CHARTA LEARNING · LESSON NOTES

The Art of Term Structure Models: Volatility and Distribution

FRM Part II

1 A story: how jumpy is the short rate?

- Kestrel's short rate today: 4.00%
- Views: as jumpy at 1% as at 8%, or moves shrink near 0
- Drift: where the rate is expected to go
- Volatility rule: how far it can wander, and when
- 1 basis point (bp) = 0.01%; volatility in bp a year

IN THE LESSON

- Kestrel's short rate today is four per cent.
- Two analysts on the desk disagree.
- One says the rate is just as jumpy at one per cent as at eight per cent.
- The other says moves shrink as rates fall towards zero.
- Who is right depends on the model's two parts.
- The drift says where the rate is expected to go.
- The volatility rule says how far it can wander, and when.
- That rule shapes the whole distribution of future rates, including whether they can turn negative.
- One unit before we start.
- A basis point is one hundredth of one per cent, and the volatility of the short rate is usually quoted in basis points a year.

2 Volatility that depends on time

dr = λ(t) dt + σ(t) dw
Example schedule: σ(t) = σ × e^{−αt}

Drift, λ(t)	Set to fit today's rate curve
Volatility, σ(t)	Known today for every date
Shock, dw	Normal shock, variance dt

IN THE LESSON

- Here is the first model.
- Write d r for the tiny change in the short rate over a tiny interval of time, d t.
- The change equals lambda of t times d t, plus sigma of t times d w.
- The second line shows a common schedule for the volatility: sigma of t equals sigma times e to the minus alpha t, so volatility decays exponentially at the rate alpha.
- Now the three parts.
- Lambda of t, the drift, is a function of calendar time, usually chosen so that the model reproduces today's term structure of rates.
- Sigma of t, the volatility, is also a function of calendar time only, so it is known today for every future date.
- And d w is a normal shock with a variance of d t.

3 What time dependence does, and does not, mean

- σ(t) depends on calendar time only
- Not on the rate level, not on the path taken
- Rate changes are normal; bp volatility can rise or fall
- No shrinking at low rates: negative rates possible
- Trap: it is not stochastic or rate-dependent volatility

IN THE LESSON

- Be precise about what time dependence means.
- Sigma of t depends on calendar time only.
- It does not depend on the level of the rate, and it does not depend on the path the rate has taken to get there.
- Because each shock is normal and its size is fixed in advance, rate changes are normally distributed, with a basis point volatility that can rise or fall as time passes.
- And because volatility does not shrink when rates are low, negative rates remain possible.
- The classic trap is to call this stochastic volatility, or rate dependent volatility.
- It is neither.
- For the exam, time dependence means a deterministic schedule of basis point volatility, nothing more.

4 Example 1: Kestrel's volatility schedule

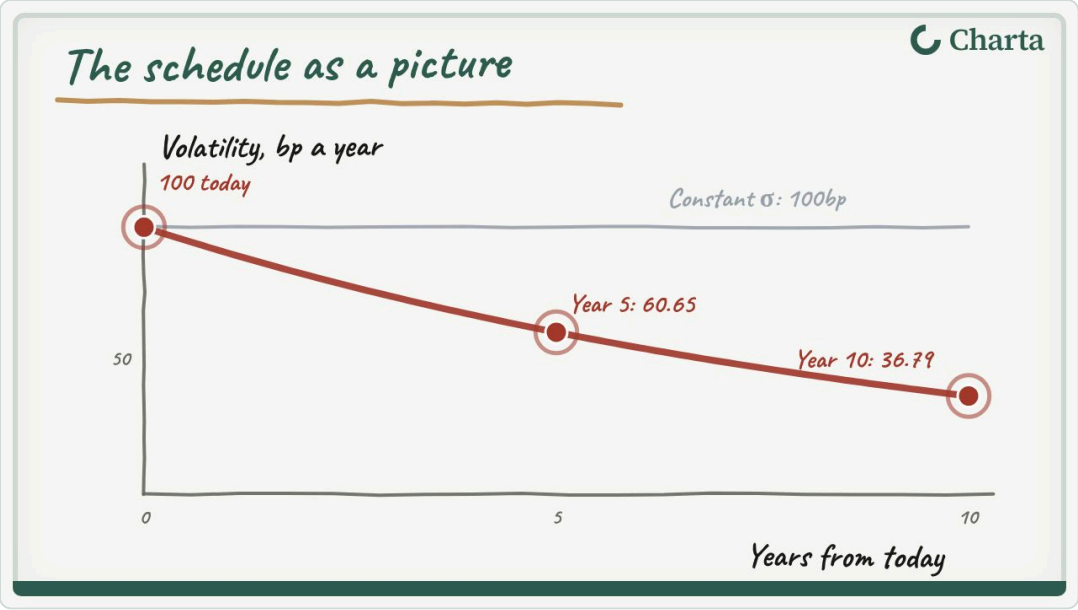
Kestrel: σ = 100bp a year, decay α = 0.10. Volatility at each date

Today, t = 0	100 × e^{0}	**100.00bp**
In 1 year	100 × e^{−0.1}	**90.48bp**
In 5 years	100 × e^{−0.5}	**60.65bp**
In 10 years	100 × e^{−1}	**36.79bp**
At a rate of 1% or 8%	the level plays no part	**unchanged**

IN THE LESSON

- Kestrel sets sigma at one hundred basis points a year and the decay rate alpha at point one zero.
- Let us find the volatility at each date.
- Today, at time zero, it is one hundred times e to the zero, so one hundred basis points.
- In one year it is one hundred times e to the minus point one, about ninety point four eight.
- In five years it is one hundred times e to the minus point five, about sixty point six five.
- In ten years it is one hundred times e to the minus one, about thirty six point seven nine.
- And whether the rate is then at one per cent or eight per cent, these figures are unchanged, because the level of the rate plays no part.

5 The schedule as a picture



IN THE LESSON

- Here is the same schedule drawn across ten years.
- The red curve is Kestrel's volatility, falling from one hundred basis points today, through sixty point six five at year five, to thirty six point seven nine at year ten.
- The grey line is a constant volatility of one hundred basis points, for comparison.
- The whole red curve is fixed today.
- The desk knows now that a rate move four years from now will be smaller than one tomorrow, whatever the market does in between.
- That is the strength and the weakness of the model in a single picture.
- A decaying schedule like this is one way a desk can make long dated rates less volatile than short dated ones.

6 One step of the model

$$\Delta r = \lambda(t) \times \Delta t + \sigma(t) \times \sqrt{\Delta t} \times \varepsilon$$
$$\text{Mean} = \lambda(t) \times \Delta t \quad \text{SD} = \sigma(t) \times \sqrt{\Delta t}$$

Expected change	Scales with Δt
Standard deviation	Scales with $\sqrt{\Delta t}$; time, not level
Shock, ε	Standard normal draw

IN THE LESSON

- On a spreadsheet, take a short but finite step, delta t.
- The change in the short rate, delta r, equals lambda of t times delta t, plus sigma of t times the square root of delta t, times epsilon.
- On the second line, the mean of the change is lambda of t times delta t, and its standard deviation is sigma of t times the square root of delta t.
- Now the three parts.
- The expected change scales with delta t itself.
- The standard deviation scales with the square root of delta t, and depends on calendar time but not on the rate level.
- And epsilon is a standard normal draw, with a mean of zero and a standard deviation of one.

7 Example 2: one month from today

Rate 4.00%, drift λ = 20bp a year, σ(0) = 100bp. One month, ε = +1.2

Time step Δt	1 month = 1/12 year	**√Δt = 0.2887**
Expected change	20 × 1/12	**1.67bp**
Standard deviation	100 × 0.2887	**28.87bp**
Change with ε = +1.2	1.67 + 28.87 × 1.2	**36.31bp**
New short rate	4.00% + 0.3631%	**4.36%**

IN THE LESSON

- Kestrel's rate is four per cent, the drift is twenty basis points a year and today's volatility is one hundred basis points.
- Take one month, with a shock epsilon of plus one point two.
- The time step is one twelfth of a year, so the square root of delta t is point two eight eight seven.
- The expected change is twenty times one twelfth, one point six seven basis points.
- The standard deviation is one hundred times point two eight eight seven, twenty eight point eight seven basis points.
- The change is one point six seven plus twenty eight point eight seven times one point two, which is thirty six point three one basis points.
- The new short rate is four per cent plus point three six three one per cent, about four point three six per cent.

8 Example 3: five years on, and a negative rate

Five years on: σ(5) = 60.65bp, drift still 20bp a year, same ε = +1.2

Standard deviation	60.65 × 0.2887	**17.51bp**
Change with ε = +1.2	1.67 + 17.51 × 1.2	**22.68bp**
Same SD at any rate level	1% or 8% alike	**17.51bp**
Today, rate 0.25%, ε = −2	1.67 − 28.87 × 2	**−56.07bp**
New short rate	0.25% − 0.5607%	**−0.31%**

IN THE LESSON

- Five years ahead the schedule gives sixty point six five basis points; keep the drift and the shock of plus one point two.
- The monthly standard deviation is sixty point six five times point two eight eight seven, seventeen point five one basis points.
- The change is one point six seven plus seventeen point five one times one point two, twenty two point six eight basis points.
- That standard deviation is the same whether the rate is then one per cent or eight per cent.
- Now go back to today, with the rate at a quarter of one per cent and a shock of minus two.
- The change is one point six seven minus twenty eight point eight seven times two, minus fifty six point zero seven.
- The new rate is minus point three one per cent.

9 Over longer horizons, variances add

Variance = ∫ σ(t)^{2} dt, from 0 to T
Constant σ: σ^{2} × T, so SD = σ × √T
Decaying σ: σ^{2} × (1 − e^{−2αT}) ÷ (2α)

Cumulative variance	Sum of each slice's σ² × Δt
Horizon SD	Square root of the variance
Volatilities	Never add or average them

IN THE LESSON

- Over a longer horizon, the variances of many small independent steps add, but their volatilities do not.
- The first line says the variance of the cumulative change, from now to time T, is the integral of sigma of t squared, d t, from zero to T: each slice adds its own sigma squared times its length.
- On the second line, with a constant sigma the variance is sigma squared times T, so the standard deviation is sigma times the square root of T.
- With Kestrel's decaying schedule, the third line, it works out to sigma squared, times one minus e to the minus two alpha T, divided by two alpha.
- Now the three boxes.
- The cumulative variance is the sum of the slices.
- The horizon standard deviation is its square root.
- And volatilities are never added or averaged.

10 Example 4: two years with two volatilities

Volatility 100bp a year in year 1, then 80bp a year in year 2. Two-year SD?

Year 1 variance	100 × 100 × 1	**10,000**
Year 2 variance	80 × 80 × 1	**6,400**
Two-year variance	10,000 + 6,400	**16,400**
Two-year SD	√16,400	**128.1bp**
Wrong: add the vols	100 + 80	**180bp (+41%)**

IN THE LESSON

- Here is a simple case.
- Volatility is one hundred basis points a year in year one, then eighty in year two.
- What is the standard deviation of the two year change?
- Year one's variance is one hundred times one hundred, times one year, ten thousand, in basis points squared.
- Year two's is eighty times eighty, six thousand four hundred.
- The two year variance is ten thousand plus six thousand four hundred, sixteen thousand four hundred.
- The two year standard deviation is its square root, about one hundred and twenty eight point one basis points.
- The wrong answer adds the volatilities: one hundred plus eighty, one hundred and eighty basis points.
- That overstates the risk by about forty one per cent.

11 Example 5: ten years of a decaying schedule

Kestrel's schedule, σ = 100bp, α = 0.10. SD of the 10-year cumulative change?

2αT	2 × 0.10 × 10	**2.0**
1 − e^{−2}	1 − 0.1353	**0.8647**
Variance	10,000 × 0.8647 ÷ 0.2	**43,235**
10-year SD	√43,235	**207.9bp**
If σ stayed at 100bp	100 × √10	**316.2bp**

IN THE LESSON

- Now Kestrel's own schedule, with sigma of one hundred basis points and alpha of point one zero, over ten years.
- First, two alpha T is two times point one zero times ten, which is two.
- Next, one minus e to the minus two is one minus point one three five three, which is point eight six four seven.
- The variance is ten thousand times point eight six four seven, divided by point two, about forty three thousand two hundred and thirty five.
- The ten year standard deviation is its square root, about two hundred and seven point nine basis points.
- If volatility had stayed at one hundred basis points, it would be one hundred times the square root of ten, three hundred and sixteen point two.

12 Mistakes to avoid with time-dependent volatility

- Scale the shock by √Δt, never by Δt
- 1 month: 100 × 0.2887 = 28.87bp, not 100 ÷ 12 = 8.33
- Over a horizon add variances, never volatilities
- Not stochastic, not rate-dependent: a fixed schedule
- Exam: annual bp vol × √Δt = vol for the period

IN THE LESSON

- Three mistakes cost marks here.
- First, scale the shock by the square root of delta t, never by delta t.
- For one month, one hundred basis points times point two eight eight seven is twenty eight point eight seven basis points, not one hundred divided by twelve, which is eight point three three.
- Second, over a horizon add variances, never volatilities, and never average them either.
- Third, do not call this stochastic or rate dependent volatility.
- It is a fixed schedule, set today.
- And the exam point to carry away: to turn an annual basis point volatility into the volatility for a period, multiply by the square root of the period's length in years.

13 Strength: fitting today's volatility curve

- Drift λ(t) fits today's term structure of rates
- σ(t) fits the term structure of volatility too
- For example cap or swaption vols across expiries
- Implied vol to expiry T: root-mean-square of σ(t), 0 to T
- So cumulative variance to T = T × (implied vol)^{2}

IN THE LESSON

- How well do time dependent volatility models work?
- Start with the strength, which is flexibility.
- The drift, lambda of t, is already chosen to fit today's term structure of rates.
- Choosing sigma of t as well lets a normal model fit the term structure of volatility too.
- The market quotes such volatilities, for example cap or swaption volatilities across expiries.
- Think of a normal implied volatility to an expiry T as a root mean square average of sigma of t over the whole period from now to T.
- Then the cumulative variance to T is T times that implied volatility squared, and the desk picks sigma of t, slice by slice, so that the model matches each expiry.

14 Example 6: fitting σ(t) to market volatilities

Market normal vols: 1-year expiry 100bp, 2-year expiry 95bp. Fit σ by year

Variance to year 1	1 × 100^{2}	**10,000**
Variance to year 2	2 × 95^{2}	**18,050**
Year 2 slice	18,050 − 10,000	**8,050**
σ in year 1	√10,000	**100.00bp**
σ in year 2	√8,050	**89.72bp**

IN THE LESSON

- Kestrel sees normal implied volatilities of one hundred basis points for a one year expiry and ninety five for a two year expiry.
- The variance to year one is one times one hundred squared, ten thousand.
- The variance to year two is two times ninety five squared, eighteen thousand and fifty.
- The year two slice is eighteen thousand and fifty minus ten thousand, eight thousand and fifty.
- So sigma in year one is the square root of ten thousand, one hundred basis points.
- And sigma in year two is the square root of eight thousand and fifty, about eighty nine point seven two basis points.
- The quoted ninety five is a root mean square average; year two alone is lower.

15 Weaknesses: describing how volatility behaves

- Tomorrow's vol curve = the rest of today's curve
- The market rarely obliges: recalibrate often
- No dependence on the rate level; never random
- One vol per expiry: no strike skew
- Rates still normal: negative rates possible

IN THE LESSON

- Now the weaknesses, which are structural.
- Because volatility is a deterministic function of calendar time, the model says that tomorrow's volatility curve is simply the remaining part of today's curve.
- The market rarely delivers that, so the parameters must be recalibrated often, which means the model contradicts itself from one day to the next.
- Volatility cannot depend on the level of rates, and it can never be random.
- There is one volatility per expiry, so the model cannot capture a strike skew, where options with different strikes trade at different volatilities.
- And rates remain normal, so negative rates are still possible.

16 Example 7: a curve no σ(t) can fit

Suppose the 2-year expiry vol were 65bp, with the 1-year still 100bp

Variance to year 1	1 × 100^{2}	**10,000**
Variance to year 2	2 × 65^{2}	**8,450**
Year 2 slice	8,450 − 10,000	**−1,550**
σ in year 2	√(−1,550)	**no real value**
Fix	a model with mean reversion	**needed**

Good at fitting today's vol curve; poor at predicting vol

IN THE LESSON

- Fitted values can also be impossible.
- Suppose the two year volatility were sixty five basis points, with the one year still one hundred.
- The variance to year one is ten thousand, as before.
- The variance to year two is two times sixty five squared, eight thousand four hundred and fifty.
- The year two slice is eight thousand four hundred and fifty minus ten thousand, which is minus one thousand five hundred and fifty.
- A variance cannot be negative, so sigma in year two would be the square root of a negative number, and has no real value.
- When implied cumulative variance falls with the horizon, no real sigma of t exists.
- A fix needs a model with mean reversion.
- So time dependent volatility is good at fitting today's volatility curve, but poor at predicting how volatility will behave.

17 Near zero: two other volatility rules

- Normal models: bp volatility constant (or a schedule)
- CIR: bp volatility proportional to √r
- Lognormal: bp volatility proportional to r
- Last two: moves shrink as the rate falls
- So the rate cannot go negative

IN THE LESSON

- Back to the story.
- Kestrel's rate has fallen to a quarter of one per cent, and the time dependent model has just produced a rate of minus point three one per cent.
- The analyst who said moves shrink near zero wants a different volatility rule.
- There are three to know.
- In normal models, basis point volatility is constant, or a fixed schedule.
- In the C I R model, it is proportional to the square root of the rate.
- In lognormal models, it is proportional to the rate itself.
- In the last two, moves shrink as the rate falls, so the rate cannot go negative.

18 The CIR model

dr = k(θ − r) dt + σ√r dw
bp volatility = σ × √r

Speed, k	Speed of mean reversion
Long-run level, θ	Level the rate is pulled to
Volatility, σ√r	High at high r, shrinks near 0

IN THE LESSON

- Here is the model of Cox, Ingersoll and Ross.
- The change in the rate, d r, equals k times theta minus r, times d t, plus sigma times the square root of r, times d w.
- The second line gives its basis point volatility: sigma times the square root of r.
- The speed k says how fast the rate is pulled back.
- Theta is the long run level it is pulled towards, so this is mean reversion.
- And the volatility term, sigma times root r, is high when rates are high and shrinks as the rate approaches zero.

19 How the CIR rate behaves

- Above θ the drift is negative; below θ positive
- Volatility rises with √r: high when rates are high
- Near zero the shock dies away: no negative rates
- If 2kθ ≥ σ^{2}, the rate cannot even reach zero

IN THE LESSON

- What does that do to the path of the rate?
- When the rate is above theta, the drift is negative and pulls it down.
- When it is below theta, the drift is positive and pulls it up.
- Volatility rises with the square root of the rate, so it is high when rates are high.
- Near zero the shock dies away, while the drift still pushes upwards, so the rate cannot go negative.
- And there is a stronger result.
- If two times k times theta is at least sigma squared, the rate cannot even reach zero.
- That condition compares the strength of the upward pull with the size of the shocks.

20 Example 8: Kestrel's CIR parameters

Kestrel's CIR: k = 0.25, θ = 5%, σ = 0.05. Rate today 4%

Drift at r = 4%	0.25 × (0.05 − 0.04)	**+25bp a year**
Drift at r = 6%	0.25 × (0.05 − 0.06)	**−25bp a year**
2kθ	2 × 0.25 × 0.05	**0.025**
σ^{2}	0.05 × 0.05	**0.0025**
2kθ ≥ σ^{2}?	0.025 ≥ 0.0025	**never hits 0**

IN THE LESSON

- Kestrel's C I R model has k of point two five, theta of five per cent and sigma of point zero five, with the rate at four per cent.
- The drift at four per cent is point two five times point zero five minus point zero four, which is plus twenty five basis points a year: the rate is pulled up.
- The drift at six per cent is point two five times point zero five minus point zero six, minus twenty five basis points a year: pulled down.
- For the zero test, two k theta is two times point two five times point zero five, point zero two five.
- Sigma squared is point zero five times point zero five, point zero zero two five.
- Since point zero two five is at least point zero zero two five, the rate never hits zero.

21 The lognormal model

dr = a r dt + σ r dw
bp volatility = σ × r; ln r is normal

Drift, a × r	Grows with the rate
Volatility, σ × r	Proportional to the rate
ln r is normal	So r is always positive

IN THE LESSON

- The second rule gives the lognormal model.
- The change d r equals ay times r times d t, plus sigma times r times d w.
- On the second line, the basis point volatility is sigma times r, and the logarithm of the rate, l n r, is normally distributed.
- Now the parts.
- The drift, ay times r, grows with the rate.
- The volatility, sigma times r, is proportional to the rate itself.
- And because the log of r is normal, r is lognormal and so it is always positive.

22 How the lognormal rate behaves

- ln r is normal, so the rate is lognormal
- Always positive: the rate never crosses zero
- Volatility proportional to r: double r, double it
- No mean reversion unless a term is added
- σ here is a % of the rate, not bp

IN THE LESSON

- How does the lognormal rate behave?
- The log of the rate is normal, so the rate itself is lognormal.
- It is always positive and can never cross zero.
- Its basis point volatility is proportional to the rate, so if the rate doubles, the basis point volatility doubles.
- Without an added term, the lognormal model has no mean reversion, so nothing pulls the rate back towards a long run level.
- And be careful with units: sigma here is a percentage of the rate, such as twenty five per cent, not a number of basis points.

23 Match each model to its volatility rule

The usual mistake is mixing up the three volatility rules

Model	bp volatility	Negative rates?
Normal, incl. time-dependent	σ or σ(t)	Possible
CIR	σ × √r	No
Lognormal	σ × r	No

IN THE LESSON

- The usual mistake is to mix up the three volatility rules, so match each model to its own.
- Read the table row by row.
- In normal models, including the time dependent one, the basis point volatility is sigma, either constant or a fixed schedule sigma of t, and negative rates are possible.
- In C I R, the basis point volatility is sigma times the square root of r, and negative rates cannot happen.
- In the lognormal model, the basis point volatility is sigma times r, and negative rates cannot happen either.
- Constant, square root, proportional: that is the sequence to remember.

24 One step of CIR and lognormal

Δr CIR = k(θ − r) × Δt + σ√r × √Δt × ε
Δr Log = a × r × Δt + σ × r × √Δt × ε

CIR bp volatility	σ × √r, per year
Lognormal bp volatility	σ × r, per year
Then × √Δt	Gives the period's SD

IN THE LESSON

- Now the discrete steps, both using epsilon, a standard normal draw.
- Over a short interval delta t, the C I R change equals k times theta minus r, times delta t, plus sigma root r times the square root of delta t, times epsilon.
- The lognormal change equals ay times r times delta t, plus sigma times r times the square root of delta t, times epsilon.
- Now the boxes.
- The annual basis point volatility is the coefficient on the shock before the square root of delta t is applied.
- For C I R it is sigma times root r.
- For the lognormal it is sigma times r.
- Then multiply by the square root of delta t for the period's standard deviation.

25 Example 9: one month under CIR

CIR: $k = 0.25$, $\theta = 5\%$, $\sigma = 0.05$, $r = 4\%$. One month, $\varepsilon = -1.5$

bp vol a year	$0.05 \times \sqrt{0.04} = 0.05 \times 0.2$	**100bp**
SD for the month	100×0.2887	**28.87bp**
Drift for the month	$25 \div 12$	**2.08bp**
Change with ε = −1.5	$2.08 - 28.87 \times 1.5$	**−41.22bp**
New short rate	$4.00\% - 0.4122\%$	**3.59%**

IN THE LESSON

- Take Kestrel's C I R model at four per cent, with a one month shock of minus one point five.
- The basis point volatility a year is point zero five times the square root of point zero four, which is point zero five times point two, point zero one, or one hundred basis points.
- Next the standard deviation for the month: one hundred times point two eight eight seven, twenty eight point eight seven basis points.
- The drift for the month is twenty five divided by twelve, two point zero eight basis points.
- The change is two point zero eight minus twenty eight point eight seven times one point five, or minus forty one point two two.
- The new rate is about three point five nine per cent, four per cent minus point four one two two.

26 Example 10: one month under the lognormal model

Lognormal: $a = 0.02$, $\sigma = 0.25$ (25% of r), $r = 4\%$. One month, $\varepsilon = +1.0$

bp vol a year	$0.25 \times 0.04 = 0.01$	**100bp**
SD for the month	100×0.2887	**28.87bp**
Drift for the month	$0.02 \times 0.04 = 8$bp; $\div 12$	**0.67bp**
Change with ε = +1.0	$0.67 + 28.87$	**29.54bp**
New short rate	$4.00\% + 0.2954\%$	**4.30%**

IN THE LESSON

- Now a lognormal model, with ay of point zero two, sigma of point two five, meaning twenty five per cent of the rate, the rate at four per cent, and a one month shock of plus one.
- The basis point volatility a year is point two five times point zero four, point zero one, or one hundred basis points.
- The standard deviation for the month is one hundred times point two eight eight seven, twenty eight point eight seven.
- The drift is point zero two times point zero four, eight basis points a year, so point six seven for the month.
- The change is point six seven plus twenty eight point eight seven, twenty nine point five four basis points.
- The new rate is about four point three zero per cent, four per cent plus point two nine five four.

27 Example 11: when the rate level changes

All three set to 100bp a year at $r = 4\%$. Now the rate moves

Normal, any rate	σ fixed	**100bp**
CIR at 8%	$0.05 \times \sqrt{0.08} = 0.01414$	**141.4bp**
Lognormal at 8%	0.25×0.08	**200bp**
CIR at 2%	$0.05 \times \sqrt{0.02}$	**70.7bp**
Lognormal at 2%	0.25×0.02	**50bp**

IN THE LESSON

- All three of Kestrel's models give one hundred basis points a year at four per cent.
- Now let the rate move.
- The normal model's sigma is fixed, so it stays at one hundred.
- Under C I R at eight per cent, it is point zero five times the square root of point zero eight, which is about point zero one four one four, or one hundred times the square root of two basis points, about one hundred and forty one point four.
- The lognormal gives point two five times point zero eight, two hundred.
- At two per cent, C I R gives point zero five times the square root of point zero two, about seventy point seven, and the lognormal gives point two five times point zero two, fifty basis points.
- So doubling the rate multiplies C I R volatility by root two, but lognormal volatility by two.

28 Three volatility rules in one picture



IN THE LESSON

- Here are the three rules side by side, against a rate from zero to ten per cent.
- The blue line is the normal model, flat at one hundred basis points whatever the rate.
- The gold curve is C I R.
- It starts at zero, rises steeply at first and then flattens, reaching point zero five times the square root of point one zero, about one hundred and fifty eight basis points, at ten per cent.
- The red line is the lognormal, a straight line from zero, reaching point two five times point one zero, two hundred and fifty basis points, at ten per cent.
- All three meet at four per cent and one hundred basis points, where Kestrel calibrated them.

29 Mistakes to avoid with CIR and lognormal

- r in %: 0.05 × √4 = 0.10 = 1,000bp, 10 times too big
- Lognormal form for CIR: 0.05 × 0.04 = 20bp, wrong
- Forgetting √Δt: a month is 28.87bp, not 100bp
- Exam: volatility coefficient first, then × √Δt

IN THE LESSON

- Three mistakes come up again and again.
- First, entering the rate in per cent instead of as a decimal.
- Point zero five times the square root of four is point one zero, which would be one thousand basis points, ten times too big.
- Use point zero four.
- Second, using the lognormal form for C I R.
- Point zero five times point zero four gives just twenty basis points, which is wrong.
- Third, forgetting the square root of delta t for a monthly change.
- A month's standard deviation here is twenty eight point eight seven basis points, not one hundred.
- The exam point is a fixed order of work: compute the volatility coefficient first, then scale it by the square root of delta t.

30 Lognormal models with a drift

Let Y = ln r
Deterministic drift: dY = a(t) dt + σ dw
Mean reversion: dY = k(ln θ − Y) dt + σ dw

Drift, a(t)	Fits today's rate curve
Pull to ln θ	Speed k: Black-Karasinski
r = e^{Y}	So r is always positive

IN THE LESSON

- A lognormal model can be given a drift in two ways.
- Let Y stand for l n r, the log of the short rate.
- With a deterministic drift, d Y equals ay of t times d t, plus sigma times d w.
- With mean reversion, d Y equals k times l n theta minus Y, times d t, plus sigma times d w.
- Now the boxes.
- The drift ay of t is chosen to fit today's term structure of rates.
- The pull towards l n theta at speed k is the mean reverting version, as in the Black-Karasinski model.
- And in both, r equals e to the Y, so the rate is always positive.

31 Deterministic drift: dispersion keeps growing

- a(t) chosen to fit today's term structure
- No mean reversion: SD of Y grows with √t
- σ = 0.20: SD of ln r is 0.20 at 1 year
- 0.40 at 4 years, 0.60 at 9 years, without limit
- Rate always positive, but its range keeps widening

IN THE LESSON

- Take the deterministic drift first.
- The function ay of t is chosen to fit today's term structure of rates.
- There is no mean reversion, so the standard deviation of Y, the log rate, grows with the square root of time.
- With sigma of point two zero, the standard deviation of l n r is point two zero at one year.
- It is point two times the square root of four, point four zero, at four years, and point two times the square root of nine, point six zero, at nine years, without limit.
- The rate stays positive, but the range of likely rates keeps widening as the horizon grows.

32 Mean reversion in ln r: median and mean

- Y is pulled toward ln θ at speed k
- Long-run variance of Y: Var = σ^{2} ÷ (2k)
- Median of the long-run rate = θ
- Mean = θ × exp(Var ÷ 2), above the median
- Traps: reverting r instead of ln r; θ as the mean

IN THE LESSON

- Now mean reversion.
- In the Black-Karasinski style model, Y is pulled toward l n theta at speed k.
- The long run variance of Y settles at sigma squared divided by two k.
- Since Y is normal, r is lognormal and always positive.
- The median of the long run rate is theta.
- But its mean is larger, because a lognormal has a long right tail.
- The mean equals theta times the exponential of the variance divided by two, so it sits above the median.
- Two traps follow.
- One is applying mean reversion to r instead of to l n r.
- The other is treating theta as the mean, when it is the median.

33 Example 12: median and mean of the long-run rate

Kestrel's Black-Karasinski: θ = 4%, k = 0.5, σ = 0.20. Long-run rate?

Variance of ln r	0.20^{2} ÷ (2 × 0.5)	**0.04**
SD of ln r	√0.04	**0.20**
Median rate	θ	**4.00%**
Mean rate	4% × e^{0.02}	**4.08%**
If σ = 0.40	Var 0.16; 4% × e^{0.08}	**mean 4.33%**

IN THE LESSON

- Kestrel's Black-Karasinski model has theta of four per cent, k of point five and sigma of point two zero.
- What does the long run rate look like?
- The variance of l n r is point two zero squared divided by two times point five, which is point zero four.
- The standard deviation of l n r is the square root of point zero four, point two zero.
- The median rate is theta, four per cent.
- The mean rate is four per cent times e to the point zero two, about four point zero eight per cent.
- If sigma were point four zero instead, the variance would be point one six, and the mean four per cent times e to the point zero eight, about four point three three per cent.

34 Try three quick questions

1. Time-dependent model, σ = 90bp a year. SD of a one-quarter change?
   Answer: 90 × √0.25 = 45bp.
2. CIR: σ = 0.04, short rate 9%. Annual bp volatility?
   Answer: 0.04 × √0.09 = 0.04 × 0.3 = 0.012 = 120bp.
3. Black-Karasinski: θ = 5%, k = 0.4, σ = 0.4. Median and mean?
   Answer: Var = 0.16 ÷ 0.8 = 0.2. Median 5%; mean 5% × e^{0.1} = 5.53%.

IN THE LESSON

- Now try three quick questions.
- Pause, think, then listen.
- Question one.
- In a time dependent volatility model, the annual basis point volatility is ninety basis points.
- What is the standard deviation of the change over one quarter?
- Multiply by the square root of a quarter, point five: forty five basis points.
- Question two.
- In a C I R model, sigma is point zero four and the rate is nine per cent.
- What is the annual basis point volatility?
- Point zero four times point three, the root of point zero nine, is point zero one two, or one hundred and twenty basis points.
- Question three.
- A Black-Karasinski model has theta of five per cent, k of point four and sigma of point four.
- Find the median and mean.
- The variance is point one six over point eight, point two.
- The median is five per cent, and the mean is five per cent times e to the point one, about five point five three per cent.