Pathwise valuation, simulation and rate models

Fixed Income

1 A story: two routes to one price

- Mira: discount each payment at today's spot rate
- Tomas: build a tree of rates, work backwards
- Plain bond: both get exactly the same price
- So when is the tree worth the effort?

IN THE LESSON

- Imagine two colleagues at a small bond desk, Mira and Tomas.
- Both are asked to price the same plain government bond.
- Mira simply discounts each payment at today's discount rate for that date.
- Tomas builds a branching tree of future interest rates and works backwards through it.
- Tomas spends a whole afternoon.
- Mira spends two minutes.
- Whose answer is better?
- The surprising truth is that if Tomas built his tree properly, their answers match to the penny.
- The tree was forced to agree with the same market prices Mira used.
- So the real question is, when is the extra afternoon worth it?

2 Quick reminder: the zero curve and no free lunch

- Spot rate = yield on a zero coupon bond
- Each bond payment is a tiny zero coupon bond
- Wrong price = risk free profit (arbitrage)
- The lattice is built to match the same prices

IN THE LESSON

- A quick reminder from Level One.
- A spot rate is the yield on a zero coupon bond for one particular maturity.
- Each payment from a bond is a little zero coupon bond, so we can discount it at its own spot rate.
- This is the arbitrage free way to value a bond with fixed payments.
- If a dealer quoted any other price, you could buy the cheap pieces, sell the dear ones, and lock in a profit with no risk.
- A binomial lattice is a second tool built on the very same idea.
- It is a tree where the one year rate can go up or down each year, built so that it prices the zero coupon bonds correctly.

3 Two ways to the same value

Spot: $V = CF_1/(1+z_1) + CF_2/(1+z_2)^2 + \ldots$
Tree: $V = \frac{1}{2}$ (V up + V down) ÷ (1 + r)

Cash flow (CF)	A coupon or the principal
Spot rate (z)	Discount rate for that date
One-year rate on the tree (r)	Rate at one node of the tree
Value	What the bond is worth

IN THE LESSON

- Here are the two calculations side by side.
- On the zero curve, the price equals each cash flow divided by one plus its spot rate, raised to the number of years, all added up.
- On the lattice, you start at the end.
- At each node, the value is the average of the two possible next values, each with its coupon added, divided by one plus the one year rate at that node.
- Let me take the terms.
- The cash flow is a coupon or the principal.
- The letter z is the spot rate for that date.
- The letter r is the one year rate at a node on the tree.
- And the last term is the value you are looking for.

4 Example 1: Mira prices off the spot curve

Spot rates 2.00%, 3.00%, 3.50%. Bond: 4% annual coupon, 3 years, face 100.

Year 1 coupon	$4 \div 1.02$	**3.92**
Year 2 coupon	$4 \div 1.03^2$	**3.77**
Year 3: coupon + face	$104 \div 1.035^3$	**93.80**
Mira's price	sum of the three	**101.49**

IN THE LESSON

- Example one.
- The government spot rates are two per cent for one year, three per cent for two years and three and a half per cent for three years.
- Mira values a three year bond paying a four per cent annual coupon, with a face value of one hundred.
- Step one, the year one coupon is four divided by one point zero two, which is three point nine two.
- Step two, the year two coupon is four divided by one point zero three squared, which is three point seven seven.
- Step three, the final payment of one hundred and four divided by one point zero three five cubed is ninety three point eight zero.
- Add them up and Mira gets one hundred and one point four nine.

5 How Tomas builds a matching tree

Build the tree

Choose volatility: 15% → Find year 1 rates → Find year 2 rates

Then test it

Price every zero correctly → Fixed payments: same value

IN THE LESSON

- Now Tomas.
- He cannot choose the tree rates freely.
- He builds them in stages, and at every stage he forces the tree to reproduce the market.
- First he picks a volatility, say fifteen per cent, which fixes how far apart the up rate and the down rate sit at each date.
- Second, he finds the year one rates so that a two year zero coupon bond is priced exactly.
- Third, he finds the year two rates so that a three year zero is priced exactly.
- Only then is the tree calibrated.
- And because it prices every zero coupon bond correctly, it must price any bond made of fixed payments correctly too.

6 Example 1: Tomas's calibrated tree

Today 2.00%. 15% volatility. Up, down equally likely.

Date and node	One-year rate
Year 1: rise, fall	4.61 / 3.42%
Year 2: top node	5.97%
Year 2: middle node	4.42%
Year 2: bottom node	3.27%

IN THE LESSON

- Here is the tree Tomas ends up with, at fifteen per cent volatility.
- Today's one year rate is two per cent.
- A year from now the rate is either four point six one per cent after a rise, or three point four two per cent after a fall.
- Notice the up rate is higher than the down rate by the same proportion at every date.
- A year after that, there are three rates: five point nine seven per cent at the top, four point four two per cent in the middle, and three point two seven per cent at the bottom.
- The middle one can be reached two ways, up then down or down then up.
- We will lean on that fact shortly.

7 Example 1: working back, and the verdict

Work backwards on the 15% tree; then change the volatility

Year 1 node values	with coupon: 102.34 and 104.71	
Average of the two	(102.34 + 104.71) ÷ 2	**103.52**
Divide by 1 + 2.00%	103.52 ÷ 1.02	**101.49**
Lattice minus spot price	at 10%, 15% or 25%	**0.00**

IN THE LESSON

- Tomas works backwards.
- At year two, the last payment of one hundred and four is discounted one year at each of the three year two rates.
- Then he steps back to year one, averaging the two values that can follow each node, discounting at that node's rate, and adding the coupon.
- That gives one hundred and two point three four at the node after a rise, and one hundred and four point seven one at the node after a fall.
- Average the two, which is one hundred and three point five two, and divide by one point zero two.
- The answer is one hundred and one point four nine.
- Exactly Mira's number.
- He tries ten and twenty five per cent volatility too.
- Still the same.

8 Example 2: a callable bond needs the lattice

3-year 4% bond, callable at 100 on years 1 and 2; tree at 15%

Option-free bond	same on spot curve and tree	**101.49**
Callable bond	value capped at call price	**101.15**
Value of the call	101.494 − 101.146	**0.35**
Who owns the call	the issuer	**investor loses**

IN THE LESSON

- Now change the bond.
- Example two.
- The same three year, four per cent bond, but the issuer may call it at one hundred on either of the first two coupon dates.
- The payments now depend on what rates do.
- The spot curve cannot say when the issuer will call, because it has no rates by state.
- The tree can.
- At each node, if the bond's value would rise above the call price, the issuer calls and the value is capped.
- On the fifteen per cent tree the callable bond is worth one hundred and one point one five.
- The option free bond is worth one hundred and one point four nine.
- So the call is worth thirty five cents to the issuer, which is thirty five cents the investor gives up.

9 Volatility changes the call, not the plain bond

Value of the call (per 100)

10% vol	**0.16**
15% vol	**0.35**
25% vol	**0.72**

The plain bond is 101.49 at all three

IN THE LESSON

- Here is what happens when volatility changes.
- The three bars show the value of the call to the issuer at three levels of volatility.
- The vertical axis shows the value in price points per hundred of face value.
- At ten per cent volatility the call is worth sixteen cents.
- At fifteen per cent it is worth thirty five cents.
- At twenty five per cent it is worth seventy two cents.
- Why does it grow?
- When rates can swing further, they can fall further, and a falling rate is exactly when the issuer exercises the call.
- Meanwhile the plain bond stayed at one hundred and one point four nine every time.
- That is the proof that a correct lattice adds no value of its own.

10 Which tool for which bond

Do the cash flows depend on where rates go?

Instrument	Right tool
Plain fixed rate bond	Spot curve or tree
Callable bond	Lattice
Putable bond	Lattice
Capped floating bond	Lattice

IN THE LESSON

- Let us pull this together.
- Look at each bond type and ask one question.
- Do its cash flows depend on where rates go?
- A plain fixed rate bond does not, so the spot curve gives the right price, and a calibrated tree gives the same.
- A callable bond depends on rates, because the issuer calls when rates are low.
- A putable bond is the mirror image, because the investor puts it back when rates are high.
- A bond with a cap on its floating coupon also depends on rates.
- For all three, you need the lattice.
- The exam trap is to think the tree gives a different price for a plain bond because it contains volatility.
- It should not, unless it is mis calibrated.

11 A story: every route to the village

- List every full route rates could take
- Discount the bond along each route
- Each route is equally likely
- Average the route values

IN THE LESSON

- Imagine walking to a village over three days.
- Each morning you reach a fork, and each time you may go uphill or downhill.
- Different routes lead to different views, and a different total effort.
- Pathwise valuation is the same idea for interest rates.
- Instead of working backwards one node at a time, you list every full route the rate could take from today to the end, value the bond along each route, and then average the results.
- Because each fork is equally likely to go up or down, each full route is equally likely too.
- So a plain average is all you need.

12 Pathwise valuation: the rule

Value = average of path values
Number of paths = 2^(n − 1)

Value of one path	Cash flows discounted on one path
Number of paths	Two to the power n − 1
Years to the cash flow (n)	Sets how many forks
Bond value	Average over all paths

IN THE LESSON

- Here is the rule.
- The bond's value equals the simple average of the values along all the paths.
- The number of paths for a cash flow at year n is two raised to the power of n minus one.
- Why n minus one?
- The first rate, today's, is already known, so only the later years offer a choice.
- For a three year bond, that is two squared, which is four paths.
- Take the terms in turn.
- The value of one path is that path's cash flows each discounted at the rates along that path.
- The number of paths is two to the power of n minus one.
- And n, the years to the cash flow, sets how many choices there are.

13 Example 3: the four paths

Rates earned in years 1, 2 and 3 on each path

Path	Rates in years 1, 2, 3
Up, up	2.00 / 4.61 / 5.97
Up, down	2.00 / 4.61 / 4.42
Down, up	2.00 / 3.42 / 4.42
Down, down	2.00 / 3.42 / 3.27

IN THE LESSON

- Example three.
- We use the same fifteen per cent tree and the same three year, four per cent bond.
- List the four paths.
- Up then up: the rates are two per cent, four point six one and five point nine seven.
- Up then down: two per cent, four point six one, four point four two.
- Down then up: two per cent, three point four two, four point four two.
- Down then down: two per cent, three point four two, three point two seven.
- Notice the middle two paths end on the same final rate, but they are different paths, because one spends year two at four point six one and the other at three point four two.
- They will not be worth the same.

14 Example 3: value the bond on each path

Each cash flow is discounted at the rates along its own path

Up, up	3.92 + 3.75 + 91.98	**99.65**
Up, down	3.92 + 3.75 + 93.34	**101.01**
Down, up	3.92 + 3.79 + 94.42	**102.13**
Down, down	3.92 + 3.79 + 95.47	**103.18**

IN THE LESSON

- Now discount along each path, using that path's own rates.
- On the up up path, the coupons are worth three point nine two and three point seven five, and the final payment is ninety one point nine eight.
- The total is ninety nine point six five.
- Up then down gives one hundred and one point zero one.
- Down then up gives one hundred and two point one three.
- Down then down gives one hundred and three point one eight.
- See the pattern.
- The more the rates rise along the way, the less the bond is worth.
- The all rising path is cheapest, and the all falling path is dearest.
- Each path has the same chance, one in four.

15 Example 3: average, then check

Average of the four path values

Sum of the four paths	99.65 + 101.01 + 102.13 + 103.18	**405.97**
Divide by four paths	each path has a chance of 1/4	**÷ 4**
Pathwise value	average of the paths	**101.49**
Backward induction	same tree	**101.49**

IN THE LESSON

- Now the payoff.
- Add the four path values: ninety nine point six five, plus one hundred and one point zero one, plus one hundred and two point one three, plus one hundred and three point one eight.
- That comes to four hundred and five point nine seven.
- Divide by four paths, and you get one hundred and one point four nine.
- That matches backward induction exactly, and it matches Mira's spot curve answer too.
- Three methods, one price.
- For a bond with fixed cash flows, they must agree.
- Pathwise valuation is therefore a handy check on a tree, and later it becomes the method we really need when payments depend on the route.

16 Why counting paths by final node goes wrong

Weight in the average (%)

Top rate 5.97%	**25%**
Middle rate 4.42%	**50%**
Bottom rate 3.27%	**25%**

By final node you would wrongly use 33.3% each

IN THE LESSON

- Here is the classic mistake, drawn as bars.
- The vertical axis shows the weight each year three rate carries in the average, in per cent.
- There are three final rates but four paths.
- The top rate, five point nine seven, is reached by one path, so it carries a quarter of the weight, twenty five per cent.
- The bottom rate is also reached by one path, twenty five per cent.
- The middle rate, four point four two, is reached by two paths, so it carries half the weight, fifty per cent.
- If you counted by final nodes you would give each a third, thirty three per cent, and the middle rate would be given too little weight.

17 Example 4: when the path matters

Bonus of 5 at year 3, paid only on the down-down path

Plain bond value	average of the four paths	**101.49**
Bonus on down, down	5 ÷ (1.02 × 1.0342 × 1.0327)	**4.59**
Weight of that path	one path in four	**× 1/4**
Note value	101.49 + 1.15	**102.64**

IN THE LESSON

- Why bother with paths at all, if the answer is the same?
- Because some instruments pay according to history.
- Example four.
- A note pays the same four per cent coupons, but also a bonus of five at maturity, and only if rates fell in both years, the down then down path.
- A single node cannot know how it was reached, so backward induction on this tree cannot handle it.
- Pathwise valuation can.
- The bonus is paid only on one path, and its present value there is five divided by the three discount factors, which is four point five nine.
- Weight it by one quarter and the bonus adds about one point one five to the value of the note.

18 Mistakes to avoid so far

- 1. Plain bond: lattice must equal the spot curve
- 2. Count routes, not final nodes: 2^(n − 1)
- 3. Discount each path at its own rates
- 4. Path-dependent cash flows: use paths

IN THE LESSON

- Let us stop and collect the mistakes that cost marks.
- One.
- Expecting the lattice to value a plain bond differently from the spot curve.
- If it does, the tree is wrong, not clever.
- Two.
- Counting paths by final nodes.
- Count routes, not endpoints, so a three year cash flow has four paths.
- Three.
- Forgetting that each path is discounted at its own rates, not at the average rate.
- Four.
- Using backward induction for an instrument whose payments depend on the path.
- A node cannot remember how it was reached, so use paths.
- Keep these four in your head as we move on to simulation.

19 A story: mortgage homeowners remember

- Mortgage borrowers repay early when rates fall
- What they do depends on the route rates took
- A small tree cannot hold enough routes
- So simulate thousands of random routes

IN THE LESSON

- Think about a homeowner with a mortgage.
- When rates drop a long way, she refinances and repays early.
- When rates are rising, she stays put.
- But here is the catch.
- If rates fell last year and she already refinanced, she cannot refinance again this year.
- So what she does now depends on the whole route rates have taken, not just today's level.
- A security backed by thousands of such mortgages therefore has cash flows that depend on history.
- A tree with only a handful of routes struggles with this.
- Monte Carlo simulation instead rolls the dice thousands of times, to produce thousands of full routes, often with a step of one month.

20 Monte Carlo: four stages

Monte Carlo valuation

Simulate paths → Calibrate with a drift → Project cash flows → Discount, then average

IN THE LESSON

- Here are the four stages of a Monte Carlo valuation, read left to right.
- Stage one, simulate.
- Using an assumed volatility and a model of how rates behave, the computer generates many random paths of future rates, often monthly.
- Stage two, calibrate.
- Before trusting the paths, test them on benchmark bonds.
- If they misprice the benchmarks, add a constant drift to every simulated rate until they price correctly.
- Stage three, project.
- On each path, work out the security's cash flows, such as the prepayments.
- Stage four, discount each path's cash flows at that path's rates, and average across all the paths.

21 Where simulation is used

- Mortgage backed securities: prepayments follow the path
- Notes whose payments depend on past rates
- Keeps history; short steps, realistic models
- Cost: computing time and sampling error

IN THE LESSON

- So where does a bank use this?
- Mostly for securities whose payments depend on the path of rates.
- The leading case is the mortgage backed security, where prepayments depend on past rates.
- Other cases include bonds whose coupons depend on the history of rates, and some structured notes.
- Why not always use a tree?
- A tree has few paths and loses the history at each node.
- A simulation keeps the history, and can have a short time step, so it can handle a messy, realistic model.
- The price is computing time, and one other thing: the answer is an average of random paths, so it carries a small sampling error that we shall measure shortly.

22 Example 5: calibrating the paths

Today 3.00%; year-2 rates on four paths; benchmark zero trades at 93.81

Year-2 rates on the paths	5.0, 4.2, 3.4, 2.6 (%)	
Model price, unadjusted	mean of 100 ÷ (1.03 × (1+r))	**93.54**
Market price	benchmark zero	**93.81**
Gap	93.81 − 93.54	**0.27**

IN THE LESSON

- Example five, the calibration step.
- A simulation gives four illustrative paths for the one year rate in year two: five point zero, four point two, three point four and two point six per cent.
- Today's one year rate is three per cent.
- A two year zero coupon benchmark trades at ninety three point eight one.
- Value it on each path: one hundred divided by one point zero three, divided by one plus that year two rate.
- The four values are ninety two point four six, ninety three point one seven, ninety three point eight nine, and ninety four point six three.
- They average ninety three point five four.
- That is below the market price, so the gap is point two seven.

23 Example 5: the drift fix

Subtract the same drift from every simulated rate

Drift adjustment	applied to every rate	**−0.30%**
Adjusted year-2 rates	4.7, 3.9, 3.1, 2.3 (%)	
Benchmark on each path	92.73, 93.44, 94.17, 94.90	
Average	equals the market price	**93.81**

IN THE LESSON

- The model price of ninety three point five four is below the market price of ninety three point eight one, so the simulated rates are a little too high.
- The fix is one constant drift, subtracted from every simulated rate.
- A drift of minus zero point three per cent works.
- The year two rates become four point seven, three point nine, three point one and two point three per cent.
- Reprice the benchmark on these paths and the four values are ninety two point seven three, ninety three point four four, ninety four point one seven, and ninety four point nine zero.
- The average is ninety three point eight one, exactly the market price.
- Now the paths are fit to use.

24 Example 6: value the security on the paths

Cash flows in years 1 and 2 on each adjusted path (illustration)

Path 1, rate 4.7%	8 and 98	**98.64**
Path 2, rate 3.9%	12 and 94	**99.49**
Path 3, rate 3.1%	20 and 85	**99.46**
Path 4, rate 2.3%	30 and 74	**99.36**
Average of the paths	security value	**99.24**

IN THE LESSON

- Now use the calibrated paths on the security itself.
- Example six.
- A small pass through security pays cash at the end of year one and year two.
- Its cash flows depend on the path, because when rates fall, borrowers repay early, so more cash arrives in year one and less in year two.
- On path one the cash flows are eight and ninety eight.
- On path two, twelve and ninety four.
- Path three, twenty and eighty five.
- Path four, thirty and seventy four.
- Discount each at that path's rates.
- The path values are ninety eight point six four, ninety nine point four nine, ninety nine point four six, and ninety nine point three six.
- The average is ninety nine point two four.

25 Sampling error and the drift

Sampling error = s ÷ √N
Adjusted rate = simulated rate + drift

Spread of path values (s)	How far path values vary
Number of paths (N)	More paths, smaller error
Sampling error	Noise in the average
Drift adjustment	Same shift on every rate

IN THE LESSON

- Two small formulas.
- First, the sampling error of the average.
- It equals the spread of the path values, divided by the square root of the number of paths.
- In example six the four path values spread by about forty cents.
- A real run might have a spread of two points and use two thousand five hundred paths.
- The error is then two divided by fifty, which is four cents.
- Second, the drift.
- The adjusted rate is the simulated rate plus the drift, and the drift is chosen so benchmark bonds are priced exactly.
- Notice what more paths do.
- They shrink the sampling error.
- They do not fix a biased model.

26 More paths: error shrinks slowly

Sampling error (100 paths = 100)

100 paths	**100**
400 paths	**50**
1,600 paths	**25**
6,400 paths	**12.5**

Half the error needs four times the paths

IN THE LESSON

- Here is how sampling error falls as you add paths.
- The vertical axis shows the error relative to a run of one hundred paths.
- With one hundred paths the error is a hundred, our starting point.
- Quadruple the paths to four hundred and the error halves to fifty.
- Quadruple again to sixteen hundred and it halves again to twenty five.
- At sixty four hundred paths it is twelve and a half.
- Each halving costs four times the work.
- That is the square root at work.
- Remember the limit of this: if the model is biased, say it misprices the benchmark bond, then ten thousand paths give you a very precise wrong answer.

27 A story: a rubber band on the short rate

- A model says how the short rate moves
- Drift: a pull back to a long-run level
- Shock: a random push each step
- Models differ in the rules they choose

IN THE LESSON

- Think of a dog on a long elastic lead, tied to a post.
- The dog wanders randomly, but the further it strays, the harder the lead pulls it back.
- The short term interest rate behaves a little like that dog.
- It jumps around randomly, but central banks and markets pull it back towards a long run level.
- A term structure model is just a rule for how the short rate moves: a pull back, called the drift, plus a random push, called the shock.
- Different models choose different rules.
- These rules matter because we need them to build the trees and the simulations that we have just used.

28 Vasicek and Cox-Ingersoll-Ross

Vasicek: $\Delta r = a(b - r)\Delta t + \sigma z$
CIR: $\Delta r = a(b - r)\Delta t + \sigma\sqrt{r}\, z$

Speed of reversion (a)	How fast rates are pulled back
Long-run rate (b)	Where rates settle
Current short rate (r)	Rate today
Volatility (σ)	Size of the random push
Random shock (z)	A random draw each step

IN THE LESSON

- Both well known equilibrium models share the same pull back.
- The change in the short rate equals the speed a, times the gap between the long run rate b and today's rate r, times the time step, plus a random shock.
- The speed a says how fast the pull works.
- The long run rate b is where rates tend to settle.
- In Vasicek, the shock is the volatility sigma times a random number, so the size of the shock is the same whatever the rate.
- In Cox-Ingersoll-Ross, the shock also has the square root of the rate in it, so as rates get small, the shocks shrink.

29 Mean reversion: the picture

Chart: Expected short rate vs Years ahead (0 to 8)

- Long-run rate 4%
- 1% today
- Pulled back up

IN THE LESSON

- Here is mean reversion, drawn.
- Time in years runs along the bottom, and the expected short rate up the side.
- The flat line is the long run rate, four per cent.
- The curve starts low, at one per cent today, and climbs towards the long run line.
- At first it rises quickly, because the gap is wide and the pull is strong.
- As the gap closes, the pull weakens and the curve flattens, getting close to four per cent but never quite touching it.
- If the rate started above four per cent, the curve would instead fall towards the line.
- This pull is common to both Vasicek and Cox-Ingersoll-Ross.

30 Example 7: one quarter, two models

Speed 0.3, long-run 4%, rate 1%, step of a quarter-year (illustration)

Pull back, both models	0.3 × 3% × 0.25	**+0.225%**
Vasicek shock	2% × √0.25	**1.00%**
Cox-Ingersoll-Ross shock	0.1 × √0.01 × √0.25	**0.50%**
Two shocks down	Vasicek vs CIR	**−0.78% | 0.23%**

IN THE LESSON

- Example seven.
- Both models share a speed of mean reversion of point three and a long run rate of four per cent.
- Today's short rate is one per cent, and we step forward one quarter.
- The pull back is point three times three per cent times a quarter, which is point two two five per cent, so the expected rate is one point two two five per cent in both models.
- The Vasicek shock has a size of two per cent times the square root of a quarter, which is one per cent.
- In Cox Ingersoll Ross, a sigma of point one gives the same two per cent when the rate is four, so the shock at one per cent is half that: point five per cent.
- Two shocks down from the expected rate, Vasicek goes below zero.
- Cox-Ingersoll-Ross stays positive.

31 Four models, side by side

Family, volatility and whether rates can go negative

Vasicek	equilibrium, constant vol.	**can go below 0**
Cox-Ingersoll-Ross	equilibrium, vol. ∝ √rate	**stays positive**
Ho-Lee	arbitrage-free, constant vol.	**can go below 0**
Kalotay-Williams-Fabozzi	arbitrage-free, log of rate	**stays positive**

IN THE LESSON

- Now put the names to the rules.
- Vasicek and Cox-Ingersoll-Ross are equilibrium models.
- They start from assumptions about how the rate behaves, and the yield curve comes out as an answer, so it may not match today's market curve.
- Vasicek has constant volatility, so negative rates are possible.
- Cox-Ingersoll-Ross has volatility that grows with the square root of the rate, so rates cannot go negative.
- The other pair are arbitrage free models.
- Ho Lee has constant volatility and a time varying drift chosen to match today's curve exactly; it can produce negative rates.
- Kalotay Williams Fabozzi models the logarithm of the rate, which keeps it positive.

32 Which family for which job

Equilibrium models

Assume rate behaviour → Curve comes out → May miss the market

Arbitrage-free models

Take today's curve → Time-varying drift → Prices derivatives consistently

IN THE LESSON

- Why have two families?
- Because they answer different questions.
- An equilibrium model starts from the economy: how fast do rates revert, where do they settle?
- Its curve is a result, and can disagree with the market.
- That is fine for studying how rates might behave.
- An arbitrage free model works the other way.
- It takes today's curve as given and adds a drift that changes with time, so the model reproduces the market curve exactly.
- That is what you need to price a derivative, such as an option on a bond, consistently with the bonds that are trading.
- If a model misprices the bonds, the option prices built from it are not trustworthy.

33 Recap: four ideas in one picture

Four valuation ideas

Lattice value	Equals spot curve for a plain bond
Pathwise value	Average over all paths
Simulation value	Thousands of paths, drift-fit
Term structure model	Equilibrium or arbitrage-free

IN THE LESSON

- Let us pull it together.
- One.
- For a plain bond, the spot curve and a calibrated lattice give the same value.
- The lattice earns its place when payments depend on rates, as with calls and puts.
- Two.
- Pathwise valuation lists all two to the power of n minus one paths, discounts each at its own rates and averages them.
- Three.
- Monte Carlo simulation does this with thousands of random paths.
- You calibrate with a drift, and more paths cut noise but not bias.
- Four.
- Term structure models describe how the short rate moves, and fall into equilibrium and arbitrage free families.

34 Try three quick questions

1. Lattice price for a plain bond differs from the spot price. Why?
   Answer: It is mis-calibrated: fixed cash flows must agree.
2. How many paths does a year-4 cash flow have?
   Answer: $2^3 = 8$ paths, since there are 2^(n − 1).
3. Which model keeps rates positive and matches today's curve?
   Answer: Kalotay-Williams-Fabozzi: models the log of the rate.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A calibrated lattice values a plain three year bond at a different price from the spot curve.
- What does that tell you?
- The answer is that the lattice is mis calibrated, because the two must agree for fixed cash flows.
- Question two.
- How many paths does a cash flow at year four have in a binomial tree?
- The answer is eight, which is two cubed.
- Question three.
- Which model keeps rates positive and matches today's curve exactly?
- The answer is Kalotay Williams Fabozzi, the arbitrage free model of the log of the rate.