g whose values are defined at every node. A process is said to be a *martingale* relative to the risk-neutral probabilities if it satisfies

$$g(t) = E_{\text{RN}}[g(t')]$$

for all $t < t'$. The risk-neutral probabilities are determined by the fact that

- $s(t)/A(t)$ is a martingale relative to the risk-neutral probabilities

where $s(t)$ is the stock price process (for a non-dividend-paying stock), or equivalently by the fact that

- $\mathcal{F}(t)$ is a martingale relative to the risk-neutral probabilities

where $\mathcal{F}(t)$ is a futures price. Options are tradeables, so the value V of any option is determined by the condition that

- $V(t)/A(t)$ is a martingale relative to the risk-neutral probabilities.

One advantage of this framework is that it makes easy contact with the continuous-time theory. The central connection is this: in continuous time, the solution of a stochastic differential equation $dy = f dt + g dw$ is a martingale exactly if $f = 0$.

We can use this insight to explain and/or confirm some results previously obtained by other means. We return here to the constant-interest-rate environment, so $A(t) = e^{rt}$, and we focus (just to be specific) on options on a non-dividend-paying stock (rather than on a futures price).

Question: why does the risk-neutral stock price process satisfy $ds = rsdt + \sigma sdw$? Answer: because the risk-neutral stock price has the property that $s(t)/A(t) = s(t)e^{-rt}$ is a martingale. Explanation: if we assume that the risk-neutral price process has the form $ds = f dt + g dw$ for some f, we easily find that

$$d(se^{-rt}) = e^{-rt} ds - re^{-rt} sdt = (f - rs)dt + e^{-rt} gdw.$$

So se^{-rt} is a martingale exactly if $f = rs$. (You may wonder why the risk-neutral stock price process has the same *volatility* as the subjective stock price process. This is because changing the drift has the effect of re-weighting the probabilites of paths, without actually changing the set of "possible" paths; changing the volatility on the other hand has the effect of considering an entirely different set of "possible paths." This is the essential content of Girsanov's theorem, which is discussed in the class Stochastic Calculus, and applied in the course Continuous Time Finance.)

Question: why does the option price satisfy the Black-Scholes PDE? Answer: because the option price normalized by $A(t)$ must be a martingale. Explanation: suppose the option price has the form $V(s(t), t)$ for some function $V(s, t)$. Then

$$\begin{aligned} d\left(V(s(t),t)e^{-rt}\right) &= e^{-rt} dV - re^{-rt} V dt \\ &= e^{-rt}(V_t dt + V_s ds + \tfrac{1}{2} V_{ss}\sigma^2 s^2 dt) - re^{-rt} V dt \\ &= e^{-rt}(V_t + rsV_s + \tfrac{1}{2}\sigma^2 s^2 V_{ss} - rV)dt + e^{-rt}\sigma s V_s dw. \end{aligned}$$

For this to be a martingale the coefficient of dt must vanish. That is exactly the Black-Scholes PDE.

Question: why does the solution of the Black-Scholes PDE give the discounted expected payoff of the option? Answer: because the option price normalized by $A(t)$ is a martingale. Explanation: suppose V solves the Black-Scholes PDE, with final value $V(s, T) = f(s)$. We have shown that $e^{-rt}V(s(t), t)$ is a martingale. Therefore

$$V(s(0), 0) = E_{\mathrm{RN}}\left[e^{-rt} V(s(t), t)\right]$$

for any $t > 0$. Bringing e^{-rt} out of the expectation and setting $t = T$ gives

$$V(s(0), 0) = e^{-rT} E_{\text{RN}} [V(s(T), T)] = e^{-rT} E_{\text{RN}} [f(s(T))]$$

as asserted.

The preceding questions and answers are, of course, simply convenient reorganizations our prior calculations connecting risk-neutral expectation to the Black-Scholes PDE.

Why, exactly, must $V(t)/A(t)$ be a martingale, if V is the price of a tradeable? In discrete time, this is true because $V(0)/A(0) = V(0)$ is the intial cost of a self-financing trading strategy that replicates the value of V at time T. In continuous time the same assertion holds. In general it is a consequence of the *martingale representation theorem*, which lies beyond the scope of this class. But for an option on a lognormal stock in a constant interest rate environment the argument reduces to our second explanation of the Black-Scholes PDE (bottom of page 7, Section 6). In fact, in that setting $V(t)/A(t) = V(t)e^{-rt}$ is a martingale because

$$d(e^{-rt} V(s(t), t)) = e^{-rt} \sigma s V_s \, dw$$

by Ito combined with the Black-Scholes PDE. This is equivalent (another application of Ito) to

$$dV(s(t), t) = \sigma s V_s \, dw + rV \, dt = V_s \, ds + r(V - sV_s) \, dt.$$

This equation is familiar: we used it in Section 6 to know that our trading strategy (holding V_s units of stock and a bond worth $V - sV_s$ at each time) was self-financing. In summary: if we knew nothing about the Black-Scholes PDE, but we knew that $V(t)e^{-rt}$ was a martingale (and a little more: we would need the coeffient of dw in the SDE satisfied by Ve^{-rt}), we could identify – by arguing as above – a trading strategy with initial cost $V(0)$ that replicates the option.

****

(c) What if you only hedge at discrete times? From the viewpoint of a basic, practical class like Derivative Securities this topic is a digression. But theoretically-minded students may find it interesting. My discussion follows the beginning of a paper by H. E. Leland, *Option pricing and replication with transaction costs*, J. Finance 40 (1985) 1283-1301 (available online through JSTOR). A thoughtful, quite readable discussion of this topic is the paper by E. Omberg, *On the theory of perfect hedging*, Advances in Futures and Options Research 5 (1991) 1-29 (not available online unfortunately). Making a choice, I'll focus on the hedging of a European option on a non-dividend-paying stock. A parallel discussion can however be given for an option on a forward price.

Suppose an investment bank sells an option and tries to replicate it dynamically, but the bank trades only at evenly spaced time intervals $j\delta t$. (Now δt is positive, not infinitesimal). The bank follows the standard trading strategy of rebalancing to hold $\phi = \partial V/\partial s$ units of stock each time it trades, where V is the value assigned by the Black-Scholes theory. As we shall see in a moment, this strategy is no longer self-financing – but it is *nearly so*, in a suitable stochastic sense, in the limit $\delta t \to 0$.

People often ask, when examining the derivation of the Black-Scholes PDE by examination of the hedging strategy, "Why do we apply Ito's lemma to $V(s(t), t)$ but not to Δ, even though the choice of Δ also depends on $s(t)$?" The answer, of course, is that the hedge portfolio is held fixed from t to $t + \delta t$. The following discussion – in which δt is small but not infinitesimal – should help clarify this point.

OK, let's return to that investment bank. The question is: how much additional money will the bank have to spend over the life of the option as a result of its discrete-time (rather than continuous-time) hedging? We shall answer this by considering each discrete time interval, then adding up the results.

The bank holds a short position on the option and a long position in the replicating portfolio. The value of its position just after rebalancing at any time $t = j\delta t$ is (by hypothesis)

$0 = -V(s(t), t) + \phi s(t) + [V(s(t), t) - \phi s(t)]$ = short option + stock position + bond position

with $\phi = \frac{\partial V}{\partial s}(s(t), t)$. The value of its position just before the next rebalancing is

$$-V(s(t + \delta t), t + \delta t) + \phi s(t + \delta t) + [V(s(t), t) - \phi s(t)]e^{r\delta t}.$$

The cost (or benefit) of rebalancing at time $t + \delta t$ is minus the value of the preceding expression. Put differently: it is the difference between the two preceding expressions. So it equals

$$\delta V - \phi\delta s - [V - \phi s](e^{r\delta t} - 1).$$

If we estimate δV by Taylor expansion keeping just the terms one normally keeps in Ito's lemma, we get (remembering that $\phi = \partial V/\partial s$)

$$\frac{\partial V}{\partial s}\delta s + \tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}(\delta s)^2 + \frac{\partial V}{\partial t}\delta t - \frac{\partial V}{\partial s}\delta s - rV\delta t + rs\frac{\partial V}{\partial s}\delta t.$$

Notice that the first and fourth terms cancel. Also notice that the substitution $(\delta s)^2 = \sigma^2 s^2 \delta t$ leads to an expression that vanishes, according to the Black-Scholes equation. Thus, the failure to be self-financing is attributable to two sources: (a) errors in the approximation $(\delta s)^2 \approx \sigma^2 s^2 \delta t$, and (b) higher order terms in the Taylor expansion. Our task is to estimate the associated costs.

Collecting the information obtained so far: if the investment bank re-establishes the "replicating portfolio" demanded by the Black-Scholes analysis at each multiple of δt then it incurs cost

$$\tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}(\delta s)^2 + \frac{\partial V}{\partial t}\delta t - rV\delta t + rs\frac{\partial V}{\partial s}\delta t$$

at each time step, plus an error of magnitude $|\delta t|^{3/2}$ due to higher order terms in the Taylor expansion. Using the Black-Scholes PDE, this cost has the alternative expression

$$\tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}[(\delta s)^2 - \sigma^2 s^2 \delta t] \quad \text{plus an error of order } |\delta t|^{3/2}.$$

It can be shown that when $ds = (\mu + \frac{1}{2}\sigma^2)s\,dt + \sigma s\,dw$,

$$\delta s = \sigma s u\sqrt{\delta t} + (\mu + \tfrac{1}{2}\sigma^2)s\delta t \quad \text{plus an error of order } |\delta t|^{3/2}$$

where u is Gaussian with mean 0 and variance 1 (this is closely related to our our discussion of Ito's lemma). Therefore

$$(\delta s)^2 = \sigma^2 s^2 u^2 \delta t \quad \text{plus an error of order } |\delta t|^{3/2}.$$

Thus neglecting the error terms, the cost of refinancing at any given timestep is

$$\tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2 (u^2 - 1)\delta t$$

where u is Gaussian with mean value 0 and variance 1. This expression is obviously random; its expected value is 0 and its standard deviation is of order δt. Moreover the contributions associated with different time intervals are independent. Notice that the distribution of refinancing costs is *not* Gaussian, since it is proportional to $u^2 - 1$ not u.

Pulling this together: since the expected value of $u^2 - 1$ is zero, the *expected cost* of refinancing at any given timestep is at most of order $|\delta t|^{3/2}$, due entirely to the "error terms." However the *actual cost* (or benefit) of refinancing is larger, a random variable of order δt. But the picture changes when we consider many time intervals. Over $n = T/\delta t$ intervals, the terms $\frac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2 (u^2 - 1)\delta t$ accumulate to a sum

$$\sum_{j=1}^{n} \tfrac{1}{2}\sigma^2 s^2(t_j)\frac{\partial^2 V}{\partial s^2}(s(t_j), t_j)(u_j^2 - 1)\delta t$$

with mean 0 and standard deviation of order $\sqrt{n\delta t^2} = \sqrt{T\delta t}$; the sum is still random, but it's small, statistically speaking, if δt is close to zero, by a sort of law-of-large-numbers. (Notice the resemblance of this argument to our explanation of Ito's lemma. That's no accident: we are in essence deriving Ito's formula all over again.) We've been ignoring the error terms – but they cause no trouble, because they too accumulate to a term of order $\sqrt{\delta t}$, because $n(\delta t)^{3/2} = T\sqrt{\delta t}$.

Final conclusion: the errors of refinancing tend to self-cancel, by a sort of law-of-large-numbers, since their mean value is 0. The net effect, when δt is small, is random but small — in the sense that its mean and standard deviation are of order $\sqrt{\delta t}$.

We have argued that the cost of refinancing tends to zero as $\delta t \to 0$. An article by A. Lo, D. Bertsimas, and L. Kogan goes further, examining the statistical distribution of refinancing costs when δt is small. (The relation between their work and the preceding discussion is like the relation between the central limit theorem and the law of large numbers.) The reference is: J. Financial Economics 55 (2000) 173-204 (available online through sciencedirect.com).

****

(d) Path-dependent options. We explained in part (b) why the value of an option is its discounted expected payoff. This is true even for path-dependent options, whose payoff isn't a function of s_T alone. As a result, Monte Carlo methods are a crucial tool for valuing path-dependent (exotic) options. There are many different varieties of path-dependent options (see Hull's Chapter 25). Some examples: an *Asian* option (whose payoff is determined by the average price of the underlying from time 0 to T); a *lookback* option (whose payoff is determined by the maximum or minimum price of the underlying from time 0 to T); or a *down-and-out barrier* option (whose payoff is 0 if the underlying ever goes below the barrier, and a specified function of the final-time value otherwise). Most path-dependent options cannot be valued by working backward in the tree, since you can't even get started: being path-dependent, the option doesn't have a well-defined value at the final-time nodes. (Barrier options are the exception – due to their special structure, valuation by working backward in the tree possible – and is actually the preferred technique, since Monte Carlo is computationally expensive.)

Serious Monte Carlo calculations must be done numerically, and they involve sampling (rather than enumerating all paths as we do below). But to illustrate the idea let's consider a two-period example that's amenable to by-hand calculation. Consider a two-year call struck at 90 on a stock with current price $S_0 = 100$. We'll take $r = 5\%$, $d = 4\%$, $\sigma = 20\%$, and we'll work with the forward price tree; notice that $F_0 = S_0 e^{(r-d)T} = 100e^{(.05-.04)2} = 102.02$. One's first thought might be to use a multiplicative tree with $\delta t = 1$ and $u = \exp(-\frac{1}{2}\sigma^2\delta t + \sigma\sqrt{\delta t}) = e^{.18}$, $d = \exp(-\frac{1}{2}\sigma^2\delta t - \sigma\sqrt{\delta t}) = e^{-.22}$, but we would rather take the risk-neutral probability of the up state be exactly 1/2. This affects the tree, as we'll see in a moment.

We already noted that our two-period forward price tree starts from $F_0 = 102.02$. The other nodes should be

$$\begin{aligned}
F_u &= 102.02e^{.18} = 122.14 \\
F_d &= 2(102.02) - 122.14 = 81.9 \quad \text{to assure that } F_0 = \tfrac{1}{2}F_u + \tfrac{1}{2}F_d \\
F_{uu} &= 102.02e^{2(.18)} = 146.23 \\
F_{ud} &= 2(122.14) - 146.23 = 98.05 \quad \text{to assure that } F_u = \tfrac{1}{2}F_{uu} + \tfrac{1}{2}F_{ud} \\
F_{dd} &= 2(81.90) - 98.05 = 65.75 \quad \text{to assure that } F_d = \tfrac{1}{2}F_{ud} + \tfrac{1}{2}F_{dd}
\end{aligned}$$

We have taken care to make F_t a martingale, since this property is crucial: it assures that the underlying is correctly priced. (Our procedure is equivalent to making the forward price tree multiplicative with $d = 2 - u$, so that $(u + d)/2 = 1$). Note that once the forward prices have been determined, so are the stock prices, e.g. $S_u = F_u e^{-.01} = 120.92$ and $S_d = F_d e^{-.01} = 81.09$ (see the figure).

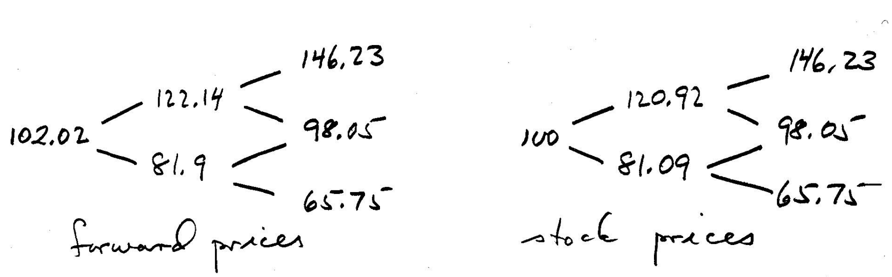

Let's use this tree to price a call struck at 90 on the average value of S_t at times 1 and 2: enumerating the cases,

- in scenario uu the payoff is $(120.92 + 146.23)/2 - 90 = 43.58$
- in scenario ud the payoff is $(120.92 + 98.05)/2 - 90 = 19.49$
- in scenario du the payoff is 0
- in scenario dd the payoff is 0

so the value of the option is

$$e^{-rT} E_{\mathrm{RN}}[\text{payoff}] = e^{-.1}[(1/4)43.58 + (1/4)19.49] = 14.27$$

As a second example, let's price a knock-out barrier call with barrier 85 and strike 90:

- in scenario uu the payoff is 56.23
- in scenario ud the payoff is 8.05
- in scenario du the payoff is 0
- in scenario dd the payoff is 0

so the value of the option is

$$e^{-rT} E_{\mathrm{RN}}[\text{payoff}] = e^{-.1}[(1/4)56.23 + (1/4)8.05] = 14.54.$$

The Asian example could not have been valued by working backward in the tree. (We couldn't have even gotten started: the final-time value depends on the path, not on the final-time node.) The barrier option could however have alternatively been valued by working backward (due to its special structure). In fact:

- value at node uu is 56.23; value at node ud is 8.05; value at node dd is 0
- value at node u is $e^{-.05}(\frac{1}{2}56.23 + \frac{1}{2}8.05) = 30.57$
- value at node d is 0
- value at initial node is $e^{-.05}(\frac{1}{2}30.57 + 0) = 14.54$.

The preceding examples used a binomial tree, but this was purely a matter of choice. It can, for example, be convenient to use a trinomial tree instead of a binomial tree. (One reason: in the binomial setting the price oscillates depending whether the number of steps is even or odd, see e.g. Hull's Figure 20.4; trinomial trees avoid this. Another reason: in pricing a barrier option it is a good idea to have a node at the barrier at each timestep; trinomial trees permit this.) Different trees correspond to different discrete approximations of the risk-neutral price process. The fact that a trinomial market is incomplete doesn't matter, since our fundamental understanding is now based on the continuous time theory. The use of trees is just a numerical approximation to that theory.

****

(e) American options revisited. We learned early in the semester how to price an American option by working backward in a tree. American options *must* be priced this way; Monte Carlo *cannot* be used. (Well, I exaggerate a little bit. American options can also be solved using a PDE free boundary problem, as explained below; but that's the continuous-time version of working backward in the tree. There are some Monte Carlo type alternatives, see Hull's Section 26.8 and the material below, but they are much more complicated than the basic Monte Carlo discussed above.)

Valuation using the discounted expected payoff. For a European option, we saw that the value assigned by the binomial tree was expressible in the form $e^{-rT}E_{\mathrm{RN}}[f(s(T))]$. A similar calculation applies to the American option – however $f(s(T))$ must be replaced by the value realized *at exercise*: the value of the option is $E_{\mathrm{RN}}\left[e^{-r\tau}f(s(\tau))\right]$ where τ is the exercise time. Once we've worked backward through the tree we know how to determine τ – for each realization of the risk-neutral process, it's the first time that realization reaches a node of the tree associated with early exercise (or T, if that realization does not reach an "early-exercise" node).

We see now why a straightforward Monte Carlo scheme cannot be used to value an American option: one would have to know τ. When we work backward in the tree we find the exercise rule as part of the calculation. Without working backward in the tree, we seem to have no idea about τ.

Actually, there is (at least conceptually) a Monte-Carlo-based way to determine τ. It uses the following theorem for American options:

$$\text{Value} = \max_{\text{exercise rules}} E_{\mathrm{RN}}\left[e^{-r\tau}f(s(\tau))\right].$$

In other words the exercise rule selected by backsolving the binomial tree is the one that maximizes the discounted expected payoff. An honest proof of this fact is far from trivial – partly because it requires formalization of what one means by an "exercise rule." But here is a rough idea how it goes. Any exercise rule determines a hedging strategy, i.e. a synthetic option that is available in the marketplace. So the max over exercise rules gives a lower bound for the value of the option. Our strategy of working backward through the tree gives an upper bound. The two bounds agree since the value obtained by working backward through the tree is associated with a special exercise rule. (Finding τ in practice using this theorem is difficult, since any algorithm for optimization requires evaluating the function being optimized many times, and each Monte Carlo evaluation is computationally expensive.)

Valuation of American options using a PDE. (This material is not in Hull; you can find a brief summary in the "student guide" by Wilmott-Howison-Dewynne.) For a European option the continuous-time limit of working backward through the tree amounts to solving the Black-Scholes PDE for $t < T$, with final data $f(s)$ at $t = T$. There is an analogous statement for an American option, however the PDE is replaced by a *free boundary problem*:

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2 + rs\frac{\partial V}{\partial s} - rV \leq 0,$$

$$V(s,t) \geq f(s),$$

and

$$\frac{\partial V}{\partial t} + \tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2 + rs\frac{\partial V}{\partial s} - rV = 0 \quad \text{or} \quad V(s,t) = f(s).$$

The logic behind the first inequality is this: in our derivation of the Black-Scholes PDE, the crucial juncture was when we saw that the choice $\phi = \partial V/\partial s$ made $d(V - \phi s)$ deterministic:

$$d(V - \phi s) = \left(\frac{\partial V}{\partial t} + \tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2\right) dt.$$

We concluded, by the principle of no arbitrage, that this must equal $r(V - \phi s)dt$. But that arbitrage argument assumed that you continued to hold the option. In the present context, where early exercise is permitted, the absence of arbitrage gives a weaker conclusion: the deterministic portfolio $(V - \phi s)$ can grow *no faster* than the risk-free rate. Thus

$$\left(\frac{\partial V}{\partial t} + \tfrac{1}{2}\frac{\partial^2 V}{\partial s^2}\sigma^2 s^2\right) \leq r\left(V - \frac{\partial V}{\partial s}s\right);$$

this is our first inequality. The logic behind the second inequality is obvious: the value is no smaller than can be realized by immediate exercise. The third relation simply says that one of the first two relations always holds – because for any given (s,t) the optimal strategy involves either holding the option a little longer (in which case the Black-Scholes equation applies) or exercising it immediately.

We call this a free-boundary problem because the value is still governed by the Black-Scholes PDE in *some* region of the (s,t) plane – the region where immediate exercise isn't optimal – however this region isn't given as data but must be found as part of the problem. Schematically:

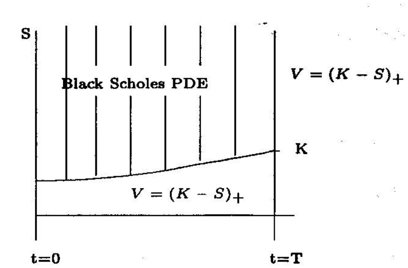

Figure 1: Schematic of the free boundary problem whose solution values an American put.

One can show that V and $\Delta = \partial V/\partial s$ are both continuous across the free boundary. Of course, on the "exercise" side of the boundary $V = f(s)$ and $\partial V/\partial s = f'(s)$ are known, giving two boundary conditions. If the domain of the PDE were known then just one boundary condition would be permitted; but the domain isn't known, and the extra boundary condition serves to fix the free boundary.