RECORD OF SOCIETY OF ACTUARIES
1995 VOL. 21 NO. 1

ARITHMETIC OF OPTION PRICING

Leader: JOSEPH D. KOLTISKO
Instructor: CRAIG MERRILL*
Recorder: JOSEPH D. KOLTISKO

This teaching session will examine option-pricing techniques, considerations involved and reasons for discounting at the short-term rate. It will also include biased probabilities, option-adjusted duration and convexity option-adjusted spread.

MR. JOSEPH D. KOLTISKO: I'm very excited about Craig Merrill's presentation. He is an assistant professor at Brigham Young University and has a Ph.D. from Wharton. He's also currently working on a monograph with David Babbel on interest rate contingent claim valuation.

I personally was very pleased to see that this was billed as the "Arithmetic of Option Pricing." If it had been the stochastic calculus of option pricing, I would have been scared away. Without any further ado I'd like to introduce our speaker, Craig Merrill.

MR. CRAIG MERRILL: As Joe mentioned, I'm working on a monograph with Dave Babbel for the Society of Actuaries on interest rate contingent claim pricing. It seemed that this might be a nice chance to present some of that material. In particular, I want to concentrate on the ideas that underlie option pricing. I'm not going to get into the calculus of it. I'd rather provide you with a conceptual framework from which it would be easier to jump into the option-pricing literature. If you have a road map, it's much easier to understand the details as you hit them.

Option pricing is quite a broad area. In fact, we hear much about derivative securities in the recent popular press, and I thought about that as I talked to my class recently about what a derivative security is. In fact, it's almost anything but a productive asset like a machine or a factory. Things as fundamental as an equity or a debt security have been valued as an option on the underlying assets of the firm. A life insurance contract covers not just mortality risk—one has to value loan privileges and the right to walk away from a policy and cash it out. All of those things are interest-rate-contingent options. When buying a car, whether to lease or purchase becomes an option pricing type of decision, because the lease gives you the option to walk away after two or three years or to buy the car at some prespecified price. In the broad area of option pricing, I've decided to focus on interest rate contingent claims because that's probably of more value within an insurance framework than equity option pricing would be, although you get the same fundamental ideas either way.

I'll give some option-pricing examples that you're probably quite familiar with. You know there are futures and options on treasury securities, such as calls and puts. You have options on futures. There are floating-rate securities. Any kind of floating-rate debt could be priced in an option-pricing framework. There are interest rate options such as caps and floors, there are options on those (captions) and options on swaps (swaptions),

---

*Mr. Merrill, not a member of the sponsoring organizations, is an Assistant Professor at Brigham Young University in Provo, UT.

and then corridors and collars and the list goes on and on. What you'll find as you dig into this area is that all of these things are built up from some very simple basics. Most of them can be traced back to a forward contract. A futures contract is just a series of forward contracts where you mark to market as you go. A put and a call held appropriately replicate a forward contract. It's just a matter of building up from simple pieces.

I want to introduce what's called a binomial lattice model as a framework here. As you know, if you ask anyone who has looked at term structures what the interest rate is, you're going to get laughed out of the room because there's no such thing as *the* interest rate. We teach our students in corporate finance how to do a present value calculation, but often we gloss over how to come up with the appropriate discount rate. To say that there's one rate that we can use for all maturities and cash-flow risk levels is an oversimplification.

What we want to do is try to capture the stochastic nature of interest rates in a simple model shown in Chart 1. The easiest way to approach this is with today's short-term interest rate. This could be the 30-day T-bill rate, or the prime rate or the London interbank offered rate (LIBOR) or any of the short interest rates. It can either go up or down over the next increment in time. That increment could be the next minute, hour, day, or month. It depends on how fine a data observation we need. The idea is that we have an interest rate right now at time t and it could either go up or down between now and time $t+1$. The probability of that is designated by $p(t)$, the probability of going up, and $\{1 - p(t)\}$ the probability of it going down. Now if we repeat that process we start to build the lattice. We go from r at time t to either r[up] or r[down] at time $t+1$ and then r[up][up], or r[up][down], or r[down][up] (you end up in the same place with the appropriate assumption) and r[down][down]. If you have cash flows that are not path dependent, that is if the cash flows are not determined by today's value of the interest rate, this recombining tree is useful.

We're going to look at a specific example of this in Chart 2 which we've called the multiplicative binomial model. The idea here is that the interest rate $r(t)$ is given today and can go up or down as follows. It could go up to $r(t)$ times $(1 + \sigma)$, or it could go down to $r(t)$ divided by $(1 + \sigma)$. And then if it goes up and comes back down, you end up back at the original $r(t)$. By going down, you'd be dividing by $(1 + \sigma)$ twice; going up twice you'd be multiplying by $(1 + \sigma)$ twice. This is a nice model in that it satisfies the need to be viable, but it's not particularly detailed or flexible. It works well for understanding the ideas that are involved, but it has some shortcomings which we'll talk about later.

As an example, let's say that today's interest rate is 5% (annualized), so this is a one-year spot rate. The one-year spot rate has a 50% chance of going up to 6.25% or down to 4%, so we've got a 25% volatility here. It could go up to 7.81%, come back down to 5%, go down and up to 5%, or down to 3.2%. We want to make use of this model, this scenario set of interest rates, to try to do some valuation. The idea here is that we're trying to create a set of possible future interest rates and we're going to see that there are some restrictions that have to be put on it to make sure that it's a viable or realistic scenario set. Given this set, we want to look at how we might value something.

CHART 1
GENERAL BINOMIAL MODEL

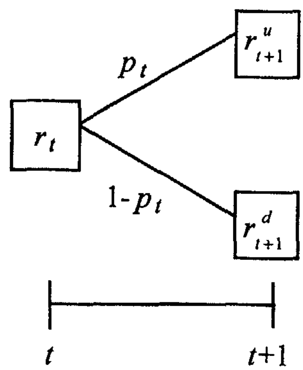

TWO-PERIOD BINOMIAL MODEL

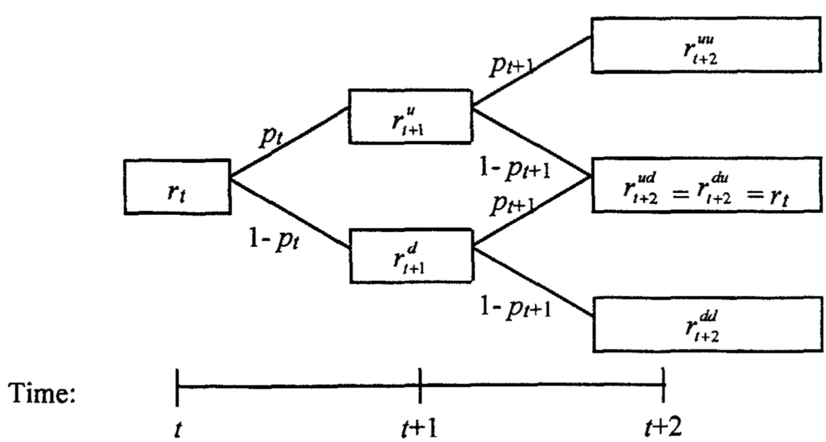

CHART 2
MULTIPLICATIVE BINOMIAL MODEL

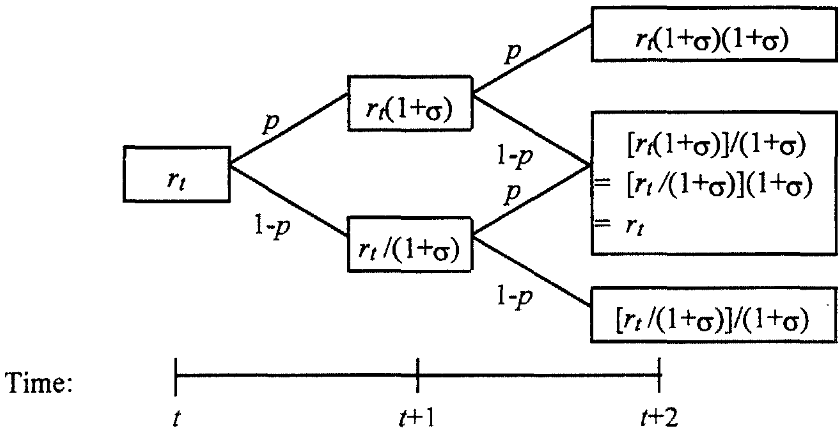

EXAMPLE OF MULTIPLICATIVE MODEL

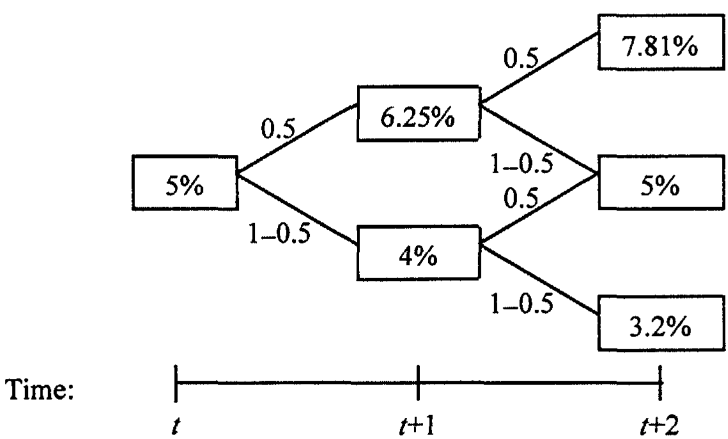

The simplest example is a one-year Treasury strip, a zero-coupon bond, as shown in Chart 3. It doesn't matter where the interest rate goes over the next period. Either way you're going to be paid a face value of $100, so I've just used $100 for both the upward or downward interest rate movement. The value of the strip is just $100 divided by 1.05, using the 5% interest rate at the first node of the tree.

CHART 3
ONE-YEAR STRIP

- $95.24 = $100/(1.05)

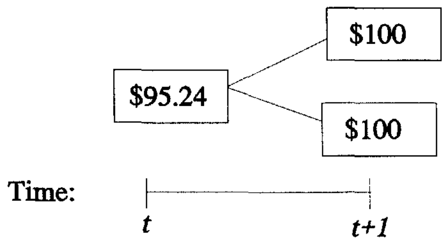

TWO-YEAR STRIP

- Discounting by paths:

$$0.5\left[\frac{\$100}{(1+0.0625)(1+0.05)}\right]+0.5\left[\frac{\$100}{(1+0.04)(1+0.05)}\right]=\$90.61$$

- Or, work through the tree one period at a time...

The next step would be a two-year strip. There are a couple of ways we might do this. One would be using what's called discounting by paths; as you can see from Chart 2, the interest rate could be 6.25% to 4% in period two. We've got a 50% chance of each of those things happening. We take the $100 and we discount it along the first path, which would be up in period two, so we discount by 1.0625 times 1.05 for the down movement. So we discount cash-flows along the two paths, weight them by their probabilities, and come up with $90.61. This is fine for a simple tree, but if you have a tree with 60 branching periods, you have many paths to try to keep track of and discount all the way through.

FROM THE FLOOR: Rather than just up and down, can you have up one, two, or three levels and down one?

MR. MERRILL: Sure, instead of a binomial you could have a trinomial or any sort of multinominal. It turns out though that the binomial model is almost good enough and the trinomial can be shown to be totally adequate to perfectly fit any current term structure.

So if we don't discount by paths, another possibility is to work our way through the lattice, through the tree. This turns out to be computationally much more tractable than trying to figure out all of the paths and discount each of them and weight them by their probabilities. In this one we just do the simple one-period discounting over and over and over again back through the tree.

Chart 4 shows this on a two-year strip. We know that there's going to be $100 at maturity on any path. We can discount those by the appropriate short rate, which in this case is our 6.25%, and we get $94.12 at the upper node at $t + 1$. If we do that same thing for the node below, we get $96.15 by discounting the $100. Now we need to discount these two cash flows, $94.12 in the case of an up movement, and $96.15 in the case of a down movement. We discount them by 5% and weight each by the probability of occurrence (which is 50% each) and we come up with $90.61. It's no surprise. This is just a different way of doing the calculation and the discounting by paths, just some algebra to rearrange the calculation, but we come up with the same result. This is a much better way, in terms of computational speed and simplicity, to work your way through a complex tree than by trying to do the discounting by paths. One thing I'd like to illustrate is that in the tree we've modeled explicitly only the movements in the one-year rate.

FROM THE FLOOR: Can you expand on why it's a much easier method computationally?

MR. MERRILL: Imagine if you have 60 periods in your tree. That means you've got 61 terminal nodes and 2^{60} possible paths. So you would need to lay out each path, find the probability of that path's occurrence and then do the probability weight at summation of the results. In this case we still have to do the same number of nodes-worth of calculation, but because each calculation is so simple, you can write just a few lines of code and do it repetitively. It's much easier to program and computes more quickly.

Explicitly, we're only modeling the evolution of one interest rate, so the assumption here is that this is what's called a one-factor model of the term structure. It's a single factor discrete time model of the term structure. Now that seems to be a big claim, but what I'm saying here is that implicit in every node on this lattice is a complete term structure, and I want to show you that in a simple case here with just two periods.

As shown in the example below, the first spot rate is 5%. That's just the first node on the tree. We can find the second spot rate because we know the price of a zero-coupon bond that has a two-period maturity. So we take 100, divided by one plus that spot rate squared and solve for the spot rate and we come up with 5.054%. If we had a larger tree we could work our way through all of the spot rates going out to the longest maturity. Those future spot rates would each be implicit at each node in the tree because the initial spot rate could be any node on the tree—it would just be a matter of drawing the lattice out from there. Similarly, we can find the forward rates, which are called implied forward rates. They're implied by an absence of arbitrage. In these models, the fundamental assumption that we're using is that there's an absence of arbitrage. There's no way to make money risklessly, no money machines.

CHART 4
TWO-YEAR STRIP

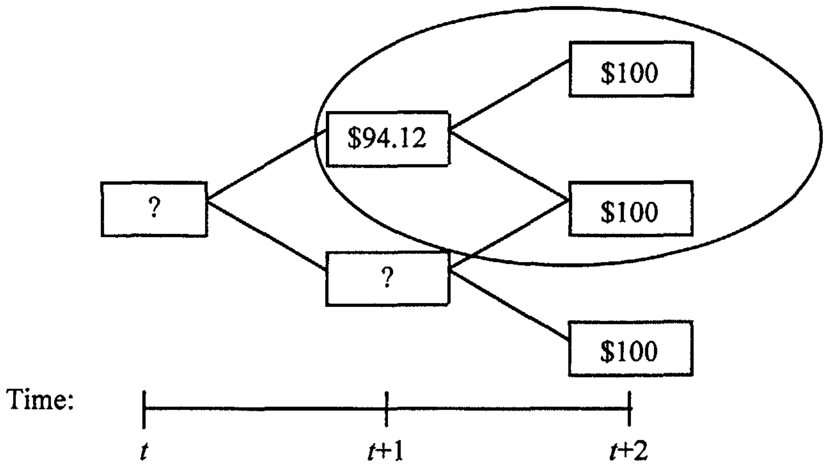

TWO-YEAR STRIP

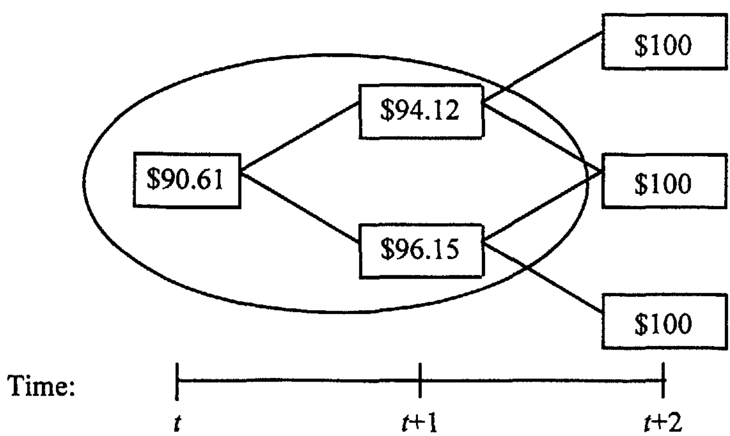

Spot rate term structure:

- One-year spot rate: 5%
- Two-year spot rate can be found using:

$$\$90.61 = \frac{\$100}{(1+s_2)^2}$$

- Solve to get the two-year spot rate:

$$s_2 = \sqrt{\frac{\$100}{\$90.61}} - 1 = 5.054\%$$

Consider two securities, asset one and asset two. Both of them pay $100 in two years. They differ in the way we value them. We value A_2 using the two-period spot, so that's $100 divided by one plus the two-year spot rate squared. We value A_1 using a series of one-period rates. Now in order to have an absence of arbitrage these two things have to have the same value. If A_1 and A_2 pay the same amount of money at the same point in time, they have to have the same price. In other words, if it walks like a duck and quacks like a duck, it must be a duck. We already know the forward rate at the initial time period, denoted by f_0, today's one-period spot rate. The forward rate at time one, which is the rate of interest from time one to time two obtainable now at time zero, is what we need to solve for.

Forward rates and arbitrage:

- Forward rates implied by absence of arbitrage.
- Consider two securities:

$$A_1 = \frac{\$100}{[1+f_0][1+f_1]} \qquad A_2 = \frac{\$100}{[1+s_2]^2}$$

- Assuming no arbitrage implies these must be equal.

Since A_1 and A_2 have to be equal, we can set them equal to each other and solve for that rate. If we do that, we come up with the example below. We can note that that's just the price of a one-period zero-coupon bond divided by the price of a two-period zero-coupon bond minus one, which is going to be the formula for a forward rate. If you've got f at some period t in the future, the implied forward rate is the price of a t period zero-coupon bond divided by the price of a t+1 period zero-coupon bond, minus one.

Implied forward rates:

- Solve for f_1:

$$f_1 = \frac{\frac{\$100}{[1+s_1]}}{\frac{\$100}{[1+s_2]^2}} - 1$$

- Note that:

$$f_1 = \frac{P_1}{P_2} - 1$$

Within our lattice is an implicit model of the movement of the entire term structure through time. This is a single-factor model. We might do things that are more complex like a two- or three-factor model, or as many arbitrary factors as we'd like to introduce. We can model each one as moving through time in some manner similar to this. That allows us more flexibility in the shape of the term structure that we can model. This particular model will allow for upward sloping, downward sloping, and single-hump term structures. So it must have some flexibility but it still has more pedagogical value than anything else.

FROM THE FLOOR: Am I reading your structure right? You're fixing probability and volatility factors which determine what interest rates will occur. What interest rate is likely to occur?

MR. MERRILL: If I was trying to calibrate this model, I would set up the lattice with arbitrary probability of movement and volatility parameters. Then I'd allow it to solve for those two to get a best fit to today's current term structure. Today's term structure would drive the selected value of the probability and the volatility parameter. It's the volatility parameters or the sigma that determines which interest rates can be realized.

Now I'd like to look at an interest rate cap in Chart 5. Let me just give a simple introduction for those who haven't run into these. An interest rate cap is a call option on the interest rate as opposed to on some tangible security like a T-bill, a bond, or a stock. It's a call option on the interest rate. It pays the maximum of the interest rate, minus some strike rate, so $\max(r - x,$ zero) times some notional principal, times the number of days in the period for the interest rate divided 360 or 360 days in a year. You want to cap your interest rate exposure at some level, so you want a security that's going to pay the excess beyond that cap if the interest rate is higher. For example, in our simple tree in Chart 2, if we wanted to cap the interest rate at 6%, then we would want a security that would pay 0.25% at the $t + 1$ upper node and that would pay 1.81% at upper node at time $t + 2$, but we wouldn't need anything to pay anywhere else which is illustrated in Chart 6. These are the cash flows for the interest rate cap.

Interest rate floors are similar instruments that would pay like a put option on the interest rate. A floor would pay when the interest rate drops below some prespecified level. We could also have interest rate collars that keep the interest rate within some level or a corridor that pushes it outside of some prespecified level. There are all kinds of variations on this. To value an interest rate cap is quite easy within this framework.

CHART 5
INTEREST RATE CAP—CASH FLOWS

$\text{Max}(r - x, 0) \times NP \times (n/360)$
6% strike rate

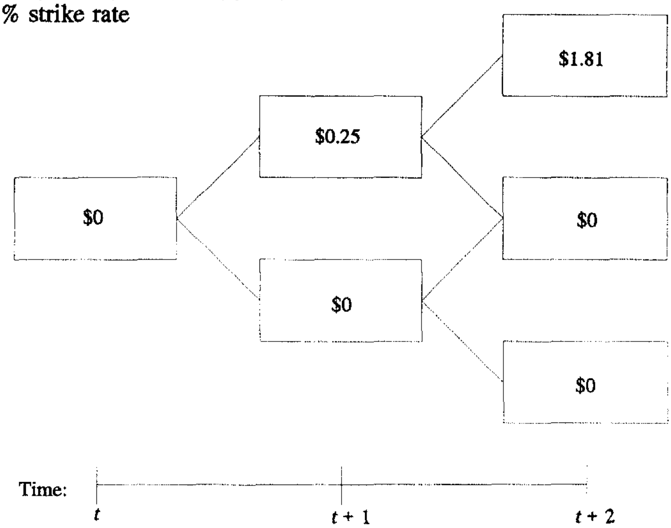

CHART 6
INTEREST RATE CAP—VALUE

$\$1.10 = 0.5(\$1.81/1.0625) + \$0.25$

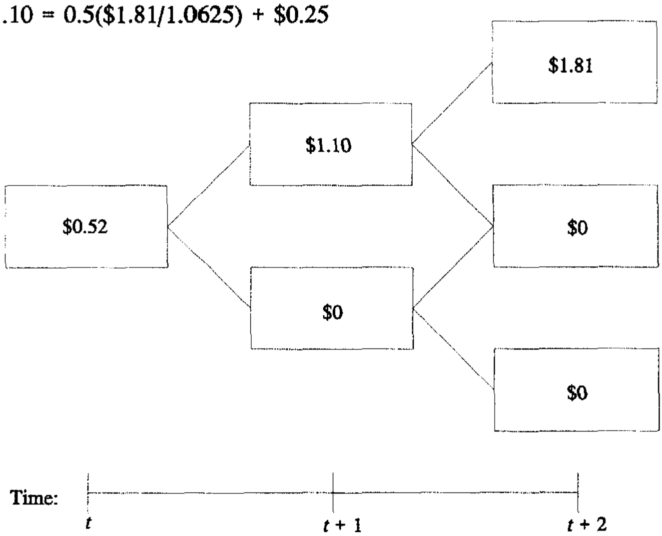

We start at time $t + 2$ with the amount this cap would pay, either $1.81 or zero, and we just discount our way back through the tree like we did before. There is a 50% chance of receiving $1.81 discounted by the 6.25% and a 50% chance of zero dollars. Besides the discounted cash flows, we add the 25 cents cash flow that occurs at $t + 1$ on the upper node, and that equals $1.10. We've discounted the final period cash flow to the intermediate period, $t + 1$, added in the cash flow that occurs there, and then we discount that $1.10 back at 5% times the 50% chance of it occurring and we get the 52 cents value for this interest rate cap at time t. Now this is for a $100 notional principal cap that pays once per year over two years.

There are a variety of features to these interest rate caps. The way I've portrayed this one is that the reset date and the cash flow occur at the same point in time. It's quite common to find caps that have a reset date at time $t + 1$ with the cash flow occurring at time $t + 2$. It can be valued within this framework quite easily by discounting the payout back at the time $t + 1$ interest rate and just valuing it at the tree at that point. Because from time $t + 1$ to time $t + 2$ there's no uncertainty about what the cash flow will be, it's quite easy to do that sort of cap as well in this framework.

I've introduced the idea that we can capture uncertainty about movement in interest rates by using a binomial lattice structure. I've introduced one possible way that we could build the lattice using this multiplicative process. Then, once the lattice is built, there is a spot-rate term structure and an implied forward-rate term structure implicit in each node of the tree. We've seen how we can attach cash flows to the nodes and discount them back to get a value. Those cash flows could be noninterest rate contingent like a zero-coupon bond in which case we might just as well use the spot rates to discount them. Or, they could be interest rate contingent as with the interest rate cap, in which case we do need to discount through the tree because the value of the cash flows are dependent upon the probability of their occurrence.

Now I want to turn again to the idea of no arbitrage. This is an important concept in financial valuation. This is just fundamental and we must have a firm belief that it's not going to be violated. That seems like a very strong statement, but if you go out looking for arbitrage opportunities it's going to be a long, laborious, and expensive hunt. There was a professor from the University of California, Berkeley who talked about how he had helped one of the investment banks set up an elaborate arbitrage. They had found a way to make $60,000 risk free, but it cost them $2 billion in trade face value to get $60,000. They needed supercomputers. They needed the full capacity of an investment bank and $2 billion in assets to do all of the trades that they needed to do to get this $60,000. These arbitrage opportunities are rare and so we can use that assumption in deriving values for these securities if there's no arbitrage with fair confidence.

I want to show you that with an appropriately structured lattice, one that does not allow for arbitrage opportunities; we can show, using an arbitrage argument, that we have the right answer. If there's time I will get into how we can tell whether there are arbitrage opportunities in the lattice. If not, at least you'll have that idea in your mind so that as you go and look at the literature you know that that's an important assumption in generating the interest rate lattice.

Let's discuss a one-year cap with the same parameters as before. We're still talking about this term structure or this evolution of the term structure and we're talking about the same

\$100 notional principal. In this case we just have one cash flow at the end of one year. If we used the lattice to discount the \$0.25 at the up movement and the zero at the down movement at 5% with a 50% probability of each occurrence, we would get an \$0.119 per \$100 notional principal price, with rounding.

By building a replicating portfolio, we will demonstrate that \$0.119 is the right price. What we're going to do is sell short delta two-year strips at \$90.6055 for \$100 of face value each. Now we're going to lend D dollars at the short rate. Selling short is just selling these things even though we might not own them. We're, in essence, borrowing them from our broker and selling them. But we're entering into the same obligation as if we had sold these things out of our own portfolio. So we've created a portfolio that's funded by selling short the two-year strips or the delta of them and investing at the short rate with D dollars. The cost of the portfolio, as shown below, is 90.6055 times delta (what we received for selling the strips) minus D, the amount of money that we need to invest. It turns out we need to invest more than the short sale has generated in order to replicate the cap, and that difference is the cost of the cap.

Replicating portfolio—cash flows:

- Cost of the portfolio:

  $\$90.6055\Delta - D$

- Cash flows in one year:

  $-94.1177\Delta + D(1.05)$ if short rate 6.25%

  $-96.1539\Delta + D(1.05)$ if short rate 4.0%

First, we want to identify the cash flows that would occur in each possible future state of the world. If the interest rate goes up to 6.25%, the two-year strip, which has now been shortened to a one-year strip, is worth \$94.1177 and we would have to pay that amount times the delta of those that we have sold. If the interest rate goes down to 4%, then the remaining one year of the strip is worth \$96.1539 and we'd have to pay that times the delta strips that we've sold. On the other hand, on our loan of D dollars, we now get back D times 1.05 in either case. What we want to do is try to replicate the interest rate cap. In order to do that we've set the value of the portfolio equal to the cash flow from the interest rate cap in each state of the world. That's reflected in the first two equations below. This is the value of the portfolio on the left-hand side and the cash flow to the interest rate cap on the right-hand side. With two equations and two unknowns we can solve for delta and D and that gives us the price of the cap. Delta is 0.1228, so we sell that many strips at \$90.6055 each to generate \$11.1243, while D equals \$11.2434. Plug those in and you get the difference of \$0.119, which we obtained just by discounting in the lattice framework. So we've confirmed that the absence of arbitrage and the lattice approach are consistent and that they give us the same price. That's not surprising given that I constructed the lattice knowing that it would not allow for arbitrage opportunities.

Replicating Portfolio—Δ and D:

- Set cash flows equal to the cap

  $-\$94.1177\Delta + D(1 + 0.05) = \$0.25$

  $-\$96.1539\Delta + D(1 + 0.05) = 0$

- Solve for D and D to get

  $D = \$11.2434$ and $\Delta = 0.1228$

- Finally, we have the price of the cap

  $C = \$90.6055\Delta - D = \$0.119$

In Chart 7, there's a tree that has six-month periods instead of one-year periods. In this case we have the annualized six-month spot rates. I've constructed the tree in such a way that the annualized current six-month spot rate is the same as the annual rate that we were using before and we end up with roughly the same terminal dispersion of interest rates. The annualized six-month rate is 6.25% if we go up and up again, while the annualized six-month rate if we go down and then down again is 4%. So we've got roughly the same spread in a one-year period as we had before, but we've got an intermediate branch. The idea is that as we add more and more intermediate branches, we get a richer and richer scenario set. The question I want to ask is, do we still get a valid answer out of this? Can we still show an absence of arbitrage?

CHART 7
TREE WITH INTERMEDIATE BRANCHING

- six-month periods
- σ = 0.118034

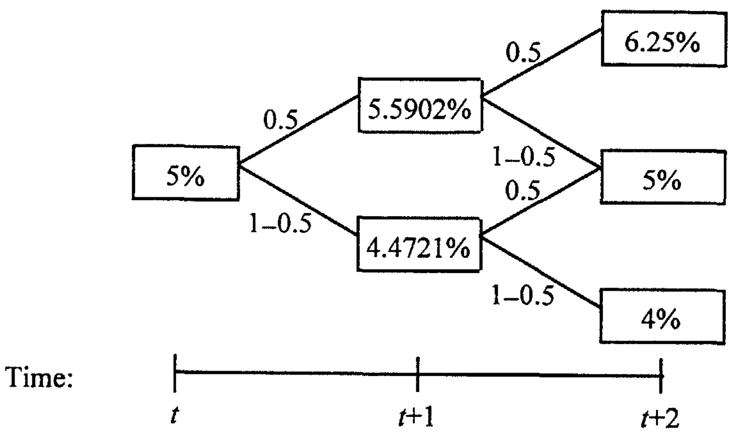

18-MONTH STRIP

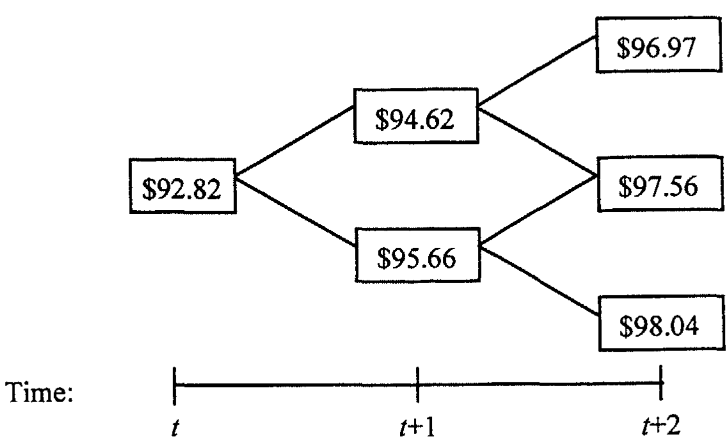

To test this I set up an 18-month strip, which is a zero-coupon security that matures one period after the end of the lattice. I haven't written in the $100 at time $t + 3$, but that's what we're doing. We discount the numbers back to t, all appropriately weighted by their probabilities and we determine the value of an 18-month strip as $92.82 and we can see the evolution of that price through the tree.

Now let's add a one-year interest rate cap with six-month reset. This is the same 6% strike rate as before, so if we get to the point where we have a 6.25% interest rate then the thing's going to pay $0.25 and so on back through.

Chart 8 shows the value of the cap. Recall the discounting is done over six-month periods, so half of the annualized rate applies. The lattice gives us $0.0593 as the value of the cap.

CHART 8
ONE-YEAR CAP

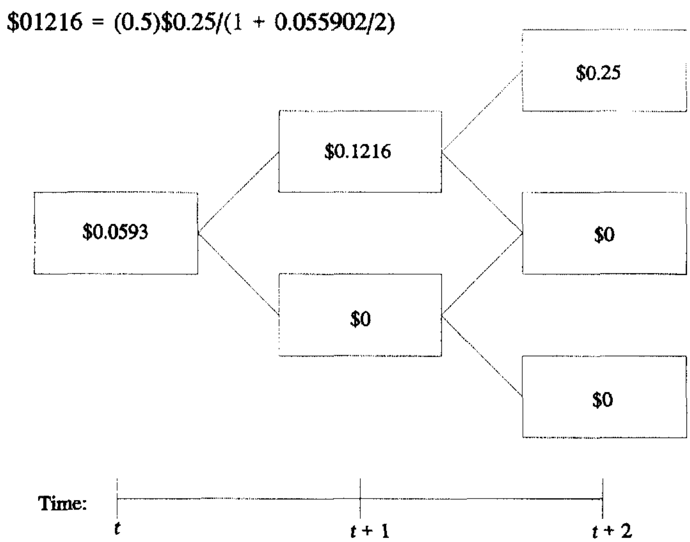

Now can we replicate that and show an absence of arbitrage? The idea here is if we can construct a portfolio that replicates the cash flows of the instrument, then it has to have the same price. We know the price of the Treasury strip and we know the discount rate to use, so we can set up a portfolio and value it and compare it to what our lattice generated. In this case we're going to try the same thing as before, sell short delta strips at $92.82, which is what we generated in our lattice, and lend D dollars at the short rate.

Now in one year the portfolio has to satisfy the four equations shown below in two unknowns and that's a bit of a problem. We can't satisfy all four of these with any particular values of delta and D. So we need to figure out a different way to do the replication. It turns out that we need to do a dynamic replication so that we work our way through this lattice one node at a time. We would start by replicating this one branch at a time in the way we did before. We'd come up with two equations in two unknowns and calculate how much of the strip to sell and how many dollars to lend. We would figure out what the replicating portfolio is to go back to the initial time period. We would have a strategy that says to sell a certain number of strips and lend a certain amount of money in the first period and then a strategy that says we change our portfolio to a different number of strips and number of dollars depending on which node we move to, to go into the last time period. That portfolio movement over time is going to have a particular value which will equal the $0.0593 that we obtained from the lattice approach.

Replicate the one-year cap

- Sell short ΔSTRIPS at $92.82 each and
- Lend D dollars at short rate
- In one year, portfolio must satisfy:

$$-\$96.97\Delta + (1+0.05/2)(1+0.055902/2)D = \$0.25$$
$$-\$97.56\Delta + (1+0.05/2)(1+0.055902/2)D = 0$$
$$-\$97.56\Delta + (1+0.05/2)(1+0.044721/2)D = 0$$
$$-\$98.04\Delta + (1+0.05/2)(1+0.044721/2)D = 0$$

Those are the two approaches to option pricing that you're going to find in any of the literature. It's either going to be generating a lattice and then pricing based on the lattice, or it's going to be generating what's called a hedge portfolio, the replicating portfolio that I demonstrated here and using that to replicate the cash flows of some instrument and finding the price that way.

As I mentioned, there is more than one approach to generating these lattices. One approach is called equilibrium and one is called arbitrage free. In an equilibrium model we specify some stochastic process that the interest rate follows, or a series of interest rates in a multifactor model, or volatility, and given that specification, we fit the model as best as we can to today's term structure. So, for example, this multiplicative binomial model would be considered an equilibrium model. We're specifying that the interest rate follows a multiplicative process over the time that the short rate does and then implicitly the whole term structure moves according to that model. As I stated, we can use the volatility parameter and the probability as our parameters, and estimate those to best fit today's term structure. We then use that to price interest rate contingent securities. If you fit a curve to a bunch of points there will be some error. You will have this term structure that you've estimated and you'll have some points that will be somewhat above and below each point in time depending on how well you can fit with your model. Let me just give you an example.

Draw a "best fitting" smooth curve through the yields for each on-the-run Treasury note. If you were to price with this curve and then actually go out in the market and trade, you would find that, right now, the prices you generate admit slight arbitrage opportunities at each maturity. What's been developed to avoid this is a different approach that has been referred to as arbitrage-free interest rate modeling. Now in this case, we would allow the parameters to vary with time. So, for example, the volatility would be a function of time, and you would estimate a different level of volatility for each time. The goal there is to fit a curve that goes exactly through each of the points that you have observations on, so that for known securities it admits no arbitrage opportunities.

Now one of them is more appealing to the academic and one of them is more appealing to the practitioner. You can guess which. If you're going to go out and trade millions or billions of dollars you want a model that is as accurate as possible today when you're trading. For that, these arbitrage-free models are, in fact, what is most often implemented. On the other hand, if you look at how they do over time, an arbitrage-free model that fits exactly to today's term structure turns out to have larger errors tomorrow than the equilibrium model will have. So if you look at how they do in fitting term structures that are revealed over the next few days, because of the exact fit nature of the arbitrage-free model, it really does need to be estimated every time it's used. The equilibrium model tends to be useful for a little bit longer period of time just because these interest rates tend to fluctuate in some range anyway.

This makes it worthwhile to look at a more general lattice model. In general, the short rate can move up or down in our binomial models. Then we can add several other securities to the model, a note and a bond, that will move up or down with it or opposite to the short rate. These can become the basis of valuation, whether it's by replication or whether it's by lattices. I may come back to this notation later.

One of the more successful lattice models is one that was done by Hull and White. It's a trinomial lattice model. In this case the interest rate can move either low, middle, or high in the next period. In this case, Chart 9 illustrates what's called the Extended Vasicek Model. This was derived by Hull and White as an extension to the Vasicek Equilibrium Model.

Recall the binomial model that I showed you earlier. Take the time increments to zero in the limit that becomes a lognormal distribution of short-term interest rate movement. So the shorter the time period becomes, the closer to a lognormal of distribution you get at the terminal time period of the lattice. That's a good thing. That does not allow for a negative interest rate and it also allows for prices that are consistent with the assumption of no arbitrage.

There was another binomial lattice model by Ho and Lee that used an additive process, in which you add or subtract some amount to or from the interest rate at each point in time. In the limit that became a normally distributed interest rate process, which allows for negative interest rates and arbitrage. It's very analytically tractable and can be implemented in such a way as to give good prices, but it's not desirable to have a model that allows for arbitrage opportunities and negative interest rates, neither of which we'd like to think exist. Hull and White have developed a trinomial lattice model that has the same failing in that it can allow negative interest rates if the parameters are not selected carefully, but it has been shown very carefully to allow absolutely no arbitrage and a perfect fit to almost any arbitrary term structure that we might see in the economy.

This is the way it works: the initial node, A, currently represents a 10% short rate. These are arbitrary lengths of time, so I'm just going to refer to them as periods. That 10% could go up, stay the same, or decrease. But notice, B, the short rate, could go way up or stay the same. There's some added flexibility in this model. Hull and White specify a drift and expect a change in the short rate. Then they compare the next nodes on the tree and what they want. The middle branch should go to the node that is as close as possible to the expected interest rate in the next period. They take the current interest rate, add to it the expected change, and then put the middle branch as close as possible to that interest rate. If you have a positive expected change you see that as the interest rate gets low it jumps up quite a bit, whereas, when it's high it doesn't jump up by quite as much. By this mechanism they keep from obtaining negative interest rates.

CHART 9
A FOUR-PERIOD ILLUSTRATION OF SHORT RATES GENERATED BY THE EXTENDED-VASICEK TRINOMIAL LATTICE MODEL

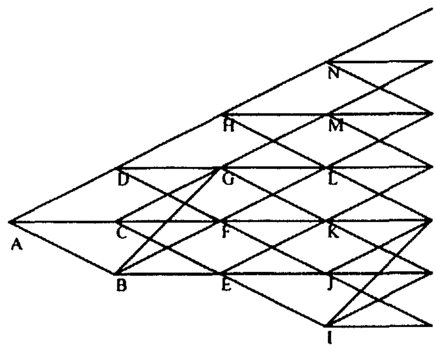

Node	A	B	C	D	E	F	G
Rate (%)	10.00	7.58	10.00	12.42	7.58	10.00	12.42
p_1	0.462	0.044	0.507	0.415	0.286	0.221	0.166
p_2	0.493	0.477	0.451	0.534	0.627	0.657	0.667
p_3	0.045	0.479	0.042	0.051	0.087	0.122	0.167
Node	H	I	J	K	L	M	N
Rate (%)	14.85	5.15	7.58	10.00	12.42	14.85	17.27
p_1	0.121	0.042	0.455	0.370	0.293	0.228	0.171
p_2	0.657	0.426	0.499	0.570	0.623	0.654	0.667
p_3	0.222	0.532	0.046	0.060	0.084	0.118	0.162

MR. LINGDE HONG: What is the expected rate?

MR. MERRILL: They specify a stochastic process that the interest rate follows, to produce the expected rate and then a volatility. I didn't bring a full description of the model, but they'll specify, for example, that the expected change in the interest rate would be some average amount, plus a time dependent parameter. They estimate that time dependent parameter at each step of the tree.

MR. HONG: It almost looks as though you expect it to always be flat. From E, given that you come through A, then B, or C, does that imply you expect to go to node J?

MR. MERRILL: Or do you go somewhere closer to J than either K or I. The expected value does not necessarily land on a node; the test is which node is it closest to. The nodes in this particular tree are obtained by taking the initial interest rate, 10%, and adding to that an integer value, positive or negative, times a delta R which they estimate using some theorems which assure that there's no arbitrage.

MR. HONG: Well, I'm just saying, by looking at this picture, it seems like there would be an upward drift effect.

MR. MERRILL: That's right. In this particular case there's always an expected increase in the interest rate. Sometimes that expected increase when you get down low is large enough to pull the three branches up one node as I've illustrated with the drawing. Sometimes it's not large enough and so it moves kind of straight through time up, across the middle, and down.

MR. HONG: Those times are determined by the initial yield curve?

MR. MERRILL: This particular model has what's called mean reversion, so that as the interest rate gets further and further away from some long run average interest rate, there's a stronger and stronger and pull back towards it. That pull is stronger sooner on the low side than on the high side because of the positive drift. If we went up much further you'd see that it would start to pull back in on that side as well.

MR. HONG: The model implies that interest rates long-term are going to stay the same. Do you buy that?

MR. MERRILL: It seems to fit history fairly well. In a stable economy like the U.S., or in post-industrial Japan, Germany, or whatever, we have not seen incredibly high interest rates last for any significant amount of time. Plus given the policies of our government it seems like the interest rate is going to be held in some range.

MR. HONG: But wouldn't that imply some kind of an arbitrage opportunity? If you know there is a target interest rate, long-term, you know what's coming and you ought to be able to trade on that.

MR. MERRILL: Right. If it were going there with certainty, yes. But because there's a random element to it and it's always bouncing around, it may be pulled in that direction, but because the random kick can send it up instead of down or down instead of up, there's not a sure profit and that's what's required for an arbitrage.

MR. HONG: So it's a better bet.

MR. MERRILL: It might be a better bet, but it's not a sure bet. So we generate a lattice that represents a viable set of interest rates, and again, I'd refer you to the Hull and White paper. It's in a 1991 issue of the *Journal of Financial and Quantitative Analysis*. Hull and White have done a whole series of papers. They have two more coming out in the next year that walk through the mechanics of doing this. Do a literature search on Hull and White, and you'll come up with a variety of papers on this topic.

MR. KOLTISKO: Their book on option pricing is part of the Society's syllabus.

MR. MERRILL: In Chart 9 we have a scenario set. We have interest rates that could go from 10% up to 12.42%, or down to 7.58%, or stay the same at 10%. From there we've got a little bit broader range and so on out. Now what we want to do is use this to price. This model has a good property. As we construct this tree we generate a bunch of state contingent prices. Now what do I mean by that? Let's look at node D, for example. I want to know what is the value of $1 if and only if I arrive at node D? Well, in order to do that I need to know the probability of node D occurring and the interest rate path or paths that may lead to it.

In Chart 10 I've generated the probabilities as moving from any node to the next. Notice that these are not constant probabilities. The probabilities vary at each point in time. That's a characteristic of an arbitrage-free model. We vary one parameter in the drift and we vary the probabilities according to that parameter in the drift at each point in time. So these probabilities apply at a given node and at a given point of time. The drift applies to every node at a different point in time. Given the probability of 0.462 of going up to node D and the discount rate being 10% for that one time period movement, the value of $1 if and only if we reach point D or node D is $0.418. Similarly, the value of $1 if we reach node H is given here as $0.153. Node U is probably the least likely to be obtained because only one branch leads there. Node O would be almost as unlikely. Node U would be the least likely and, accordingly, it has the lowest value for $1 if and only if we reach that node.

One good thing about this model is we have to generate these prices anyway to produce the model, so we can store them in an array and use them for pricing interest rate contingent securities. Remember the idea of an interest rate contingent security is that we put cash flows on the tree where they will occur and then we find their value.

If we know the value of $1 anywhere that it occurs on the tree, if we just assign the cash flows there and weight them by their state contingent prices, these Qs, and add those up, we have the value of the security. Note here the "Qs" are like present value factors, not mortality decrements. We get the same number we would come up with if we discounted our way back through the tree. So, for example, if we want to value an interest rate cap, with an 11% strike rate, $100,000 notional principal, and interpreting this as a one-year per time period, we can lay out the cash flows on this tree. Finding a price means merely multiplying the cash flow at node D, for example, times the Q that we generated for node D, plus the cash flow at node G times the Q we generated for the node G, and then add those up all the way through.

MR. STEPHEN ADAM BEKER: The probabilities don't seem to add to one. For instance, look at the Q's movement from node B to either E, F, or G.

MR. MERRILL: Those aren't probabilities. The Qs are the values of $1 that get paid if and only if you reach that node. So, for example, let's look at node F. There are one, two, or three different ways that we could arrive at node F. The value of arriving there is determined by the interest rate along those paths and the probabilities of those paths being followed. Given those Qs, it's very easy to value an interest rate contingent security because all we need to do is lay the cash flows on the tree which is shown in Chart 11. If we wanted to just kind of make the point, we could just lay the one tree on top of the other and multiply the Q times the cash flow in each case. The cash flows here are zero in many places and positive in other places by the nature of the interest rate cap. So wherever there's a zero cash flow, there's no chance of receiving any money, so the Q is zeroed out.

CHART 10
A FOUR-PERIOD ILLUSTRATION OF $Q(ij)$s GENERATED BY THE EXTENDED-VASICEK TRINOMIAL LATTICE MODEL

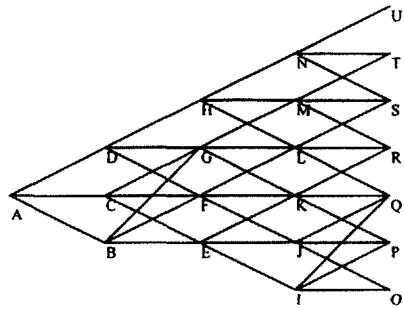

Each node (i,j) represents $Q(i,j)$ which is the value of a security that pays one dollar if and only if node (i,j) is reached and zero otherwise. For example, the node identified by the label G in this figure represents $Q(2,1)$.

Node	A	B	C	D	E	F	G
$Q(i,j)$	1.000	0.041	0.446	0.418	0.035	0.219	0.403
Node	**H**	**I**	**J**	**K**	**L**	**M**	**N**
$Q(i,j)$	0.153	0.003	0.045	0.199	0.311	0.146	0.016
Node	**O**	**P**	**Q**	**R**	**S**	**T**	**U**
$Q(i,j)$	0.003	0.033	0.145	0.752	0.165	0.038	0.002

CHART 11
CASH FLOWS FOR A FOUR-YEAR CAP WITH AN 11% STRIKE RATE, \$100,000 NOTIONAL AMOUNT AND ANNUAL PAYMENTS GENERATED BY THE EXTENDED-VASICEK TRINOMIAL LATTICE MODEL

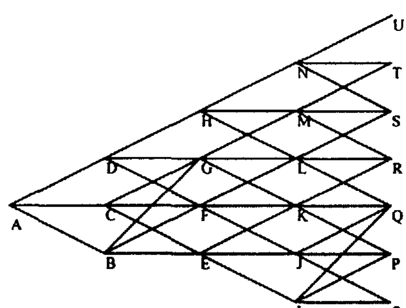

Each cash flow is given by $(r_0 + j\Delta r - 11\%)(\$100,000)$

Node	A	B	C	D	E	F	G
Cash Flow	0	0	0	1424.87	0	0	1424.87
Node	**H**	**I**	**J**	**K**	**L**	**M**	**N**
Cash Flow	3849.74	0	0	0	1424.87	3849.74	6274.61
Node	**O**	**P**	**Q**	**R**	**S**	**T**	**U**
Cash Flow	0	0	0	1424.87	3849.74	6274.61	8699.49

But, for example, node G, if we obtain it, has a cash flow of $1,424.87. That has a value of 0.403 times the $1424.87. We add all of those up and we come up with the value of our contingent payment. Now if you find this of interest and you go to the Hull and White papers, what I would suggest that you do is generate a lattice following their instructions that looks like this. You can use these to check your work because these are ones that have been calculated and you can use them to validate your code or your spreadsheet. If you do that you'll find that whether you discount back one time period at a time, like we did in the binomial lattice, or you use the Qs as state contingent claim prices, you'll come up with the same price for your interest rate contingent claim.

Now once you have this lattice generated you can do all sorts of interesting things with it. For example, let's say that you have a callable bond. We need to be careful here because there are two types of call options. One allows you to call anytime during a period of time. Another would allow you to call only at a certain time. These are American and European options, respectively.

A European-style option is very easy to do in the lattice framework because we know at one particular date if the bond is worth more called than not called it will be called and we can calculate what the cash flow would be if it is called. We can put those cash flows on the tree where they might occur if that point is reached. The decision to call or not call is going to be dependent upon the level of interest rates on the call date.

FROM THE FLOOR: What about expenses?

MR. MERRILL: Those could be factored in, since they are known amounts. What the procedure allows us to do is analyze the impact of any particular provision of an interest rate contingent claim. We can separate out the value of the call option from the value of the underlying bond by analyzing the option-related cash flows, or just the bond-related cash flows, or the two together. Furthermore, not only can we put a value on those things, but we can figure out a duration impact or a convexity impact.

There's some disagreement about what the best way to calculate duration is, based on these models. The original MacCauley duration involves shifting the term structure a fixed amount at every point in time, a parallel shift so to speak. But there are some in the finance community who think that duration ought to mean jiggling the interest rate that's driving the whole term structure and letting the term structure move with it so that you get a nonparallel shift. So, for example, in a single-factor model of the term structure like the one we're using here, you would move the short rate 50 basis points, and let the model regenerate so that you get a nonparallel shift in your current term structure.

Alternately you could just shift every node 50 basis points and get a parallel shift in the term structure. I will not attempt to answer which is the right form of duration. Probably the more traditional answer would be to just add 50 basis points to every node in the lattice and reprice your securities. Then subtract 50 basis points from every node in the lattice from the original values and value your securities. Then use a first difference for duration and a second difference for convexity as if this were a finite difference approximation to a derivative. In that way you can calculate the impact of various features of a contract in terms of value, duration, and convexity.

Of particular interest to me anyway is the option-like characteristics of a life insurance policy. In those cases we have something that is more like an American option. The right to surrender my policy does not exist on certain days, but at any point in time if I'm rational after I vest it in the cash value of the policy. That becomes an American-style option. Now with European-style options it turns out that we don't always have to use the lattice framework, but sometimes there's actually a known analytical closed-form solution for the value of the security or the option. When we move into American-style options—no one has yet found a way to get in a closed-form solution—lattice style approaches are the only things that seem to work.

There are other variations on the lattice. For example, if you're doing continuous time models you might use a finite difference approximation to a partial differential equation, but in operation you're doing exactly what you're doing here. You're discretizing a process and working your way through the discretized time space to discount the cash flows by their probability of occurrence. This is where the real power of the lattice approach lies. You can take the tools that have been used in the mortgage-backed security market and put values on the option-like characteristics of life insurance or annuity products and understand exactly what you're exposing yourselves to in terms of financial value by having those options in your contract. In some cases knowing the value may just tell you how much the market's driving you to give the options away. It may aid in accurately pricing the instruments that are being sold.

MR. KOLTISKO: One of the things that struck me when listening to your talk, Craig, is how you might effectively use a lattice by sampling or bundling certain paths in the lattice. Some practitioners have undertaken that approach for American-style option pricing. Have you found that practical and have you had success with sampling strategies?

MR. MERRILL: The need arises particularly in American-style options because of what's called path dependency. Take node K, for example, in the trinomial lattice. If you're looking at an instrument that's like a callable bond, it may get called or it may not get called by the time you reach node K, depending on whether interest rates go up and then down to K or down and then up to K. So you need to value the arrival at node K as if it could have come from either direction. You now are back to the point where you have to keep track of every possible path. We lose some of the ability to just discount through the tree and that's quite a bit to keep track of. How many paths do we have going to node K? There are at least four different paths (definitely more as I look at it), but you have to keep track of each and every one of those. Now what Joe pointed out is this idea of sampling paths. In this case you would randomly generate a path through the tree and then save the results of that path somewhere in a variable, a range of cells in the spreadsheet, whatever. Then you would randomly generate another path through the tree and save the results somewhere. Along with the path you need the probability of its occurrence, but since those are generated that's easy to save. You generate enough of those and you have a good quality random sampling of the universe's paths that could be followed through the tree. In practice the results I've seen suggest that somewhere around 1,000 or 1,500 paths out of a 60–100-period lattice are more than adequate to get a good approximation.

MR. DANIEL A. ANDERSON: I'd be interested in your comments about how this mechanism gives us the capability to calculate and how it's driven by the probability structures and so on. The probability structures you were talking about can be developed based on what today's current term structure is. So implicit in that is the assumption that today's rates somehow are predictive of future rates. If you take a different focus, you realize that some people are uncomfortable with the idea of relating the concepts of probabilities to the concept of interest rates. It's different from mortality-type probabilities where you do have a universe of people out there and they die at different rates and you can look at the law of large numbers. I've heard the expression the law of one number, meaning the interest rates. So if an actuary was to go back to management and say, "I'm going to price this thing for the company and I've got all these probability structures," I'm interested in whether he or she has to recognize that there's only going to be one path. The price they set isn't going to be a safe price necessarily because they're not going to get a dispersion of things in the future. Only one thing is going to happen and if interest rates do, in fact, go up (and that's the bad news), it's hard to argue to somebody that the probability of interest rates rising wasn't that high.

MR. MERRILL: You raised a couple of issues. Let me see if I can address them. First, there's this idea of a probability structure for interest rate movement. The note that I've included derives what's called an equivalent martingale measure or a risk neutral probability structure. That is what we use in finance to price securities. It's not actually like a probability of an event occurring. It's not as if we all got together and said how likely is an upward movement, how likely is a downward movement. It's not a consensus, or any subjective measurement, or estimate of what the probability of the price movement is. Rather it is the probability structure that is consistent with there being no arbitrage in the economy and with people behaving as if they're risk neutral. That is, they're concerned only about expected value and not about variance. Now it has a particular relationship to any sort of subjective or objective probability estimate that we may make and that's derived there. So whatever the probability structure out there might be, we can transform it into this risk neutral structure and use that for pricing. Now, again, the pricing is assuming no arbitrage which is very important. That means that in expectation you'll always meet a risk-adjusted rate of return. There's no guaranteed way to beat that risk-adjusted rate of return.

Now the question of setting prices. There are two ways that it can be done and I think the one that's chosen is quite clear and the reasons are quite clear. As I would present to students in a Principles in Insurance course, there's the actuarially fair value of a policy (the expected law, so to speak), and then there's the price that you actually pay which includes risk loading, expenses, etc. That risk loading represents the extra amount that is charged to ensure, if a worse-than-expected scenario were to arrive, that we don't go bankrupt. So another possible application of the lattice is to try to get a probability distribution of possible outcome and then pricing to make sure that we're covered for X% of those outcomes, say 80% or 90%, depending on the risk tolerance of our company.

FROM THE FLOOR: I was interested in your remark about the tendency to regress toward some sort of a mean. Interest rates are really composed of at least two components; one of those is an inflationary expectation. The other is a real interest rate which is added onto that. You hedged your remark by saying, in a stable economy, rates revert to the mean, but if you eliminate that characteristic then, of course, inflation rates may go anywhere. If we include volatile inflation, interest rates may go to 50% or higher as we've seen recently in the turmoil in Mexico. I wondered if any work has been done on this lattice relationship in connection with splitting these two apart and predicting the likelihood and severity of the impact of inflation on interest rates.

MR. MERRILL: Yes, one of the most commonly cited and used single-factor models is one by Cox, Ingersoll, and Ross. Their model is stated in continuous time, but there are discrete time analogs to it that could be done in the lattice framework. A variety of practitioner and academic researchers have tried to test the ability of that model to fit the term structure over time because the model assumes constant parameters. One of the biggest problems they've run into is exactly what you identified, there are two components to the interest rate, the real rate and the nominal rate which has the inflation component in it. The theoretical models are stated in terms of real interest rates and in order to capture the inflation component as well, you really need a two-factor model where the interest rate that you're modeling is driven by two stochastic processes, one for the real rate and one for the inflation rate. We've seen instability in both. We saw real rate movement in the early 1980s and a great deal of nominal rate movement besides that that's attributable to changes in inflation expectations. So at best any single-factor model is only going to be an approximation to what a two-factor model could do in terms of modeling nominal rates.

MR. BEKER: This is a question that's really more targeted to the audience. I'm curious to know how many companies are using straight caps and some of these other options. Second, how are you going about valuing them and pricing them?

MR. KOLTISKO: We've seen a number of major insurers quite active in hedging the risk that they already take in their interest-sensitive lines. These are principally the large, well-capitalized and sophisticated companies that have both developed proprietary pricing systems and purchased them from vendors. That has been my experience.

MR. MERRILL: I have a friend who has worked in the capital management division of a couple of large insurance companies and one of the main things he's been doing is interest rate models in order to do both scenario testing for a GIC portfolio and valuation for hedging instruments.

One of the difficulties that seems to exist is that derivative securities, like options and futures, are perceived as speculative or risky instruments. So when we try to use them in a hedging capacity to reduce risk, it's difficult to do given that there's a fear that they may be used in a speculative capacity. It's difficult, as Barings Brothers will tell you. It's difficult to set up adequate controls so that they get used only for hedging and not for speculative purposes. That has slowed the acceptance of these instruments in the insurance industry.

MR. KOLTISKO: I have a comment on the question of long tails for interest rate models. Lognormal models have finite variances which limit the range of outcomes. In a Stable Paretian Model we fatten the tails a little bit which makes it more likely that those extreme scenarios will occur. Have you seen that implemented in a lattice framework?

MR. MERRILL: Not in the literature. That would be quite a cutting edge interest rate model and would likely be held for proprietary uses at this stage. But it is doable. If you can specify an interest rate process that can be discretized and has as its limit an arbitrary excess kurtosis type of distribution, it can then be implemented in this format.

I'm new to the Society. I'm excited about the monograph that we're working on and I'm also excited about this particular area. This is I think a very valuable set of tools for especially this industry. This industry has more assets under management than any other financial services sector; therefore, there's a larger set of challenges and opportunities for the use of these tools.