CHARTA LEARNING · LESSON NOTES

Industry analysis and choosing a valuation method

Equity Investments

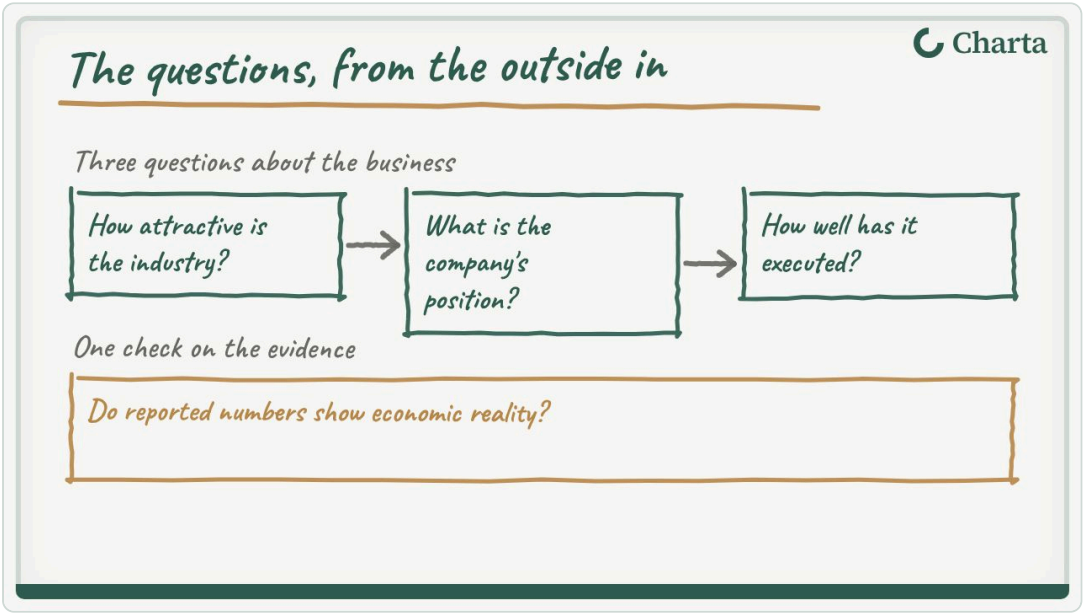
1 A story: the coffee cart by the station

- A cart grew 20% a year with no rivals nearby
- A line through the past says: more of the same
- Then a chain opens next door
- Always ask why the past happened

IN THE LESSON

- Imagine a coffee cart outside a busy railway station.
- For three years it has been the only place to buy coffee there, and its sales have grown twenty per cent every year.
- An analyst draws a straight line through those years and forecasts more of the same for another decade.
- Then the station owner lets a famous chain open a shop right next to the cart.
- The growth was never about the cart.
- It was about having no competition.
- The forecast looked at the past and forgot to ask why the past happened.
- That is the lesson of this first section.
- Before you forecast a single number, understand the industry, the company's place in it, and whether its advantage will last.

2 The questions, from the outside in



IN THE LESSON

- Here are the questions an industry and competitive analysis should answer, in order.
- Start on the left with the industry.
- Is it a place where companies can earn good returns over the long run, or does competition compete profits away?
- Next, the company's competitive position.
- Does it win by being the cheapest, by being different, or by serving a small niche well?
- Then execution.
- Has the company actually delivered its strategy, judged against its own history and against its rivals?
- Finally, the bottom row.
- A cross-check on all three: can you trust the financial statements?
- Reported earnings are the raw material for everything else, and they may not show the economic reality.

3 Question one: how attractive is the industry?

- Rivalry between existing competitors
- Threat of new entrants
- Threat of substitutes
- Power of suppliers, and power of buyers

IN THE LESSON

- A quick reminder from Level One, then we apply it.
- An industry is attractive when the businesses in it can earn returns above their cost of capital for a long time.
- Five forces decide that.
- Rivalry between existing competitors: many similar firms and slow growth push prices down.
- The threat of new entrants: if it is easy to get in, high profits attract newcomers who compete them away.
- The threat of substitutes: a different product that does the same job caps what you can charge.
- And the bargaining power of suppliers, who can raise your costs, and of buyers, who can push down your prices.
- The stronger these forces, the weaker the long-run profits for everyone in the industry.

4 Scoring a made-up industry

Harlow Carriers: pressure from each force

Force	Pressure
Rivalry	High
New entrants	High
Substitutes	Medium
Suppliers, buyers	Low, High

IN THE LESSON

- Let's score an industry the way a vignette might.
- Harlow Carriers is a regional road freight business.
- Rivalry is high: dozens of carriers sell almost the same service.
- New entrants are a high threat, because a used lorry and a licence are enough to start.
- Substitutes are a medium threat, since some customers can switch to rail.
- Supplier power is low, because fuel and drivers come from many sources.
- Buyer power is high, because big retailers re-tender contracts every year and switching costs them very little.
- Read the column: four forces lean against the carriers.
- However well Harlow is run, this industry will struggle to deliver strong long-run returns.
- A forecast of widening margins needs a very good reason.

5 Questions two and three: position and execution

- Low-cost producer: wins on price or margin
- Differentiator: customers pay extra for it
- Focused niche: one segment served very well
- Execution: judged against its history and rivals

IN THE LESSON

- Industry attractiveness sets the ceiling.
- Inside it, companies differ.
- The second question is the company's competitive position, and there are three classic answers.
- The low-cost producer wins by having a lower cost per unit than anyone, so it can undercut rivals or keep better margins.
- The differentiator offers something customers value enough to pay extra for, such as quality, brand or service.
- The focused player concentrates on one niche and serves it better than the generalists.
- The third question is execution.
- A strategy on paper proves nothing.
- The analyst checks whether margins, market share and returns have really improved, comparing the company with its own history and with its rivals.
- Good strategy and poor delivery is a warning sign.

6 Forecasting when the advantage fades

E(n) = E0 × (1 + g1)^t × (1 + g2)^(n − t)

Earnings today	Starting point
Growth while it lasts	High rate, while protected
Years the advantage lasts	Then growth drops
Growth after it fades	Lower rate, for the rest

IN THE LESSON

- Now let's turn the idea into arithmetic.
- Suppose a company has an advantage that will last for a number of years, and after that it fades.
- We forecast earnings in two stages.
- Start from today's earnings.
- Grow them at the high rate for as many years as the advantage lasts.
- Then grow them at a lower, fade rate for the remaining years.
- The formula multiplies today's earnings by one plus the high rate, raised to the number of years the advantage lasts, and then by one plus the fade rate, raised to the years that remain.
- The fade rate is itself a judgement, and another analyst may choose a different one.
- What matters is that the analyst has asked the question at all.

7 Example 1: extrapolate or fade?

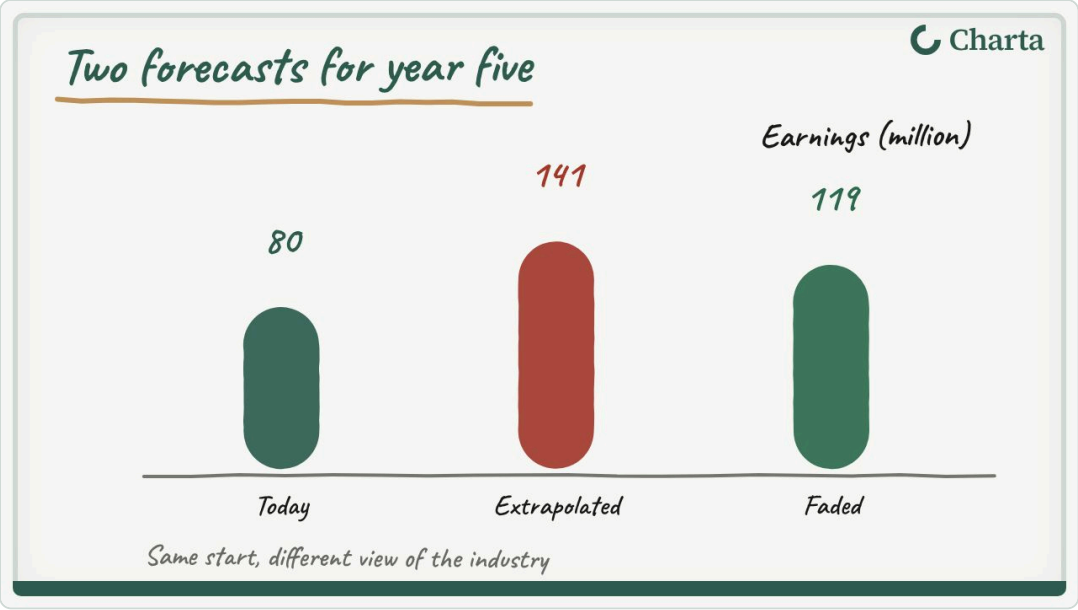
Ferrowick Pharma: earnings 80m, growing 12%

Extrapolated, year 5	80×1.12^5	**141m**
Advantage ends	patent expires in year 3	**year 3**
Faded, year 5	$80 \times 1.12^3 \times 1.03^2$	**119m**
Overstatement	141 ÷ 119 − 1	**18%**

IN THE LESSON

- Our first example.
- Ferrowick Pharma makes a drug protected by a patent.
- Earnings today are eighty million and have grown twelve per cent a year.
- The patent expires in three years, and then generic rivals may enter.
- A lazy forecast extrapolates twelve per cent for five years.
- Eighty times one point one two to the fifth power is about one hundred and forty-one million.
- An analyst who has done the industry work asks what changes.
- Entry is no longer blocked.
- So she grows earnings at twelve per cent for three years and at three per cent for the last two.
- That gives about one hundred and nineteen million.
- The extrapolated forecast is eighteen per cent too high.

8 Two forecasts for year five



IN THE LESSON

- Here are the three numbers as bars.
- The first bar is earnings today, eighty million.
- The tall bar in the middle is the extrapolated forecast for year five, about one hundred and forty-one million.
- The bar on the right is the faded forecast, about one hundred and nineteen million.
- Both forecasts start from the same eighty and both use twelve per cent for the early years.
- The gap between the two tall bars comes entirely from one judgement: whether the advantage survives the patent.
- Since a share's value comes from future earnings, an eighteen per cent error here flows straight through into the valuation.
- The cost of skipping the industry question is a model that looks precise and is simply wrong.

9 The fourth check: are the numbers real?

- Profit rising while cash flow falls is a warning
- Early revenue, or costs pushed onto the balance sheet
- Repeating one-off items
- Trap: extrapolating growth without asking why

IN THE LESSON

- There is one more part of the analysis, and it runs alongside the other three.
- The analyst checks the quality of the financial statements.
- Everything above, from margins to growth to returns, was judged from reported figures.
- If those figures flatter the company, every conclusion is built on sand.
- So the analyst looks for warning signs.
- Revenue recognised early.
- Costs moved onto the balance sheet.
- Profit rising while operating cash flow falls.
- Large one-off items that always seem to repeat.
- The trap is the one from our story: extrapolating past growth without asking whether the advantage that produced it will survive.
- Check the industry, the position, the execution, and the evidence, in that order, before you trust a forecast.

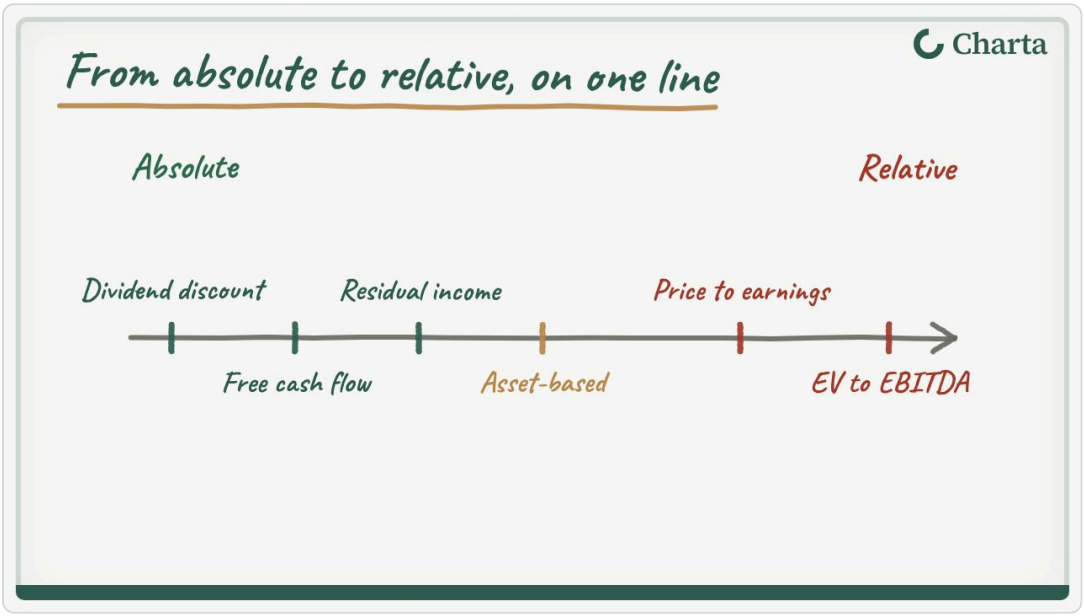
10 A story: two ways to price a flat

- Method one: what will the flat itself earn?
- Method two: what did similar flats sell for?
- One is absolute, one is relative
- They can disagree, and the gap teaches you

IN THE LESSON

- Now to the second outcome.
- Suppose you want to know what a flat is worth.
- There are two natural ways.
- The first is to look at the flat itself: how much rent will it earn over the years, and what is that stream of rent worth today, after allowing for risk?
- The second is to look next door: what did similar flats in the same street sell for recently?
- The first method values the asset from its own earnings.
- The second prices it against comparable assets.
- Finance uses the same two families.
- The first is called absolute valuation, the second relative valuation.
- They often give different answers, and understanding why is more useful than either answer alone.

11 From absolute to relative, on one line



IN THE LESSON

- Here is the family of models on a line.
- On the left is absolute valuation, which estimates intrinsic value from the asset's own characteristics.
- The first three points are the classic examples.
- The dividend discount model discounts expected dividends.
- The free cash flow model discounts the cash the business could pay out.
- The residual income model adds the present value of profits earned above the charge for equity capital to the book value.
- The fourth point, the asset-based model, values a company at the market value of its assets less its liabilities.
- Many students file it under relative, but it is absolute.
- On the right is relative valuation: price to earnings, and enterprise value to E B I T D A, which price a share against similar companies.

12 What each family can and cannot tell you

- Absolute: value from its own cash flows or assets
- Relative: price against similar assets, via multiples
- Asset-based models are absolute: assets less liabilities
- Relative answers only: cheap or dear against the peers

IN THE LESSON

- Let's be precise about each.
- An absolute model builds intrinsic value from the asset's own expected cash flows, or from its assets less its liabilities.
- Its answer stands on its own, but it depends on estimates of growth and the required return, and small changes in those can move the answer a lot.
- A relative model rests on one idea: similar assets should sell at similar prices.
- You take a multiple from comparable companies and apply it to your company.
- It is quick and grounded in the market, but it carries a built-in weakness.
- It can only say a share is cheap or expensive compared with the benchmark.
- If the whole peer group is overpriced, a share can look cheap and still be dear.

13 The simplest absolute model: constant growth

V0 = D1 ÷ (r − g)

Next year's dividend	Cash received next year
Required return	What investors demand
Growth, forever	Must be below r
Value today	What the share is worth

IN THE LESSON

- To compare the families we need one absolute formula, and the simplest is the constant growth dividend model.
- The value today equals next year's dividend, divided by the required return minus the growth rate.
- Take the pieces.
- The dividend is cash the owner expects to receive next year.
- The required return is what investors demand for bearing this risk.
- The growth rate is how fast the dividend is expected to rise, forever.
- The formula only works when the required return is bigger than the growth rate.
- It also reveals how sensitive the answer is.
- When the gap between the two rates is small, the denominator is small and the value jumps.
- We will see that in the next example.

14 Example 2: cheap against peers, dear on its own

Harbor Textiles: price 36, EPS 2.40, peers on 18×

Forward P/E	36 ÷ 2.40	**15×**
Relative value	18 × 2.40	**43.20**
Absolute value	1.20 ÷ (0.09 − 0.04)	**24.00**
Justified P/E	24 ÷ 2.40	**10×**

IN THE LESSON

- Our second example.
- Harbor Textiles trades at thirty-six.
- Next year's earnings are forecast at two forty a share, so it trades at fifteen times forward earnings.
- Its peers trade at eighteen times.
- The relative answer: eighteen times two forty is forty-three twenty, so the share looks twenty per cent cheap.
- Now the absolute answer.
- Harbor pays half its earnings, a dividend of one twenty.
- Growth is four per cent and the required return is nine per cent.
- Value equals one twenty divided by five per cent, which is twenty-four.
- That is a third below the price, and it implies a fair multiple of only ten times earnings.
- The peers at eighteen times look expensive, not Harbor cheap.

15 One share, three numbers



IN THE LESSON

- Here are the three numbers side by side.
- The left bar is the market price, thirty-six.
- The middle bar is the relative value, forty-three twenty, which says buy.
- The right bar is the absolute value, twenty-four, which says sell.
- Which is right?
- Both are right about what they measure.
- The relative model tells you Harbor is priced below its peers.
- It cannot tell you whether the peers are sensibly priced.
- The absolute model supplies that reference point, but it leans on a growth rate and a required return that are estimates too.
- Notice that if growth were six per cent instead of four, the value would jump to forty.
- That is the sensitivity we warned about.
- When the two models disagree, go looking for the cause.

16 Use them together

- Relative says cheap: ask what the peers' multiple assumes
- Absolute says dear: check your growth and required return
- A disagreement is a prompt to investigate
- Trap: cheap against peers is not cheap in absolute terms

IN THE LESSON

- So which family should you trust?
- The usual answer is that they complement each other.
- Analysts often use one to cross-check the other.
- If a relative model says a share is cheap, ask what growth and risk the peers' multiple implies, and whether you believe it.
- If an absolute model says a share is worth far less than its price, ask whether your growth and required return are too gloomy, or whether the market is genuinely too optimistic.
- A disagreement is not a failure, it is a prompt to investigate.
- The trap is to read a relative result as an absolute one.
- A share that is cheap relative to peers has not been shown to be cheap.
- It has only been shown to be cheaper than its neighbours.

17 Mid-lesson check: four mistakes so far

- 1. Extrapolating growth without asking why
- 2. Judging the industry by how well the firm is run
- 3. Cheap against peers read as cheap, full stop
- 4. Asset-based models filed under relative valuation

IN THE LESSON

- Let's pause and collect the mistakes that cost marks so far.
- One.
- Extrapolating past growth without asking whether the advantage that produced it will survive.
- The coffee cart and the patent are the same error.
- Two.
- Judging an industry by the company's quality, when the five forces set the ceiling whoever runs the company.
- Three.
- Reading a share that looks cheap against its peers as cheap in absolute terms, forgetting that the whole peer group may be overpriced.
- Four.
- Filing the asset-based model under relative valuation.
- It is absolute, because it values the company from its own assets less liabilities.
- If you can avoid these four, the first two outcomes are safe.
- Next, a group made of several businesses.

18 A story: three businesses, one share price

- A group owns software, packaging and haulage
- One multiple fits none of them well
- Value each part on its own, then add
- Compare the total with the group's share price

IN THE LESSON

- Now the third outcome.
- Picture a group that owns a software business, a packaging factory and a road haulage firm.
- These are three different companies sharing one share price and one head office.
- If you valued the group with a single multiple, you would apply one number to a mix that no single number fits.
- Software grows fast and deserves a high multiple.
- Haulage is slow and risky and deserves a low one.
- The cleaner way is to value each business separately, using the model that suits it, then add the pieces together.
- That is sum of the parts valuation.
- And very often the total comes out above where the whole group trades.
- That gap is the conglomerate discount.

19 Sum of the parts, step by step

Equity = sum of segments − net debt − other claims
Discount = 1 − price ÷ sum-of-parts value

Segment values	Each at its own multiple
Net debt	Debt less cash
Other claims	Pension deficit, minorities
Equity value	What shareholders own
Conglomerate discount	Gap below the total

IN THE LESSON

- Here is the method as a formula.
- First, value each segment with a multiple or a cash flow model suited to it, and add them up.
- Because the segment values are enterprise values, belonging to lenders and shareholders together, you then deduct net debt and any other corporate claims that rank ahead of the shareholders, such as a pension deficit or a minority stake held by outsiders.
- Unallocated head office costs are valued too, as a negative item.
- What is left is equity value.
- Divide by the number of shares for a value per share.
- The conglomerate discount is then one minus the market price divided by the sum of the parts value.
- It is the share of that value the market refuses to pay for.

20 Example 3: adding up Vantor Group

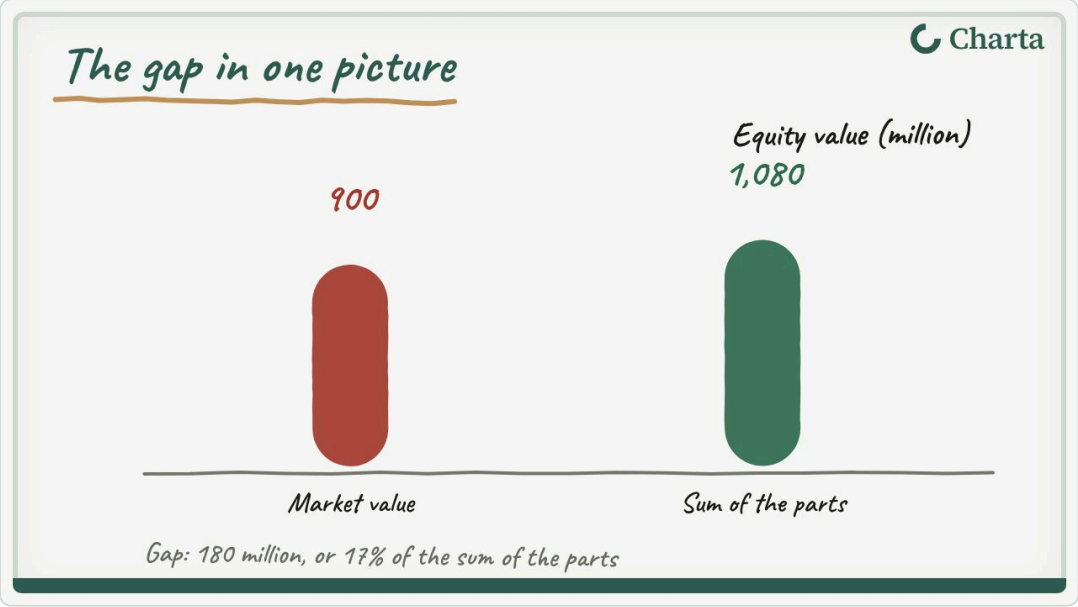
Vantor Group, in millions of dollars

Software	60 × 14	**840**
Packaging	120 × 6	**720**
Haulage and head office	90 × 5 − 30 × 6	**270**
Equity	1,830 − 700 − 50	**1,080**
Per share; discount	18.00 vs 15.00	**17%**

IN THE LESSON

- Our third example.
- Vantor Group has three segments.
- Software has E B I T D A of sixty million and deserves fourteen times, so it is worth eight hundred and forty.
- Packaging has one hundred and twenty million at six times, which is seven twenty.
- Haulage has ninety million at five times, which is four fifty.
- Head office costs thirty million a year, valued at six times, a negative one hundred and eighty.
- The total enterprise value is eighteen thirty.
- Now take away net debt of seven hundred and a pension deficit of fifty.
- Equity is ten eighty.
- With sixty million shares that is eighteen dollars a share.
- The shares trade at fifteen, so the market is paying seventeen per cent less than the parts are worth.

21 The gap in one picture



IN THE LESSON

- Here is the same result as two bars.
- The left bar is what the market pays for the equity: fifteen dollars a share times sixty million shares is nine hundred million.
- The right bar is the sum of the parts equity, one thousand and eighty million.
- The difference, one hundred and eighty million, is the conglomerate discount in money, which is seventeen per cent of the right-hand bar.
- Notice what drives the sum of the parts.
- The software multiple of fourteen times is far higher than anything the group trades on as a whole.
- Had we used one average multiple, we would have undervalued software and overvalued haulage.
- Also notice that every deduction mattered.
- Forget the pension deficit and equity is overstated by exactly fifty million.

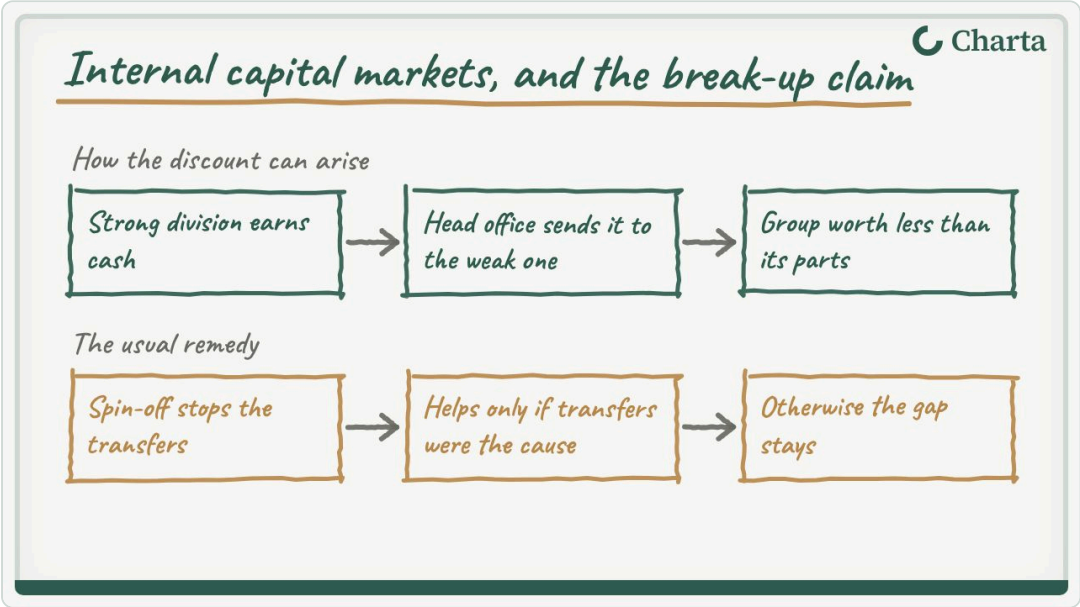
22 Why would the market pay less than the parts?

- Inefficient internal capital markets: strong subsidise weak
- Endogenous factors: weak firms diversify to escape trouble
- Research measurement errors: flawed peer comparisons
- Each explanation points to a different conclusion

IN THE LESSON

- Why does the discount exist?
- There are three proposed explanations.
- First, inefficient internal capital markets.
- Head office decides where the group's cash goes.
- If it lets strong divisions subsidise weak ones, instead of backing the best projects, value is lost.
- Second, endogenous factors.
- The causation may run the other way.
- Struggling companies often diversify to escape trouble, so they would have been cheap anyway.
- The discount describes the businesses, not the structure.
- Third, research measurement errors.
- The discount may be an illusion created by flawed comparisons, for example peer multiples taken from companies that are not truly comparable, or from companies that are in better shape.
- Keep all three in mind, because each leads to a different conclusion.

23 Internal capital markets, and the break-up claim



IN THE LESSON

- Let's look at the first explanation as a chain of cause and effect.
- On the top row, a strong division earns good cash.
- Head office, instead of paying it out or investing in the best opportunity, sends it to a weak division with poor prospects.
- The group as a whole is then worth less than the strong and weak parts valued separately.
- That is a real way to destroy value.
- On the bottom row is the usual remedy.
- A spin-off stops the transfer.
- But read the middle box carefully.
- It only helps if the transfers were the cause of the discount.
- If the discount came from the other two explanations, splitting the group changes little, and it costs money to do.
- A break-up does not release value automatically.

24 Example 4: how much would a break-up release?

Vantor Group: the 3.00 gap per share, split by cause

Cross-subsidies	fixable by a break-up	**1.20**
Measurement error	peer multiples too high	**1.00**
Endogenous	weak units bought to escape	**0.80**
Gain after costs	1.20 − 0.30	**0.90**
New price; discount left	15.90 vs 18.00	**12%**

IN THE LESSON

- Our fourth example, back with Vantor.
- The gap is three dollars a share: eighteen less fifteen.
- Suppose the analyst's investigation attributes it as follows.
- One twenty is from head office cross-subsidies.
- One dollar is from peer multiples that were too generous, a measurement error.
- And eighty cents reflects the fact that the weakest businesses were bought to escape trouble.
- A spin-off can only fix the first.
- After separation costs of thirty cents a share, the gain is ninety cents, so the shares might reach fifteen ninety.
- The remaining discount is still about twelve per cent.
- The headline discount of seventeen per cent is the start of the argument, not the conclusion.

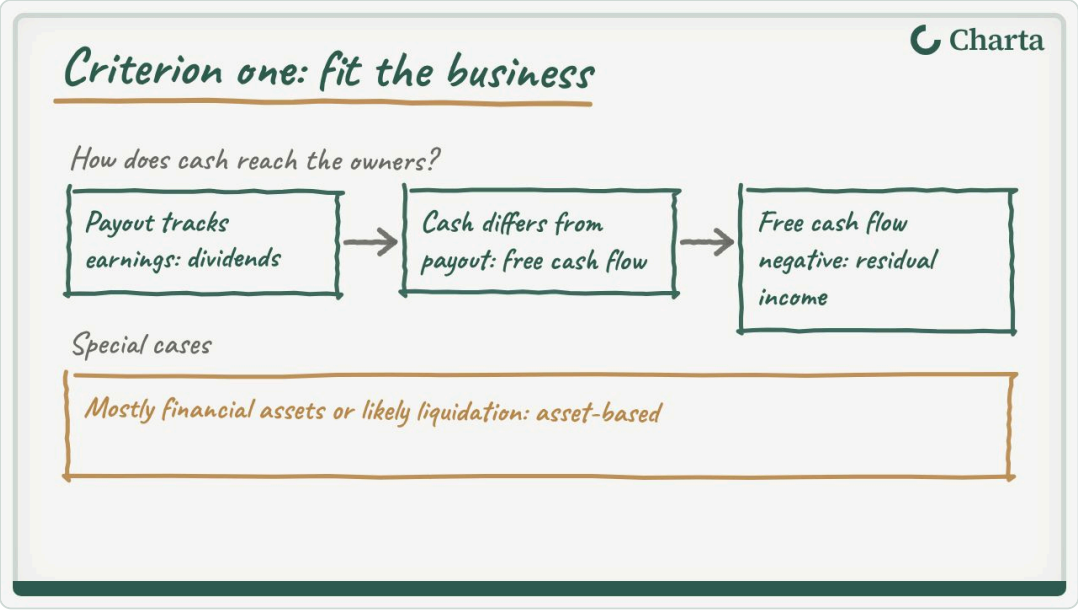
25 A story: choosing the right tool

- A hammer is wrong for a screw
- Criterion one: fit the business
- Criterion two: the quality of the data
- Criterion three: the purpose of the valuation

IN THE LESSON

- Now the fourth outcome.
- A carpenter does not use one tool for every job.
- A hammer is perfect for nails and terrible for screws.
- Valuation models are the same.
- The skill is choosing the right one for the job in front of you, and that choice often matters more than any single input.
- There are three criteria.
- First, the model should fit the business: how does this company generate cash, and how does it hand that cash to its owners?
- Second, it should suit the quality of the data, because a model is only as reliable as its inputs.
- Third, it should match the purpose of the valuation, because different buyers want different answers.
- Let's take them one at a time.

26 Criterion one: fit the business



IN THE LESSON

- The top row sorts companies by how cash reaches the owners.
- If dividends are steady and move with earnings, the dividend discount model is a fair measure of what holders receive.
- If the company generates cash that differs from what it pays out, or pays nothing at all, use a free cash flow model, which measures what it could pay.
- And if free cash flow will be negative for years while the company invests heavily, residual income can still work, because it is built from earnings and book value.
- The bottom row is for the special case.
- Where the assets are mostly financial, or the company is likely to be liquidated, an asset-based approach fits best.
- Each choice comes with a condition attached.

27 Four companies, four models

Match the company to the model

Company	Model
Steady utility	Dividend
Cash-rich, no payout	Free cash flow
Fast-growing retailer	Residual income
Holder of shares	Asset-based

IN THE LESSON

- Let's match some made-up companies to models.
- A utility that has paid about seventy per cent of its earnings for a decade: the dividend discount model, because dividends track earnings.
- A software firm that pays no dividend but produces strong, positive free cash flow: a free cash flow model, because the cash is generated even though it is not paid out.
- A retailer expanding so fast that free cash flow will be negative for five years: residual income, as long as its accounting is sound, since it starts from book value and earnings.
- And a holding company whose assets are mainly listed shares: an asset-based approach, valuing the holdings at market prices less liabilities.
- None of these choices is made by habit.

28 Criterion three: the purpose changes the answer

Calder Foods: r = 10%, g = 4%

Dividend next year	set by current board	**1.20**
Free cash flow to equity	cash the firm generates	**2.10**
Minority view	1.20 ÷ (0.10 − 0.04)	**20.00**
Control view	2.10 ÷ (0.10 − 0.04)	**35.00**

IN THE LESSON

- Our fifth example shows purpose.
- Calder Foods pays a dividend of one twenty next year, but earns far more: free cash flow to equity is two ten a share.
- Required return is ten per cent and growth is four per cent.
- A minority investor cannot change the payout, so the dividend model fits: one twenty divided by six per cent is twenty.
- A buyer taking control can set the payout, so what matters is the cash the business generates.
- Two ten divided by six per cent is thirty-five.
- The difference is not an error.
- The two investors hold different rights, so they value different cash flows.
- We kept growth equal to be simple.
- In practice the buyer would also reassess growth.

29 Use more than one model, then reconcile

Calder Mills: three models, one reconciliation

Free cash flow model	from the cash it generates	**30**
Residual income model	book value plus excess profit	**34**
Multiples approach	peer-based	**38**
Average	(30 + 34 + 38) ÷ 3	**34**
Widest gap from average	4 ÷ 34	**12%**

IN THE LESSON

- Analysts rarely rely on a single model.
- They apply more than one and reconcile the results.
- Imagine the analyst values a different firm, Calder Mills, with three methods.
- A free cash flow model gives thirty dollars, a residual income model gives thirty-four, and a multiples approach gives thirty-eight.
- The average is thirty-four.
- The widest value is four away from the average, which is twelve per cent of it.
- That is a fairly tight range, so the analyst gains confidence.
- If the range had been fifty per cent, the right response would not be to average blindly.
- It would be to find out why.
- Is one model fed a bad input, or is the data too poor for any of them?
- Reconciling is investigating, not splitting the difference.

30 Recap of the whole lesson

- Industry, position, execution, statement quality
- Absolute from within, relative from peers: use both
- Sum of the parts less claims; discount has three causes
- Pick the model by business, data and purpose

IN THE LESSON

- Let's pull it together.
- Industry and competitive analysis asks four things: how attractive is the industry, how strong is the company's position, how well has it executed, and can you trust the statements.
- Never extrapolate without asking why.
- Absolute models, like dividends, free cash flow, residual income and assets less liabilities, value from within.
- Relative models use multiples from similar companies, and answer only relative to the benchmark.
- Use them to cross-check each other.
- Sum of the parts values each segment, subtracts net debt and other claims, and compares the result with the price.
- The discount may come from internal capital markets, endogenous factors or measurement errors, so a break-up is no guarantee.
- Choose a model by the business, the data, and the purpose.

31 Try three quick questions

1. Earnings 100, then 20% for one year, then 5% for two. Earnings after three years?
   Answer: $100 \times 1.2 \times 1.05^2 \approx 132.3$
2. Segments 500 and 300, net debt 200, 40m shares, price 12. Discount?
   Answer: Equity 600 → 15.00 a share; 1 − 12 ÷ 15 = 20%
3. No dividend, free cash flow negative for years, reliable accounts. Which model?
   Answer: Residual income, built from earnings and book value.

IN THE LESSON

- Now try three quick questions yourself.
- Pause after each one.
- Question one.
- A company's earnings are one hundred and it has an advantage for one more year, growing at twenty per cent, then five per cent for two years.
- What are earnings after three years?
- One hundred times one point two times one point zero five twice is about one hundred and thirty-two point three.
- Question two.
- Two segments are worth five hundred and three hundred, net debt is two hundred, there are forty million shares, and the price is twelve.
- What is the discount?
- Equity is six hundred, fifteen a share, so the discount is twenty per cent.
- Question three.
- No dividend, negative free cash flow for years, reliable accounts.
- Which model?
- Residual income.