Financial Econometrics | Lecture 1

Economic Questions and Data, and a Review of Probability

In Class Group Activity

· SJTU, Summer School 2026 · Fatih Kansoy

How this block runs. You work in your home group of three or four, textbook open. Roles rotate today: **Scribe** (writes the answer), **Reporter** (speaks for the group), **Sceptic** (must raise one objection), **Timekeeper**. All money today is in US dollars ($) and pounds sterling (£).

- **Phase 1, Concept checks (~10 min).** For each question you vote on your own first. Then you talk for about 90 seconds with the people next to you, and vote again. Do not talk before the first vote.
- **Phase 2, Desk problem (~20 min).** One problem in parts, on a small set of stock returns. We solve it together, part by part.
- **Phase 3, Case (~25 min).** You read one real story. The ending is hidden on purpose. You answer four questions and take a position. Then we reveal what happened.

A short word list for this block. A **random variable** is a number whose value is decided by chance, such as next month's return. Its **expected value** (or mean) μ is its long-run average; its **variance** σ^2 measures spread, and the **standard deviation** σ is the square root of the variance, back in the original units. The **covariance** of two variables measures whether they move together; the **correlation** ρ puts that on a fixed scale from -1 to $+1$. To **standardise** a value is to subtract its mean and divide by its standard deviation, giving a Z-score. A **causal** question asks whether one thing makes another happen; a **predictive** question only asks whether one thing helps forecast another. A **confounder** is a third factor that drives two others and so creates a correlation between them without one causing the other.

Phase 1 – Concept checks

~10 MIN

ConcepTest 1. A pension fund notices that, across 200 listed firms last year, the firms whose chief executive was paid more also tended to have higher share returns. The fund's board asks: "Will paying our own chief executive more *cause* our share price to rise?" Which statement is the best answer?

(a) The correlation across the 200 firms answers the board's question directly.

(b) This is a causal question, but the 200 firms are observational data, so the correlation alone cannot answer it. Firm quality is a likely confounder: better firms both pay more and perform better.

(c) Because we have data on 200 firms, this is effectively a randomised controlled experiment.

(d) It is only a forecasting question, so causation does not matter.

ConcepTest 2. Stock and Watson note that a good way to forecast whether it is raining is to watch whether pedestrians are carrying umbrellas. A quant finds that the **term spread** (long minus short interest rates) reliably falls before US recessions, and uses it to forecast GDP. Which reading is best?

(a) Because the term spread forecasts recessions well, a narrowing spread must *cause* recessions.

(b) A variable can be a useful predictor without being a cause. The term spread forecasts GDP whether or not it causes anything.

(c) A forecast is only valid once we have proven the causal mechanism behind it.

(d) Umbrellas cause rain, so the analogy is wrong.

ConcepTest 3. From a random sample of 30 monthly returns on a fund you compute a sample average of +0.8% per month. A classmate says: "So the fund's true expected monthly return is 0.8%." The most accurate response is:

(a) Correct: the sample average equals the population mean.

(b) The 0.8% is one estimate. A different random sample of 30 months would give a different average, so 0.8% is itself a draw from a distribution.

(c) It is correct only if the returns are normally distributed.

(d) We should report the variance, 0.8%, instead, because variance is in the same units as the return.

Phase 2 – The desk problem

~20 MIN

Two stocks, six months

The setting. You are on a junior desk looking at two London-listed stocks over the last six months: **AeroJet plc** (an airline, Asset A) and **GridCo** (a utility, Asset B). The table gives each month's total return in per cent. Treat these six months as the whole population, so you divide by $n = 6$ (not $n - 1$) and the arithmetic stays clean.

Month	A: AeroJet return (%)	B: GridCo return (%)
1	+4	+1
2	−2	+3
3	+6	−1
4	+1	+2
5	−3	+4
6	+6	+3

(a) Mean, variance and standard deviation of Asset A. Fill in the working table (the first row is done for you), then compute the mean μ_A, the variance $\sigma_A^2 = \frac{1}{n}\sum(A - \mu_A)^2$, and the standard deviation σ_A.

Month	A	$A - \mu_A$	$(A - \mu_A)^2$
1	+4	+2	4
2	−2		
3	+6		
4	+1		
5	−3		
6	+6		

When you have σ_A^2, say in one sentence why the variance has awkward units but the standard deviation does not.

(b) Covariance and correlation of A and B. You may take $\mu_B = 2.0\%$, $\sigma_B^2 = 2.667$ and $\sigma_B = 1.63\%$ as given. Compute the covariance $\text{cov}(A, B) = \frac{1}{n}\sum(A - \mu_A)(B - \mu_B)$, then the correlation $\rho = \text{cov}(A, B)/(\sigma_A \sigma_B)$. What does the sign of the correlation tell you about how the two stocks move?

(c) Standardise and use the normal table. Forget the two stocks for a moment. The whole London market index has monthly returns that are approximately normal with mean $+1\%$ and standard deviation 4%, that is $N(1, 16)$ (remember the second number is the variance, $4^2 = 16$).

(1) What is the probability the index loses more than 7% in a month (return below -7%)? Standardise first.

(2) What is the probability it gains more than 9%?

(3) Between which two return levels does the index sit 95% of the time? Use $\mu \pm 1.96\,\sigma$.

(d) The portfolio twist. You put half your money in AeroJet and half in GridCo, so the portfolio return is $P = 0.5\,A + 0.5\,B$. Using

$$\text{var}(aX + bY) = a^2\sigma_X^2 + 2ab\,\sigma_{XY} + b^2\sigma_Y^2,$$

find the mean and the variance of P. Compare σ_P with σ_A and σ_B. What would the portfolio variance have been if the two stocks were uncorrelated? In one sentence, state what this shows about diversification.

Phase 3 – The case

~25 MIN

Twenty standard deviations: the day the model said could not happen

How to run this. Read the story with your group. The ending is hidden on purpose. Your group must take a position on the four questions and hold your decision. Do not search the internet. Use only what you know from the lecture. I will reveal what happened only after every group has answered.

The story. It is the summer of 1987. A London risk desk runs a simple, respectable model: daily moves in the Dow Jones Industrial Average (an average of 30 large US firms) are treated as **normally distributed**. The desk has years of clean data. The average daily change is about $+0.04\%$, and the standard deviation of daily changes is about 1.12%. That is the desk's whole picture of risk: a bell curve with that mean and that spread.

On most days the index drifts up or down by around 1%. A 3% fall would be a genuinely bad day. The desk sets its limits, its margin buffers and its "worst plausible loss" using this normal model, because that is what every textbook of the day recommends.

Then comes **Monday, 19 October 1987.** In a single trading day the Dow falls by 22.6%.

The story stops here.

Four discussion questions (answer all four in your group; take a position):

1. **How far out?** Using the desk's own numbers (mean 0.04%, standard deviation 1.12%), how many standard deviations below the mean is a -22.6% day? Show the arithmetic.
2. **How often should that happen?** Under the normal model, very roughly, how often should

a move of 20 or more standard deviations occur: once a decade, once since the dinosaurs, or essentially never in the age of the universe? Commit to one.

3. **Model or world?** The crash happened anyway. Which does the evidence support, and why: (a) the day was a freak we should ignore; (b) the normal distribution is the wrong model for stock returns because real returns have fat tails; (c) the standard deviation was simply mismeasured?

4. **Your call.** You run a London desk today. In one or two sentences, write what you would change about a "worst plausible loss" number that is built on a normal distribution. Commit before the reveal.

End of activity handout. Solutions are distributed after the block.