Derivative Securities – Fall 2007 – Section 10

Notes by Robert V. Kohn, extended and improved by Steve Allen.
Courant Institute of Mathematical Sciences.

Options on interest-based instruments: pricing of bond options, caps, floors, and swaptions. This section provides an introduction to valuation of options on interest rate products. We focus on two approaches: (i) Black's model, and (ii) trees. Briefly: we're taking the same methods developed earlier this semester for options on a stock or forward price, and applying them to interest-based instruments.

The material discussed here can be found in chapters 26, 27, and 28 of Hull; I'll also take some examples from the book *Implementing Derivatives Models* by Clewlow and Strickland, Wiley, 1998. (Steve Allen's version of these notes includes a few pages about how a Monte Carlo scheme based on the Heath-Jarrow-Morton theory works. I don't attempt that here.)

****

Black's model. Recall the formulas derived in Section 5 for the value of a put or call on a forward price:

$$\begin{aligned} c[\mathcal{F}_0, T; K] &= e^{-rT}\left[\mathcal{F}_0 N(d_1) - K N(d_2)\right] \\ p[\mathcal{F}_0, T; K] &= e^{-rT}\left[K N(-d_2) - \mathcal{F}_0 N(-d_1)\right] \end{aligned}$$

where

$$\begin{aligned} d_1 &= \frac{1}{\sigma\sqrt{T}}\left[\log(\mathcal{F}_0/K) + \tfrac{1}{2}\sigma^2 T\right] \\ d_2 &= \frac{1}{\sigma\sqrt{T}}\left[\log(\mathcal{F}_0/K) - \tfrac{1}{2}\sigma^2 T\right] = d_1 - \sigma\sqrt{T}. \end{aligned}$$

Black's model values interest-based instruments using almost the same formulas, suitably interpreted. One important difference: since the interest rate is no longer constant, we replace the discount factor e^{-rT} by $B(0,T)$.

The essence of Black's model is this: consider an option with maturity T, whose payoff $\phi(V_T)$ is determined by the value V_T of some interest-related instrument (a discount rate, a term rate, etc). For example, in the case of a call $\phi(V_T) = (V_T - K)_+$. Black's model stipulates that

(a) the value of the option today is its discounted expected payoff.

No surprise there – it's the same principle we've been using all this time for valuing options on stocks. If the payoff occurs at time T then the discount factor is $B(0,T)$ so statement (a) means

$$\text{option value} = B(0,T)E_*[\phi(V_T)].$$

We write E_* rather than E_{RN} because in the stochastic interest rate setting this is *not* the risk-neutral expectation; we'll explain why E_* is different from the risk-neutral expectation later on. For the moment however, we concentrate on making Black's model computable. For this purpose we simply specify that (under the distribution associated with E_*)

(b) the value of the underlying instrument at maturity, V_T, is lognormal; in other words, V_T has the form e^X where X is Gaussian.

(c) the mean $E_*[V_T]$ is the forward price of V (for contracts written at time 0, with delivery date T).

We have not specified the variance of $X = \log V_T$; it must be given as data. It is customary to specify the "volatility of the forward price" σ, with the convention that

$$\log V_T \text{ has standard deviation } \sigma\sqrt{T}.$$

Notice that the Gaussian random variable $X = \log V_T$ is fully specified by knowledge of its standard deviation $\sigma\sqrt{T}$ and the mean of its exponential $E_*[e^X]$, since if X has mean m then $E_*[e^X] = \exp(m + \frac{1}{2}\sigma^2 T)$.

Most of the practical examples involve calls or puts. For a call, with payoff $(V_T - K)_+$, hypothesis (b) gives

$$E_*[(V_T - K)_+] = E_*[V]N(d_1) - KN(d_2)$$

where

$$d_1 = \frac{\log(E_*[V_T]/K) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\log(E_*[V_T]/K) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.$$

This is a direct consequence of the lemma we used long ago (in Section 5) to evaluate the Black-Scholes formula. Using hypotheses (a) and (c) we get

$$\text{value of a call} = B(0,T)[F_0 N(d_1) - KN(d_2)]$$

where F_0 is the forward price of V today, for delivery at time T, and

$$d_1 = \frac{\log(F_0/K) + \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}}, \quad d_2 = \frac{\log(F_0/K) - \frac{1}{2}\sigma^2 T}{\sigma\sqrt{T}} = d_1 - \sigma\sqrt{T}.$$

A parallel discussion applies for a put.

It is by not obvious (at least not to me) that Black's formula is correct in a stochastic interest rate setting. We'll justify it a little later, for options on bonds. But here is a rough, heuristic justification. Since the value of the underlying security is stochastic, we may think of it as having its own lognormal dynamics. If we treat the risk-free rate as being constant then Black's formula can certainly be used. Since the payoff takes place at time T, the only reasonable constant interest rate to use is the one for which $e^{-rT} = B(0,T)$, and this leads to the version of Black's formula given above.

Black's model applied to options on bonds. The following examples is taken from Clewlow and Strickland (section 6.6.1). Let us price a one-year European call option on a 5-year discount bond. Assume:

- The current term structure is flat at 5 percent per annum; in other words $B(0,t) = e^{-.05t}$ when t is measured in years.

- The strike of the option is 0.8; in other words the payoff is $(B(1,5)-0.8)_+$ at time $T=1$.
- The forward bond price volatility σ is 10 percent.

Then the forward bond price is $F_0 = B(0,5)/B(0,1) = .8187$ so

$$d_1 = \frac{\log(.8187/.8000) + \frac{1}{2}(0.1)^2(1)}{(0.1)\sqrt{1}} = 0.2814, \quad d_2 = d_1 - \sigma\sqrt{T} = 0.2814 - 0.1\sqrt{1} = .1814$$

and the discount factor for income received at the maturity of the option is $B(0,1) = .9512$. So the value of the call option now, at time 0, is

$$.9512[.8187N(.2814) - .8N(.1814)] = .0404.$$

Black's formula can also be used to value options on coupon-paying bonds; no new principles are involved, but the calculation of the forward price of the bond must take into account the coupons and their payment dates; see Hull's Example 26.1.

One should avoid using the same σ for options with different maturities. And one should never use the same σ for underlyings with different maturities. Here's why: suppose the option has maturity T and the underlying bond has maturity $T' > T$. Then the value V_t of the underlying is known at both $t=0$ (all market data is known at time 0) and at $t=T'$ (all bonds tend to their par values as t approaches maturity). So the variance of V_t vanishes at both $t=0$ and $t=T'$. A common model (if simplified) model says the variance of V_t is $\sigma_0^2 t(T'-t)$ with σ_0 constant, for all $0 < t < T'$. In this case the variance of V_T is $\sigma_0^2 T(T'-T)$, in other words $\sigma = \sigma_0\sqrt{T'-T}$. Thus σ depends on the time-to-maturity $T'-T$. In practice σ – or more precisely $\sigma\sqrt{T}$ – is usually inferred from market data.

Black's model applied to caps. A cap provides, at each coupon date of a bond, the difference between the payment associated with a floating rate and that associated with a specified cap rate, if this difference is positive. The ith caplet is associated with the time interval (t_i, t_{i+1}); if $R_i = R(t_i, t_{i+1})$ is the term rate for this interval, R_K is the cap rate, and L is the principal, then the ith caplet pays

$$L \cdot (t_{i+1} - t_i) \cdot (R_i - R_K)_+$$

at time t_{i+1}. Its value according to Black's formula is therefore

$$B(0, t_{i+1}) L \Delta_i t [f_i N(d_1) - R_K N(d_2)].$$

Here $\Delta_i t = t_{i+1} - t_i$; $f_i = f_0(t_i, t_{i+1})$ is the forward term rate for time interval under consideration, defined by

$$\frac{1}{1 + f_i \Delta_i t} = \frac{B(0, t_{i+1})}{B(0, t_i)};$$

and

$$d_1 = \frac{\log(f_i/R_K) + \frac{1}{2}\sigma_i^2 t_i}{\sigma_i\sqrt{t_i}}, \quad d_2 = \frac{\log(f_i/R_K) - \frac{1}{2}\sigma_i^2 t_i}{\sigma_i\sqrt{t_i}} = d_1 - \sigma_i\sqrt{t_i}.$$

The volatilities σ_i must be specified for each i; in practice they are inferred from market data. The value of a cap is obtained by adding the values of its caplets.

A floor is to a cap as a put is to a call: using the same notation as above, the ith floorlet pays

$$L\Delta_i t(R_K - R_i)_+$$

at time t_{i+1}. Its value according to Black's formula is therefore

$$B(0, t_{i+1})L\Delta_i t[R_K N(-d_2) - f_i N(-d_1)]$$

where d_1 and d_2 are as above. The value of a floor is obtained by adding the values of its floorlets.

The market convention is to quote a single volatility for a cap or floor which is then applied to each of the constituent caplets or floorlets – but this is just a convention to make it easy to communicate. To actually price a cap or floor one must evaluate each individual FRA option (caplet or floorlet) at the appropriate volatility, then sum the resulting prices to arrive at the price of the cap or floor. Then one can solve for a single volatility that, applied to each individual FRA option, would give the same price.

Here's an example, taken from Section 26.3 of Hull. Consider a contract that caps the interest on a 3-month, \$10,000 loan one year from now; we suppose the interest is capped at 8% per annum (compounded quarterly). This is a simple caplet, with $t_1 = 1$ year and $t_2 = 1.25$ years. To value it, we need:

- The forward term rate for a 3-month loan starting one year from now; suppose this is 7% per annum (compounded quarterly).
- The discount factor associated to income 15 months from now; suppose this is .9169.
- The volatility of the 3-month forward rate underlying the caplet; suppose this is 0.20.

With this data, we obtain

$$d_1 = \frac{\log(.07/.08) + \frac{1}{2}(0.2)^2(1)}{0.2\sqrt{1}} = -0.5677, \quad d_2 = d_1 - 0.2\sqrt{1} = -0.7677$$

so the value of the caplet is, according to Black's formula,

$$(.9169)(10,000)(1/4)[.07N(-.5677) - .08N(-.7677)] = 5.162 \text{ dollars.}$$

Black's model applied to swaptions. A swaption is an option to enter into a swap at some future date t (the maturity of the option), with a specified fixed rate c and frequency f, which lasts until a specified later time T. (This is evidently a $T - t$-year swap.) For example, if it's an option to pay the fixed rate, and if $c = 5.50\%$, $t = 3$ years and $T = 8$ years, then option can be exercised in 3 years to enter into a 5 year swap with a 5.50% coupon.

To value it, recall that if the payment times are t_j then the value of the swap at time t will be

$$V_{\text{swap}}(t) = L\left[\sum \frac{c}{f} B(t, t_j) + B(t, T) - 1\right],$$

and the par swap rate at time t will be the value of c that makes the right hand side equal to 0:

$$R_{\text{swap}}(t) \sum \frac{1}{f} B(t, t_j) = 1 - B(t, T).$$

Of course $R_{\text{swap}}(t)$ isn't known now, because it dependes on discount rates for lending at time t. But we get the *forward swap rate* F_{swap} by replacing $B(t, t_j)$ above by the forward rate $B(0, t_j)/B(0, t)$: after some arithmetic,

$$F_{\text{swap}}(0) \cdot \sum \frac{1}{f} B(0, t_j) = B(0, t) - B(0, T).$$

If the coupon is set to F_{swap}, then the swap has no value at time 0. (We could alternatively have reached the same conclusion from the formula at the top of page 8 in Section 9.) The forward swap rate can be calculated at any time $0 \leq \tau \leq t$ of course: arguing as above, it is

$$F_{\text{swap}}(\tau) \cdot \sum \frac{1}{f} B(\tau, t_j) = B(\tau, t) - B(\tau, T)$$

and it agrees with R_{swap} at $\tau = t$.

The swap will be in the money if the proposed coupon rate c is better than the swap rate R_{swap} when the option matures (time t). For a swap to receive floating rate and pay fixed rate, this occurs if $R_{\text{swap}} > c$. If it is exercised, the value of the swap at exercise is

$$(R_{\text{swap}} - c)\frac{L}{f} \sum B(t, t_j),$$

i.e. the exercised swap has the same value as a stream of payments of $(R_{\text{swap}} - c)\frac{L}{f}$ at each coupon date t_j. Black's Black's formula gives the time-0 value of the jth payment as

$$B(0, t_j)\frac{L}{f}[F_{\text{swap}} N(d_1) - cN(d_2)]$$

where F_{swap} is the forward swap rate and

$$d_1 = \frac{\log(F_{\text{swap}}/c) + \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}}, \quad d_2 = \frac{\log(F_{\text{swap}}/c) - \frac{1}{2}\sigma^2 t}{\sigma\sqrt{t}} = d_1 - \sigma\sqrt{t}.$$

To get the value of the swap itself we sum over all i:

$$\text{value of swap} = A[F_{\text{swap}} N(d_1) - cN(d_2)] \quad \text{where } A = \tfrac{L}{f} \textstyle\sum_{i=1}^{N} B(0, t_t).$$

Here σ is the volatility of forward swap rate F_{swap} (which would normally be determined by calibrating the predictions of the model to market prices).

The option to enter into a swap that receives the floating rate and pays the fixed rate uses the call option formula. By entirely similar reasoning, an option to enter into a swap that pays the floating rate and receives the fixed rate uses the put option formula. There is another type of interest rate option that consists of a swap that can be cancelled at a given point in time. The right to cancel a swap to receive the floating rate and pay the fixed rate is equivalent to having the option to enter into a swap paying the floating rate and receiving the fixed rate, thereby offsetting the existing swap. Therefore, the option to cancel a swap receiving floating and paying fixed uses the put option formula. Similarly, an option to cancel a swap paying floating and receiving fixed uses the call option formula.

Options on swaps can be either cash-settlement or settled by delivery. Settlement by delivery involves actually entering into the swap specified. Cash settlement means that the value of entering into this swap at the market rate prevailing at the time of settlement is calculated and then a cash payment is made of this value. Market convention is that caps and floors are always cash-settled.

When we calculate the value of an option on a swap, we are only looking at the value of the fixed rate bond portion of the swap (we have implicitly been assuming that we are always valuing options for dates on which the swap has just made a coupon payment, so that the floating rate bond portion is worth par). Therefore, options on fixed rate bonds can be valued using the exact same formula as options on swaps.

Here's an example, taken from Clewlow and Strickland section 6.6.1. Suppose the yield curve is flat at 5 percent per annum (continuously compounded). Let us price an option that matures in 2 years and gives its holder the right to enter a one-year swap with semiannual payments, receiving floating rate and paying fixed term rate 5 percent per annum. We suppose the volatility of the forward swap rate is 20% per annum.

The first step is to find the forward swap rate F_{swap}. It satisfies

$$F_{\text{swap}}(1/2)\sum_{i=1}^{2} B(0,t_i) = (B(0,t) - B(0,t_2))$$

with $t = 2$, $t_1 = 2.5$, and $t_2 = 3.0$. Since the yield curve is flat at 5% compounded continuously, we have $B(0,t) = e^{-.05t}$ for all t. A bit of arithmetic gives gives $F_{\text{swap}} = .0506$, in other words 5.06%. Now

$$d_1 = \frac{\log(.0506/.0500) + \frac{1}{2}(0.2)^2(2)}{0.2\sqrt{2}} = 0.1587, \quad d_2 = d_1 - 0.2\sqrt{2} = -.0971,$$

and

$$\sum_{i=1}^{2} B(0,t_i)(t_i - t_{i-1}) = \tfrac{1}{2}(e^{-(.05)(2.5)} + e^{-(.05)(3)}) = .8716,$$

so the value of the swaption is

$$.8716L[.0506N(.1587) - .05N(-.0971)] = .0052L$$

where L is the notional principal of the underlying swap.

****

When can Black's model be used? Why is it correct? Black's model is widely-used and appropriate for pricing European-style options on interest-based instruments. It has two key advantages: (a) simplicity, and (b) directness. By simplicity I mean not that Black's model is easy to understand, but rather that it requires just one parameter (the volatility) to be inferred from market data. By directness I mean that we model the underlying instrument directly – the basic hypothesis of Black's model is the lognormal character of the underlying.

Black's model *cannot* be used, however, to value American-style options, i.e. options with whose exercise date is not fixed in advance. Many bond options permit early exercise – sometimes American-style (permitting exercise at any time) but more commonly Bermudan (permitting exercise at a list of specified dates, typically coupon dates). Such options are best modelled using a tree (much as we did a few weeks ago for American options on equities). When working with interest rates, the tree models the risk-neutral interest-rate process, which can then be used to value bonds of all types and maturities, and American as well as European options on these bonds. Interest-rate trees are not "simple" in the sense used above: to get started we must calibrate the entire tree to market data (e.g. the yield curve). And they are not "direct" in the sense used above: we are modeling the risk-neutral interest rate process, not the underlying instrument itself; thus there are two potential sources of modeling error: one in modeling the value of the underlying instrument, the other in modeling how the option's value depends on that of the underlying instrument.

The simplicity and directness of Black's model are also responsible for its disadvantages. Black's model must be used separately for each class of instruments – we cannot use it, for example, to hedge a cap using bonds of various maturities. For consistent pricing and hedging of multiple instruments one must use a more fundamental model such as an interest rate tree.

Now we turn to the question of *why* Black's model is correct. The following explanation involves "change of numeraire". (The word numeraire refers to a choice of units.)

Up to now our numeraire has been cash (dollars). Its growth as a function of time is described by the money-market account introduced in Section 7. The money-market account has balance is $A(0) = 1$ initially, and its balance evolves in time by $A_{\text{next}} = e^{r\delta t} A_{\text{now}}$. We are accustomed to finding the value f of a tradeable instrument (such as an option) by working backward in the tree using the risk-neutral probabilities. At each step this amounts to

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

where q and $1-q$ are the risk-neutral probabilities of the up and down states. As we noted in Section 7, this can be expressed as

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

and it can be iterated in time to give

$$f(t)/A(t) = E_{\text{RN}}[f(t')/A(t')] \quad \text{for } t < t'.$$

This is captured by the statement that "$f(t)/A(t)$ is a martingale relative to the risk-neutral probabilities."

But sometimes the money-market account is not the convenient comparison. In fact we may use *any* tradeable security as the numeraire – though when we do so we must also change the probabilities. Indeed, for any tradeable security g there is a choice of probabilities on the tree such that

$$\frac{f_{\text{now}}}{g_{\text{now}}} = \left[q_* \frac{f_{\text{up}}}{g_{\text{up}}} + (1 - q_*) \frac{f_{\text{down}}}{g_{\text{down}}} \right].$$

This is an easy consequence of the two relations

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

which hold (using the risk-neutral q) since both f and g are tradeable. A little algebra shows that these relations imply the preceding formula with

$$q_* = \frac{q g_{\text{up}}}{q g_{\text{up}} + (1 - q) g_{\text{down}}}.$$

(The value of q_* now varies from one binomial subtree to another, even if q was uniform throughout the tree.) Writing E_* for the expectation operator with weight q_*, we have defined q_* so that

$$f_{\text{now}}/g_{\text{now}} = E_*[f_{\text{next}}/g_{\text{next}}].$$

Iterating this relation gives (as in the risk-neutral case)

$$f(t)/g(t) = E_*[f(t')/g(t')] \quad \text{for } t < t';$$

in other words "$f(t)/g(t)$ is a martingale relative to the probability associated with E_*." In particular

$$f(0)/g(0) = E_*[f(T)/g(T)]$$

where T is the maturity of an option we may wish to price.

Let us apply this result to explain why Black's formula is valid. For simplicity we focus on options whose maturity T is also the time the payment is received. (This is true for options on bonds, not for caplets or swaptions. The justification of Black's formula for caplets and swaptions is a little different; see Hull.) The convenient choice of g is then

$$g(t) = B(t, T).$$

Since $g(T) = 1$ this choice gives

$$f(0) = g(0) E_*[f(T)] = B(0, T) E_*[f(T)].$$

We shall apply this twice: once with f equal to the value of the underlying instrument, what we called V_t on page one of these notes; and a second time with f equal to the value of the option. The first application gives

$$E_*[V_T] = V_0/B(0, T)$$

and we recognize the right hand side as the *forward price* of the instrument. For this reason the probability distribution associated with this E_* is called *forward risk-neutral.* The second application gives

$$\text{option value} = B(0,T)E_*[\phi(V_T)]$$

where $\phi(V_T)$ is the payoff of the option – for example $\phi(V_T) = (V_T - K)_+$ if the option is a call.

This explains Black's formula, except for one crucial feature: the hypothesis that V_T is lognormal with respect to the distribution associated with E_* (the forward-risk-neutral distribution). This is of course only asserted in the continuous-time limit, and only if the risk-neutral interest rate process is itself lognormal. The assertion is most easily explained using continuous-time (stochastic differential equation) methods, and we will not attempt to address it here.

****

Interest rate trees and American style interest rate options. Thus far we have only discussed European-style interest rate options – ones where there is a single option exercise date. But there are interest rate derivatives which permit multiple exercise dates: a particularly popular product in the US Dollar interest rate market is the Bermudean swaption, an option to enter into a particular swap on any of a series of payment dates for the swap. For example, the underlying swap might be one to pay 6 month LIBOR and receive a fixed coupon of 6% payable semiannually through Sept. 30, 2017 (the "maturity date" of the swap). The Bermudean option could be exercisable every 6 months starting Sept. 30, 2008 and ending Sept. 30, 2013 – the maturity date of the swap stays fixed at Sept. 30, 2017.

Valuation of American style interest rate options is almost always performed using a binomial or trinomial tree. This is conceptually the same as what we did for American-style options on forwards. However, there is an important difference. In pricing options on equities, we specified a tree for market price of a forward, which had to be a martingale under the risk-neutral measure. (We used this to determine the risk-neutral measure.) For pricing interest-based instruments, we will specify a tree for the *short-term interest rate under the risk-neutral measure.* The short-term interest rate is not the price of an asset, so there is no reason for it to be a martingale.

[Let's pause for a brief digression. You might wonder why we don't build a tree based on the price of an asset, an interest rate forward or a swap, rather than on an interest rate. There are basically two reasons: (1) if we made the convenient assumption that the price of the asset follows a lognormal process, it might become so high on some nodes that it would imply a negative interest rate and negative interest rates only occur in very extraordinary circumstances; (2) as a swap gets close to maturity, its price must go towards par (its duration is getting shorter and shorter) so it certainly can't be described by a martingale).]

How does an interest rate tree work? The basic idea is shown in the figure: each node of the tree is assigned a risk-free rate, different from node to node; it is the one-period risk-free rate for the binomial subtree just to the right of that node.

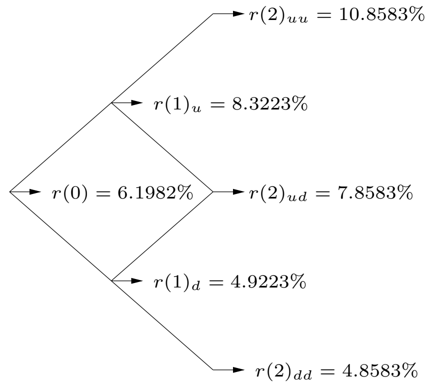

What probabilities should we assign to the branches? It might seem natural to start by figuring out what the subjective probabilities are. But why bother? All we really need for option pricing are the risk-neutral probabilities. When we discussed equities we used the volatility and drift of the forward price to establish a tree, then used its nodes to find the risk-neutral probabilities. But recall that when we considered the continuum limit $\delta t \rightarrow 0$, the risk-neutral probabilities were very close to 1/2. For an interest rate tree we have the right to choose the branching probabilities as we please; the usual practice is to make them exactly 1/2.

But we don't have complete freedom. The discount rates $B(0,T)$ are known at time 0 for all maturities T. To be useful, our interest rate tree must agree with this market data. In other words, it must be *calibrated* to the present term structure in the marketplace. In summary: for interest rate trees we

- restrict attention to the risk-neutral interest rate process.
- assume the risk-neutral probability is $q = 1/2$ at each branch, and
- choose the interest rates at the various nodes so that the long-term interest rates associated with the tree match those observed in the marketplace.

The last bullet – calibration of the tree to market information – is the most subtle one. We'll discuss it only briefly (this is a major focus of the class Interest and Credit Models).

But first let's just be sure we understand how such a tree determines long-term interest rates. As an example let's determine $B(0,3)$, the value at time 0 of a dollar received at time 3, for the tree shown in the figure. (Put differently: $B(0,3)$ is the price at time 0 of a zero-coupon bond which matures at time 3.) We take the convention that $\delta t = 1$ for simplicity.

Consider first time period 2. The value at time 2 of a dollar received at time 3 is $B(2,3)$; it has a different value at each time-2 node. These values are computed from the fact that

$$B(2,3) = e^{-r\delta t}[\tfrac{1}{2}B(3,3)_{\text{up}} + \tfrac{1}{2}B(3,3)_{\text{down}}] = e^{-r\delta t}$$

since $B(3,3) = 1$ in every state, by definition. Thus

$$B(2,3) = \begin{cases} e^{-r(2)_{uu}} = .897104 & \text{at node } uu \\ e^{-r(2)_{ud}} = .924425 & \text{at node } ud \\ e^{-r(2)_{dd}} = .952578 & \text{at node } dd. \end{cases}$$

Now we have the information needed to compute $B(1,3)$, the value at time 1 of a dollar received at time 3. Applying the rule

$$B(1,3) = e^{-r\delta t}[\tfrac{1}{2}B(2,3)_{\rm up} + \tfrac{1}{2}B(2,3)_{\rm down}]$$

at each node gives

$$B(1,3) = \begin{cases} e^{-r(1)_u}(\frac{1}{2}\cdot .897104 + \frac{1}{2}\cdot .924425) = .838036 & \text{at node } u \\ e^{-r(1)_d}(\frac{1}{2}\cdot .924425 + \frac{1}{2}\cdot .952578) = .893424 & \text{at node } d. \end{cases}$$

Finally we compute $B(0,3)$ by applying the same rule:

$$\begin{aligned} B(0,3) &= e^{-r\delta t}[\tfrac{1}{2}B(1,3)_{\rm up} + \tfrac{1}{2}B(1,3)_{\rm down}] \\ &= e^{-r(0)}[\tfrac{1}{2}\cdot .838036 + \tfrac{1}{2}\cdot .893424] = .8137. \end{aligned}$$

Valuing an option using an interest rate tree is easy: just work backward (the option is a tradeable). Hedging is easy too: each binomial submarket is complete, so a risky instrument can be hedged using any pair of interest-based insruments (for example, two distinct zero-coupon bonds).

Here is a toy (two-period) example to communicate the idea of calibration. Recall that if $A(t)$ is the balance of a money-market fund with value 1 at time 0, then the value of a European option with maturity T is $E_{\rm RN}[\text{payoff}/A(T)]$. For a zero-coupon bond with maturity T, the payoff is 1, so the value of the option is $E_{[}1/A(T)]$. Suppose that the present marketplace yields are $y(0,1) = 5\%$ and $y(0,2) = 6\%$, so $B(0,1) = e^{-.05} = .9512$ and $B(0,2) = e^{-.06*2} = .8869$. Suppose you have guessed that the interest rate tree (with branching probabilities 1/2) has nodal yields 5% initially, branching to 4% and 6% after one year as shown in the left side of the next figure. Such a tree gives the price of the discount bound as

$$B(0,2) = \tfrac{1}{2}\left[\frac{1}{1.06*1.05} + \frac{1}{1.04*1.05}\right] = \tfrac{1}{2}(.8985 + .9158) = .9071$$

which is far off from the observed value .8869. There is a systematic algorithm for correcting this (see e.g. sections 8.1-8.4 of Clelow and Strickland). The output of such an algorithm might for example be the tree shown on the right hand side of the figure. In fact, it gives the proper value for $B(0,2)$, since its prediction is

$$B(0,2) = \tfrac{1}{2}\left[\frac{1}{1.088*1.05} + \frac{1}{1.0592*1.05}\right] = \tfrac{1}{2}(.874 + .8992) = .8869.$$

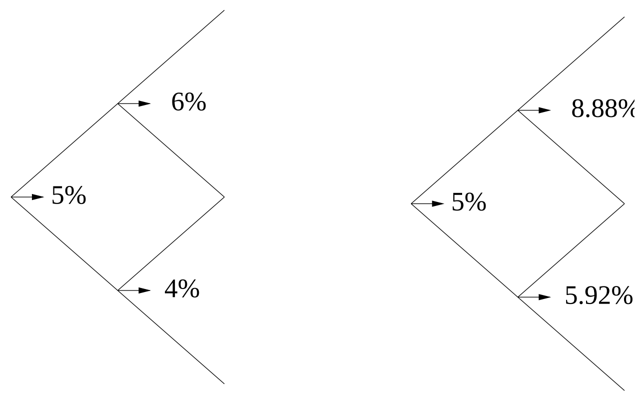

You may be disturbed by this example – how could a 5% interest rate evolve to either 8.88% or 5.92%? Not only is the average of these rates 7.40%, not equal to 5%, but even the lower of the two rates is greater than 5%. Don't we have an arbitrage? No we don't, because an interest rate is not the price of a tradable asset.

In fact, the rates in a tree are not related to one another by any theory. The 5% rate was the rate that applied to the first year; the 8.88% and 5.92% rates where the ones that applied in the second year. The observable quantity that's most closely related to 8.88% and 5.92% to is the current 1 year forward rate, which is 7% since $e^{-.06*2} = e^{-.05*2}e^{-.07*2}$.

A key issue in building interest rate trees is how much *mean reversion* to build into the tree. Mean reversion refers to a negative correlation between rate levels in one period and rate levels in the immediately succeeding period. When you have mean reversion, the expected rate for the following period will be lower than the current rate for higher rate nodes and higher than the current rate for lower rate nodes. Mean reversion is an issue that does not arise for trees that are being used to model tradable assets; tradable asset prices are martingales (more precisely: their present-valued prices are martingales), so that every node must have the current price equal to the expectation of the price at the following period. This is inconsistent with mean reversion. But interest rates are not tradable assets. One reason mean reversion is an economically reasonable assumption for interest rates is that the central bank tends to act as a counterbalancing force in the economy and pull rates back towards target levels. The reasonableness of mean reversion assumptions can be seen from historical data, by noting that the volatility of forward rates with later starting dates is lower than those with shorter starting dates.

Any interest-based instrument can be priced on a tree. To price a swap, you'll need access to more than a single year's worth of forward prices. But you can easily access this by calculations on the tree. For example, let's say you are at a node where you need to calculate the price of an annual pay swap with 3 years remaining that pays a fixed coupon of 6%. Let's say the 1 year rates at your node and its immediate branches look like the following figure: The one year discount factor is $e^{-.0725} = .9301$. The two year discount factor is the average of $e^{-.0725}e^{-.0895} = .8504$ and $e^{-.0725}e^{-.065} = .8715$, which is .8610. The three year discount factor is the average of $e^{-.0725}e^{-.0895}e^{-.0925} = .7753$, $e^{-.0725}e^{-.0895}e^{-.0775} = .7870$, $e^{-.0725}e^{-.065}e^{-.0775} = .8065$, and $e^{-.0725}e^{-.065}e^{-.0525} = .8270$, which is .8464. So the value

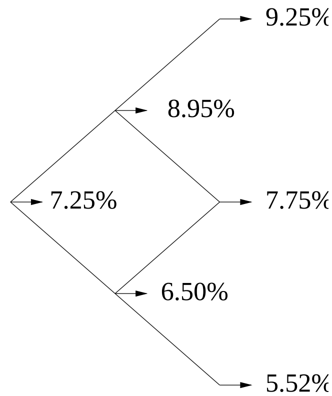

of the swap is $100 * .8464 + 6(.9301 + .8610 + .8464) - 100 = .4650$.

You may have noticed that we are implicitly assuming perfect correlation between movements in different segments of the rate curve by valuing the American-style swaptions on a tree. Most people in the financial industry believe that this method provides sufficient accuracy for valuing American-style swaptions, since mean reversion can be used to produce the same sort of relationship between European swaptions prices and American swaptions prices as is produced by assuming less-than-perfect correlation.

****

Convexity. Options models may be required for interest rate products that do not at first glance seem to be structured as options. Consider the following example: let.s say that a conventional forward rate agreement to receive the 3 month LIBOR resetting on April 30, 2010 is currently 4.50%. So a conventional FRA with a coupon rate of 4.50% will be priced at 0. What about an unconventional FRA with a 4.50% coupon rate to receive the 3 month LIBOR resetting on April 30, 2010 that , instead of making its settlement payment on July 30, 2010 (April 30 + 3 months), makes its settlement payment on April 30, 2010. At first it might seem that it too should be priced at 0, since it just represents the discounted value of a set of payments whose expected value is 0. But the rate at which this discounting will be done is directly tied to the rate at which the contract settles. If the 3 month LIBOR on April 30, 2010 turns out to be 6.50%, the discounted value the LIBOR receiver is owed will be 2.00%/(1+6.50%/4) = 1.96802, while if this LIBOR turns out to be 2.50%, the discounted value the LIBOR receiver owes will be 2.0%/(1+2.50%/4) = 1.98758. This asymmetry, known as *convexity*, will always disadvantage the LIBOR receiver and will be larger the greater the change in LIBOR away from 4.50%, so it has the same economic effect as an option on the forward LIBOR sold by the LIBOR receiver to the LIBOR payer. Since the unconventional FRA eliminates this asymmetry, it needs to have a higher coupon rate than the conventional FRA, to compensate the LIBOR payer for having lost this advantage.

Some other examples of convexity:

- *A LIBOR futures contract.* When forward prices and interest rates are uncorrelated, futures contracts and forward contracts for the same terms should have equal prices (the argument is very similar to the one from the Appendix to Hull's Chapter 5 that we discussed at the beginning of the semester). But for a futures contract based on an interest rate, there will be a correlation between price and interest rates. When rate levels rise, the receiver of LIBOR will get a cash payment reflecting the higher value of the future, which can then be reinvested till the settlement date at the now higher interest rates. When rate levels fall, the receiver of LIBOR will make a cash payment reflecting the lower value of the future, which can be borrowed till the settlement date at the now lower interest rates. So the futures rate needs to be higher than the forward rate to compensate the payer of LIBOR for this asymmetry that favors the LIBOR receiver.

- *A swap that pays a fixed amount times the change in swap rate.* We know from Section 9 that the value of a conventional swap is

$$L \sum \frac{c - R_{\text{swap}}}{f} B(0, t_j)$$

  So any change in the swap rate is being multiplied by the sum of a set of discount factors, which will be higher when rate levels are lower; this is favorable to the fixed rate receiver. If an unconventional swap pays based on the change in swap rates multiplied by a fixed amount, it will be more favorable to the floating rate receiver than the conventional swap, so the break-even coupon rate for an unconventional swap must be set higher than the break-even coupon rate for a conventional swap.

- *The cheapest-to-deliver option on a Treasury bond future.* See below.

Treasure bond futures and the cheapest-to-deliver option. As we.ve noted, most interest rate products are tied to deposit index rates rather than government rates. There are exceptions, almost all of which can be handled using the exact same models as we.ve discussed for LIBOR rates (just use the current par coupon rate on governments in place of R_{swap} and use discount factors based on the government bond market rather than the LIBOR market).

One product that is quite popular and has some special characteristics are futures based on government bond prices. We.ll focus on one of the most popular of these products, the future on long US Treasury bonds, but futures on government bonds for other maturities and countries have similar characteristics.

In designing a futures product for long US Treasury bonds, the desire was to have settlement based on actual delivery of a bond, rather than cash settlement based on a published price. Partly, this might have been due to the fact that methodology for deriving par coupon rates from observed prices were not as well developed at the time this contract was introduced as they are now. Partly, there is always a bias in favor of settlement through actual delivery, since it avoids issues of possible manipulation in quoted prices or questions about how large a transaction the quoted price applies to. At the same time, it was clearly necessary to allow several different Treasury bond issues to be deliverable in settlement of the futures contract, since designating only a single bond would increase the likelihood that idiosyncratic factors applying to that one bond (e.g., shortage of bonds available for borrowing due to technical factors or manipulation) could distort the settlement. The decision was made to allow any Treasury bond with remaining maturity greater than 15 years eligible for delivery. This also had the desirable effect of reducing the degree of idiosyncratic differences between different Treasury issues, since the ability to hedge different issues with a single contract served to help unify pricing.

But different longer term bonds clearly trade at very different prices (for example, they could have very different coupon rates). To allow all of them to be deliverable against a single futures contract, there has to be some formula that determines the relative prices between them. You could just look at market price quotes, but that would return us to the same issues about applicability of quoted prices that created the desire for physical settlement rather than cash settlement. The solution was to create a table of fixed conversion factors between bonds that apply to the futures contract settlement. The fixed conversion factors were chosen based on a simple algorithm . discounting all cash flows at a permanently fixed rate of 6.00%. But when real market rates are currently higher than 6%, this rule will tend to overvalue longer duration bonds relative to shorter duration bonds, since longer cash flows are being overvalued by discounting at only 6%. When real market rates are below 6%, this rule has the opposite effect, overvaluing shorter duration bonds relative to longer duration bonds.

The party that is short the future must be given the choice of which bond to deliver (otherwise having a wide range of bonds to deliver does not get around the problems associated with a single bond eligible for delivery). This party will take advantage of this by choosing to deliver the bond that is currently most overvalued by the fixed conversion factors. This is known as exercising the cheapest-to-deliver option. It should be fairly obvious that the further market rates move away from 6% (either up or down), the more this option will be worth. So prices of Treasury bond futures need to adjust for this convexity cost by selling at lower prices than they would without the option, and the degree of adjustment will be higher in higher volatility environments. Note that this effect will be present despite the fact that very few contracting parties actually hold their futures contracts all the way to settlement. Any variance in pricing from what is implied by this reasoning will be taken advantage of by enough arbitrage traders to bring relationships back into theoretical line.