CHARTA LEARNING · LESSON NOTES

Regression Hedging and PCA

FRM Part II

1 A story: the hedge that made a bet

- Kestrel Bank's rates desk: long $100m of a 10-year bond
- Hedge: sell a 5-year note, sized so the DV01s match
- Quiet weeks: the hedge seems to work
- Then the yields move unequally and P&L swings
- Lesson: matching DV01s is not enough

IN THE LESSON

- Kestrel Bank's rates desk buys one hundred million dollars of face value of a ten-year bond.
- If yields rise, the bond loses value, so the desk wants a hedge.
- It sells a five-year note, sized so that both positions have the same D V zero one, which means the same dollar change in value for a one basis point move in yield.
- In quiet weeks the hedge seems to work, and the profit and loss, or P and L, barely moves.
- Then the two yields start moving by different amounts, and the P and L swings by hundreds of thousands of dollars.
- The lesson: matching D V zero ones is not enough.

2 The DV01-neutral hedge

DV01 = $ change in value for a 1 bp move in yield
Hedge = Target × (DV01_T ÷ DV01_H)

Target face	Face of the bond being hedged
DV01 of target	Per 100 of face
DV01 of hedge	Per 100 of face
Hedge face	Traded on the opposite side

IN THE LESSON

- D V zero one is the change in a position's value, in dollars, for a one basis point change in its yield.
- A D V zero one neutral hedge sizes the offsetting position so that its D V zero one equals the target's.
- The face of the hedge equals the face of the target, times the target's D V zero one divided by the hedge's D V zero one.
- Look at the four boxes.
- The target face is the bond being hedged.
- Its D V zero one, and then the hedge's D V zero one, are both quoted per one hundred of face.
- The hedge face is what you trade, on the opposite side.
- And the catch, in the note at the bottom: this is exact only if both yields move by the same number of basis points.

3 Example 1: sizing the DV01-neutral hedge

Long $100m of the 10-year, DV01 0.08 per 100; 5-year note DV01 0.05 per 100

Target DV01, per bp	100m × 0.08 ÷ 100	**$80,000**
DV01 ratio	0.08 ÷ 0.05	**1.6**
DV01-neutral hedge face	100m × 1.6	**$160m**
Hedge DV01, per bp	160m × 0.05 ÷ 100	**$80,000**
Side	long the bond, so sell	**Sell 5-year**

IN THE LESSON

- The ten-year bond has a D V zero one of point zero eight per one hundred of face, and the five-year note has point zero five.
- First, the target's D V zero one in dollars: one hundred million times point zero eight, divided by one hundred, is eighty thousand dollars per basis point.
- Second, the D V zero one ratio: point zero eight divided by point zero five is one point six.
- Third, the hedge face: one hundred million times one point six is one hundred and sixty million.
- Check it: one hundred and sixty million times point zero five, over one hundred, is again eighty thousand dollars.
- Last, the side.
- The desk is long the bond, so it sells the five-year note.

4 Example 2: yields that do not move one for one

Hedge $160m. The 5-year yield rises 4 bp; the 10-year rises only 3 bp

Long 10-year bond	−80,000 × 3	**−$240,000**
Short 5-year note	+80,000 × 4	**+$320,000**
Net P&L	320,000 − 240,000	**+$80,000**
Face that matched the moves	160m × 3 ÷ 4	**$120m**
Extra $40m: DV01 $20,000	20,000 × 4 bp	**$80,000**

If both yields fall 4 and 3 bp instead: net −$80,000

IN THE LESSON

- Now suppose that one day the five-year yield rises four basis points but the ten-year yield rises only three.
- The long bond loses eighty thousand times three, two hundred and forty thousand dollars.
- The short note gains eighty thousand times four, three hundred and twenty thousand.
- Net, the desk makes eighty thousand dollars.
- With moves of three against four, only three quarters of the hedge was needed: one hundred and sixty million times three over four is one hundred and twenty million.
- The extra forty million of face carries a D V zero one of twenty thousand dollars, and twenty thousand times four basis points is exactly the eighty thousand.
- If both yields had fallen by those amounts instead, the desk would have lost eighty thousand dollars.

5 Example 3: when the curve twists

Same $160m hedge. The 5-year yield falls 2 bp; the 10-year yield rises 3 bp

Long 10-year bond	−80,000 × 3	**− $240,000**
Short 5-year note	−80,000 × 2	**− $160,000**
Net P&L	−240,000 − 160,000	**− $400,000**
No hedge, for comparison	−80,000 × 3	**− $240,000**

IN THE LESSON

- Suppose the curve twists: the five-year yield falls two basis points while the ten-year yield rises three.
- Again the long bond loses two hundred and forty thousand dollars, eighty thousand times three.
- The short note also loses, because its yield fell and its price rose: eighty thousand times two, one hundred and sixty thousand.
- Net, the desk loses four hundred thousand dollars.
- With no hedge, the bond alone would have lost only two hundred and forty thousand.
- On this day the hedge added to the loss instead of reducing it.
- A D V zero one neutral hedge protects only against the parallel move it assumes, and in a twist the two yields move in opposite directions.

6 Why a DV01-neutral hedge falls short

- DV01-neutral is exact only for a parallel shift
- Not one for one: say 0.7 bp of 10y per 1 bp of 2y
- The curve also steepens, flattens and twists
- Hedged yield moves less: hedge too big, a bet on the hedge
- Exam: the error comes from unequal yield changes

IN THE LESSON

- A D V zero one neutral hedge is exact only for a parallel shift, where both yields move by the same number of basis points.
- Real yields at different maturities do not move one for one.
- For instance, a ten-year yield might move only point seven of a basis point for each basis point of the two-year yield.
- The curve also steepens, flattens and twists.
- When the hedged yield tends to move less than the hedging yield, the D V zero one neutral hedge is too large, and the position takes on the risk of the hedging instrument itself, as Kestrel's extra forty million did.
- The exam point: the hedging error comes from unequal yield changes, which are not one for one.

7 The idea of a regression hedge

- Use the history of yield changes, not a one-for-one guess
- Dependent: change in the target's yield
- Explanatory: change in the hedge's yield
- β scales the DV01-neutral hedge; β = 1 is DV01-neutral
- β = 0.8: hedge with 80% of the DV01-neutral amount

IN THE LESSON

- A regression hedge starts from the observed history of yield changes instead of assuming the two yields move one for one.
- The desk runs a regression, which fits a straight line through past data.
- The dependent variable, the one being explained, is the change in the target's yield.
- The explanatory variable is the change in the hedge's yield.
- The slope of the fitted line, called beta, then scales the D V zero one neutral hedge.
- A beta of one gives back the D V zero one neutral hedge, so that hedge is just a special case.
- A beta of point eight means hedging with eighty per cent of the D V zero one neutral amount.

8 The regression, and the risk it removes

Δy_T = α + β × Δy_H + ε
Var(hedged P&L) = (1 − R^2) × Var(unhedged P&L)

Beta β	bp of target per bp of hedge
Error ε	What stays unhedged
R squared	Share of variance removed

IN THE LESSON

- The change in the target's yield equals alpha, a constant, plus beta times the change in the hedge's yield, plus epsilon, an error term.
- Read the three boxes.
- Beta is the slope: how many basis points the target yield tends to move per basis point of the hedge yield.
- Epsilon is the part of the target's move that the hedge yield does not explain, and that is exactly what stays unhedged.
- R squared is the share of the variance of target changes that the regression explains, so it is the share of the variance that the hedge removes.
- That gives the second line.
- The variance of the hedged P and L equals one minus R squared, times the variance of the unhedged P and L.

9 Example 4: how much risk each hedge leaves

Daily changes: σ_T = 5 bp, σ_H = 6 bp, ρ = 0.9; the regression gives β = 0.75

No hedge: P&L SD	80,000 × 5 bp	**$400,000**
R^2 of the regression	0.9^2	**0.81**
Regression hedge: P&L SD	400,000 × √(1 − 0.81)	**$174,356**
DV01-neutral: P&L SD	80,000 × √(25 + 36 − 54)	**$211,660**

Variance removed: regression 81%, DV01-neutral 1 − 7 ÷ 25 = 72%

IN THE LESSON

- The ten-year yield has a standard deviation, or volatility, of five basis points a day, and the five-year six.
- Their correlation is point nine, and the regression gives a beta of point seven five.
- Without a hedge, the daily P and L standard deviation is eighty thousand times five, four hundred thousand dollars.
- The R squared is point nine squared, point eight one.
- The regression hedge leaves four hundred thousand times the square root of one minus point eight one, about one hundred and seventy four thousand three hundred and fifty six dollars.
- The D V zero one neutral hedge leaves eighty thousand times the square root of twenty five plus thirty six minus fifty four, where fifty four is two times point nine times five times six.
- That is about two hundred and eleven thousand six hundred and sixty dollars.
- The regression hedge removes eighty one per cent of the variance, the D V zero one neutral hedge only one minus seven over twenty five, seventy two per cent.

10 Working out beta

β = Cov(Δy_T, Δy_H) ÷ Var(Δy_H)
β = ρ × (σ_T ÷ σ_H)

Covariance	How the two changes co-move
Variance	Of the hedge changes only
Correlation ρ	√R^2, with the slope's sign
Volatilities σ	Of changes, in bp, same horizon

IN THE LESSON

- Beta is the slope of the regression of target yield changes on hedge yield changes.
- The first line says beta equals the covariance of the two changes, divided by the variance of the hedge changes.
- The second line says the same thing using the correlation, rho, and the two volatilities: beta equals rho times sigma T over sigma H.
- Now the four boxes.
- The covariance measures how the two changes move together.
- The variance is that of the hedge changes only.
- The correlation can come from R squared, by taking its square root and keeping the sign of the slope.
- And the volatilities are of yield changes, in basis points, measured over the same horizon, such as a day or a week.

11 Example 5: beta from Kestrel's statistics

Daily changes: σ_T = 5 bp, σ_H = 6 bp, ρ = 0.9

Covariance	0.9 × 5 × 6	**27 bp^2**
Variance of hedge changes	6^2	**36 bp^2**
Beta: Cov ÷ Var	27 ÷ 36	**0.75**
Check: ρ × σ_T ÷ σ_H	0.9 × 5 ÷ 6	**0.75**
Given R^2 = 0.81 instead	ρ = √0.81, slope's sign	**0.9**

IN THE LESSON

- Let us compute Kestrel's beta.
- Sigma T is five basis points, sigma H six, and rho point nine.
- The covariance is rho times the two volatilities: point nine times five times six, twenty seven basis points squared.
- The variance of hedge changes is six squared, thirty six.
- Beta is twenty seven divided by thirty six, point seven five.
- Check: point nine times five over six is also point seven five.
- So the ten-year yield tends to move three quarters of a basis point for each basis point of the five-year yield.
- Suppose instead you were given only an R squared of point eight one.
- Its square root, point nine, is the correlation, taking a positive sign because the slope is positive.

12 Example 6: beta from five days of data

A toy sample of five days of changes, in bp; each series averages zero

Hedge yield changes	−8, −4, 0, 4, 8	**mean 0**
Target yield changes	−6, −4, 1, 4, 5	**mean 0**
Sum of products	48 + 16 + 0 + 16 + 40	**120**
Sum of hedge squares	64 + 16 + 0 + 16 + 64	**160**
Slope β	120 ÷ 160	**0.75**

IN THE LESSON

- Take a toy sample of five days; real work uses hundreds.
- The hedge yield changes are minus eight, minus four, zero, four and eight basis points.
- The target changes on the same days are minus six, minus four, one, four and five.
- Both average zero, which keeps the arithmetic short, because then covariance and variance are just sums divided by the same count, and the count cancels.
- The sum of products is forty eight plus sixteen plus zero plus sixteen plus forty, one hundred and twenty.
- The sum of the hedge changes squared is sixty four plus sixteen plus zero plus sixteen plus sixty four, one hundred and sixty.
- The slope is one hundred and twenty over one hundred and sixty, point seven five.

13 Slips in computing beta

σ_T = 5 bp, σ_H = 6 bp, ρ = 0.9, R^2 = 0.81. Three common slips

Inverted volatility ratio	0.9 × 6 ÷ 5	**1.08**
Correlation left out	5 ÷ 6	**0.83**
R^2 used in place of ρ	0.81 × 5 ÷ 6	**0.675**
Correct	0.9 × 5 ÷ 6	**0.75**

IN THE LESSON

- Kestrel's inputs are sigma T of five basis points, sigma H of six, rho of point nine and R squared of point eight one.
- Three slips cost marks on beta.
- The first slip inverts the volatility ratio: point nine times six over five gives one point zero eight.
- Remember that the target's volatility goes on top, because the target is the variable being explained.
- The second slip leaves out the correlation: five over six gives about point eight three, which would be right only if the two changes were perfectly correlated.
- The third slip uses R squared where the correlation belongs: point eight one times five over six gives point six seven five.
- The correct answer is point nine times five over six, point seven five.

14 The face amount of a regression hedge

Hedge = Target × (DV01_T ÷ DV01_H) × β
Hedge DV01 = β × Target DV01

Target face	The position being hedged
DV01 ratio	Same face basis, per 100
Beta β	Multiply by β, never divide
Hedge face	Opposite side: sell if long

IN THE LESSON

- The regression hedge face equals the target face, times the D V zero one ratio, times beta.
- The second line says the same thing in risk terms: the hedge's D V zero one should equal beta times the target's D V zero one.
- The four boxes explain each part.
- The target face is the position being hedged, such as Kestrel's long bond.
- The D V zero one ratio uses both D V zero ones on the same face basis, for instance per one hundred of face.
- Beta multiplies, and never divides.
- And the hedge face is traded on the opposite side: sell the hedge if you are long the target.

15 Example 7: Kestrel's regression hedge

Target $100m; DV01 0.08 and 0.05 per 100; β = 0.75

DV01-neutral face	100m × 0.08 ÷ 0.05	**$160m**
Regression hedge face	160m × 0.75	**$120m**
Hedge DV01	120m × 0.05 ÷ 100	**$60,000**
Check: β × target DV01	0.75 × 80,000	**$60,000**
Typical day: 5y +4, 10y +3 bp	−240,000 + 60,000 × 4	**$0 net**

Sell $120m of the 5-year note against the long bond

IN THE LESSON

- The target is one hundred million of face, the D V zero ones are point zero eight and point zero five per one hundred, and beta is point seven five.
- The D V zero one neutral face is one hundred million times point zero eight over point zero five, one hundred and sixty million.
- Multiply by beta: one hundred and sixty million times point seven five is one hundred and twenty million.
- The hedge D V zero one is one hundred and twenty million times point zero five over one hundred, sixty thousand dollars per basis point.
- Check it: beta times the target D V zero one, point seven five times eighty thousand, is also sixty thousand.
- On a typical day the five-year yield rises four basis points and the ten-year three.
- The bond loses two hundred and forty thousand, and the hedge gains sixty thousand times four, the same amount, so the net is zero.
- The desk sells one hundred and twenty million of the five-year note.

16 Slips in sizing the hedge

Target $100m; DV01 0.08 vs 0.05 per 100; β = 0.75. Three common slips

DV01 ratio inverted	100m × (0.05 ÷ 0.08) × 0.75	**$46.9m**
Divided by β	160m ÷ 0.75	**$213.3m**
β ignored	the DV01-neutral amount	**$160m**
Correct	160m × 0.75	**$120m**

IN THE LESSON

- Take the same inputs: a one hundred million dollar target, D V zero ones of point zero eight and point zero five, and beta of point seven five.
- Sizing has its own slips.
- Inverting the D V zero one ratio gives one hundred million times point zero five over point zero eight, times point seven five, about forty six point nine million, far too small.
- Dividing by beta instead of multiplying gives one hundred and sixty million over point seven five, about two hundred and thirteen point three million, far too large.
- Ignoring beta altogether gives the D V zero one neutral amount, one hundred and sixty million, which we have seen over hedges.
- The correct trade is one hundred and sixty million times point seven five, one hundred and twenty million.

17 Two explanatory variables

$\Delta y = \alpha + \beta_2 \times \Delta S_2 + \beta_{30} \times \Delta S_{30} + \varepsilon$
$\text{Face}_i = F \times (\text{DV01}_T \div \text{DV01}_i) \times \beta_i$

Beta on 2-year	Holding the 30-year rate fixed
Beta on 30-year	Holding the 2-year rate fixed
Swap DV01	Per 100 of face, each swap

IN THE LESSON

- Kestrel now tries hedging the ten-year bond with two swaps, a two-year and a thirty-year.
- It regresses the change in the bond's yield on both swap rate changes at once.
- The first line is the regression: the yield change equals alpha, plus beta two times the change in the two-year swap rate, plus beta thirty times the change in the thirty-year swap rate, plus an error.
- The second line sizes each swap on its own: its face equals the target face, times the target's D V zero one over that swap's D V zero one, times that swap's own beta.
- In the boxes: beta two holds the thirty-year rate fixed, beta thirty holds the two-year rate fixed, and each swap's D V zero one is per one hundred of face.

18 Example 8: hedging with two swaps

Target $100m, DV01 0.08. Swaps: 2y 0.02, 30y 0.16 per 100. β_2 = 0.30, β_30 = 0.60

2-year swap face	100m × (0.08 ÷ 0.02) × 0.30	**$120m**
2-year swap DV01	120m × 0.02 ÷ 100	**$24,000**
30-year swap face	100m × (0.08 ÷ 0.16) × 0.60	**$30m**
30-year swap DV01	30m × 0.16 ÷ 100	**$48,000**
Check: 0.30 and 0.60 × 80,000	24,000 + 48,000	**$72,000**

Pay fixed on both swaps: a payer swap gains when rates rise

IN THE LESSON

- The target is still the one hundred million bond, D V zero one point zero eight.
- Kestrel's regression gives beta two of point three and beta thirty of point six.
- The two-year swap has a D V zero one of point zero two per one hundred, and the thirty-year swap point one six.
- The two-year face is one hundred million times point zero eight over point zero two, times point three: one hundred and twenty million.
- Its D V zero one is one hundred and twenty million times point zero two over one hundred, twenty four thousand dollars.
- The thirty-year face is one hundred million times point zero eight over point one six, times point six: thirty million.
- Its D V zero one is thirty million times point one six over one hundred, forty eight thousand.
- Check: point three and point six times eighty thousand give the same two figures, seventy two thousand in all.
- To gain when rates rise, Kestrel pays fixed on both swaps.

19 Partial slopes, and two traps

- Betas are partial: each holds the other rate fixed
- Correlated rates: partial slopes differ from simple ones
- They need not sum to 1: here 0.30 + 0.60 = 0.90
- Trap: a one-variable slope for each swap
- Trap: 0.90 on each swap gives $360m and $45m

IN THE LESSON

- The betas here are partial slopes.
- Each one measures the effect of its own swap rate holding the other one fixed.
- Swap rates are correlated, so a one-variable slope on the two-year rate alone would also pick up part of the thirty-year rate's effect, and the two partial slopes differ from the simple ones.
- The betas also need not sum to one.
- Here they sum to point nine.
- There are two traps.
- The first is using a one-variable slope for each swap.
- The second is applying the sum of the betas to each swap: point nine on each would give three hundred and sixty million of the two-year and forty five million of the thirty-year, far too much.

20 Level regressions and change regressions

Same yields, two kinds of regression

	Level regression	**Change regression**
Regresses	Yield on yield	Change on change
R^2	Very high, flattering	Lower, more honest
Errors	Serially correlated	Closer to independent
Best use	Rich or cheap	Sizing a hedge

IN THE LESSON

- There are two ways to regress one yield on another.
- A level regression regresses the yield level of one security on the yield level of another.
- A change regression regresses the change in one yield on the change in the other, over a fixed interval such as a day or a week.
- The R squared of a level regression is very high and flattering, while that of a change regression is lower and more honest.
- The errors of a level regression are serially correlated, meaning today's error is close to yesterday's, while change regression errors are much closer to independent.
- A level regression is best used to judge whether a security is rich or cheap, and a change regression to size a hedge.

21 What they share, and why levels mislead

- Both fit an intercept, a slope and an error term
- Both can drift between sample periods
- Levels trend and persist, so R^2 overstates the fit
- Hedging is about P&L from rate moves: use changes
- Kestrel: levels R^2 = 0.98; changes R^2 = 0.81

IN THE LESSON

- Both regressions fit an intercept, a slope and an error term.
- And both can drift between sample periods.
- So why do levels mislead?
- Yield levels trend and are persistent.
- Two trending series line up even when their day to day moves are only loosely linked, so R squared overstates how well one yield explains the other.
- Hedging is about the P and L from moves in rates, so it needs changes.
- Kestrel's data show the gap: the level regression has an R squared of point nine eight, the change regression point eight one.

22 Example 9: rich or cheap, and the wrong hedge

Level fit: 10y yield = 0.30% + 1.05 × 5y yield.
Today: 5y 4.00%, 10y 4.58%

Fitted 10-year yield	0.30% + 1.05 × 4.00%	**4.50%**
Actual minus fitted	4.58% − 4.50%	**+8 bp**
Verdict	yield above fit, price below	**10-year cheap**
Hedge sized on level slope	160m × 1.05	**$168m: wrong**
Hedge sized on change beta	160m × 0.75	**$120m**

IN THE LESSON

- Kestrel's level fit says the ten-year yield equals point three per cent plus one point zero five times the five-year yield.
- Today the five-year yield is four per cent, and the ten-year is four point five eight.
- The fitted ten-year yield is point three plus one point zero five times four, four point five per cent.
- Actual minus fitted is plus eight basis points.
- The bond yields more than the fit, so its price is lower than the fit implies: the ten-year looks cheap.
- But the level slope must not size the hedge.
- One hundred and sixty million times one point zero five would give one hundred and sixty eight million: wrong.
- The change beta gives one hundred and twenty million.

23 Regressing the other way round

Reverse slope = ρ × (σ_H ÷ σ_T)
Implied ratio = σ_T ÷ (ρ × σ_H) = β ÷ ρ^2
β × reverse slope = ρ^2 = R^2

Beta β	Target on hedge: the right one
Reverse slope	Hedge on target
Implied ratio	One over the reverse slope

IN THE LESSON

- Now regress the other way round, the hedge yield change on the target yield change.
- The first line is the slope of that reverse regression: rho times sigma H over sigma T.
- Taking its reciprocal gives another candidate hedge ratio.
- The second line shows it: sigma T divided by rho times sigma H, which equals beta divided by rho squared.
- The last line says beta times the reverse slope equals rho squared, which is R squared.
- In the boxes: beta, target on hedge, is the right slope.
- The reverse slope is hedge on target.
- And the implied ratio is one over the reverse slope.
- Since rho squared is below one unless rho is plus or minus one, this ratio is bigger than beta: an over hedge.

24 Example 10: Kestrel regresses the wrong way

σ_T = 5 bp, σ_H = 6 bp, ρ = 0.9; the correct β = 0.75

Reverse slope	0.9 × 6 ÷ 5	**1.08**
Implied hedge ratio	1 ÷ 1.08 = 0.75 ÷ 0.81	**0.926**
Hedge face, reverse	160m ÷ 1.08	**$148.1m**
Hedge face, correct	160m × 0.75	**$120m**
Product of the slopes	0.75 × 1.08	**0.81 = R^2**

Over-hedge: 148.1m − 120m ≈ $28m

IN THE LESSON

- Sigma T is five basis points, sigma H six and rho point nine, so the correct beta is point seven five.
- Suppose an analyst regresses the five-year yield change on the ten-year yield change.
- The reverse slope is point nine times six over five, one point zero eight.
- Its reciprocal, one over one point zero eight, is about point nine two six.
- The same number comes from beta over rho squared: point seven five over point eight one.
- The hedge face becomes one hundred and sixty million divided by one point zero eight, about one hundred and forty eight point one million, against the correct one hundred and twenty million.
- The product of the slopes, point seven five times one point zero eight, is point eight one, the R squared.
- That leaves about twenty eight million of over hedge.

25 Example 11: why the direction matters

Daily P&L SD = 80,000 × √(25 + h^2 × 36 − 2h × 27), h = hedge ratio

h = 0.75: regression	25 + 20.25 − 40.5 = 4.75	**$174,356**
h = 1 ÷ 1.08: reverse	25 + 30.86 − 50 = 5.86	**$193,729**
h = 1: DV01-neutral	25 + 36 − 54 = 7	**$211,660**
h = 0: no hedge	25	**$400,000**

Only target-on-hedge minimises the variance of hedged P&L

IN THE LESSON

- For a hedge ratio h, Kestrel's daily P and L standard deviation is eighty thousand times the square root of twenty five plus h squared times thirty six, minus two h times twenty seven, because the hedged P and L follows the target change minus h times the hedge change.
- At point seven five, the variance is four point seven five, a standard deviation of about one hundred and seventy four thousand three hundred and fifty six dollars.
- At one over one point zero eight, five point eight six, about one hundred and ninety three thousand seven hundred and twenty nine.
- At one, seven, about two hundred and eleven thousand six hundred and sixty.
- At zero, four hundred thousand.
- Why is target on hedge the right direction?
- Ordinary least squares minimises the squared errors in the dependent variable.
- So minimising hedged P and L variance is exactly the regression of target on hedge.
- The position being hedged is the dependent variable.

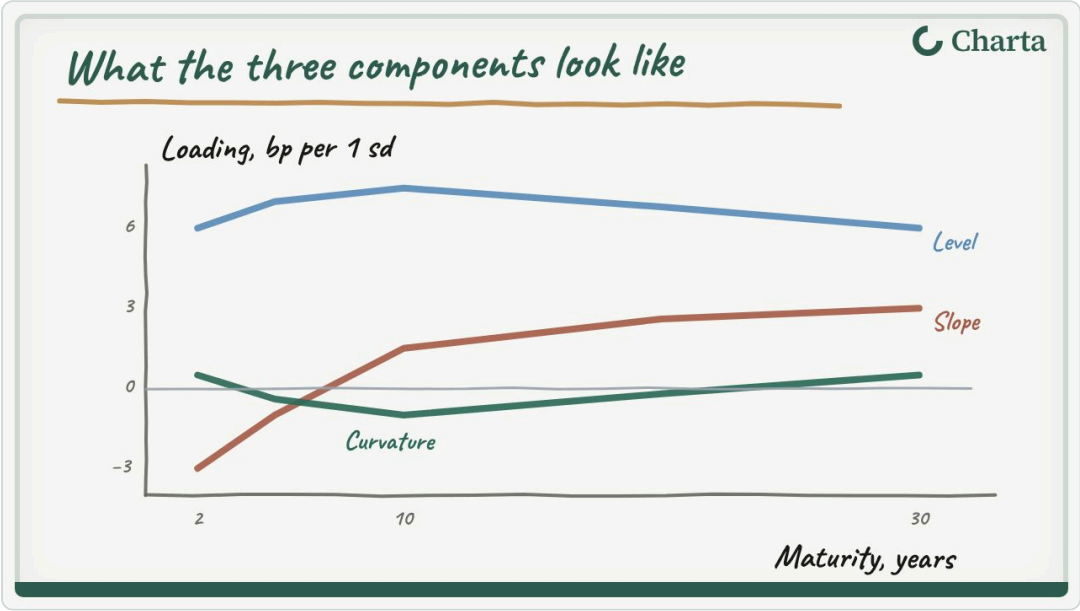
26 Principal component analysis

- Start: covariance matrix of changes in many rates
- Rewrite as uncorrelated components, ranked by variance
- Component variances add up to the total variance
- First three: level, slope and curvature
- Made-up sample: 80% + 12% + 4% = 96% of variance

IN THE LESSON

- A real book is exposed to the whole curve.
- Principal component analysis, or P C A, starts from the covariance matrix of changes in many rates.
- It rewrites that matrix as a set of uncorrelated factors, the principal components, ranked by the variance each one explains.
- The component variances add up to the total variance.
- For yield curves, the first three are usually seen as a level shift, a slope change and a curvature change.
- In a made up sample, level might explain eighty per cent of the variance, slope twelve and curvature four, ninety six per cent together.

27 What the three components look like



IN THE LESSON

- This chart plots, against maturity, how many basis points each yield moves when a component moves by one standard deviation.
- These numbers are called the loadings.
- The level curve is high and fairly flat: about six basis points at two years, seven and a half at ten, and six at thirty, so every yield moves the same way.
- The slope curve rises from minus three at two years, through one and a half at ten, to plus three at thirty: the short end falls while the long end rises, so the curve steepens.
- The curvature curve is a dip: plus a half at two and thirty years, minus one at ten, so the middle moves against the ends.
- The grey line marks zero, the left axis marks minus three, three and six, and the maturities, two, ten and thirty years, run along the bottom.

28 Example 12: the bond's exposure to each component

Loadings at 2y, 10y, 30y. Level 6, 7.5, 6; slope −3, 1.5, 3; curvature 0.5, −1, 0.5

Level: 10y moves +7.5 bp	−80,000 × 7.5	**−$600,000**
Slope: 10y moves +1.5 bp	−80,000 × 1.5	**−$120,000**
Curvature: 10y moves −1 bp	−80,000 × (−1)	**+$80,000**
Unhedged SD, three factors	√(600^2 + 120^2 + 80^2) × 1,000	**$617,090**

IN THE LESSON

- The loadings at two, ten and thirty years are: level six, seven and a half and six; slope minus three, one and a half and three; curvature a half, minus one and a half.
- A one standard deviation level move lifts the ten-year yield by seven and a half basis points, so the bond loses eighty thousand times seven and a half, six hundred thousand dollars.
- A slope move lifts it by one and a half basis points: a loss of one hundred and twenty thousand.
- A curvature move lowers it by one basis point, so the bond gains eighty thousand.
- Because the components are uncorrelated, their variances simply add.
- The unhedged standard deviation, over these three components, is the square root of six hundred squared plus one hundred and twenty squared plus eighty squared, in thousands: about six hundred and seventeen thousand and ninety dollars.

29 Example 13: hedging level and slope with two swaps

Pay fixed on 2y and 30y swaps; their DV01s are D_2 and D_30 per bp

Level nets to zero	6 D_2 + 6 D_30 = 600,000	**sum 100,000**
Slope nets to zero	−3 D_2 + 3 D_30 = 120,000	**gap 40,000**
Solve	D_30 = 70,000; D_2 = 30,000	**2 swaps, 2 factors**
2-year swap face	30,000 ÷ (0.02 ÷ 100)	**$150m**
30-year swap face	70,000 ÷ (0.16 ÷ 100)	**$43.75m**

IN THE LESSON

- To hedge two components, Kestrel needs two instruments.
- It pays fixed on a two-year and a thirty-year swap, with D V zero ones D two and D thirty, so each gains its D V zero one for every basis point its rate rises.
- For level to net to zero, six D two plus six D thirty must equal six hundred thousand, so the two sum to one hundred thousand.
- For slope, minus three D two plus three D thirty must equal one hundred and twenty thousand, so D thirty exceeds D two by forty thousand.
- Solving, D thirty is seventy thousand and D two is thirty thousand dollars.
- With D V zero ones of point zero two and point one six per one hundred, the faces are one hundred and fifty million and forty three point seven five million.

30 Example 14: what risk is left

After the hedge: level and slope net to zero; curvature does not

Level P&L per sd	−600,000 + 6 × 100,000	**$0**
Slope P&L per sd	−120,000 − 90,000 + 210,000	**$0**
Curvature P&L per sd	80,000 + 15,000 + 35,000	**+$130,000**
Residual SD	√(130,000^2): uncorrelated	**$130,000**
Before the hedge	from Example 12	**$617,090**

IN THE LESSON

- After the hedge, level and slope net to zero but curvature does not.
- For level, the bond loses six hundred thousand and the swaps gain six times one hundred thousand, so the net is zero.
- For slope, the bond loses one hundred and twenty thousand, the two-year swap loses three times thirty thousand, ninety thousand, and the thirty-year swap gains three times seventy thousand, two hundred and ten thousand: net zero again.
- For curvature, the bond gains eighty thousand, and the swaps gain half of thirty thousand and half of seventy thousand, fifteen and thirty five thousand: one hundred and thirty thousand in all.
- Because the components are uncorrelated, the residual risk is the sum of the variances of the unhedged components, here a standard deviation of one hundred and thirty thousand, down from six hundred and seventeen thousand and ninety dollars.
- Notice the swaps raise the curvature exposure, from eighty thousand to one hundred and thirty thousand; hedging it would need a third instrument.

31 Mistakes to avoid

- DV01-neutral is exact only for parallel shifts
- Regress target changes on hedge changes, not the reverse
- β = ρ × σ_T ÷ σ_H — no inverting, no dropping ρ, not R^2
- Multiply DV01-neutral face by β; estimate from changes
- PCA: one instrument per factor; the rest stays as risk

IN THE LESSON

- First, a D V zero one neutral hedge is exact only for parallel shifts, and real yields do not move one for one.
- Second, regress the target's yield changes on the hedge's yield changes, never the reverse, because the position being hedged is the dependent variable.
- Third, beta is rho times sigma T over sigma H: do not invert the ratio, drop the correlation, or use R squared in its place.
- Fourth, multiply the D V zero one neutral face by beta, and estimate beta from changes, not levels.
- Fifth, in P C A, you need one hedging instrument for each factor you hedge, and what you leave unhedged stays as risk.

32 Try three quick questions

1. R^2 = 0.64, positive slope; σ_T = 6 bp, σ_H = 8 bp. Beta?
   Answer: ρ = √0.64 = 0.8. β = 0.8 × 6 ÷ 8 = 0.6.
2. Long $50m, DV01 0.09 per 100; hedge DV01 0.06; β = 0.8. Hedge?
   Answer: 50m × (0.09 ÷ 0.06) × 0.8 = $60m, sold.
3. β = 0.6 and ρ = 0.8. Ratio implied by the reverse regression?
   Answer: 0.6 ÷ 0.8^2 = 0.6 ÷ 0.64 = 0.9375: an over-hedge.

IN THE LESSON

- Now try three quick questions.
- Pause the video, think, then listen to the answers.
- Question one.
- A regression of target yield changes on hedge yield changes has an R squared of point six four and a positive slope.
- The target's volatility is six basis points and the hedge's is eight.
- What is beta?
- The correlation is the square root of point six four, point eight, so beta is point eight times six over eight, point six.
- Question two.
- You are long fifty million of face with a D V zero one of point zero nine per one hundred, the hedge has point zero six, and beta is point eight.
- What hedge?
- Fifty million times point zero nine over point zero six, which is one point five, times point eight: sixty million, sold.
- Question three.
- With beta point six and correlation point eight, what ratio does the reverse regression imply?
- Point six over point six four, point nine three seven five, an over hedge.