CHARTA LEARNING · LESSON NOTES

Term Structure Models: Drift

FRM Part II

1 A story: one rate drives the whole curve

- Short rate: the rate for the next short period
- Term structure model: how the short rate moves
- Drift: the expected change. Volatility: the random part
- Drawn as a tree: up or down at every step
- Prices are worked backward through the tree

IN THE LESSON

- Halden's desk holds bonds and interest rate options.
- To price them all consistently, the desk wants one model of the short rate, which is the interest rate for borrowing over the next instant, or in practice over the next short period.
- A term structure model describes how that short rate moves.
- It has two parts.
- The drift is the expected change, the part we can predict.
- The volatility is the random part.
- In practice the model is drawn as a tree.
- Each step is a short slice of time, and from every node the rate moves up or down.
- Prices are then worked backward through the tree.

2 The normal model, with and without drift

dr = λ × dt + σ × dw
Node move = λ × dt ± σ × √dt

Drift, λ: zero or a constant	Zero, or a constant
σ: annual volatility in bp	Annual volatility, in bp
dt: the time step, in years	The time step, in years

IN THE LESSON

- The simplest model says the rate changes by a random amount that does not depend on the level of the rate.
- With no drift, d r equals sigma times d w.
- Here sigma is the annual volatility, quoted in basis points, and d w is a random shock with mean zero.
- With a constant drift lambda, we add lambda times d t, where d t is the time step in years.
- In a tree, each node moves to a new rate.
- The move is lambda times d t, plus or minus sigma times the square root of d t, each with probability one half.
- When the drift is zero, the first term simply drops out.

3 Example 1: Halden's zero-drift tree

r0 = 4.00%, σ = 100 bp a year, quarterly steps (dt = 0.25), no drift

Move per quarter	1.00% × √0.25	**±0.50%**
Date 1 nodes	4.00 ± 0.50	**4.50%, 3.50%**
Date 2 nodes	up-down = down-up	**5.00, 4.00, 3.00%**
Expected rate at date 2	¼×5 + ½×4 + ¼×3	**4.00%**
Nodes after 40 steps	n + 1, not 2^n	**41, not 2^{40}**

IN THE LESSON

- Halden's short rate today, r nought, is four per cent, volatility is one hundred basis points a year, and steps are quarterly.
- The move per quarter is one per cent times the square root of point two five, which is plus or minus point five zero per cent.
- At date one the nodes are four point five zero and three point five zero per cent.
- At date two, up then down equals down then up, so the nodes are five, four and three per cent.
- The expected rate at date two is a quarter of five plus a half of four plus a quarter of three, still four per cent.
- Because it recombines, after forty steps there are forty one nodes, not two to the forty, about one point one trillion.

4 Example 2: adding a constant drift

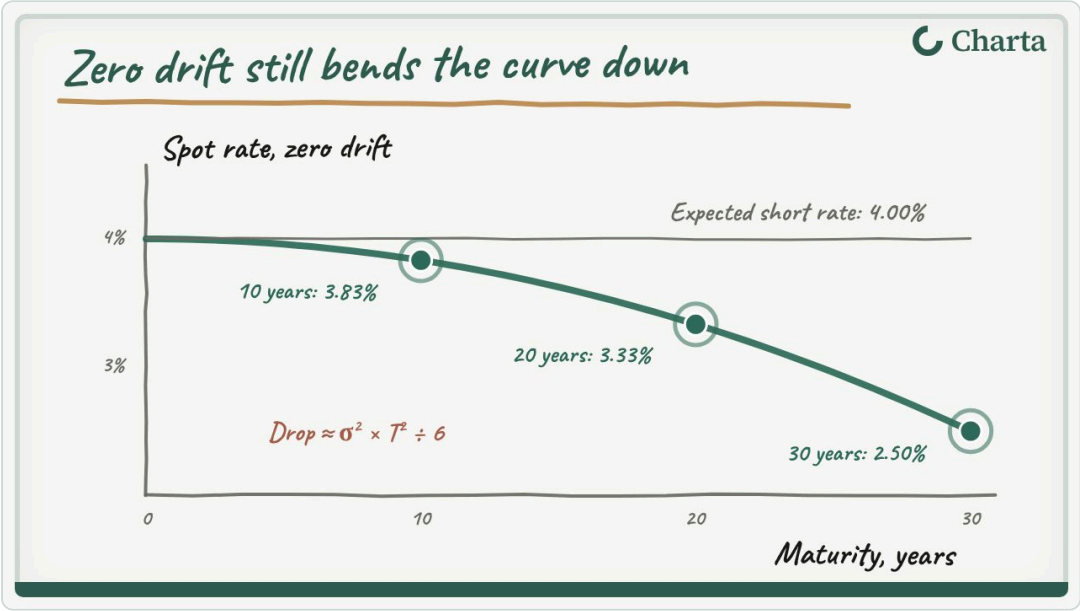
Same tree, plus a constant drift λ = 0.40% a year

Drift per quarter	0.40% × 0.25	**+0.10%**
Date 1 nodes	4.00 + 0.10 ± 0.50	**4.60%, 3.60%**
Date 2 nodes	4.20 ± 1.00	**5.20, 4.20, 3.20%**
Quarter: drift vs vol	40 × 0.25 vs 100 × √0.25	**10 vs 50 bp**
4 years: drift vs vol	40 × 4 vs 100 × √4	**160 vs 200 bp**

IN THE LESSON

- Now add a constant drift of point four zero per cent a year: per quarter, point four zero times point two five, or point one zero per cent.
- At date one the nodes are four plus point one zero, plus or minus point five zero: four point six zero and three point six zero per cent.
- At date two the centre is four plus two drifts, four point two zero, plus or minus one per cent: five point two zero, four point two zero and three point two zero per cent.
- Over a quarter the drift is ten basis points, against volatility of one hundred times root point two five, fifty basis points.
- Over four years the drift is one hundred and sixty basis points, against two hundred.
- The exam point: drift scales with time itself, volatility with the square root of time.

5 Zero drift still bends the curve down



IN THE LESSON

- With zero drift the expected short rate stays at four per cent forever, the grey line.
- Yet the spot curve the model produces, in green, slopes down.
- The reason is convexity.
- A bond price is a convex function of rates, so when rates are uncertain, the average bond price is higher than the price at the average rate.
- A higher price means a lower yield.
- The effect is roughly sigma squared times T squared, divided by six.
- At ten years it takes off about seventeen basis points, giving three point eight three per cent.
- At twenty years the spot rate is three point three three per cent, and at thirty years only two point five zero.

6 Example 3: one drift fits one point

Market spots: 2-year 4.30%, 10-year 4.50%. Fit one constant drift λ

Model spot at T	r0 + λT/2 − σ^{2}T^{2}/6	**approximate**
Convexity at 10 years	0.01^{2} × 100 ÷ 6	**0.167%**
Drift that fits 10 years	(4.50 − 4.00 + 0.167) ÷ 5	**λ = 0.133%**
Model 2-year spot	4.00 + 0.133 − 0.007	**4.13%**
Miss at 2 years	4.13 − 4.30	**−17 bp**

IN THE LESSON

- Suppose the market's two year spot rate is four point three zero per cent and its ten year spot rate four point five zero.
- Approximately, the model spot rate at maturity T is r nought, plus lambda times T over two, minus sigma squared times T squared over six.
- The convexity term at ten years is point zero one squared times one hundred, over six, which is point one six seven per cent.
- To fit, lambda is four point five zero minus four, plus point one six seven, all divided by five: point one three three per cent a year.
- Now the model's two year spot is four, plus point one three three times two over two, minus the convexity at two years, point zero one squared times four over six, or point zero zero seven: about four point one three per cent.
- Four point one three minus four point three zero: the model is seventeen basis points too low.

7 How well do these simple trees work?

- Recombines: n + 1 nodes, quick to build and price
- Zero drift: spot curve slopes down by about σ^{2}T^{2}/6
- One drift fits one point: neither version fits a real curve
- Rates can go negative; volatility fixed in bp
- Exam: volatility grows with √t, drift with t

IN THE LESSON

- On the plus side, the tree recombines, so it has n plus one nodes after n steps, and it is quick to build and to price on.
- On the minus side there are three problems.
- With zero drift, convexity makes the spot curve slope down, by about sigma squared T squared over six.
- A constant drift lifts the curve, but one drift can fit only one point, as Halden just found, so neither version fits a real curve.
- And rates can go negative, while volatility stays fixed in basis points whatever the rate level.
- The exam point: volatility grows with the square root of time, drift with time itself.

8 How far can the rate move?

SD over one step = σ × √dt
SD over T years = σ × √T

σ: annual volatility in bp	Per year, in basis points
dt: the time step, in years	Month 1/12; trading day 1/252
Variances add, not SDs	So SD grows with √time

IN THE LESSON

- With zero drift, the change in the short rate over one step is normally distributed with mean zero.
- Its standard deviation is sigma times the square root of d t.
- Over a horizon of T years, the standard deviation is sigma times the square root of T.
- Sigma is quoted per year in basis points, so T and d t must be in years too.
- A monthly step is one twelfth of a year, and a trading day one two hundred and fifty second.
- Why a square root?
- Because independent changes add in variance, not in standard deviation.
- Variance grows in line with time, so standard deviation grows with its square root.

9 Example 4: scaling Halden's volatility

Halden: σ = 100 bp a year. Expected change zero; variances add

One month	100 × √(1/12)	**28.87 bp**
One trading day	100 × √(1/252)	**6.30 bp**
One quarter	100 × √0.25	**50.00 bp**
One year, then	100 × √1, then	**100, then**
four years	100 × √4	**200 bp**
Five years	100 × √5	**223.61 bp**

IN THE LESSON

- Take Halden's one hundred basis points a year.
- The expected change is zero, and variances add.
- Over a month, d t is one twelfth, and one hundred times the square root of one twelfth is twenty eight point eight seven basis points.
- Over one of two hundred and fifty two trading days a year, it is one hundred times the square root of one over two hundred and fifty two, six point three zero basis points.
- Over one quarter, the square root of point two five is point five, so fifty basis points.
- Over one year it is one hundred, and over four years one hundred times two, which is two hundred: quadrupling the horizon only doubles the standard deviation.
- Over five years it is two hundred and twenty three point six one basis points.

10 Example 5: the usual scaling errors

Scaling Halden's 100 bp annual volatility to a month, and to a day

Right: scale by √dt	100 × √(1/12)	**28.87 bp**
Wrong: scale by dt	100 × 1/12	**8.33 bp**
Wrong: quote the variance	100^2 × 1/12 = 833 bp^2	**not an SD**
Wrong: annual σ for a day	100 ÷ 6.30 = √252	**15.87 times**

IN THE LESSON

- Now scale Halden's one hundred basis points to a month and to a day.
- The right answer scales by the square root of d t: one hundred times the root of one twelfth, twenty eight point eight seven basis points.
- The first error scales by d t itself: one hundred divided by twelve, eight point three three basis points, far too small.
- The second error quotes the variance as if it were the standard deviation.
- One hundred squared over twelve is eight hundred and thirty three basis points squared, not a standard deviation.
- The third error mixes daily and annual volatility.
- Using the annual one hundred basis points for a single day overstates the daily figure of six point three zero by a factor of the square root of two hundred and fifty two, about fifteen point eight seven.

11 How likely is a negative rate?

Z = −r ÷ (σ × √T)
Probability of a negative rate = N(Z)

r: today's short rate	Today's short rate
σ: annual volatility in bp	Annual volatility
T: the horizon, in years	Horizon, in years
N: the standard normal CDF	Standard normal CDF

IN THE LESSON

- A normal model has no lower bound.
- Every future date carries a positive probability that the short rate is below zero.
- For a zero drift model, the distance below zero, measured in standard deviations, is Z.
- It equals minus r, divided by sigma times the square root of T.
- Here r is today's short rate, sigma the annual volatility and T the horizon in years.
- The probability of a negative rate at that horizon is N of Z, where N is the standard normal cumulative distribution function.
- The probability is small when the starting rate is high and volatility low, but it grows as the horizon lengthens.

12 Example 6: Halden's chance of a negative rate

Zero drift, σ = 100 bp a year. Chance the short rate is below zero at T

r = 4%, T = 1 year	Z = −4 ÷ (1 × √1) = −4.00	**0.003%**
r = 4%, T = 5 years	Z = −4 ÷ √5 = −1.79	**3.7%**
r = 4%, T = 10 years	Z = −4 ÷ √10 = −1.26	**10.3%**
r = 1%, T = 1 year	Z = −1 ÷ 1 = −1.00	**15.9%**

Small for high rates and short horizons; grows with T

IN THE LESSON

- With zero drift and one hundred basis points of volatility, what is the chance of a negative rate?
- Starting from four per cent, over one year Z is minus four divided by one, minus four point zero zero, and the probability is just point zero zero three per cent.
- Over five years Z is minus four over the square root of five, minus one point seven nine, and the probability rises to three point seven per cent.
- Over ten years Z is minus four over the square root of ten, minus one point two six, and the probability is ten point three per cent.
- Now suppose rates start at only one per cent.
- Over a single year Z is minus one, and the probability is fifteen point nine per cent.

13 Three ways to cope

- No lower bound: a negative rate always has some chance
- Accept it when the probability is small
- Floor the tree at zero: simple, but no longer normal
- Rate-dependent volatility: lognormal or square root
- Mean reversion alone does not remove negative rates

IN THE LESSON

- How do term structure models cope?
- With no lower bound, a negative rate always has some chance.
- The first answer is to accept it, when the probability is small enough not to matter for the securities being priced.
- The second is to floor the tree at zero, replacing every negative node with zero.
- That is simple, but the rate is no longer normal, and the floor lifts the expected rate above the unfloored tree's.
- The third is to change the volatility so that it depends on the rate, as in a lognormal or square root model.
- Note what does not work: mean reversion alone does not remove negative rates.

14 Example 7: what a floor does

Low-rate tree: r0 = 0.50%, σ = 100 bp, one annual step (dt = 1)

Unfloored nodes	0.50 ± 1.00	**1.50%, −0.50%**
Unfloored expected rate	½ × (1.50 − 0.50)	**0.50%**
Floored nodes	max(rate, 0)	**1.50%, 0.00%**
Floored expected rate	½ × (1.50 + 0)	**0.75%**
Lift from the floor	0.75 − 0.50	**+25 bp**

IN THE LESSON

- Take a low rate tree with one annual step: the rate starts at point five zero per cent, and the move is a full one per cent.
- Without a floor, the nodes are point five zero plus or minus one: one point five zero and minus point five zero per cent.
- The expected rate is half of one point five zero minus point five zero, still point five zero per cent.
- With a floor, each node becomes the larger of the rate and zero, so the nodes are one point five zero and zero.
- The expected rate is now half of one point five zero, which is point seven five per cent.
- The floor has lifted the expected rate by twenty five basis points, and the rate is no longer normal.

15 Volatility that depends on the rate

Each set to 100 bp at a 4% rate; rate-dependent volatility shrinks as the rate falls

Model	Volatility rule	bp vol at 4%	bp vol at 1%
Normal	σ = 1.00%	100	100
Lognormal	25% × r	100	25
Square root	0.05 × √r	100	50

IN THE LESSON

- Each model is set to give one hundred basis points at a four per cent rate, and rate dependent volatility shrinks as the rate falls.
- In the normal model sigma is one per cent: one hundred at either rate.
- In a lognormal model, volatility is proportional to the rate: twenty five per cent of the rate.
- That is one hundred basis points at four per cent, but only twenty five at one per cent.
- In a square root model, volatility is point zero five times the square root of the rate.
- At four per cent that is point zero five times point two, one hundred basis points; at one per cent, point zero five times point one, fifty.
- Near zero the moves shrink, so rates cannot go negative: the lognormal model stays above zero, the square root model at or above it.
- The cost: basis point volatility now moves with the rate level.

16 The Ho-Lee model: drift that depends on time

dr = λ(t) × dt + σ × dw
λ(t) ≈ forward-curve slope + σ^2 × t

λ(t): shifts the whole tree	Set date by date to fit prices
σ: sets the spread of the tree	Constant, in basis points
dt: the time step, in years	The time step, in years

IN THE LESSON

- We saw that one constant drift cannot fit the curve.
- The Ho-Lee model keeps normally distributed rates and constant volatility, but lets the drift depend on time.
- The rate change is lambda of t times d t, plus sigma times d w.
- Lambda of t is chosen date by date, so that the tree reprices the observed zero coupon bonds.
- In continuous time, lambda of t is roughly the slope of the forward curve plus sigma squared times t.
- The second term is the convexity correction we met earlier.
- The exam point is the division of labour: the drift shifts the whole tree up or down at each date, while volatility sets how far the tree spreads.

17 Example 8: a Ho-Lee tree for Halden

Ho-Lee, quarterly: r0 = 4.00%, σ√dt = 0.50%, λ1dt = +0.20%, λ2dt = −0.10%

Date 1 nodes	4.00 + 0.20 ± 0.50	**4.70%, 3.70%**
Up then down	4.70 − 0.10 − 0.50	**4.10%**
Down then up	3.70 − 0.10 + 0.50	**4.10%**
Date 2 nodes	4.10 ± 1.00	**5.10, 4.10, 3.10%**
Rule: j ups at date n	r0 + Σ drifts + (2j − n)σ√dt	**2 ups: 5.10%**

IN THE LESSON

- Halden fits a Ho-Lee tree with quarterly moves of point five zero per cent and fitted drifts per quarter of plus point two zero, then minus point one zero per cent.
- At date one the nodes are four plus point two zero, plus or minus point five zero: four point seven zero and three point seven zero per cent.
- Up then down is four point seven zero, minus point one zero, minus point five zero, which is four point one zero.
- Down then up is three point seven zero, minus point one zero, plus point five zero, also four point one zero.
- The drift depends only on the date, so the tree recombines; the date two nodes are four point one zero plus or minus one per cent: five point one zero, four point one zero and three point one zero per cent.
- The rule: r nought plus the drifts so far, plus two j minus n, times sigma root d t; two ups gives five point one zero.

18 Example 9: fitting the first drift

New annual tree (dt = 1), annual compounding: r0 = 4%, σ = 1%. Market 2-year zero 92.02

Price, discounting by 1 ÷ (1 + r)	avg(100 ÷ (1 + r1)) ÷ 1.04	**annual compounding**
Try λ1 = 0	rates 5.00% and 3.00%	**92.46**
Compare with the market	92.46 > 92.02: rates too low	**raise λ1**
Try λ1 = 0.50%	rates 5.50% and 3.50%	**92.02**
Fitted drift	matches the market price	**λ1 = 0.50%**

Next: hold λ1, choose λ2 to price the 3-year zero

IN THE LESSON

- To fit a drift, switch from the quarterly tree to a new one: annual steps and annual compounding, with r nought four per cent and sigma one per cent.
- A two year zero coupon bond trades at ninety two point zero two.
- With annual compounding, each step discounts at one over one plus the rate, so the price is the average of one hundred over one plus each date one rate, divided by one point zero four.
- First try a drift of zero.
- The date one rates are five and three per cent, and the price is ninety two point four six.
- That is above the market, so rates are too low and the drift must rise.
- Try a drift of point five zero per cent.
- The rates become five point five zero and three point five zero, and the price is ninety two point zero two, matching the market.
- The fitted first drift is point five zero per cent; the next is then chosen to price the three year bond.

19 What arbitrage-free models are for

- Takes today's market prices as inputs
- Parameters fitted so benchmark bonds reprice exactly
- Use: price derivatives and off-the-run bonds consistently
- Not a forecast of future rates
- Cannot call fitted benchmarks rich or cheap: they match

IN THE LESSON

- Ho-Lee is an arbitrage free model.
- Such a model takes today's market prices as inputs.
- Its parameters, such as the Ho-Lee drift at each date, are fitted so the model reprices the benchmark bonds exactly.
- What is it for?
- Pricing other securities, mainly derivatives and off the run bonds, consistently with those benchmarks, with no arbitrage against them.
- An off the run bond is an older issue that trades less than the newest benchmark.
- What is it not for?
- It is not a forecast of future rates.
- And it cannot call the fitted benchmarks rich or cheap, since they match the model by construction.

20 Example 10: noise in the fitted drift

Same tree. The 2-year input is 5 cents higher: 92.07, from bid-ask noise

Fitted drift at 92.02	from Example 9	**0.50%**
Refit at 92.07	solve again on the tree	**0.445%**
Check	rates 5.45% and 3.45%	**92.07**
Change in the drift	0.445 − 0.500	**−5.5 bp**
Effect on the next date	λ2 tends to offset it	**zigzag**

IN THE LESSON

- Matching is not free.
- Suppose the two year quote sat five cents above the mid price, so the input is ninety two point zero seven instead of ninety two point zero two.
- The fitted drift was point five zero per cent; solving again, it is point four four five per cent.
- Check it: the date one rates are five point four five and three point four five per cent, and the price is ninety two point zero seven.
- The fitted drift has moved by about minus five and a half basis points because of five cents of noise.
- If the three year price is clean, the next drift swings back to price it, so fitted drifts tend to zigzag.

21 Why matching the market is hard

- One parameter per date: many prices, many parameters
- Fitted values absorb noise, such as bid-ask, and swing
- Bond prices say nothing about volatility
- Volatility is calibrated to option prices separately
- Exam: a perfect fit does not validate the dynamics

IN THE LESSON

- Here are the difficulties of making a model match observed prices.
- First, each date needs its own parameter, so many prices need many parameters.
- Second, the fitted values absorb noise, such as bid ask spreads and stale quotes, and so they can swing erratically, as we just saw.
- Third, a drift fitted to bond prices says nothing about volatility.
- Volatility must be calibrated separately, to the prices of options such as caps and swaptions.
- And the exam point: a perfect fit to today's bond prices does not validate the model's dynamics.
- Any model with enough free parameters can fit today; the test is whether it moves rates realistically.

22 The Vasicek model: drift that depends on the rate

$dr = k \times (\theta - r) \times dt + \sigma \times dw$
$Gap = 2 \times k \times \sigma \times dt \times \sqrt{dt}$

k: speed of mean reversion	Speed of mean reversion
θ: the long-run rate	The long-run rate
σ: annual volatility in bp	Annual volatility, in bp

IN THE LESSON

- The Vasicek model lets drift depend on the current rate.
- The rate change is k times theta minus r, times d t, plus sigma times d w.
- Here k is the speed of mean reversion, theta is the long run rate, and sigma is the annual volatility in basis points.
- The sign of theta minus r sets the direction: below theta the drift pulls up, above it pulls down.
- Because the drift now differs from node to node, up then down no longer lands exactly where down then up does.
- The gap between the two is two times k times sigma times d t times the square root of d t, small for short steps.

23 Example 11: the first Vasicek step

Vasicek, quarterly: $r0 = 4\%$, $\theta = 5\%$, $k = 0.35$, $\sigma\sqrt{dt} = 0.50\%$

Drift at 4.00%	$0.35 \times (5 - 4) \times 0.25$	**+0.0875%**
Up node	$4.00 + 0.0875 + 0.50$	**4.5875%**
Down node	$4.00 + 0.0875 - 0.50$	**3.5875%**
Drift at the up node	$0.35 \times (5 - 4.5875) \times 0.25$	**+0.0361%**
Drift at the down node	$0.35 \times (5 - 3.5875) \times 0.25$	**+0.1236%**

IN THE LESSON

- Halden now fits Vasicek with theta five per cent and k point three five, keeping quarterly moves of point five zero per cent.
- At four per cent, the drift is point three five times five minus four, times point two five, which is point zero eight seven five per cent.
- The up node is four plus point zero eight seven five plus point five zero, four point five eight seven five per cent.
- The down node is four point zero eight seven five minus point five zero, three point five eight seven five.
- At the up node it is point three five times five minus four point five eight seven five, times point two five, point zero three six one per cent.
- At the down node it is point one two three six per cent.
- The higher rate gets a smaller pull up.

24 Example 12: the gap, and making it recombine

Second quarter, from the up node (4.5875%) and the down node (3.5875%)

Up then down	4.5875 + 0.0361 − 0.50	**4.1236%**
Down then up	3.5875 + 0.1236 + 0.50	**4.2111%**
Gap	4.2111 − 4.1236	**8.75 bp**
Formula check	2 × 0.35 × 1% × 0.25 × 0.5	**0.0875%**
Middle node: the average	(4.1236 + 4.2111) ÷ 2	**≈ 4.167%**

Outer nodes 5.1236%, 3.2111%. Averaging shifts E[r] ≈ 2 bp

IN THE LESSON

- Now take the second quarter.
- Up then down is four point five eight seven five, plus point zero three six one, minus point five zero: four point one two three six per cent.
- Down then up is three point five eight seven five, plus point one two three six, plus point five zero: four point two one one one.
- The gap is four point two one one one minus four point one two three six, point zero eight seven five per cent, or eight point seven five basis points.
- The formula agrees: two times point three five times one per cent times point two five times point five.
- To recombine, the middle node is set to one value, such as the average, four point one six seven per cent.
- The outer nodes are five point one two three six and three point two one one one.
- The average is a simplification: it shifts each date one node's expected rate by about two basis points, so in practice the probabilities or outer nodes are re-solved.

25 Vasicek: where the rate is expected to go

E[r_T] = θ + (r0 − θ) × e^{−kT}
Half-life = ln(2) ÷ k

r0: today's short rate	Today's short rate
θ: the long-run rate	The long-run rate
k: speed of mean reversion	Speed of mean reversion

IN THE LESSON

- Over a single step, the expected change in the Vasicek rate is the drift, k times theta minus r, times d t.
- Its standard deviation is the same as in the zero drift model, sigma times the square root of d t.
- Over a longer horizon the gap between the rate and its long run mean shrinks by the factor e to the power minus k T.
- So the expected rate T years from now is theta plus r nought minus theta, times e to the minus k T.
- The rate approaches theta without overshooting it.
- The time for the expected gap to halve is the half-life, the natural log of two divided by k.

26 Example 13: Halden's Vasicek numbers

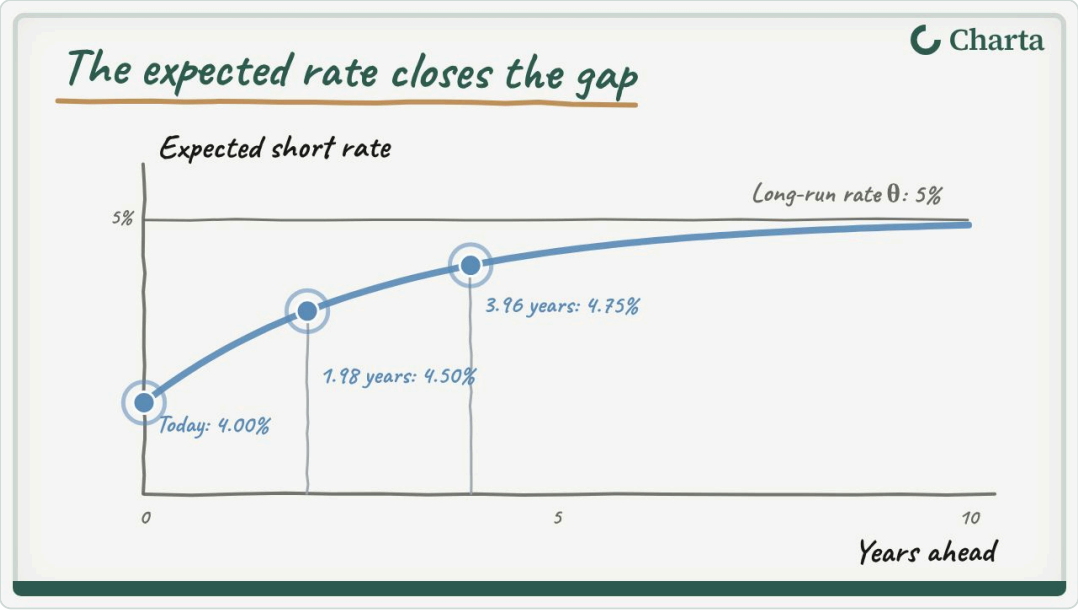
Halden: r0 = 4%, θ = 5%, k = 0.35, σ = 100 bp a year

Drift, one quarter	0.35 × (5 − 4) × 0.25	**+8.75 bp**
SD, one quarter	100 × √0.25	**50 bp**
Expected rate in 2 years	5 − 1 × e^{−0.70}	**4.503%**
Expected rate in 5 years	5 − 1 × e^{−1.75}	**4.826%**
Half-life	ln(2) ÷ 0.35	**1.98 years**

IN THE LESSON

- With Halden's r nought of four per cent, theta five, k point three five and sigma one hundred basis points, over one quarter the expected change is point three five times one times point two five, eight point seven five basis points.
- The standard deviation over the quarter is one hundred times the square root of point two five, fifty basis points, so the random part dominates in the short run.
- In two years, the expected rate is five minus one times e to the minus point seven zero, four point five zero three per cent.
- In five years it is five minus e to the minus one point seven five, four point eight two six per cent.
- The half-life is the natural log of two over point three five, one point nine eight years.

27 The expected rate closes the gap



IN THE LESSON

- Here is Halden's expected short rate over ten years.
- It starts at four per cent and rises towards the long run rate of five per cent, the grey line, but never crosses it.
- The path closes a fixed fraction of the remaining gap each year.
- After one half-life, one point nine eight years, half of the one percentage point gap remains, so the expected rate is four point five zero per cent.
- After two half-lives, three point nine six years, a quarter remains, and the expected rate is four point seven five.
- The exam point: after one half-life half of the gap to theta remains, after two a quarter.

28 Example 14: three classic mistakes

Same Halden numbers: r0 = 4%, θ = 5%, k = 0.35, quarterly steps

Half-life as 1 ÷ k	1 ÷ 0.35	**2.86 years, wrong**
Right half-life	ln(2) ÷ 0.35	**1.98 years**
Linear decay, 2 years	5 − 1 × (1 − 0.35 × 2)	**4.70%, wrong**
Right: exponential decay	5 − 1 × e^{−0.70} = 5 − 0.497	**4.503%**
Drift scaled by √dt	0.35 × 1 × √0.25	**17.5 bp, wrong**

IN THE LESSON

- The first mistake is taking the half-life as one over k: one over point three five is two point eight six years, which is wrong.
- The right half-life is the natural log of two over point three five, one point nine eight years.
- Second, linear decay.
- If the gap of one percentage point shrank by k times T, after two years it would be one times one minus point seven, leaving an expected rate of four point seven zero per cent, which is wrong.
- Exponential decay leaves e to the minus point seven zero, point four nine seven of the gap, giving four point five zero three.
- Third, scaling the drift by the square root of d t: point three five times one times root point two five gives seventeen point five basis points for a quarter, double the right eight point seven five.

29 Vasicek dispersion has a ceiling

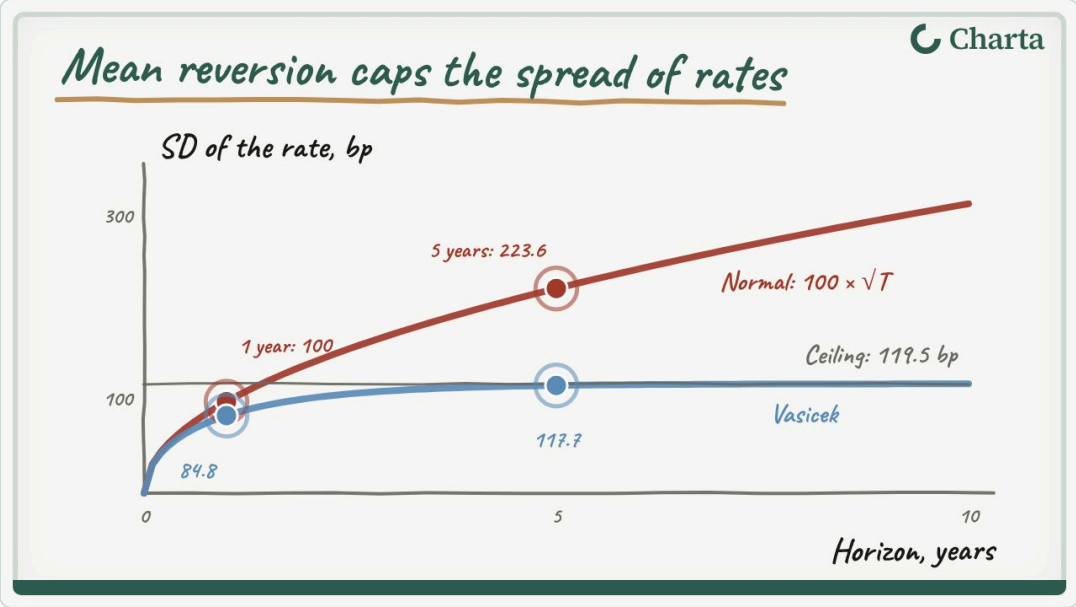
$$SD_T = \sigma \times \sqrt{(1 - e^{-2kT})} \div (2k)$$
$$\text{Ceiling as T grows} = \sigma \div \sqrt{(2k)}$$

σ: annual volatility in bp	Annual volatility, in bp
k: speed of mean reversion	Speed of mean reversion
T: the horizon, in years	Horizon, in years

IN THE LESSON

- How well does Vasicek perform?
- Start with its strength: mean reversion bounds the dispersion of rates far into the future.
- The standard deviation of the rate at horizon T is sigma times the square root of one minus e to the minus two k T, all over two k.
- Here sigma is the annual volatility and k the speed of mean reversion.
- As T grows, e to the minus two k T falls to zero, so the standard deviation tends to a ceiling of sigma divided by the square root of two k.
- In the normal model without mean reversion, by contrast, it grows with the square root of T forever.

30 Mean reversion caps the spread of rates



IN THE LESSON

- For Halden, the red curve is the normal model with no mean reversion: one hundred times the square root of T.
- The blue curve is Vasicek.
- At one year the normal model gives one hundred basis points, and Vasicek eighty four point eight.
- At five years the normal model gives two hundred and twenty three point six, but Vasicek only one hundred and seventeen point seven.
- The grey line is the ceiling, one hundred divided by the square root of point seven, one hundred and nineteen point five basis points.
- Vasicek's curve flattens towards it, while the normal model keeps climbing.

31 Where Vasicek does well

- Mean reversion bounds long-horizon dispersion
- Halden: SD ceiling 100 ÷ √0.70 = 119.5 bp
- Long-maturity yields less volatile than short ones
- A simple tree and formulas; realistic rate dynamics

IN THE LESSON

- So the first strength is bounded long horizon dispersion.
- For Halden, the standard deviation can never exceed one hundred over root point seven, one hundred and nineteen point five basis points, however far ahead we look.
- Second, also from mean reversion, long maturity yields are less volatile than short ones, as in the data, because a shock to the short rate is expected to fade before long bonds mature.
- And the model gives a simple tree and clear formulas.
- In short, Vasicek offers realistic rate dynamics.

32 Where Vasicek falls short

- Constant parameters: a limited family of curve shapes
- Does not fit market bonds, so not arbitrage-free
- Rates can go negative: long run Z = −5 ÷ 1.195 = −4.18
- Constant volatility; one factor moves all rates
- Fix the fit: time-varying θ (extended Vasicek, Hull-White)

IN THE LESSON

- With constant parameters, the model can generate only a limited family of curve shapes.
- So it does not fit the observed term structure, and it cannot price relative to market bonds: it is not arbitrage free.
- Normal rates can still go negative.
- For Halden, the long run distribution has mean five per cent and standard deviation one point one nine five per cent, so Z is minus four point one eight, a probability of about point zero zero one four per cent: tiny, but not zero.
- Volatility is constant, and one factor moves all rates together.
- The repair is to let theta depend on time, as in the extended Vasicek or Hull-White model.
- The exam point: realistic dynamics, but not arbitrage free.

33 Try three quick questions

1. Zero drift, σ = 120 bp a year. SD of the change over a quarter? Over 9 years?
   Answer: 120 × √0.25 = 60 bp; 120 × √9 = 360 bp.
2. Vasicek: r0 = 5%, θ = 3%, k = 0.20. Expected rate in 3 years? Half-life?
   Answer: 3 + 2 × e^{−0.6} ≈ 4.10%; ln(2) ÷ 0.2 ≈ 3.47 years.
3. Zero drift, r = 3%, σ = 120 bp, T = 4. Chance of a negative rate?
   Answer: Z = −3 ÷ (1.2 × 2) = −1.25; N(Z) ≈ 10.6%.

IN THE LESSON

- Pause the video, think, then listen to the answers.
- Question one.
- A zero drift normal model has volatility of one hundred and twenty basis points a year.
- Standard deviation over one quarter, and over nine years?
- One hundred and twenty times the square root of point two five is sixty basis points, and one hundred and twenty times the square root of nine is three hundred and sixty.
- Question two.
- In a Vasicek model, r nought is five per cent, theta three per cent and k point two zero.
- Expected rate in three years, and half-life?
- Three plus two times e to the minus point six, about four point one zero per cent; and the natural log of two over point two, about three point four seven years.
- Question three.
- Zero drift, a three per cent rate, volatility one hundred and twenty basis points, four years.
- Chance of a negative rate?
- Z is minus three over one point two times two, minus one point two five: about ten point six per cent.