CHARTA LEARNING · LESSON NOTES

Enterprise value multiples and comparables

Equity Investments

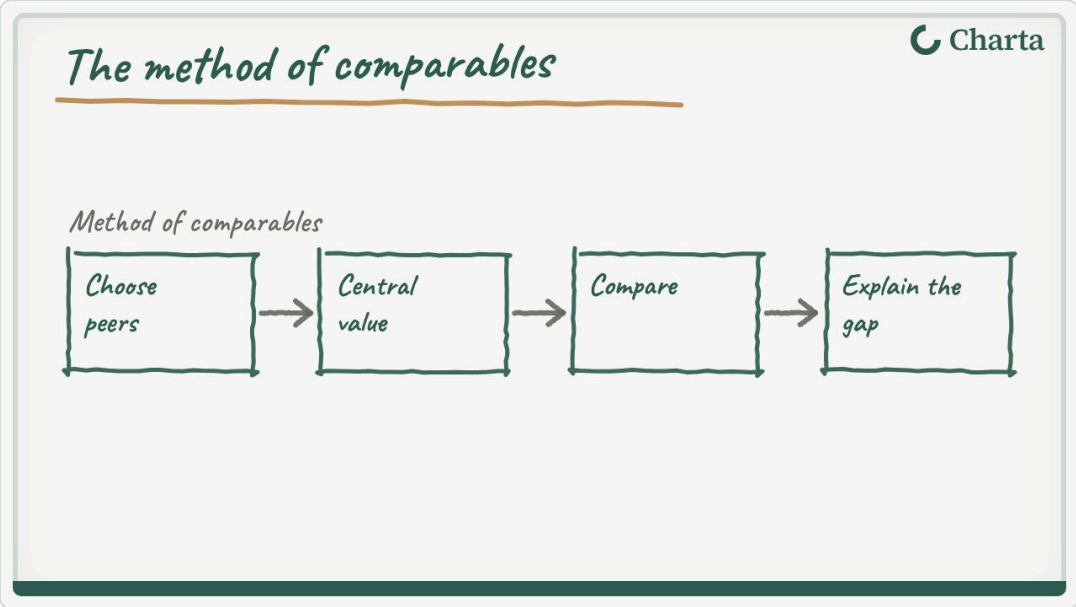
1 A story: the cheap house on the street

- A low price per square metre may be a bargain
- Or the house may simply be worse
- Shares work the same way
- Check the fundamentals before you celebrate

IN THE LESSON

- Imagine you are buying a house.
- A friend tells you about one that costs less per square metre than every other house on the street.
- A bargain?
- Maybe.
- But before you rush, you walk round it.
- It has a damp cellar, a noisy main road behind it and a roof that needs replacing.
- Suddenly the low price per square metre looks fair, even generous.
- Shares are no different.
- A share on a low multiple may be a bargain, or it may simply be a weaker business that the market has priced correctly.
- The method of comparables is the habit of checking the cellar before celebrating the price.
- It has four steps, and the last one is the one that matters.

2 The method of comparables



IN THE LESSON

- Here are the four steps, drawn left to right.
- Step one, choose the comparable companies.
- Usually they are in the same industry and are similar in size, growth and risk.
- Step two, calculate each peer's multiple, such as price to earnings, and then a central value.
- The usual central value is the median, and we will see why later.
- Step three, compare the subject company's multiple with that central value.
- That is only a subtraction.
- Step four is the judgement.
- Ask whether differences in fundamentals explain the gap.
- If the subject grows more slowly, carries more risk or earns a lower return on equity than its peers, then a lower multiple is exactly what fair pricing looks like.

3 Example 1: a subject and five peers

P/E: peers 14, 16, 18, 22, 35 (growth 7, 8, 9, 10, 12%); subject 15 (6%)

Peer mean	(14+16+18+22+35) ÷ 5	**21**
Peer median	middle of five	**18**
Subject vs median	15 ÷ 18 − 1	**−17%**
Subject vs mean	15 ÷ 21 − 1	**−29%**

The 35 inflates the mean; the median ignores it

IN THE LESSON

- Example one.
- Meet an analyst called Priya.
- She is valuing a company I will call the subject.
- Five peers trade on price to earnings ratios of fourteen, sixteen, eighteen, twenty two and thirty five, with expected growth of seven, eight, nine, ten and twelve per cent a year.
- The subject trades on fifteen times earnings, with expected growth of six per cent.
- Risk is about the same across the group.
- First, the central value.
- The mean is fourteen plus sixteen plus eighteen plus twenty two plus thirty five, which is one hundred and five, divided by five, so twenty one.
- The median is the middle peer once sorted, which is eighteen.
- That thirty five pulls the mean up and leaves the median alone.

4 Allowing for growth: the PEG ratio

PEG = (P/E) ÷ expected growth rate (in % a year)

Price to earnings (P/E)	Price paid per unit of earnings
Growth rate (% a year)	Expected earnings growth
P/E to growth (PEG)	Price paid per point of growth

IN THE LESSON

- To bring growth into the comparison, analysts often use the P E G ratio.
- It is the price to earnings ratio divided by the expected growth rate, where the growth rate is written as a whole number of per cent, not a decimal.
- Take each term.
- The price to earnings ratio is the price you pay for one unit of earnings.
- The growth rate is how fast those earnings are expected to rise each year.
- And the P E G is the answer: how much you pay for each point of growth.
- A lower P E G means you pay less for the growth you get.
- It is a quick way to adjust a multiple for one fundamental, though it ignores risk, so use it as a first check only.

5 Example 1 continued: is the subject really cheap?

PEG = P/E ÷ growth. Peers: 2.0, 2.0, 2.0, 2.2, 2.9; subject growth 6%

Peer median PEG	middle of five	**2.0**
Subject PEG	15 ÷ 6	**2.5**
Raw P/E verdict	17% below median	**looks cheap**
Growth adjusted verdict	2.5 against 2.0	**looks dear**

The discount reversed once growth was allowed for

IN THE LESSON

- Now Priya applies step four.
- Each peer's growth adjusted ratio is its price to earnings divided by its growth.
- Fourteen over seven is two.
- Sixteen over eight is two.
- Eighteen over nine is two.
- Twenty two over ten is two point two.
- And thirty five over twelve is two point nine.
- The median is two.
- The subject is fifteen divided by six, which is two point five.
- So the discount of seventeen per cent on the raw multiple has turned into a premium.
- The subject costs more per point of growth than four of its five peers.
- The market has priced in its slow growth, and then some.
- Calling it cheap would have been a mistake.

6 Fundamentals, history and the trap

- Higher growth, lower risk, higher ROE: higher multiple
- History is a fair yardstick only if conditions match
- Trap: calling a gap mispricing when fundamentals explain it
- Ask what the market already knows

IN THE LESSON

- Which fundamentals matter?
- Higher expected growth, lower risk and a higher return on equity all justify a higher price to earnings or price to book multiple.
- So a gap that these differences explain is not mispricing.
- Analysts can also compare a company with its own past multiples.
- That works only if interest rates and business conditions are similar to those in the past.
- A multiple set when rates were low is not a fair yardstick for a world of high rates.
- The trap, and a favourite of questions, is inferring mispricing from a gap that fundamental differences fully explain.
- A cheap looking share can be exactly fairly priced.

7 Cash flow: which cash flow?

- No single definition of cash flow
- Each one measures at a different point of the pipe
- Closer to the shareholder: more honest, less stable
- Know what each leaves out

IN THE LESSON

- Earnings can be a shaky base for a multiple, because they depend on accounting choices.
- So analysts also use cash flow multiples, such as price to cash flow.
- The trouble is that there is no single cash flow.
- Think of a water pipe running from the company's earnings down to the shareholders' pocket.
- At every joint along the pipe some water leaks away: depreciation, a rise in working capital, spending on new equipment, repayment of loans.
- Each definition of cash flow measures the water at a different joint.
- The nearer the shareholders' pocket you measure, the more honest the number, but the less steady it is.
- Let's walk down the pipe with a worked example, per share.

8 Example 2: from net income to FCFE

Per share: net income 3.00, D&A 1.20, working capital +0.50

Earnings + non-cash charges	3.00 + 1.20	**4.20**
Cash flow from operations	4.20 − 0.50	**3.70**
Free cash flow to equity	3.70 − 2.40 + 0.40	**1.70**

Capex 2.40 and net borrowing 0.40 come off after CFO

IN THE LESSON

- Example two.
- A company has net income of three dollars a share.
- First joint: add back non cash charges, depreciation and amortisation, of one dollar twenty.
- That gives four dollars twenty, earnings plus non cash charges.
- Second joint: take off a rise in working capital of fifty cents.
- That gives three dollars seventy, which is cash flow from operations, here with no unusual items to remove.
- Third joint: take off two dollars forty of capital spending, and add forty cents of net borrowing.
- We arrive at one dollar seventy.
- That is free cash flow to equity, the cash that could actually be paid out to shareholders after the business has been kept running.

9 Where each definition falls short

Cash flow definition: its weakness

Definition	Falls short because
Earnings + non-cash	ignores WC and capex
Adjusted CFO	still ignores capex
FCFE	volatile, can be negative
EBITDA	pre-debt; ignores tax, capex

IN THE LESSON

- Now the shortfalls, one by one.
- Earnings plus non cash charges is simple, but it ignores working capital and capital spending.
- Adjusted cash flow from operations strips out one off items, but it still ignores the spending needed to replace equipment.
- Free cash flow to equity is theoretically the best, because it measures cash left for shareholders after investment.
- But it is volatile and can even be negative in a heavy investment year, which makes the ratio meaningless.
- E B I T D A is a different animal.
- It belongs to every provider of capital, lenders as well as shareholders, so it is not a cash flow for share price at all.
- It ignores tax and capital spending too.

10 Example 2 continued: ratios at a price of 45

Share price 45. Same company, three definitions

P / (earnings + non-cash)	45 ÷ 4.20	**10.7**
P / CFO	45 ÷ 3.70	**12.2**
P / FCFE	45 ÷ 1.70	**26.5**

Peers must use the same definition

IN THE LESSON

- Let's price the same company at forty five dollars a share.
- Divide by earnings plus non cash charges, four twenty, and the share is on ten point seven times cash flow.
- Divide by cash flow from operations, three seventy, and it is on twelve point two.
- Divide by free cash flow to equity, one seventy, and it is on twenty six point five.
- One share, three very different answers, depending only on the definition.
- That is why you must never compare a price to cash flow figure from one source with another until you know which cash flow is underneath.
- Peers must be measured the same way, or the comparison is worthless.

11 EBITDA: a flow for everyone, not just shareholders

- EBITDA = before interest, tax, D&A
- Belongs to all capital providers, not just owners
- Trap: price per share ÷ EBITDA per share
- Interest paid is classified differently: IFRS vs US GAAP

IN THE LESSON

- There is one more trap in this outcome.
- E B I T D A stands for earnings before interest, tax, depreciation and amortisation.
- Because it is measured before interest, it is cash available to every provider of capital, lenders and shareholders together.
- Share price covers only the shareholders' slice.
- Pairing the two is like dividing a whole pizza's worth of toppings by a single slice.
- Suppose the same company has E B I T D A of six dollars fifty a share.
- Forty five divided by six fifty gives six point nine, which looks cheap, but the debt has been ignored.
- Also note that interest paid sits in different places in the cash flow statement under I F R S and U S G A A P, so cash flow figures are not always comparable between the two.

12 Enterprise value: the price of the whole business

- Buy the whole business, not just the shares
- Pay for equity, debt, preferred, minority owners
- Get the cash back
- Fair across different financing mixes

IN THE LESSON

- So what is the right partner for E B I T D A?
- Enterprise value.
- Imagine you wanted to buy a whole company.
- You would pay the shareholders for their shares.
- You would also take on the company's debt.
- You would honour the preferred shares, and you would owe something to outside owners of any subsidiaries the company only partly owns, called non controlling interests.
- But you would also get the company's cash and short term investments, so you could use it to pay down some of the bill.
- Enterprise value is that net price of the whole operating business.
- Because debt and equity are both counted, it does not depend on how the company is financed, which is exactly what makes it fair for comparing companies with different amounts of borrowing.

13 The enterprise value formula

EV = equity + debt + preferred + NCI − cash
EV/EBITDA

Equity value	Share price × shares
Debt, preferred, NCI	Added: other claims
Cash and investments	Subtracted: not operating
Enterprise value (EV)	Price of whole business

IN THE LESSON

- Here is the formula.
- Enterprise value equals the market value of equity, plus debt, plus preferred shares, plus non controlling interests, less cash and short term investments.
- Take the terms in turn.
- Equity value is the share price times the number of shares.
- The second box groups the other claims on the business: debt, preferred shares and non controlling interests.
- All of these are added.
- The cash box is subtracted, because cash is not part of operations and an acquirer would in effect get it back.
- The result is enterprise value.
- Divide it by E B I T D A and you have the ratio at the heart of this lesson, which is the price paid for each dollar of earnings available to all capital providers.

14 Example 3: enterprise value of two firms

$ million. Firm A: equity 3,000, debt 1,500, pref 200, NCI 100, cash 400

EV, Firm A	3,000+1,500+200+100−400	**4,400**
Firm B: equity 4,000	debt 200, cash 300	
EV, Firm B	4,000 + 200 − 300	**3,900**
Equity: B larger. EV: A larger		

IN THE LESSON

- Example three.
- Firm A has equity worth three thousand million dollars, debt of fifteen hundred, preferred shares of two hundred, non controlling interests of one hundred, and cash of four hundred.
- Add the first four, which gives forty eight hundred, and take off the four hundred of cash.
- Enterprise value is forty four hundred million.
- Firm B has equity of four thousand, debt of two hundred and cash of three hundred.
- Four thousand plus two hundred, less three hundred, is thirty nine hundred.
- Firm B has a much bigger market value of equity, but Firm A is the bigger business once borrowing is counted.
- Equity value alone would have given the wrong picture.

15 Example 3 continued: EV/EBITDA and the cash trap

EBITDA: Firm A 550, Firm B 520 ($ million)

Firm A: EV/EBITDA	4,400 ÷ 550	**8.0**
Firm B: EV/EBITDA	3,900 ÷ 520	**7.5**
Firm A, cash forgotten	4,800 ÷ 550	**8.7**

Forgetting the cash overstates the multiple

IN THE LESSON

- Now the multiples.
- Firm A has E B I T D A of five hundred and fifty.
- Enterprise value over E B I T D A is forty four hundred over five fifty, which is eight times.
- Firm B has E B I T D A of five hundred and twenty, so thirty nine hundred over five twenty is seven point five times.
- The two look almost the same.
- Now the classic slip.
- Suppose you forgot to deduct Firm A's cash.
- You would use forty eight hundred, and get eight point seven times.
- That overstates the multiple and makes the firm look dearer than it is.
- Always deduct the cash.
- And never divide enterprise value by an equity level number such as net income.
- The top and bottom must match.

16 Example 3 continued: what EBITDA hides

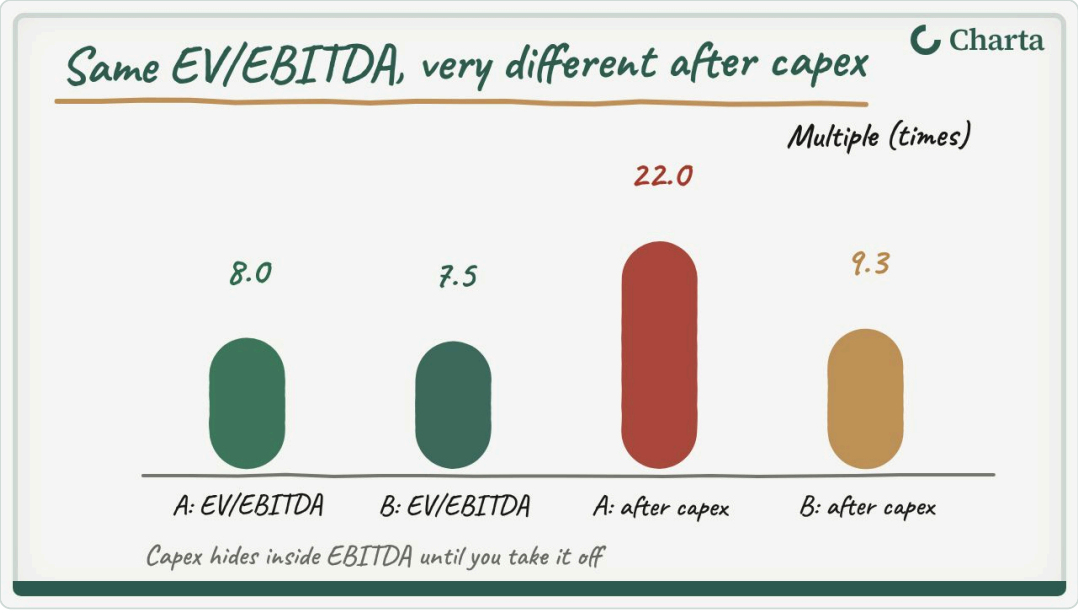
Capital spending: Firm A 350, Firm B 100 ($ million)

A: EBITDA − capex	550 − 350	**200**
B: EBITDA − capex	520 − 100	**420**
A: EV / (EBITDA − capex)	4,400 ÷ 200	**22.0**
B: EV / (EBITDA − capex)	3,900 ÷ 420	**9.3**

IN THE LESSON

- Here is the weakness.
- E B I T D A ignores capital spending.
- Firm A spends three hundred and fifty million on equipment each year, and Firm B only one hundred.
- That is a capital intensive business against a light one.
- Take capital spending off E B I T D A.
- Firm A is left with two hundred, and Firm B with four hundred and twenty.
- Now divide the same enterprise values.
- Firm A is forty four hundred over two hundred, twenty two times.
- Firm B is thirty nine hundred over four twenty, nine point three times.
- So the two firms that looked alike on E B I T D A are priced very differently once reinvestment is counted.
- Firm A's earnings need much more spending just to stay in place.

17 Same EV/EBITDA, very different after capex



IN THE LESSON

- Here is the same story as bars.
- Read from the left.
- The first two bars are enterprise value over E B I T D A: Firm A at eight times and Firm B at seven point five.
- Nearly level.
- The next two bars take capital spending off first.
- Firm A jumps to twenty two times and Firm B only to nine point three.
- The height of each bar is the number of dollars of enterprise value you pay for each dollar of earnings.
- The lesson of the picture is that a ratio that ignores capital spending can hide a heavy reinvestment burden.
- Enterprise value to sales is another enterprise value multiple, but it ignores profitability altogether, so use it with care.

18 Is EV/EBITDA a good multiple?

- For: suits different leverage
- For: positive when net income is not
- For: less affected by depreciation policy
- Against: ignores capex and working capital

IN THE LESSON

- Let's weigh it up.
- In its favour, it suits comparisons of companies with different borrowing, because enterprise value and E B I T D A are both measured before the lenders are paid.
- E B I T D A is usually positive even when net income is negative, so the ratio still works for a company having a poor year.
- And it is less distorted by different depreciation policies.
- Against it, E B I T D A ignores capital spending, which matters most for capital intensive firms.
- It ignores working capital needs.
- And it can flatter a company with heavy debt service, because the interest and repayments are not deducted.
- A sensible analyst uses it, then checks capital spending alongside.

19 Mistakes to avoid so far

- 1. Gap in multiples is not proof of mispricing
- 2. Outlier-inflated mean as benchmark
- 3. Mixed cash flow definitions across peers
- 4. EBITDA with price; EV with net income
- 5. Forgetting to subtract cash in EV

IN THE LESSON

- Time to stop and collect the common mistakes.
- One.
- Calling a share cheap from a low multiple without asking whether growth, risk or return on equity explain it.
- Two.
- Using the mean of the peer multiples when one outlier is inflating it.
- Three.
- Comparing a price to cash flow ratio across firms that define cash flow differently.
- Four.
- Pairing a firm level flow, such as E B I T D A, with the price per share.
- Five.
- Forgetting to subtract cash when you build enterprise value.
- Six.
- Dividing enterprise value by net income, an equity level number.
- In every case the top and the bottom of the ratio must belong to the same group of investors.

20 Across borders: a cheaper share?

- Foreign P/E a fifth lower: cheaper?
- Check 1: do they count earnings the same way?
- Check 2: do their economies justify different multiples?
- Inflation, rates, tax, governance, liquidity, growth

IN THE LESSON

- Now suppose Priya looks abroad.
- A company listed overseas trades on a price to earnings ratio a fifth lower than its domestic peer.
- Cheaper?
- Two other explanations must be ruled out before she says yes.
- The first is accounting.
- The two companies may count their earnings by different rules, so the earnings number underneath the ratio is not comparable.
- The second is the economy.
- Countries differ in inflation, interest rates, tax, corporate governance, market liquidity and growth prospects.
- All of these affect the return investors demand and so the multiple that is justified.
- High inflation, for example, tends to depress price to earnings ratios where companies cannot pass rising costs on to customers.

21 Accounting differences that move multiples

Difference: effect on the multiple

Area	Effect
Inventory (LIFO vs FIFO)	earnings, book value
Development costs	earnings today
Asset revaluation	book value
Fix: cash-based multiples	less accrual exposure

IN THE LESSON

- Here are the accounting differences that matter most.
- Inventory methods: I F R S forbids last in, first out, but U S G A A P allows it.
- When costs are rising, last in first out lowers reported earnings and inventory.
- Development costs: under I F R S some development spending can be capitalised as an asset, which lifts earnings today, while U S G A A P generally expenses it.
- Revaluation of assets: I F R S allows some assets to be revalued upward, which changes book value and so affects price to book.
- Because accruals are the source of these differences, multiples less exposed to accruals, such as price to cash flow or enterprise value to E B I T D A, can reduce the problem, though they do not remove it.

22 Example 4: putting two firms on one basis

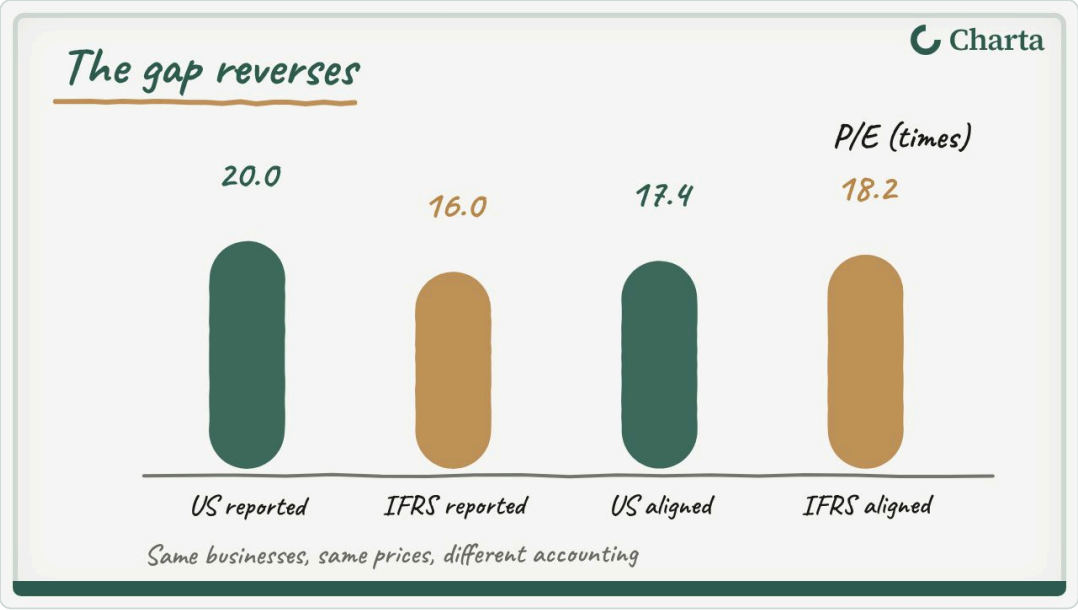
Both pay 25% tax. US peer at 60, IFRS company at 40

Reported P/Es	60 ÷ 3.00; 40 ÷ 2.50	**20.0 vs 16.0**
US peer on FIFO	3.00 + 0.60 × 0.75	**3.45**
US P/E, aligned	60 ÷ 3.45	**17.4**
IFRS EPS, expensed	2.50 − 0.40 × 0.75	**2.20**
IFRS P/E, aligned	40 ÷ 2.20	**18.2**

IN THE LESSON

- Example four.
- A US company is priced at sixty and reports earnings of three dollars a share using last in, first out.
- Its price to earnings is twenty.
- An I F R S company is priced at forty and reports two fifty, a ratio of sixteen.
- The foreign share looks twenty per cent cheaper.
- Now align them.
- On first in first out, the US company's profit before tax would be sixty cents higher.
- After twenty five per cent tax, that adds forty five cents, so earnings are three forty five, and the ratio is seventeen point four.
- The I F R S company capitalised forty cents of development spending.
- Expensed, that costs thirty cents after tax, so earnings are two twenty, and the ratio is eighteen point two.

23 The gap reverses



IN THE LESSON

- Here is the result as bars.
- Read the left pair first.
- As reported, the I F R S company is on sixteen times earnings, against twenty for its American peer.
- A screen would flag it as cheap.
- Now the right pair, after both are put on the same accounting basis.
- The I F R S company is on eighteen point two, and the American peer on seventeen point four.
- The gap has reversed.
- Nothing about either business changed, only the rules used to count earnings.
- The foreign share's apparent discount came entirely from accounting.
- The trap, then, is to conclude that a foreign share is cheap only because its price to earnings is below a domestic peer's.

24 Momentum: which way is the evidence moving?

- Valuation: where should the ball settle?
- Momentum: is the ball still moving?
- Earnings surprises and relative strength
- Popular screens, but not valuation theory

IN THE LESSON

- Everything so far asked what a share is worth.
- Momentum indicators ask a different question: which way is the evidence moving?
- They rest on the observation that trends, in earnings news and in prices, tend to persist for a while.
- Think of a rolling ball.
- Valuation tells you where the ball ought to settle.
- Momentum tells you it is still moving.
- Two families matter.
- The first is about earnings surprises: how far reported earnings beat or missed what analysts expected.
- The second is relative strength: how a share's return compares with its own past or with a benchmark.
- Honest warning: these indicators are not grounded in valuation theory, yet they are widely used to screen for shares.

25 Unexpected earnings and SUE

UE = actual EPS − expected EPS
SUE = UE ÷ SD of past UE

Unexpected earnings	Actual less expected EPS
Std deviation of past surprises	How erratic surprises are
SUE	Surprise in units of normal

IN THE LESSON

- Here are the measures.
- Unexpected earnings, also called the earnings surprise, is reported earnings per share minus expected earnings per share.
- The expectation is usually the average of analysts' forecasts.
- A raw surprise can mislead, so we scale it.
- One way divides by the standard deviation of analysts' forecasts.
- The other, called standardised unexpected earnings, or S U E, divides by the standard deviation of the company's own past surprises.
- Take the terms.
- Unexpected earnings is the gap.
- The standard deviation measures how erratic past surprises have been.
- And S U E is the answer: how big this surprise is compared with what is normal for that company.

26 Example 5: a surprise and its strength

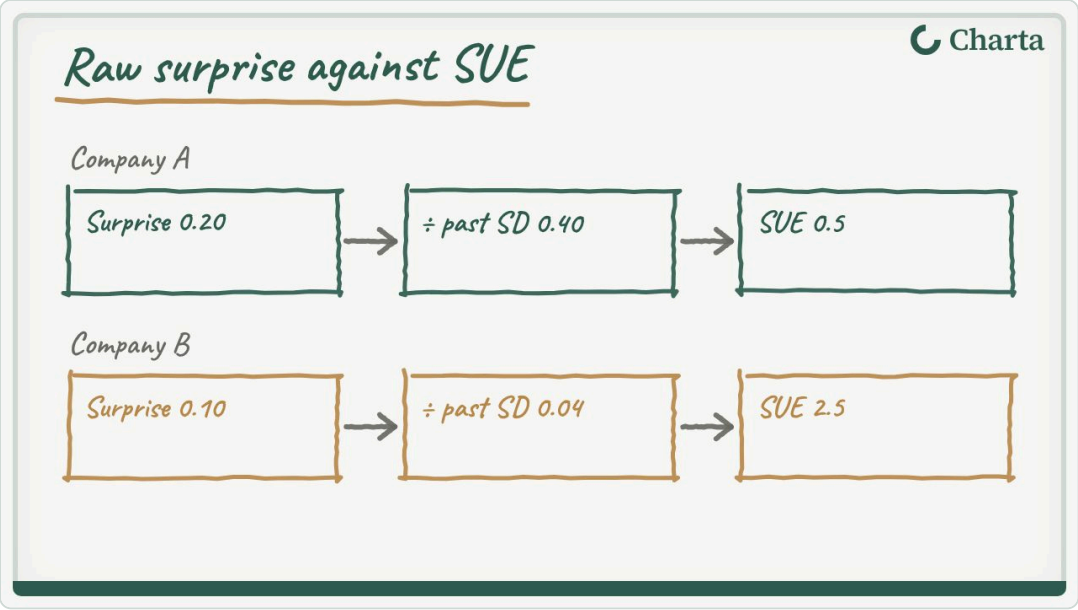
B: EPS 1.10 vs consensus 1.00; SD forecasts 0.05; SD past 0.04

Unexpected earnings	1.10 − 1.00	**0.10**
Scaled by forecast SD	0.10 ÷ 0.05	**2.0**
SUE	0.10 ÷ 0.04	**2.5**
Relative strength	1.18 ÷ 1.06	**1.11**

IN THE LESSON

- Example five.
- Company B reports earnings per share of one dollar ten against a consensus of one dollar.
- Unexpected earnings are ten cents.
- The standard deviation of analysts' forecasts was five cents, so the scaled surprise is two.
- The standard deviation of B's own past surprises is only four cents, so S U E is ten over four, two and a half.
- That is large for this company.
- Now relative strength.
- Over the last year B's shares returned eighteen per cent, against six per cent for the index.
- Divide one point one eight by one point oh six, and you get one point one one: eleven per cent ahead of the index.
- A strong surprise and a strong trend, so B would screen well on momentum.

27 Raw surprise against SUE



IN THE LESSON

- Now the trap, as a picture.
- Compare the two rows.
- Company A beat expectations by twenty cents, twice as much as Company B's ten.
- On the raw numbers A looks the more exciting.
- But A's past surprises have a standard deviation of forty cents, so a twenty cents beat is well within its normal noise.
- Divide, and its S U E is only a half.
- Company B's surprises are usually tiny, with a standard deviation of four cents.
- So its ten cents beat is two and a half times its usual surprise, a real signal.
- Comparing raw surprises across companies whose earnings differ greatly in volatility is the mistake a question is waiting for.

28 Averaging multiples: which average?

- One giant in the room bends the average
- A depressed-earnings P/E can be huge
- Four averages: mean, median, harmonic, weighted
- Each reacts differently to the outlier

IN THE LESSON

- The last question.
- Once you have a group of peer multiples, which average do you take?
- Think of a classroom of five pupils whose heights are all about the same, plus one giant who is a basketball professional visiting for the day.
- The arithmetic mean height jumps, and describes nobody.
- Multiples have giants too.
- A company with temporarily depressed earnings can show a price to earnings of forty nine, because the bottom of the ratio is tiny.
- That one number can bend the average.
- So there are four ways to summarise: the arithmetic mean, the median, the harmonic mean and the weighted harmonic mean.
- Each reacts differently to the giant.
- Let's meet them.

29 Harmonic and weighted harmonic mean

Harmonic mean = n ÷ Σ(1/x)
Weighted = 1 ÷ Σ(w/x)

Number of multiples (n)	How many multiples
Each multiple (x)	A peer's multiple
Weight (w)	Holding weight, sums to 1
Harmonic mean	The average multiple

IN THE LESSON

- Here are the two new formulas.
- The harmonic mean is the number of multiples divided by the sum of their reciprocals.
- A reciprocal is one over the multiple, and for a price to earnings ratio that is the earnings yield.
- So the harmonic mean averages the earnings yields and turns the answer back into a multiple.
- Big multiples have tiny reciprocals, so they get little weight.
- The weighted version replaces the plain count with weights.
- It is one divided by the sum of each weight over its multiple.
- Use the weights equal to the portfolio's holdings, and you get the correct multiple for a portfolio or an index.

30 Example 6: three averages of five peers

Peer P/Es: 10, 12, 14, 15, 49

Arithmetic mean	(10+12+14+15+49) ÷ 5	**20.0**
Median	middle of five	**14.0**
Sum of reciprocals	0.1+0.0833+0.0714+0.0667+0.0204	**0.3418**
Harmonic mean	5 ÷ 0.3418	**14.6**

IN THE LESSON

- Example six.
- Five peers trade on price to earnings ratios of ten, twelve, fourteen, fifteen and forty nine.
- The forty nine reflects a year of depressed earnings.
- The arithmetic mean is one hundred divided by five, which is twenty.
- That is higher than four of the five peers, and seven and a quarter above the mean of the other four, which is twelve and three quarters.
- The median is fourteen, the middle value, which pays no attention to how extreme the forty nine is.
- For the harmonic mean, the reciprocals add to point three four one eight.
- Five divided by that is fourteen point six.
- The forty nine added only point oh two oh four to the sum.

31 Example 6 continued: a fund's own multiple

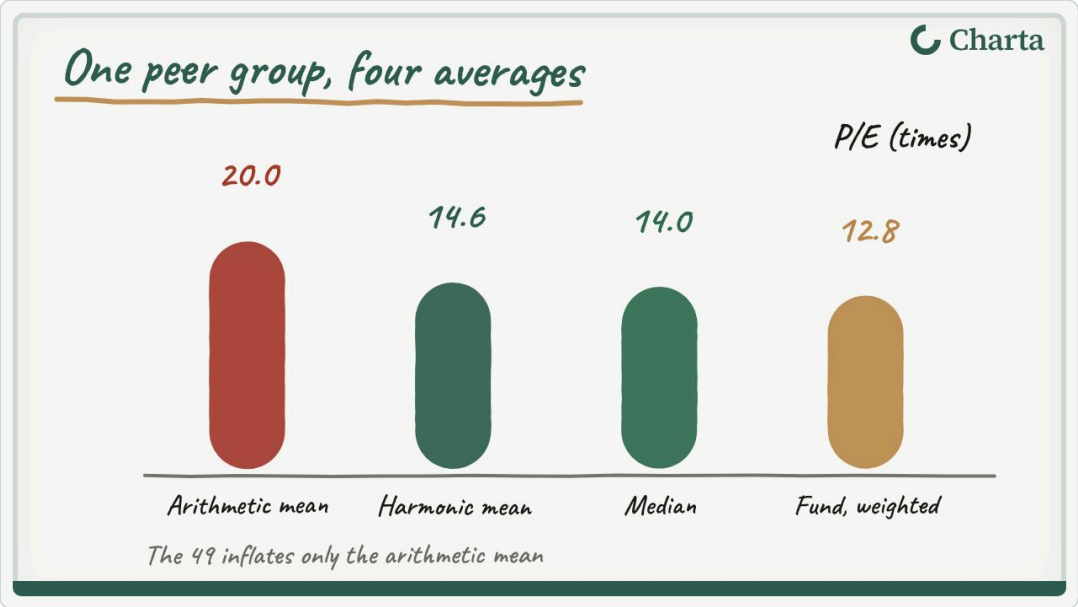
Weights: 30%, 30%, 20%, 10%, 10%

Σ(w ÷ x)	0.03+0.025+0.0143+0.0067+0.0020	**0.0780**
Fund P/E, weighted harmonic	1 ÷ 0.0780	**12.8**
Share at 16 vs mean 20	16 ÷ 20 − 1	**−20%**
Share at 16 vs median 14	16 ÷ 14 − 1	**+14%**

IN THE LESSON

- Now a fund holds these five shares in weights of thirty, thirty, twenty, ten and ten per cent.
- What is the fund's price to earnings?
- Do not take the weighted arithmetic average.
- Compute each weight over its multiple: point three over ten, point three over twelve, point two over fourteen, point one over fifteen and point one over forty nine.
- These add to point oh seven eight.
- One divided by that is twelve point eight.
- The fund's multiple is lower because it holds more of the cheaper shares.
- Finally the judgement.
- A share on sixteen looks twenty per cent cheap against the arithmetic mean of twenty.
- Against the median or the harmonic mean it looks dear.
- The verdict came from the average, not the share.

32 One peer group, four averages



IN THE LESSON

- Here are all four answers on one board.
- Read from left to right and watch the bars fall.
- The arithmetic mean treats the forty nine like any other number, and the average is pulled up to twenty.
- The harmonic mean averages the earnings yields, where forty nine is a tiny two per cent, so it gives fourteen point six.
- The median gives fourteen.
- The fund's own multiple, weighted by holdings, is twelve point eight.
- The harmonic mean can be distorted the other way, by a very low multiple, so it is wise to report the median beside it.
- The trap is averaging price to earnings ratios arithmetically when one or two very high multiples inflate the benchmark and make every other share look cheap.

33 Recap: six outcomes in one board

Six ideas for valuing with multiples

Price to earnings (P/E)	Peers, then explain the gap
Debt, preferred, NCI	Match flow to claim
Enterprise value (EV)	EV = all claims − cash
SUE	Scale the surprise
Harmonic mean	Median or harmonic

IN THE LESSON

- Let's pull it together.
- Comparables: choose peers, find a central value, compare, then explain the gap with growth, risk and return on equity.
- Cash flow: know what each definition leaves out, and never pair E B I T D A with price per share.
- Enterprise value: equity plus debt, preferred and non controlling interests, less cash, over E B I T D A, and watch capital spending.
- Cross border: accounting and economies, so align the earnings first.
- Momentum: unexpected earnings, scaled, and S U E, plus relative strength.
- And averages: the arithmetic mean is pulled up by outliers, the median resists them, the harmonic mean favours small values, and the weighted harmonic mean suits a portfolio.

34 Try three quick questions

1. Lowest growth, lowest P/E in the group. Cheap?
   Answer: Not necessarily: fundamentals may explain the gap.
2. Equity 800, debt 300, cash 100, EBITDA 100: EV/EBITDA?
   Answer: EV = 1,000, so 10.0 times.
3. Harmonic mean of P/Es of 10, 20 and 40?
   Answer: 3 ÷ 0.175 = 17.1, below the arithmetic 23.3.

IN THE LESSON

- Now try three quick questions.
- Pause after each one.
- Question one.
- A subject trades on a lower price to earnings than its peer median, but its growth is the lowest in the group.
- Is it cheap?
- The answer is not necessarily.
- Fundamentals may explain the gap, and a growth adjusted ratio could show it as dear.
- Question two.
- Equity is eight hundred, debt three hundred, cash one hundred, E B I T D A one hundred.
- What is enterprise value over E B I T D A?
- Enterprise value is eight hundred plus three hundred less one hundred, a thousand, so the ratio is ten.
- Question three.
- Three shares trade on ten, twenty and forty times earnings.
- What is the harmonic mean?
- Three over point one seventy five, which is seventeen point one.