

Risk Structure, Term Structure, and Market Efficiency

Financial Institutions and Markets · MF50004.01

Where we left off

Week 2 in one line

A bond's price is the discounted list of its cash flows, duration says how much that price moves when the discount rate does, and the rate itself is a price set by supply and demand.

This week is about why yields differ

Week 2 priced everything off a single yield. But no two bonds trade at the same yield. Two bonds maturing on the same day differ because of the borrower, the market and the tax code; the same borrower pays different rates at different maturities. This week takes those two cuts in turn, then asks how to measure the risk behind them and whether the market prices it correctly.

This week

MEETING 1 · WHAT MAKES YIELDS DIFFER

- 1 The risk structure — same date, different yields
- 2 The term structure — same issuer, different dates

MEETING 2 · MEASURING RISK, AND PRICING IT

- 1 Measuring return and risk
- 2 Market efficiency and its limits

Mishkin Chapters 5 and 6. Two ways of cutting the yield curve, then how to measure risk, then whether prices reflect it.

MEETING 1

Session 1 · The risk structure of interest rates

Session 2 · The term structure of interest rates

90 minutes · why yields differ across issuers, then across dates · closes on a curve that inverted for two years and was wrong

Same maturity, different yields

What a yield gap is made of

Hold the maturity fixed and two bonds still trade at different yields. The gap between them breaks into three parts: **default risk**, **liquidity**, and **tax treatment**. We take them one at a time.

A credit spread is a price, and these three are what it pays for. Anything else you can think of — issuer size, sector, covenant strength — works through one of them.

Default risk

The risk premium is the difference between the yield on a defaultable bond and the yield on a default-free bond of the same maturity:

$$\text{risk premium} = i_{\text{corporate}} - i_{\text{government}}$$

The demand-and-supply logic from Week 2, applied twice

News that raises default risk shifts the demand for corporate bonds left — their price falls and their yield rises. The same news shifts demand for government bonds right — a flight to quality. Both movements widen the spread, and neither requires a default to occur.

This is why spreads widen before any default happens: a spread prices what lenders expect to lose, and expectations move first.

Liquidity and taxes

LIQUIDITY

How cheaply and quickly you can turn the bond back into cash. A less liquid bond must offer a higher yield to be held at all.

Treasuries are the most liquid long-term bonds there are, which is part of why they yield less than corporates of identical credit quality.

TAXES

US municipal bonds are exempt from federal income tax, so investors accept a lower pre-tax yield — sometimes below Treasuries, despite carrying default risk.

A yield comparison that ignores tax treatment compares two different things.

COLD CALL · 3 MIN

A municipal bond can default and still yield less than a Treasury. Is that an arbitrage?

How the rating scales work

What it means	Moody's	S&P / Fitch
Highest quality	Aaa	AAA
High quality	Aa1 – Aa3	AA+ – AA–
Upper medium	A1 – A3	A+ – A–
Lower medium — investment grade ends here	Baa1 – Baa3	BBB+ – BBB–
Speculative	Ba1 – Ba3	BB+ – BB–
Highly speculative	B1 – B3	B+ – B–
Substantial risk, then default	Caa – C	CCC – D

What a rating is

An opinion about how likely a borrower is to miss a payment, relative to others. Not a probability, and not a price. The issuer pays for it.

Why the line at BBB– matters

Regulation and fund mandates are written on it. Many insurers and pension funds may not hold bonds below BBB–, so a downgrade across the line forces selling by investors with no view on the company.

A rating ranks borrowers. What that ranking is worth is settled by the price, and where the two disagree, trust the price: someone had to buy at it.

Predict before we look

PREDICT BEFORE REVEAL · 4 MIN

A Chinese AA-rated one-year corporate note. How much more than government debt does it yield?

- A. About 250 basis points
- B. About 150 basis points
- C. About 80 basis points
- D. About 40 basis points

For scale: the average **investment-grade** US corporate bond pays **78bp** over Treasuries, and a BBB one — the lowest investment grade — pays **95bp**.

What a credit spread looks like today

78bp

US investment-grade corporate bonds over Treasuries — 95bp for BBB, the lowest investment grade

270bp

US high-yield (below investment grade) — against a record 2,182bp in December 2008

41bp

China 1-year AA notes over comparable government debt — the tightest since 2007

A Chinese AA note pays less than the average US investment-grade bond


AA is an investment grade in both places, so the two are comparable in name. But a Chinese firm rated AA at home is rated around BBB– by a global agency — the bottom of investment grade, where US bonds pay **95bp**. The Chinese note pays **41bp**, less than half.

So either Chinese AA borrowers default less often than US borrowers two grades above them, or the 41bp is paying for something other than default risk.

ICE BofA option-adjusted spreads via FRED, 17 September 2026 (US corporate index, BBB index, high-yield index). China: Bloomberg, 15 May 2025. Spreads are over government debt of comparable maturity in each market.

Almost every Chinese issuer is rated in the top two grades

Share of onshore credit-bond issuers, %



On the S&P scale, two US companies are AAA

Microsoft and Johnson & Johnson — the whole list, against more than 60 US corporates thirty years ago. And only **49%** of S&P-rated corporate issuers reach investment grade at all. Half the universe sits below BBB–.

59% at AA+ or above

Among more than 6,000 onshore credit-bond issuers, **27%** are rated AAA and **32%** AA+.

Why this leaves the letters uninformative

The point of a rating is to separate borrowers. If three in five are in the top two grades, the letters carry almost no information, and investors have to look somewhere else — which is what the 41bp spread was telling us.

China: end-Q1 2024, 6,000+ onshore credit-bond issuers — AAA 27%, AA+ 32% (The Standard, on the PBoC's measures against the AAA concentration). US: S&P Global Ratings — 49% of rated corporate issuers are investment grade (78% of rated debt); AAA list as of 2026.

The same firm, rated twice

7 notches

how much higher a Chinese firm's domestic rating sits than its global one — on both the median and the average

AA / Aa2 at home, **BBB- / Baa3** abroad — the bottom rung of investment grade

Why the two disagree

They weigh the same financials differently. Domestic agencies treat **size** as a bigger positive; global agencies penalise **leverage** harder and weigh **profitability** and **state ownership** more. Add issuer-pays and a domestic panel competing for the same clients, and the scale drifts up.

It is not simply a different scale

If the gap were a constant offset, shifting one scale down seven notches would line the two up. Do that and about **24%** of jointly rated firms still disagree: the two are ranking firms differently, not just labelling them differently.

A Chinese AA and a US AA do not describe the same borrower — part of why 41bp looks so tight.

Jiang & Packer, BIS Working Paper 648 (2017), via VoxChina — 2015 data, about 50 jointly rated firms.

Same issuer, different dates

Three facts any theory has to explain

1. Yields on different maturities **move together** over time
2. When short rates are low the curve is usually **upward sloping**; when they are high it is more often inverted
3. The curve is **almost always upward sloping**

Three theories try to account for all three facts. They differ in one assumption: whether bonds of different maturities are substitutes for each other.

Pure expectations theory

A long rate is the average of the short rates the market expects over the same horizon:

$$i_{nt} = \frac{i_t + i_{t+1}^e + i_{t+2}^e + \cdots + i_{t+(n-1)}^e}{n}$$

It assumes bonds of different maturities are perfect substitutes

If they are, no maturity can offer a better expected return than another, and arbitrage forces the long rate to the average.

EXPLAINS

- Fact 1 — a rise in the short rate raises the whole average
- Fact 2 — low short rates imply expected rises, so the curve slopes up

FAILS

- Fact 3 — it predicts the curve should slope down about as often as up

Two more theories, and the scorecard

Theory	Core assumption	Explains	Fails
Pure expectations	Maturities are perfect substitutes	Facts 1 and 2	Fact 3
Market segmentation	Maturities are <i>not</i> substitutes at all; each has its own clientele	Fact 3	Facts 1 and 2
Liquidity premium	Substitutes, but imperfect — investors need paying to hold duration	All three	see below


$$i_{nt} = \frac{i_t + i_{t+1}^e + \cdots + i_{t+(n-1)}^e}{n} + \ell_{nt}, \quad \ell_{nt} > 0 \text{ and rising in } n$$

The catch: nobody observes the premium

It is the expectations theory plus one extra term, ℓ_{nt} , and that term is never measured: it is defined as whatever is left once you subtract expected short rates from the long rate. Any movement the expectations part cannot explain is attributed to it instead, which makes the theory very hard to prove wrong.

And it is large: with the 30-year at 5.29% against a 3.2% neutral rate, more than two points of that yield is premium.

What the curve implies about future short rates



THINK – PAIR – SHARE · 7 MIN

Fed funds is 3.75–4.00% and the 2-year is 4.67%. Under pure expectations, what is the market saying about the next two years — and what does the 30-year at 5.29% add?

The neutral rate, and what it says about the long end

The **neutral rate** is where the policy rate settles once inflation is at target; the FOMC's median member puts it at **3.2%**. Read the 30-year's 5.29% as a pure forecast of policy and you are claiming the Fed sits two points above that for thirty years. The gap is the term premium.

Federal Reserve H.15, 17 September 2026 — 3m 3.97%, 2Y 4.67%, 10Y 4.94%, 30Y 5.29%. Neutral rate: FOMC projections, 16 September 2026.

Why China's curve is so flat



Each theory can produce this curve

Expectations: the short rate is expected to rise slowly, averaging about 2.1% over 30 years against 1.2% today.

Liquidity premium: if short rates stay near 1.2%, the 30-year term premium is about 95bp, against roughly two points in the US.

Segmentation: insurers and banks need long bonds and there are too few, so their buying holds long yields down.

What would tell them apart

Look at what moves the long end. If long yields move with news about the future policy rate, expectations are doing the work. If they move with the supply of long bonds, or with insurers' and banks' demand for them while the policy outlook stays put, the term premium and segmentation are.

All three theories fit this curve. Telling them apart needs evidence from outside the curve itself.

US: Federal Reserve H.15, 17 Sep 2026. China: ChinaBond, 18 Sep 2026 — 3m 1.18%, 10Y 1.68%, 30Y 2.13%, a 95bp span against the US 132bp. Deposit and financing figures: Week 2.

Case — the longest inversion in 45 years, and no recession

- SITUATION** The 3-month/10-year curve inverted on **25 October 2022** and stayed inverted until **13 December 2024 — 26 months**, the longest stretch in 45 years. On the textbook reading that is the most reliable recession signal in US macro-finance, and every US recession since the late 1960s was preceded by one. No recession followed. The curve is now normal again, at **+97bp**.
- FRICTION** An inverted curve is not a forecast of recession. Under the expectations theory it carries one piece of information: **investors expect short rates to be lower later than they are now**. Whether that is bad news depends entirely on *why* rates are expected to fall — and the curve shape alone cannot tell you which reason you are looking at.
- DISCUSS** Give **two different reasons** the market might expect short rates to fall, one that ends in a recession and one that does not. Both invert the curve. Which one fits 2022–24, and what evidence would separate them?

THINK – PAIR – SHARE · 9 MIN

Two minutes alone, three with your neighbour, then we build both stories on the board.

Inversion dates: U.S. Bank; current spread from Federal Reserve H.15, 17 September 2026.

MEETING 2

Session 3 · Measuring return and risk

Session 4 · Market efficiency and its limits

90 minutes · how to measure risk, then whether prices reflect it · closes on China's quant funds

Measuring return and risk

What this session covers

1. Expected return and variance of a single asset
2. Covariance and correlation, and what happens to risk when assets are held together
3. The Capital Asset Pricing Model: which part of an asset's risk earns a return

The main result: part of each asset's risk cancels out when it is held with other assets, and investors are paid only for the part that does not.

Expected return and variance

Expected return

$$E(R) = \sum_s p_s R_s$$

Variance and standard deviation

$$\sigma^2 = \sum_s p_s (R_s - E(R))^2$$

Why variance, and not something else

It is symmetric — it treats an upside surprise and a downside surprise as equally bad, which is plainly wrong as a description of how people feel. It is still used because it is easy to work with and because it aggregates: the variance of a portfolio has a closed form in the weights, which is what makes portfolio optimisation possible.

Most criticisms of mean-variance analysis start here. The alternatives describe risk more realistically, but they are much harder to work with.

Worked example: two assets, four states

State	Probability	Asset A	Asset B
Boom	0.25	+28%	−2%
Expansion	0.35	+14%	+9%
Slowdown	0.25	−2%	+13%
Recession	0.15	−18%	+4%

Covariance and correlation

$$\sigma_{AB} = \sum_s p_s (R_{A,s} - E(R_A)) (R_{B,s} - E(R_B)), \quad \rho_{AB} = \frac{\sigma_{AB}}{\sigma_A \sigma_B}$$

Covariance is positive when the two returns tend to be above their means in the same states, and negative when one tends to be above its mean while the other is below. Correlation divides by the two standard deviations, so it always lies between −1 and +1.

$$E(R_A) = 8.70\%, \sigma_A = 15.44\% \quad \cdot \quad E(R_B) = 6.50\%, \sigma_B = 5.63\% \quad \cdot \quad \sigma_{AB} = -43.8, \text{ so } \rho_{AB} = -0.50.$$

A pays more and swings more, but the two tend to move in opposite directions: their correlation is −0.50.

Now mix them

Portfolio	Expected return	Standard deviation
All A	8.70%	15.44%
Half and half	7.60%	6.76%
21% A, 79% B — minimum variance	6.96%	3.98%
All B	6.50%	5.63%

Read the third row against the fourth

Adding the *riskier* asset to the safer one raises expected return from 6.50% to 6.96% **and** cuts standard deviation from 5.63% to 3.98%. Holding B alone is dominated.

Asset A did not become safer. The gain comes from A doing well in the states where B does badly, which you can only see by looking at the two together.

Covariance is what survives diversification

For a portfolio with weights w_i , the variance is a double sum over every pair:

$$\sigma_p^2 = \sum_i \sum_j w_i w_j \sigma_{ij} \quad \text{where} \quad \rho_{ij} = \frac{\sigma_{ij}}{\sigma_i \sigma_j}$$

Equally weighted, n assets

$$\sigma_p^2 = \frac{1}{n} \overline{\sigma^2} + \left(1 - \frac{1}{n}\right) \overline{\sigma_{ij}}$$

As n grows the first term vanishes and the second does not. Individual variances disappear; average covariance remains.

This is why only covariance with the market can be priced. The rest was free to remove, so nobody pays you for holding it.


Predict before we look

QUICK POLL · 3 MIN

How many randomly chosen stocks does it take to capture most of the available risk reduction?

- A. About 5
- B. About 20
- C. About 100
- D. There is no limit — more is always meaningfully better

How portfolio risk falls as holdings are added



Most of the benefit arrives in the first twenty holdings, and none of it removes the floor. That floor is average pairwise covariance — the systematic risk a portfolio cannot shed.

Computed from the equally-weighted expression above, with an average standard deviation of 40% and average pairwise correlation of 0.20 and 0.40. Illustrative — read the shape of the curve and treat the levels as arbitrary.

The Capital Asset Pricing Model

If only undiversifiable risk is priced, then expected return is linear in exposure to the market:

$$E(R_i) = R_f + \beta_i (E(R_M) - R_f), \quad \beta_i = \frac{\sigma_{iM}}{\sigma_M^2}$$

What beta measures

How much of an asset's movement is *shared* with the market — a covariance, scaled so that the market itself has $\beta = 1$. A very volatile asset can still have a low beta.

A very volatile stock with zero market covariance has a beta of zero and, under CAPM, earns the risk-free rate.

Worked example: a beta, and what it is worth

$$E(R_i) = R_f + \beta_i(E(R_M) - R_f) = 3.97\% + 1.3 \times 5\% = 10.5\%$$

Input	Value	Where it comes from
Risk-free rate	3.97%	Observed: the 3-month US Treasury bill yield, 17 September 2026.
Beta	1.3	Estimated: regress the stock's excess returns on the market's, often over five years of monthly data; the slope is beta. Another period or index gives a different number.
Market risk premium	5%	Assumed: the extra return investors expect from the market over the risk-free rate. It cannot be observed; estimates from historical returns or forward-looking models range from about 3% to 7%.

The answer depends on the assumed premium

With a premium of 4%, 5% or 6%, the required return is 9.2%, 10.5% or 11.8%. Report all three, so that the reader can see how much the answer depends on the assumption.

COLD CALL · 3 MIN

A gold miner's share price swings wildly but barely moves with the index. Under CAPM, what return should it offer?

Beta, by company

LOW BETA — DEFENSIVE		
General Mills	Food	0.20
Consolidated Edison	Utility	0.28
Newmont Mining	Gold	0.32
Wal-Mart	Supermarket	0.35
Coca-Cola	Drinks	0.54

HIGH BETA — CYCLICAL		
Walt Disney	Entertainment	1.21
Apple	Computer	1.26
PulteGroup	Construction	1.28
Marriott	Hotel	1.48
Tiffany	Apparel	1.80

Beta sorts by *industry*: whatever people put off buying in a downturn has a high beta. Gold mining is volatile and defensive at the same time, which is why its beta is 0.32 while its share price swings more than Tiffany's.

Estimated over 2007–2012; levels shift with the window, the ordering by industry does not.

Market efficiency and its limits

The efficient market hypothesis: prices reflect all available information, so a strategy using that information cannot earn more than its risk justifies — no **alpha**. Three forms, each including the one above.

Form	Information already in prices	So this cannot earn alpha	Evidence
Weak	Past prices and trading volume	Trading on past price patterns	Largely holds in major markets
Semi-strong	All public information: earnings, news, announcements	Analysing public information	Mostly holds, with exceptions: the anomalies
Strong	All information, including private information	Even trading on inside information	Fails: insiders earn alpha, hence insider-trading laws

To call a return abnormal you need a model of the normal return, such as the CAPM, so every test of efficiency is also a test of that model.

What an anomaly is

A pattern the model cannot explain

A group of stocks, chosen by some characteristic, that has earned more than its beta says it should. The extra return is the group’s alpha. If the market were efficient and the CAPM were right, it would be zero.


Anomaly	The pattern	First documented
Size	Small firms have earned more than large firms	Banz (1981)
Value	Cheap stocks, with low prices relative to book value, have earned more than expensive ones	Fama and French (1992)
Momentum	Stocks that rose over the past year have kept rising for a few months	Jegadeesh and Titman (1993)

An anomaly means either that the market misprices those stocks or that the model has left out a risk. Researchers have now published more than 400 of them.

Banz, Journal of Financial Economics 1981; Fama and French, Journal of Finance 1992; Jegadeesh and Titman, Journal of Finance 1993.

Research Spotlight

Depending on the statistical method, between 35% and 82% of published anomalies survive replication



Hou, K., Xue, C. & Zhang, L. (2020). Replicating Anomalies. *Review of Financial Studies* 33(5). Jensen, T. I., Kelly, B. T. & Pedersen, L. H. (2023). Is There a Replication Crisis in Finance? *Journal of Finance* 78(5).

Alpha, beta and smart beta

Three ways to invest

Alpha investing tries to earn more than the risk justifies, by trading on mispricing such as the anomalies — active funds, hedge funds, quant funds.

Beta investing holds the market through an index fund and accepts the market’s return, at very low cost.

Smart beta is an index fund that follows fixed rules to tilt towards a known factor — value, size, momentum, low volatility. It is a published anomaly, sold as a cheap product.

–58%

average fall in the returns of 97 anomalies after the research was published

7.6% → 0.8%

abnormal return to joining the S&P 500, 1990s against the last decade

2023 / 2024

years in which passive funds overtook active funds — in the US by assets, in China by A-share holdings

Once an anomaly is known and traded, its alpha shrinks. Beating the market after fees has become harder, and money has moved to beta and smart beta.

McLean and Pontiff, Journal of Finance 2016; Greenwood and Sammon, Journal of Finance 2025; Morningstar, January 2024 (US); Bloomberg, 28 October 2024 (China: ¥3.16tn passive against ¥2.84tn active, end-September 2024).

Predict before we look

How well a market prices information also depends on who is trading in it.

SHOW OF HANDS · 4 MIN

In the A-share market about 90% of trading is done by individual investors. Is such a market more efficient or less efficient than one dominated by institutions?

- A. Less efficient — individual traders are less informed
- B. More efficient — more traders means more information in prices
- C. It depends on what information we mean
- D. The question is not well posed

Who sets prices in the A-share market

Retail share of turnover, %



¥405tn

A-share turnover in 2025 — about US\$57.6tn, the first year above ¥400 trillion

3–4%

of market capitalisation held by foreign investors

What this means for efficiency

News gets into prices quickly, because so many people react to it. What a company is worth gets in more slowly, because few individual traders do that analysis. So the answer to the vote is C: it depends on what information we mean.

*Retail turnover shares: HSBC estimate reported by CNBC, 20 January 2026 — an estimate, not an official exchange statistic.
Turnover: Shanghai Stock Exchange, December 2025.*

Closing case — China's quant funds in the small-cap sell-offs of 2024 and 2026

SITUATION

Quant funds pick stocks with statistical models. In China their assets more than doubled in under a year, to over **¥2.6 trillion** by mid-2026, and 71 managers ran more than ¥10 billion each. Twice they lost heavily when small-cap stocks fell. On **19 February 2024** one fund sold ¥2.57 billion of shares in a single minute, and the exchanges suspended its trading for three days. In the week to **17 July 2026**, 73 small-cap index-enhanced quant products lost an average of **14%**.

FRICTION

Many quant funds trade on similar signals and hold similar stocks. In normal times their trading narrows mispricing and adds liquidity. When their signals fail at the same time they sell at the same time, and borrowed money makes some of them sell faster.

DECIDE

Do quant funds make the A-share market more efficient, or less stable? As the regulator, would you limit their borrowing, their trading speed or their size — or leave them alone?

THINK – PAIR – SHARE · 10 MIN

Take one side, then say what evidence would settle it.

Assets: Bloomberg, July 2026; managers above ¥10bn: AMAC via Caixin, 27 May 2026. February 2024: Shanghai and Shenzhen stock exchanges, 20 February 2024. July 2026: Bloomberg, reported 23 July 2026.

Main takeaways

What we covered

1. The risk structure: why two bonds maturing on the same day trade at different yields
2. The term structure: three facts about the yield curve, and the three theories that try to explain them
3. Measuring return and risk: variance, covariance, diversification, and the Capital Asset Pricing Model
4. Market efficiency, anomalies, alpha and beta investing, and who sets prices in the A-share market

The points to remember

- A yield gap is **default risk, liquidity and tax treatment** — only the first is about the borrower
- A US BBB bond pays **95bp** over Treasuries, a Chinese AA note **41bp** — and the domestic-to-global rating gap is **7 notches**
- A long rate is expected short rates plus a term premium. The curve was inverted **26 months** and no recession followed
- At a correlation of **-0.50**, **21%** in the riskier asset beats the safer one alone; about **20** holdings capture most of the diversification
- Beta **1.3** at a 5% premium gives **10.5%**; at 4% or 6% it is 9.2% or 11.8%
- Between **35%** and **82%** of published anomalies survive replication, and their returns fall **58%** once published

Questions to Think About

1. A Chinese AA note yields 41bp over government debt. Write down two hypotheses that would produce that number, and one observation that would distinguish them.
2. The liquidity premium theory explains all three empirical facts. Is that a strength or a warning sign?
3. If diversification removes all but the covariance term, why do most households hold fewer than ten stocks?
4. Name an A-share anomaly you would expect to disappear as institutional ownership rises, and say through what mechanism.

Before Next Class

Required

- Mishkin & Eakins — Chapters 5 and 6
- Start Chapter 7 — Week 4 opens on why financial institutions exist at all

Optional

- Jensen, Kelly & Pedersen (2023), *Journal of Finance* 78(5) — today's Spotlight; read the introduction and look at Figure 1
- Greenwood & Sammon (2025), *Journal of Finance* — the index-inclusion anomaly and its disappearance

Week 4 is different: there are no formulas for two weeks, and the question becomes why financial intermediaries exist at all.

References

- Banz, R. W. (1981). The Relationship between Return and Market Value of Common Stocks. *Journal of Financial Economics* 9(1): 3–18.
- Fama, E. F. & French, K. R. (1992). The Cross-Section of Expected Stock Returns. *Journal of Finance* 47(2): 427–465.
- Greenwood, R. & Sammon, M. (2025). The Disappearing Index Effect. *Journal of Finance* (NBER w30748).
- Greenwood, R. & Vayanos, D. (2014). Bond Supply and Excess Bond Returns. *Review of Financial Studies* 27(3): 663–713.
- Hou, K., Xue, C. & Zhang, L. (2020). Replicating Anomalies. *Review of Financial Studies* 33(5): 2019–2133.
- Jensen, T. I., Kelly, B. T. & Pedersen, L. H. (2023). Is There a Replication Crisis in Finance? *Journal of Finance* 78(5): 2465–2518.
- Jegadeesh, N. & Titman, S. (1993). Returns to Buying Winners and Selling Losers. *Journal of Finance* 48(1): 65–91.
- Jiang, X. & Packer, F. Credit Ratings of Chinese Firms by Domestic and Global Agencies. *BIS Working Paper* 648.
- McLean, R. D. & Pontiff, J. (2016). Does Academic Research Destroy Stock Return Predictability? *Journal of Finance* 71(1): 5–32.
- Mishkin, F. S. & Eakins, S. G. (2018). *Financial Markets and Institutions*, 9th Global Edition. Pearson — Chapters 5 and 6.

Market data on every slide carries its own source line and date. All values checked August 2026.

APPENDIX

Derivations

The portfolio-variance limit, the term-structure algebra, and the security market line. Covered in class only if time allows.

Appendix — the diversification limit

For an equally weighted portfolio of n assets, split the double sum into the n variance terms and the $n(n - 1)$ covariance terms:

$$\sigma_p^2 = \frac{1}{n^2} \sum_i \sigma_i^2 + \frac{1}{n^2} \sum_i \sum_{j \neq i} \sigma_{ij} = \frac{1}{n} \overline{\sigma^2} + \frac{n-1}{n} \overline{\sigma_{ij}}$$

Take the limit

$$\lim_{n \rightarrow \infty} \sigma_p^2 = \overline{\sigma_{ij}}$$

The average variance is divided by n and vanishes. The average covariance keeps its full weight however many assets you add. Everything the CAPM says follows from this one limit.

Appendix — the minimum-variance weight

For two assets with weight w in A, portfolio variance is

$$\sigma_p^2 = w^2\sigma_A^2 + (1 - w)^2\sigma_B^2 + 2w(1 - w)\sigma_{AB}$$

Differentiate with respect to w , set to zero, and solve:

$$w^* = \frac{\sigma_B^2 - \sigma_{AB}}{\sigma_A^2 + \sigma_B^2 - 2\sigma_{AB}}$$

On the session's numbers

$\sigma_A^2 = 238.5$, $\sigma_B^2 = 31.8$, $\sigma_{AB} = -43.8$, so $w^* = 75.6/357.9 = 0.21$ — 21% in the risky asset.

Note that $w^* > 0$ whenever $\sigma_{AB} < \sigma_B^2$. A negative covariance is not required; the riskier asset earns a place in the portfolio as long as it is not too closely tied to the safer one.

Appendix — forward rates from the curve

Under pure expectations, a two-period investment must pay the same either way — roll two one-period bonds, or hold one two-period bond:

$$(1 + i_{2t})^2 = (1 + i_t)(1 + i_{t+1}^e) \quad \implies \quad i_{t+1}^e = \frac{(1 + i_{2t})^2}{1 + i_t} - 1$$

Generalising, the implied one-period forward rate between $n - 1$ and n is

$$f_n = \frac{(1 + i_{nt})^n}{(1 + i_{n-1,t})^{n-1}} - 1$$

Under the liquidity premium theory the forward rate is the expected future short rate *plus* the term premium, so a forward rate sits above the expected future short rate by exactly the size of that premium.

Appendix — from the covariance limit to the security market line

Add asset i in small weight w to the market portfolio. The marginal contribution to portfolio variance is

$$\left. \frac{\partial \sigma_p^2}{\partial w} \right|_{w=0} = 2(\sigma_{iM} - \sigma_M^2)$$

so what matters for the portfolio is σ_{iM} , not σ_i^2 . Scaling by the market variance gives $\beta_i = \sigma_{iM}/\sigma_M^2$, and requiring that every asset offer the same reward per unit of that exposure gives

$$E(R_i) = R_f + \beta_i(E(R_M) - R_f)$$

The security market line is a no-arbitrage condition on marginal contributions to portfolio variance. Volatility never enters it.