



Deciding Under Uncertainty with *AI at Your Side*

AI at Your Side · Lecture 5
Sebastian Galiani · Raul A. Sosa


Kyiv School of Economics



Book website

Today: One Company, One Portfolio, One Decision

The whole lecture answers a single question, and it is not what Ferrari is worth



❓ Ferrari traded at **\$352** in April 2026. We are not going to ask what it is worth. We are going to ask **what you would have to believe** for \$352 to be a sensible price, and then whether you believe it.

💡 Last lecture the data were ours. Today they are real, and every line in a dark terminal comes from **one session with Claude Code**, the coding agent from Lecture 4. If you have never opened an agent, step zero, right after the next slide, shows how.

A Point Estimate Is Often the Least Interesting Output of a Model

The judgment lives in everything around the number

✗ What a point estimate says

Ferrari is worth \$352 per share.

It invites false confidence, hides every choice that produced it, and cannot be argued with.

✓ What a distribution says

Ferrari trades around \$352 but the 5th-percentile case is \$210, the 95th-percentile case is \$500, and here is what drives the spread.

The book's example.

It invites judgment, and it tells you where to spend it.

🕒 This lecture is about choosing what goes into the model, and finding which parts carry the weight.

Step Zero · Opening the Agent

Where it is, what you give it, and a first request whose answer you already know

What an Agent Does That a Chat Cannot

The model can be the same. The agent is allowed to touch your files



A chat

Returns text. If the text is code, you copy it, paste it, run it and bring the error back. It does not run anything for you.



An agent


Works in a folder on your computer: reads and writes files, installs what it lacks, runs the program and fixes the error. That is why it asks permission.



The agent does the work. **You decide what it gets asked, how far it may go, and what you accept from what comes back.** Lecture 1 put it this way: the unit of work goes from a reply to an assignment.

Where It Lives: Code

In Claude's desktop app, at the top left




 After the arrows, two icons. The speech bubble is the chat, and **</> is Code**, the agent. The red box marks it.

 Code needs a paid plan (Pro). The chat view also shows a folder button and a mode: for this course, click **</>** first. These screens are Windows. On a Mac it looks almost the same, and in Codex some buttons have other names, but the requests are the same. The web version cannot see your files: use the desktop app, as in the installation guide.

What You Give It: One Folder

Above the box at the bottom: Local, your own computer. Then the folder button, and Open folder




Here the folder is already chosen.

📁 The most important decision of the session.
Make a new, empty folder, here `ferrari` on the Desktop. Click the folder button, then **Open folder** and pick it.

The agent works inside that folder and has to ask before it changes anything outside it. To give it a file, copy the file into the folder. Never give it your whole Documents folder.

And a Mode That Asks First

Under the box, the word that starts as Auto. Click it and pick Manual, as here



 **The app recommends Auto**, where Claude decides by itself what to allow. For your first sessions choose **Manual**: it asks before it changes anything, and that is how you learn what it does.

In Auto you would not see the dialogs on the next slides.

A First Request Whose Answer You Know

Small on purpose: if it goes wrong here, it costs nothing

Create a file `notes.txt` that says "ready". Then tell me which folder you are working in and what files you see.

Paste it in the box at the bottom and press Enter. It stops and asks:

The screenshot shows a Windows security dialog box titled "Allow Claude to write notes.txt?". It displays the file path "C:\Users\aguss\Desktop\Ferrari\notes.txt" and the content "ready". At the bottom, there are three buttons: "Deny" (with a keyboard shortcut of 1 and Esc), "Always allow" (with a keyboard shortcut of 2 and Ctrl Shift), and "Allow once" (with a keyboard shortcut of 3 and Ctrl). The "Allow once" button is highlighted in black.

✓ **Path first:** is it inside your folder? Then what it will write, the green line. Then click **Allow once**, not Always allow, until you trust it. If the path is outside your folder, Deny.

It Did What You Asked

The folder, the file and its one word, and you can check all three

Created notes.txt, ran a command +1 -0 >

I created `notes.txt`, and it contains the word `ready`.

Working folder: `C:\Users\aguss\Desktop\ferrari`

Files I see: only `notes.txt` (5 bytes). The folder had nothing else in it, including hidden files, and it isn't a git repository.

 `notes.txt`

Undo +1 -0 >

✓ **Check it yourself:** open the folder on your Desktop, and `notes.txt` is there. The gray line above the box says it also ran a command, and that one did not ask. A few read-only commands run without a dialog. Others ask even when they only read, as on the next slide. Anything that changes a file asks.

Now It Wants to Run Something

When a command needs permission, you see all of it before it runs

Which version of Python is installed on this computer? Check it, do not guess.



On this laptop it answered Python 3.9.7. Yours should say 3.14, from the installation guide.

✂ If you cannot tell what a command does, Deny and ask it to explain first. If something breaks, do not fix it by hand: ask it what happened, and paste the error only if it came from something you ran yourself.

From here on, the Ferrari session, run with the same agent from a terminal, is shown in dark windows.

The Words of a Valuation

You do not need to know them already. Each one is a question for the tutor of Lecture 3

DCF, discounted cash flow. The value of a firm as the present value of the cash it will generate.

Free cash flow, FCF. Operating profit after tax, plus depreciation, minus capex, minus the increase in working capital.

EBITDA. Earnings before interest, taxes, depreciation and amortization: profit before those four are subtracted.

Capex, working capital. Spending on plant and equipment, and the cash tied up in inventories and in what customers still owe, net of what the firm owes suppliers.

Enterprise value, EV/EBITDA. The value of the whole firm, its shares plus its net debt, and how many years of EBITDA the market pays for it. Trailing means over the last twelve months.

WACC. The discount rate: what lenders and shareholders require, weighted by each one's share of the firm's market value.

Terminal value, g . The value of everything after the last forecast year, with cash flow growing at g forever.

CAPM, beta, risk-free rate (r_f). The return shareholders require is the risk-free rate plus beta times the equity premium. Beta says how much the stock moves with an index, and it depends on which index and which years.

bp, basis point. A hundredth of a percentage point.

Long. Owning a stock, a bet that it rises.

Block 1 · What a Valuation Is Made Of

A DCF in minutes, then what it assumes

The Decision, Not the Calculation

Ferrari, ticker RACE. About 13,600 cars a year, and a multiple that says it is not a carmaker

The company

13,663 cars in 2023, 13,752 in 2024, **13,640 in 2025**. It keeps volume below demand on purpose and publishes no ceiling. EBITDA margin went from 34% in 2019 to 38.8% in 2025.

The price

\$352 a share, about **23x EV/EBITDA** on a data provider's EBITDA, as for the peers (19x on the one Ferrari reports). Hermès 24x. Richemont 18x. LVMH 13x. Mercedes 10x. **Trailing, April 2026.**

The question

Building the DCF is the easy part. The market prices Ferrari close to the house that defines European luxury, far above any carmaker. Would you?

 That comparison is the thesis, and no spreadsheet settles it. You cannot mass-produce a Ferrari any more than a Birkin, and Ferrari has racing heritage that Hermès lacks. But Ferrari's industry is going through a technology transition, and handbags have not been disrupted in a century.

DCF: the firm as its future cash, brought back to today. **EV/EBITDA:** its price in years of EBITDA, profit before interest, taxes, depreciation and amortization.

What the Agent Builds, and What It Says It Chose

The five fields in the first request, and the list that Verification buys

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 1 of 7

condensed · it took 57 seconds

> **Context:** [...] Ferrari traded at \$352 a share in April 2026. I suspect that price needs more than these inputs can deliver.

Task: Build a five-year DCF for Ferrari. Revenue EUR 7.1B in FY2025 (year 0), growing 8% a year. EBITDA margin 38%. Depreciation 8% of revenue, capex 10% of revenue, change in working capital 1% of revenue. Tax 23%. WACC 7.5%. Terminal growth 2.5%. [...]

Constraints: Python, one script, every input at the top. Use only my numbers.

Output: [...] value per share in euros and in dollars at 1.17 per euro [...], next to \$352.

Verification: flag every number or convention you had to choose yourself. I will check free cash flow as a share of revenue by hand.

● [...] It gives EUR 208.30 a share, or \$243.71 at 1.17. [...]

● Choices I made that your inputs did not specify:


- | | | |
|--------------------------|---------------------------|------------------------|
| 1. Tax base. [...] | 2. Working capital. [...] | 3. Discounting. [...] |
| 4. Valuation date. [...] | 5. Terminal value. [...] | 6. Depreciation. [...] |
| 7. Net debt. [...] | 8. Currency. [...] | 9. Share count. [...] |

✓ **The check we promised.** $(38 - 8) \times (1 - 0.23) + 8 - 10 - 1 = 20.1\%$ of revenue. The nine choices are defensible, and none was visible until the request asked for them.

What It Says: \$244 Against \$352

🔍 The input nobody typed

The base case, in euros, converted at 1.17 dollars per euro, the April 2026 average



✖ **Resist the reflex** to call Ferrari 44% overvalued. The model says something narrower: these assumptions do not reach the price. That gap, and what would close it, is the output worth having.

⚠ **The book runs this DCF and gets \$225:** it converts at 1.08, the 2023–24 average, a rate the book never prints. That one rate moved the gap from 56% to 44%.

The tag in the corner marks an input that moved the answer and that nobody typed.

Four Fifths of the Answer Is the Part You Did Not Forecast

You forecast five years in detail. They are 19.5% of the value

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 1 of 7

condensed

[...] so capex stays above depreciation forever. That is conservative, [...] Setting capex equal to depreciation would raise terminal FCF materially. [...] Because 80.5% of EV sits in the terminal value, the result is very sensitive to WACC and g . [...]

80.5% of the value is terminal value

$$TV = \frac{FCF_T (1 + g)}{WACC - g}$$

The firm from year six on, forever, in one number. WACC is the return investors require, g the growth forever. At 7.5% and 2.5%, the denominator is **0.05**, so year six's cash flow is multiplied by twenty.

 Growth after year five is worth something only while new capital earns more than it costs, which is a question about how long Ferrari's advantages last.

The alternative the agent names, capex at depreciation, is revenue growing on a plant that never grows.

 So read a DCF upside down. The five forecast years matter mostly as the launch point of the terminal value. The discount rate and the perpetual growth rate set its multiplier, and they carry the most judgment.

The Assumption Nobody Wrote Down

🔍 The input nobody typed

Revenue per car and the exit multiple, implied by the inputs we typed



The revenue per car

Ferrari sold **13,640** cars and publishes no ceiling, so assume **15,000**. Reaching it in five years is **1.92% of units a year**. So revenue growth *is* an assumption about what each car brings in:

revenue +8%	needs per car +5.97%
revenue +10%	needs per car +7.93%

Since 2019 revenue per car has grown **about 6% a year**, mostly from mix. Our 8% sits right on that history. **Type 10% instead** and you ask for more, without ever writing it down.



The exit multiple

Divide the terminal value by year-5 EBITDA and the model is paying **10.8x**. On revenue, **4.1x**.

Ferrari traded at 23x. Hermès 24x, Richemont 18x, **LVMH 13x**, Mercedes 10x.

So this DCF ends up pricing Ferrari between Mercedes and LVMH. Part of that is mechanical, since any multiple falls when growth slows to 2.5%. The rest is a strategic claim, and it arrived through the back door: two rates and the reinvestment.

🔍 **Which assumptions are doing the work?** The list of inputs rarely says, because those assumptions are often not on it. The move that finds them is arithmetic: take an input, **ask what else it forces to be true**, and check that against history.

Block 2 · Which Assumptions Do the Work

Three ways to ask which inputs matter, then the question turned around

Move One Thing at a Time


Each row moves one input by a percentage point and holds the rest. Our run of the same model

Input, moved ± 1 pp	down	up	swing \$/sh
WACC	306.05	202.19	103.86
Terminal growth	209.31	295.32	86.01
Capex over revenue	255.87	231.56	24.32
Revenue growth	233.46	254.34	20.88
EBITDA margin	234.35	253.08	18.72
Tax rate	247.36	240.07	7.29

© The WACC moves the answer 14.2 times more than the tax rate does. And the two on top are the discount rate and the perpetual growth rate. Every other input here can be argued from a filing or a peer.

Now Each One Over Its Own Defensible Range

A point of WACC and a point of tax are not equally likely, so stop treating them that way




🔑 A sensitivity ranking depends on the ranges you choose. Terminal growth was *second* on the one-point table and it is **last** here. Nothing about Ferrari changed between these two slides. What changed is how wide each input may move. The discount rate is first both times. Alone, no range reaches \$352. All five at their optimistic end reach \$387.

The ranges. WACC and terminal growth from the book. Revenue from Ferrari's 2030 plan (4.7%) to the unit ceiling with 8% per car (10.1%). Margin 35% to 40%, against 34% to 38.8% since 2019. Capex wide, for the electric transition. Tax stays at 23%.

The Two Judgment Calls, Against Each Other

Eight of sixty-three cells land within 10% of the market price, and they lie on a diagonal



The boxed cells lie on a diagonal. A lower discount rate and a higher perpetual growth rate buy the same thing, so the pair is what matters and neither number means much alone. **Color is clipped at 500** because the top right corner reaches 1,174 and would otherwise flatten the whole table.

Run It Backwards

Stop asking what it is worth and ask what the price already assumes

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 2 of 7

condensed

> Now run it backwards. Hold every other input at base and solve for the single value of each of these that makes the price exactly \$352: WACC, terminal growth, revenue growth, EBITDA margin. For the revenue growth answer, tell me what growth in revenue per car it needs if units go from 13,640 in 2025 to 15,000 in year 5. Verify each one by re-running the model to \$352.

● [...] Every solved value re-runs through the full model to \$352.000000. [...]

Input	Base	Solved to hit \$352
WACC	7.50%	5.99% (-151 bp)
Terminal growth	2.50%	4.22% (+172 bp)
Revenue growth, years 1-5	8.00%	16.86% a year
EBITDA margin	38.0%	49.57% (+11.6 pts)

● [...] The solved 16.86% needs revenue per car to grow 14.66% a year [...]

✓ The last line of the request is the habit worth keeping.

Four Ways to \$352

🔍 The input nobody typed

Three ask Ferrari for something it has never done. The fourth asks for its Italian-index beta

	needed	up against
Revenue per car	14.7%/yr	about 6% a year since 2019
EBITDA margin	49.6%	38.8%, Ferrari's best year (2025)
Terminal growth	4.22%	forever, above 3.1% euro rf and 3.3% nominal euro-area growth
WACC	5.99%	CAPM, book's beta 0.6: 6.1%. Broad indices, about 0.9: 7.6%

CAPM: required return = rf (3.1%) + beta × equity premium (5%). With almost no net debt, the WACC is close to it.

🔑 So the gap between our model and the price is one number: the beta. The book's 0.6 matches Ferrari's beta against the Italian index to 2024 (0.58). Against broad indices it is 0.8 to 1.0, then and now, and 0.9 gives about our 7.5%. Like the exchange rate, it came without index or date.

✍️ **My bet, for the pre-mortem on slide 38 to attack:** *At \$352 I am betting that Ferrari is as safe as its Italian-index beta of 0.6 says.*

Block 3 · From a Number to a Distribution

Ten thousand futures, and the one line of code that decides whether they mean anything

The Portfolio, and Where Its Numbers Come From

The same discipline on a portfolio: 40% Apple, 30% Coca-Cola, 30% Shell, with dividends

 118 months, 2016-11 to 2026-08

These are past means, not expected returns.

	mean/mo	sd/mo
Apple	+2.45%	7.89%
Coca-Cola	+1.00%	4.74%
Shell	+1.16%	7.05%

 And how they move together

	AAPL	KO	SHEL
AAPL	1.00	0.23	0.08
KO	0.23	1.00	0.26
SHEL	0.08	0.26	1.00

0.23 means they tend to move together, but weakly. None reaches 0.27, and that is why the coming mistake is hard to see.

 This window is a choice and it is doing work. In these ten years Apple returned **12.2 times** your money, **29.0%** a year compounded. That is history. Next year's mean is a guess, and with ten years of data its standard error is still about **six points** a year.

If You Have Never Seen a Covariance

The last slide used correlations, and the next builds a covariance matrix. Ask this first



Claude Sonnet 5 · clean session · Oct 7, 2026

condensed

Explain covariance and correlation to me intuitively, assuming I know no statistics. Then: if I simulate a portfolio of three stocks and ignore the correlations between them, why does the risk change while the average return barely moves? Keep it short, with one small example.



Claude

[...] **Covariance**: do two things tend to move together? If stock A goes up on days stock B goes up (and down together too), covariance is positive. [...] **Correlation**: the same idea, but rescaled to always sit between -1 and +1. [...]

Average return of a portfolio is just a weighted average of each stock's average return. It doesn't care *how* the stocks move relative to each other [...] **Risk (volatility)** depends on whether the stocks' ups and downs *cancel out* or *pile up*. [...]

Tiny example: Two stocks, each with $\pm 10\%$ swings. [...] If they're **negatively correlated** [...] combined portfolio might only swing $\pm 3\%$. [...] If you **assume zero correlation** (independence) when they're actually highly positively correlated, you'll underestimate how much they swing *together*, and the real portfolio swings $\pm 9\%$, not $\pm 3\%$. [...]



The idea is right, and the example slips: with no correlation, two $\pm 10\%$ stocks swing about $\pm 7\%$, not $\pm 3\%$. Getting $\pm 3\%$ takes stocks that move against each other. Check a tutor's numbers too.

What the Agent Writes, and the Two Lines That Matter

Two lines, and Lecture 4's questions: what it does, why, what if we cut it, how to test it

```
Claude Code · Sonnet 5 · Sep 23, 2026 · turn 3 of 7 condensed · portfolio_sim.py

27 cov = corr_in * np.outer(sd, sd)          # = sample covariance matrix
[...]
```

```
33 shocks = rng.multivariate_normal(mu, cov, size=(N_PATHS, N_MONTHS)) # (paths, months, stocks)
34 port_monthly = shocks @ WEIGHTS          # (paths, months)
35 annual = np.prod(1 + port_monthly, axis=1) - 1          # compound 12 months
```

● [...] The largest gap is 0.0047, in line with sampling error [...] The simulated monthly means and standard deviations also match the file's to three decimals. [...]

🔍 The two worth stopping on

Line 27, the correlations inside the covariance. What if the shocks ignored them? *Hold that question. It is the next slide.*

Line 35, prod and not sum. Twelve monthly returns compound. How could I test it? Twelve months of +10% must give **213.8%**, not 120%.

⚠️ **One warning from the book that is not about Python.** Ask a spreadsheet assistant for a correlated simulation and it will often hand back the independent version, labeled correlated. The agent got this line right, and checked the means and volatilities without being asked. The next slide shows what the wrong line costs.

Two Simulations, One Assumption Different

One of them is right, and the summary table alone cannot tell you which

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 4 of 7

condensed

> What would happen if this simulation ignored the correlations [...]? Run it both ways with the same seed and show me what changes, including your correlation check.

[...] with corr no corr diff (no - with)

mean 21.46% 21.28% -0.18%

std 19.26% 16.97% -2.30%

p1 -17.97% -13.01% 4.97%

p5 -7.77% -4.44% 3.34%

p50 20.02% 20.26% 0.23%


[...]

- [...] the two matrices give differently signed and ordered factors. The same seed therefore does not give the two runs matched shocks. The differences in the table also include ordinary sampling noise. [...]

✗ The mean and median barely move, while the standard deviation and the tails do. Only the data's matrix says which run is right. The same seed did not give both runs the same draws, so part of each gap is luck (slide 30).

What the Two Runs Look Like

The chart it saved in the same turn, with its own reading



Blue, with correlations. Orange, ignoring them. Brown is where they overlap. Is the tail gap just noise? We reran the comparison with 500 seeds, and the tails move every time: about 3 points at the 5th percentile and 4 at the 1st.

Claude Code · turn 4 of 7 condensed

- [...] **Middle of the distribution.** The mean and median barely move. The mean gap is 0.18 points, which is within Monte Carlo noise. [...]
- **Tails.** Both tails are thinner without correlation, and the lower tail matters most: a 1-in-100 year loses 13% instead of 18%. [...]

Plausible Output Is Not Validation

Stack the simulated shocks, correlate them, and compare against the data's matrix

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 4 of 7

condensed

```
[...] === WITH correlations ===
```

```
[...] Simulated correlation [...]:
```

	aapl_return	ko_return	shel_return
aapl_return	1.0000	0.2334	0.0805
ko_return	0.2334	1.0000	0.2654
shel_return	0.0805	0.2654	1.0000

```
Max abs difference: 0.0047
```

```
[...]
```

```
[...] === IGNORING correlations ===
```

```
[...] Simulated correlation [...]:
```

	aapl_return	ko_return	shel_return
aapl_return	1.0000	0.0033	-0.0002
ko_return	0.0033	1.0000	0.0003
shel_return	-0.0002	0.0003	1.0000

```
Max abs difference: 0.0033
```

```
[...]
```

✅ **With them.** Largest gap against the data: **0.005**, so it passes.

❌ **Ignoring them.** Its check: **0.003** from what that run was told, zeros, so it passes. Our check against the data's matrix: **0.26**, so it fails.

🔑 **A check against the instructions cannot catch a wrong instruction.** Put the check in the request, anchored in the data: the simulated matrix next to *the data's*, and the volatilities too.

What That Check Cannot Catch

🔍 The input nobody typed

Three assumptions walk straight through it, and the heaviest is the one nobody questions

🎯 The mean, taken as data

Ten years of data give **21% a year, give or take six**. The tails move about seven tenths of a point with each point of it.

At a 12% mean, the 1st percentile goes from -18% to -24.5% .

📊 Shape and one matrix

Draw whole past months instead of normals and the 1st percentile gets about **1.7** points worse. Swap in the bad-month correlations and it gets about **0.6** worse.

They are hard to separate: bad months bring lower means, higher correlation and a more volatile portfolio, all at once.

📈 The one that grows

With n equal assets,
 $\sigma_p^2 = \sigma^2[1/n + (1 - 1/n)\rho]$, σ each asset's standard deviation, ρ the correlation of any pair.

The first term shrinks as you add assets. The second stays.

Bad months are the S&P 500's worst fifth, 24 of the 118. With their correlations rather than the other 94 months', the formula gives a standard deviation 14% higher on three assets.

🔑 One standard error of doubt about the mean moves the tails more than the other two together. The matrix is the one that grows with the portfolio, so the person who should worry most about it is the one who feels most diversified.

Whose Numbers Are These?

The same code, fed two different means, returns two different answers



The book against the agent

	book	agent	12% mean
mean	low-mid teens	21.5%	12.1%
sd	high teens	19.3%	17.9%
p5	low-mid teens loss	-7.8%	-15.0%
p1	mid-20s loss	-18.0%	-24.5%

We gave the same code a 12% mean, and it landed on every range in the book. The difference is one input.

Agent: Claude Code, Sonnet 5, Sep 23, 2026.



Four checks before any of it

- 1. Grain.** How many rows did you expect, and how many are there? Merging two tables on the wrong column runs perfectly.
- 2. Missing.** Two months here return exactly zero. We checked the daily closes: real, not gaps.
- 3. Distribution before average.** The mean is above the median because compounding is convex. The average path is not the typical investor's path.
- 4. Subgroups.** Does the pattern survive inside them?



The book's ranges need a 12% mean, and no window of four years or more inside these ten years gives one. The number that belongs in a report is the percentile with its mean and its window attached.

Block 4 · What You Show, and What You Refuse to See

The last two moves are the same move, pointed in opposite directions

Ten Thousand Numbers, One Figure

Producing the chart costs nothing now. Choosing which one to show costs what it always did

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 5 of 7

condensed · chart on the next slide

> **Context:** [...] I suspect the bad years are worse than the summary table suggests.

Task: One chart that shows how bad the bad years get.


[...] **Verification:** if the data do not show what I suspect, say so before you draw anything.

- Your suspicion is only half right. [...]
- Supports you: the 5th percentile is only where the bad years start. Once a year is that bad, the average bad year is -13.8%, almost twice the marked figure. [...]
- Does not support you: the left tail is thinner than a normal distribution with the same mean and standard deviation. [...] This comes from the model: normal monthly shocks compound into a right-skewed annual return. It does not come from the data. [...]

✓ Context states the suspicion, and Verification lets the data test it. Here the test split it in two: the thin tail comes from the model, and the 5th percentile undersells a bad year.

The Chart It Saved

The same turn, with its title, its marker and its axes read out on the right



What it drew

Its title: *"1 year in 20, the portfolio loses 7.8% or more."*

The black line is the 5th percentile, -7.8%.

Across, the one-year return of the portfolio. Up, how many of the 10,000 simulated years land in each bar.

The agent said on slide 35 that the years left of the line average -13.8%. The title leaves that out.

The Trap: The Chart That Confirms

🔍 The input nobody typed

It looks like the last request, with the verdict written in advance

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 6 of 7 condensed

```
> Make me a chart showing Ferrari has room to run.  
● I didn't draw the chart you asked for, because your own  
  model doesn't support it. [...] There is no honest chart  
  showing room to run. [...]  
● Title: "$352 a share needs a WACC near 6%, not 7.5%". [...]
```

✗ It drew chart 2 below instead. Nobody told it our 7.5% was the right rate.

📄 From the same model, three defensible charts

1. The base case against the market price. *An overvaluation thesis.*
2. The WACC and growth pairs that reach \$352. *Reachable. It drew this one, titled to argue the opposite.*
3. The implied revenue per car against its history. *A this-has-never-happened thesis.*

✓ **The antidote is also a request:** "give me three comparisons of this output, the conclusion each supports, and which is most honest with the whole model." Then you choose, and the choice is on the record.

The Pre-Mortem: Asking for the Case Against You

Assume the position already failed, and ask for the autopsy

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 7 of 7

reply on the next slide

> I am long Ferrari at \$352, one-year target \$430. A year has passed and the trade is down 40%. As post-mortem analyst, write a three-paragraph autopsy. One paragraph each: (a) a macro environment I did not anticipate, (b) a company-specific event, (c) a structural industry shift I did not take seriously. Be specific and plausible.

🔄 **Before reading it**, bring back the sentence from slide 24. *At \$352 I am betting that Ferrari is as safe as its Italian-index beta of 0.6 says.* Then mark which paragraphs that sentence had already priced in, and which ones it had not. The ones it had not are the whole reason to run a pre-mortem. The rest of the autopsy tells you what you already knew.

🔑 Ask it outright for the case against you and it will make it, with none of the social cost a colleague would pay.

What Came Back, and What It Costs the Thesis

Three futures, and one line before them worth reading twice

Claude Code · Sonnet 5 · Sep 23, 2026 · turn 7 of 7

condensed

- This is a hypothetical scenario. None of the events below is a forecast or a fact. [...] your own DCF gave \$244 against \$352 at entry, and a 40% loss puts the price at \$211. About 70% of the fall, measured in logs, is the market closing the gap the model already showed. [...]
- (a) [...] The year brought a bond-market repricing, not a recession. [...] the model is a rate bet. Moving WACC from 7.5% to 8.0% takes your DCF from \$244 to \$221 [...] The euro also slid from 1.17 to about 1.08 [...] A dollar-denominated long [...] then lost roughly another 8% in translation [...]
- (b) [...] The first fully electric model launched with weaker order conversion than the historical waiting lists implied. [...] An EBITDA margin [...] slipped to about 35%. [...]
- (c) [...] Performance stopped being scarce. [...]

✎ My bet was on Ferrari's risk, and none of the three paragraphs is about it. **Ferrari can be exactly as safe as I bet and still lose on rates, the euro, the margin or the scarcity.**

✖ **Read the 70% line again.** It is my sentence failing: the market pricing Ferrari at its broad-index beta instead of the Italian one. The autopsy measured that failure without naming it.

What You Take With You Today

Three moves, and all three are the same instinct: go looking for what would embarrass you

Run it backwards

Ask what this price already assumes, and whether you believe it. It works on any model with a price attached.

Find the input nobody typed

Take an assumption, ask what else it forces to be true, check it against history, and ask for its date. The book's exchange rate and its beta were two.

Put the check in the request

The data's matrix against the simulated one, volatilities included. Once it is in the request, it runs without your remembering it.

 **On your own, nothing to hand in.** Take `inflation_latam.csv` or `countries_2023.csv`, both in `data_lecture5.zip`, write the comparison you want in one line, chart it yourself with no AI, then ask for it twice: **blind** ("plot this data" plus the file, nothing else) and **in the five fields** (Context, Task, Constraints, Output, Verification), with the Verification line from slide 35. Compare what the three of you decided.

In a room where everyone has the same instrument, you are the variable

Next lecture: nothing to build.
Four papers of ours, and how they were actually built



Book website