COMPUTATIONAL FINANCE & FINTECH

Risk, Return and Uncertainty

Historical evidence, diversification and systematic risk

Fatih Kansoy

Worcester College, University of Oxford

Day 4 · Summer 2026

LECTURE OUTLINE

Risk, return and uncertainty

From realised market evidence to the risk that remains after diversification.

PART 1 OF 4

Evidence from market history

What does a long record of realised returns tell us about reward and downside risk?

Compare reinvested wealth, drawdowns and rolling holding periods before introducing statistical measures of risk.

An investment decision is made before returns are known

An investor choosing today does not know which asset will finish with the highest wealth.

TODAY

Choose an investment

Expectations, horizon and capacity to bear losses

→

HOLDING PERIOD

Returns are realised

Gains, losses and cash distributions arrive over time

→

LATER

Observe final wealth

One realised outcome, known only after the fact

Market history is useful evidence, but hindsight must not be mistaken for information available at the decision date.

How the historical comparison is constructed

To compare assets fairly, each series starts with the same 100 and includes all cash distributions.

$$R_t = \frac{P_t - P_{t-1} + D_t}{P_{t-1}}, \qquad W_t = 100 \prod_{s=1}^{t} (1 + R_s).$$

P_t: end-of-year price; D_t: dividend or interest received during year t; R_t: total return; W_t: accumulated wealth with distributions reinvested.

Common starting value
Every asset begins at 100 at the end of 1927.

Reinvested income
Dividends and interest buy more units of the same asset.

Logarithmic scale
Equal vertical distances represent equal proportional changes.

The continuous 1928–2025 U.S. sample compares large stocks, small stocks, Baa corporate bonds, Treasury bills and consumer prices. Taxes and transaction costs are excluded; CPI is a purchasing-power benchmark, not an investable asset.

Long-run wealth and drawdowns, 1928–2025

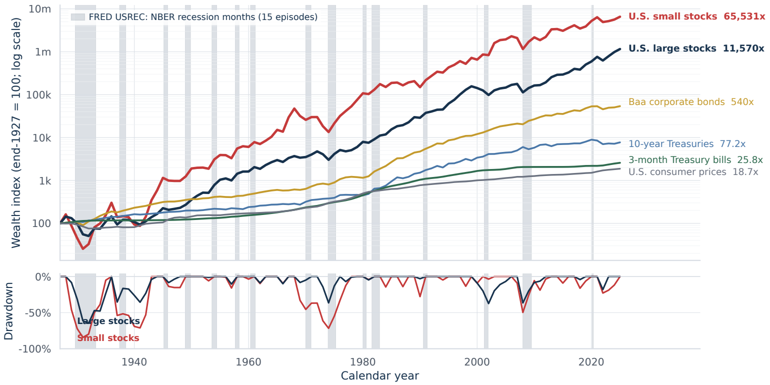

Stocks produced the highest ending wealth, but their deepest losses occurred around recessions, when income, credit and liquidity could also be weak.

Drawdown at t: $W_t / \max_{s \leq t} W_s - 1$. A value of -40% means wealth is 40% below its previous peak.

[2–4]

How one holding-period marker is calculated

Choose a starting year t and a holding period of h years.

$$\boxed{W_{t,h} = £100 \prod_{s=0}^{h-1} (1 + R_{t+s})}$$

R_{t+s}: total return in year $t+s$; $W_{t,h}$: terminal wealth from investing £100 at the start of year t for h years.

Horizontal coordinate

Plot the marker at the starting year t.

Vertical coordinate

Plot its height at the calculated terminal wealth $W_{t,h}$.

The next marker restarts with a new £100 in year $t+1$. Adjacent multi-year windows overlap; the markers are not one continuous investment path.

[2–4] Appendix A1

One-year outcomes vary across assets and starting years

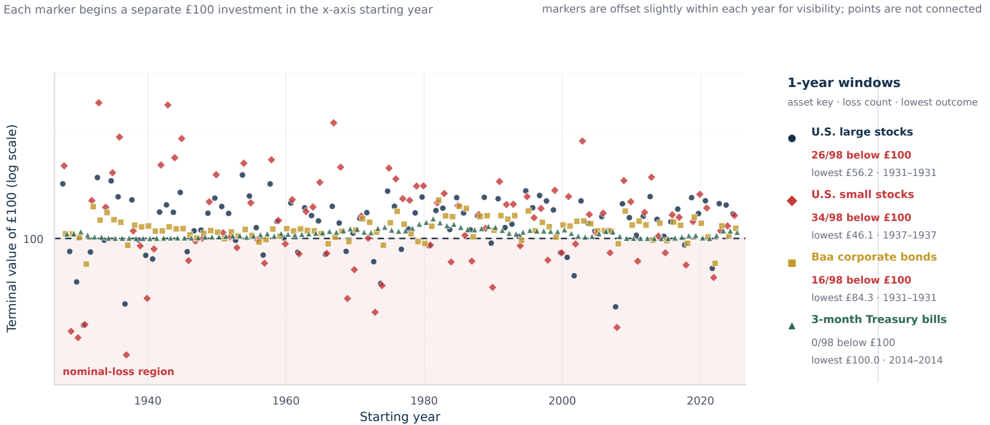

Nominal losses occurred in 34 small-stock, 26 large-stock and 16 Baa-bond windows; Treasury bills stayed at or above £100.

[2–4]

Five-year stock losses persisted

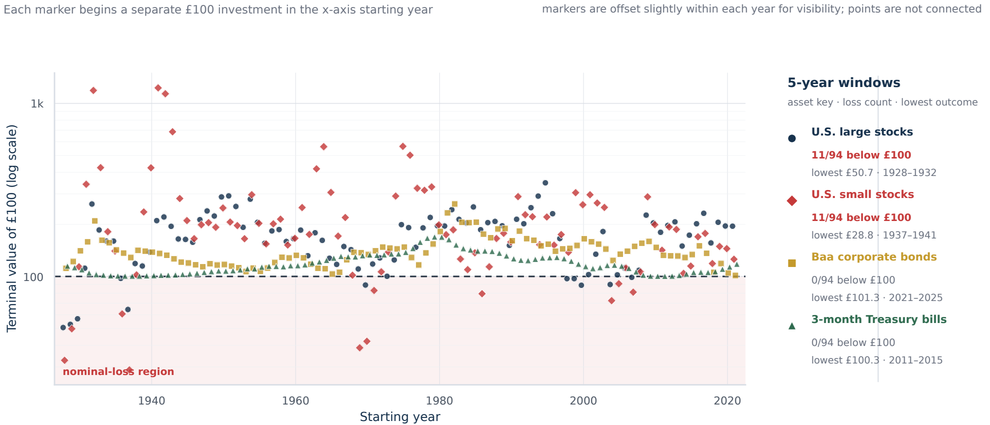

Each stock portfolio still had 11 nominal-loss windows; no five-year Baa-bond or Treasury-bill window fell below £100.

[2–4]

Ten-year stock losses were rare, not impossible

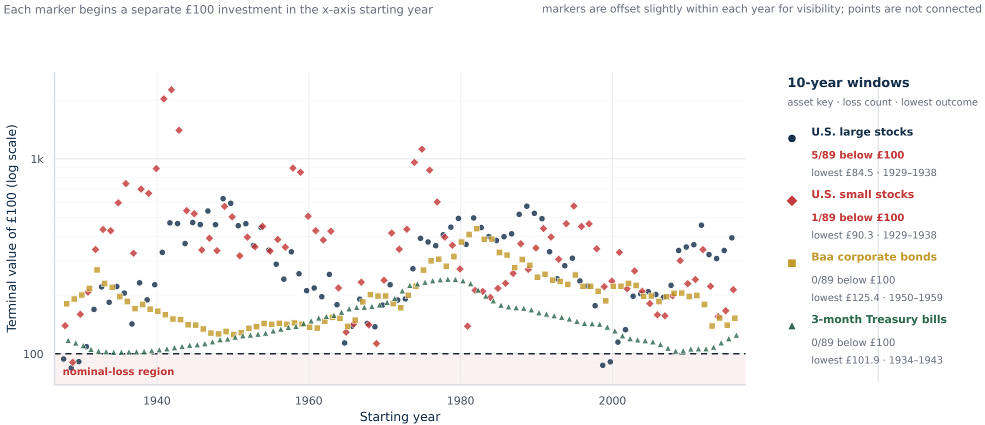

Only 5 large-stock and 1 small-stock window lost nominal principal; Baa bonds and Treasury bills remained above £100.

[2–4]

All four assets avoided twenty-year nominal losses here

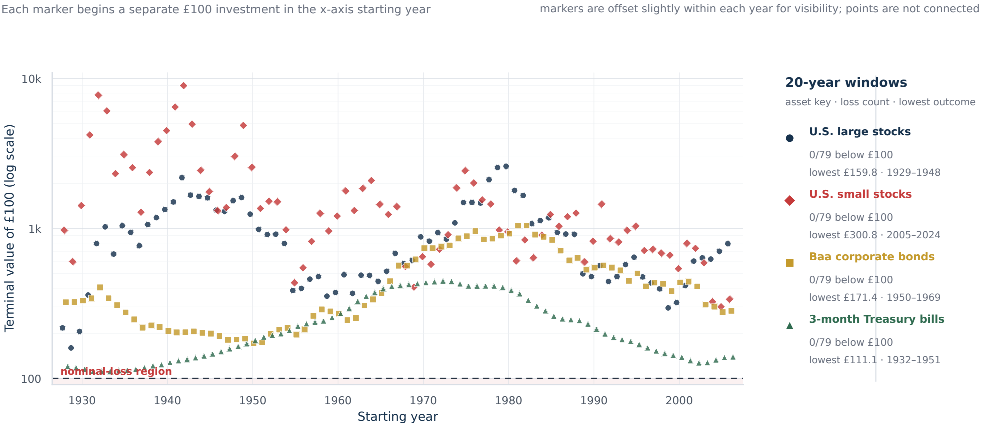

Every twenty-year window ended above £100 for all four assets in this sample. This historical result is evidence, not a guarantee.

[2–4]

How the evidence changes with the horizon

Horizon	Large stocks	Small stocks	Baa bonds	Treasury bills
1 year	26/98	34/98	16/98	0/98
5 years	11/94	11/94	0/94	0/94
10 years	5/89	1/89	0/89	0/89
20 years	0/79	0/79	0/79	0/79

- Stock outcomes were wider, and their nominal losses persisted over longer horizons.
- Baa-bond losses appeared only at one year; Treasury bills had no nominal-loss window.
- No nominal loss does not rule out an interim drawdown, a forced sale, or lagging a safer asset.

Cells report windows ending below the initial £100. The windows overlap and come from one U.S. history; they are not independent observations or a guarantee.

[2–4]

What market history can and cannot tell us

WHAT THE EVIDENCE SHOWS

- Stocks earned much higher realised long-run returns.
- Their losses were larger and sometimes persistent.
- Outcomes varied substantially with the starting date.

WHAT IT CANNOT IDENTIFY

- The exact expected return for the next period
- Which component of risk earns compensation
- The expected return of a particular security next period

The record is consistent with a reward for bearing severe losses; it does not reveal what return investors expected beforehand.

Measuring return and uncertainty

How can we summarise an unknown return distribution, and what can realised data tell us about it?

Define expected return and volatility, estimate them from history, and measure the uncertainty in those estimates.

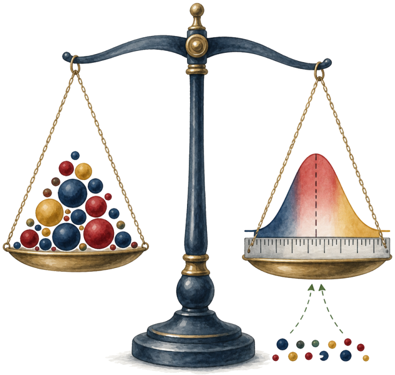

A probability distribution describes possible returns

Before the year, a distribution describes what might happen; after the year, only one return is realised. Share A costs £100 today.

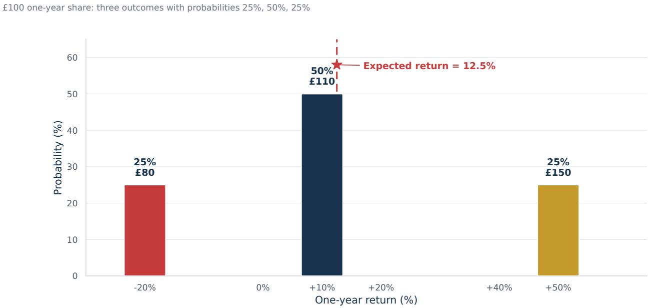

A **probability distribution** lists each possible return R_j and its probability p_j, where $p_j \geq 0$ and $\sum_j p_j = 1$.

j: possible state; R_j: one-year total return in state j; p_j: probability assigned to that state.

[1]

Expected return is a probability-weighted average

The expected return locates the centre of the stated distribution.

$$\boxed{\mathbb{E}[R] = \sum_j p_j R_j = 12.5\%}$$

What the number means

More probable outcomes receive more weight. Across many repetitions, the average realised return would approach 12.5%.

What it does not mean

It is not a promised return. Here, 12.5% is not even one of the three possible outcomes.

$\mathbb{E}[R]$: expected one-year total return; the probability-weighted centre of the distribution.

[1] Appendix A2

Uncertainty is the spread of possible returns

The centre alone does not reveal how far the realised return may land from it.

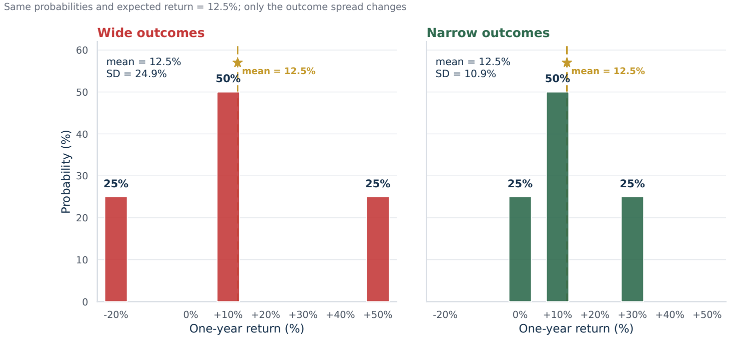

Both shares have $\mathbb{E}[R] = 12.5\%$, but Share A is more uncertain because its possible returns lie farther from that centre.

[1]

Volatility turns dispersion into one number

Standard deviation summarises the distance of possible returns from the expected return.

$$\text{deviation in state } j : d_j = R_j - \mathbb{E}[R],$$
$$\text{variance} : \operatorname{Var}(R) = \sum_j p_j d_j^2,$$
$$\text{volatility} : \boxed{\sigma = \operatorname{SD}(R) = \sqrt{\operatorname{Var}(R)}}.$$

Why square the deviations?
Positive and negative deviations cannot cancel.

Why take the square root?
Volatility returns to the same percentage units as R.

For Share A, the calculation gives $\sigma = 24.9\%$.

$\operatorname{Var}(R)$: variance, measured in squared-return units; $\operatorname{SD}(R)$: standard deviation; σ (sigma): the symbol used for population volatility.

[1] Appendix A3

How to interpret a volatility number

Suppose a one-year return distribution has $\mathbb{E}[R] = 12.5\%$ and $\sigma = 24.9\%$.

WHAT σ TELLS US

- Possible returns are widely dispersed around 12.5%.
- A larger σ means a wider distribution, holding its centre fixed.
- σ is measured in percentage-return units.

WHAT σ DOES NOT TELL US

- It is not the probability of losing money.
- It is not a worst-case loss or a guaranteed range.
- It treats equally distant gains and losses symmetrically.

Volatility measures dispersion; it does not by itself describe downside risk.

A realised return includes cash income and the price change

Suppose Share A's high outcome occurs: $P_0 = £100$, dividend $D_1 = £2$, and $P_1 = £148$.

$$R_1 = \frac{D_1 + P_1 - P_0}{P_0} = \underbrace{\frac{£2}{£100}}_{\text{dividend yield } = 2\%} + \underbrace{\frac{£148 - £100}{£100}}_{\text{capital-gain rate } = 48\%} = 50\%.$$

$$£100 \xrightarrow[+£48 \textbf{ price}]{+£2 \textbf{ cash}} £150$$ ending price plus cash received

P_0: beginning price; P_1: ending price; D_1: dividend paid during the holding period; R_1: the return that actually occurred. For a bond, replace the dividend with the coupon or other cash payment received.

[1] Appendix A4

Multi-period returns compound rather than add

Reinvesting after each period means that each return applies to the wealth carried into that period.

$$\boxed{1 + R_{0\to m} = \frac{W_m}{W_0} = \prod_{q=1}^{m}(1 + R_q)}.$$

Two-period example

$$\begin{aligned} £100 \times 1.18 \times 0.82 &= £96.76, \\ R_{0\to 2} &= -3.24\%. \end{aligned}$$

- After the $+18\%$ return, wealth is £118.
- The -18% return is then applied to £118, not £100.
- Equal percentage gains and losses do not cancel.

W_0 and W_m: wealth at dates 0 and m; R_q: return during period q; $R_{0\to m}$: the cumulative return from date 0 to date m, equal to $W_m/W_0 - 1$.

[1] Appendix A5

Historical returns form an empirical distribution

Each bar counts the number of calendar years in a return range; each dashed line is that asset's sample mean.

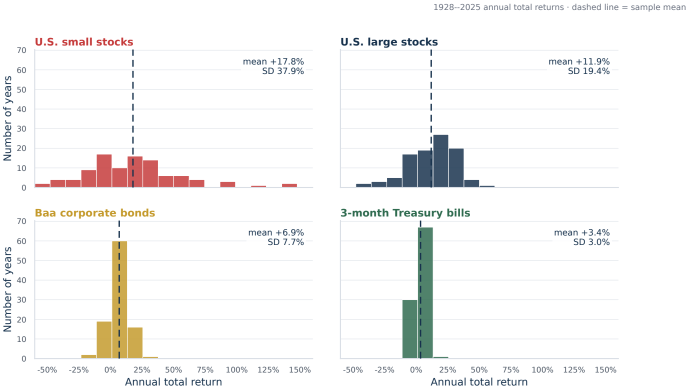

An empirical distribution summarises one realised sample. Treating it as a forecast requires an additional assumption that the underlying return distribution is sufficiently stable.

[2,3]

Historical mean and volatility are sample estimates

With T annual observations, we estimate the centre and dispersion from the realised sample.

$$\bar{R} = \frac{1}{T}\sum_{t=1}^{T} R_t, \qquad s^2 = \frac{1}{T-1}\sum_{t=1}^{T}(R_t - \bar{R})^2, \qquad s = \sqrt{s^2}.$$

U.S. series, 1928–2025	Arithmetic mean $\bar{R}$	Sample volatility s
Large stocks	11.9%	19.4%
Small stocks	17.8%	37.9%
Baa corporate bonds	6.9%	7.7%
Treasury bills	3.4%	3.0%

R_t: realised return in year t; $T = 98$; $\bar{R}$: arithmetic sample mean; s: sample standard deviation. The denominator $T-1$ accounts for estimating $\bar{R}$ from the same data.

[2,3] Appendix A6

How precisely does history estimate expected return?

We observe realised returns, not the return investors expected at the start of each year. The sample mean therefore remains an estimate.

$$\begin{aligned}\mathrm{SE}(\bar{R}) &= \frac{s}{\sqrt{T}} \\ &= \frac{19.40\%}{\sqrt{98}} = 1.96\ \text{pp}, \\ \bar{R} \pm 2\,\mathrm{SE}(\bar{R}) &= 11.85\% \pm 3.92\% \\ &\approx [7.9\%, 15.8\%].\end{aligned}$$

Even 98 annual observations leave a wide range for the underlying mean.

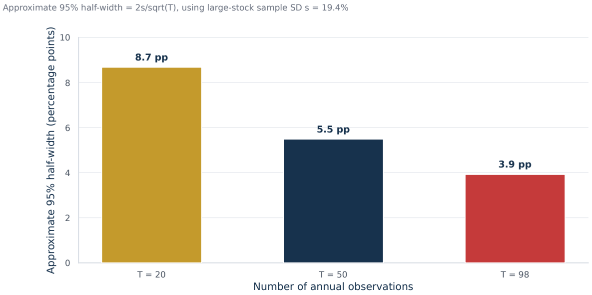

This approximate 95% interval assumes independent annual observations from one stable distribution. Serial dependence or structural change can make it too narrow.

[2,3] Appendix A7

Which average answers which question?

The same return history can support two different tasks. The appropriate average depends on what we are trying to measure.

FORECAST ONE FUTURE PERIOD

$$\bar{R} = \frac{1}{T}\sum_{t=1}^{T} R_t$$

Use the **arithmetic mean** to estimate a one-period expected return, provided the historical distribution is a relevant guide.

REPORT THE REALISED WEALTH PATH

$$g = \left[\prod_{t=1}^{T}(1+R_t)\right]^{1/T} - 1$$

Use the **compound annual growth rate** to report the constant annual rate that reproduces actual terminal wealth.

$\bar{R}$: arithmetic sample mean; g: compound annual growth rate (CAGR), also called the geometric mean return. Neither statistic answers both questions.

[1] Appendix A8

Equal arithmetic means can produce different compound growth

Both paths have an arithmetic mean return of 5% per year. Their volatility and terminal wealth differ.

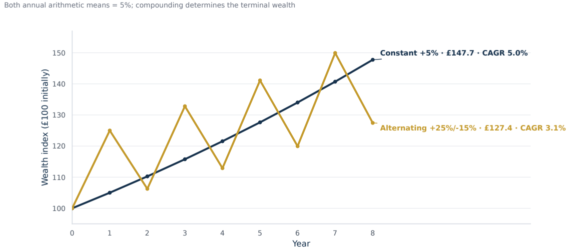

What the figure shows

The constant path compounds at 5.0%; the alternating path compounds at only 3.1%.

Why the paths separate

Percentage returns multiply wealth. The -15% loss acts on the wealth left by the preceding $+25\%$ gain.

Conclusion: equal arithmetic means do not imply equal terminal wealth or equal realised performance.

[1] Appendix A9

A historical average cannot pin down expected return

A realised historical average is not a sufficiently precise forecast of a particular security's expected return.

1. **Sampling error:** the underlying expected return is estimated imprecisely.
2. **Stability:** future valuations, institutions and economic conditions may differ from the historical sample.
3. **Portfolio relevance:** standalone volatility does not show how the security moves with risks investors already hold.

The next question is how combining assets changes the uncertainty an investor bears.

Risk inside a portfolio

Which risks disappear when holdings are combined, and which risks remain?

Separate independent from common shocks, trace how portfolio volatility changes, and derive the pricing implication.

Broad portfolios show a historical risk-return ordering

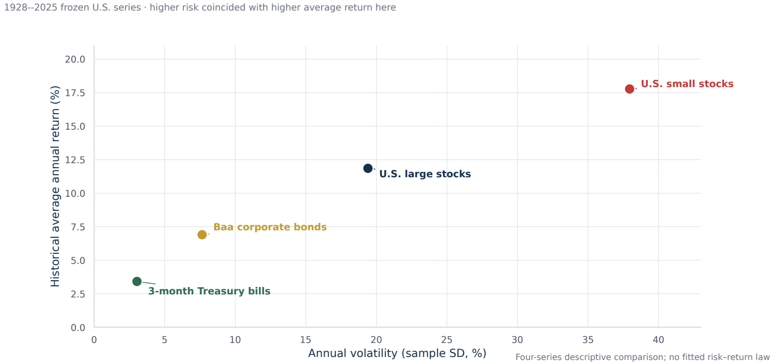

Higher realised average returns coincided with higher volatility across these four broad portfolios.

The pattern describes one historical sample; it is not a rule that every volatile security earns more.

[2,3]

Individual-stock volatility does not rank average returns

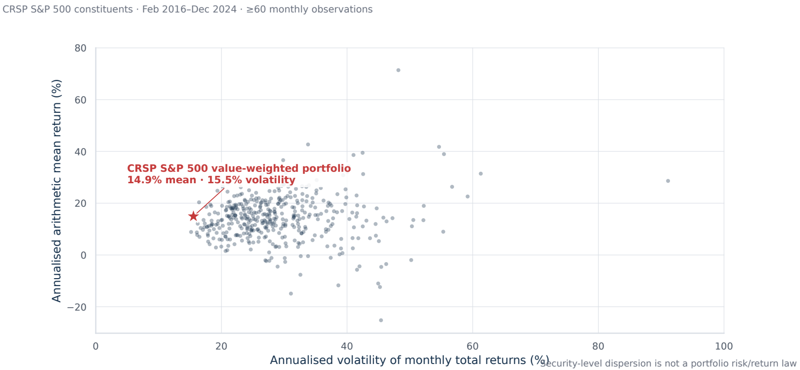

Individual stocks with similar volatility produced very different average returns.

The broad index was less volatile than most constituent stocks because firm-specific movements offset.

[5]

Equal expected default rates can imply different portfolio risk

Consider 400 loans. Each loan has a 4% probability of default. What matters for the portfolio is whether defaults occur separately or together.

INDEPENDENT DEFAULTS

- One default does not change another loan's probability.
- Defaults occur across different borrowers.
- The portfolio default rate stays near 4%.

ONE COMMON SHOCK

- The same event affects every borrower.
- The portfolio default rate is either 0% or 100%.
- Adding more exposed loans does not average out the event.

Both portfolios have an expected default rate of 4%; only the independent portfolio has a tightly concentrated outcome.

[1]

Why averaging independent defaults reduces uncertainty

Let $X_k = 1$ when loan k defaults and $X_k = 0$ otherwise. The portfolio default rate is the average of the indicators:

$$\bar{X} = \frac{1}{N}\sum_{k=1}^{N} X_k, \qquad \mathbb{E}[\bar{X}] = p, \qquad \boxed{\mathrm{SD}(\bar{X}) = \frac{\sqrt{p(1-p)}}{\sqrt{N}}}.$$

The centre does not change

The expected portfolio default rate remains $p = 4\%$, whatever N is.

The spread becomes narrower

Under independence, doubling N does not halve risk; risk falls in proportion to $1/\sqrt{N}$.

X_k: default indicator for loan k; $\bar{X}$: fraction of loans that default; p: default probability; N: number of independent loans. Standard deviations are percentage-point uncertainty in the default rate.

[1] Appendix A10

The 400-loan calculation shows the size of diversification

$$\mathrm{SD}(X_k) = \sqrt{0.04(0.96)} = 19.60 \text{ pp.} \qquad \mathrm{SD}(\bar{X}) = \frac{19.60}{\sqrt{400}} = 0.98 \text{ pp.}$$

Independent defaults reduce uncertainty by a factor of 20; a perfectly common shock retains the full 19.60 percentage points.

[1] Appendix A10

Stock-return uncertainty also has two sources

	Firm-specific news	Market-wide news
Example	A product recall or the departure of a key executive	A recession or an economy-wide interest-rate shock
Across firms	Different firms receive different news	Many firms are affected at the same time
In a broad portfolio	Positive and negative surprises can offset	The common movement remains
Name	Idiosyncratic or diversifiable risk	Systematic or market risk

- Firm-specific news can materially change expected cash flows and value.
- It does not create a separate risk premium for an investor who can diversify it.
- Market-wide news changes outcomes across many holdings and cannot be averaged away.

Portfolio size removes only the idiosyncratic component

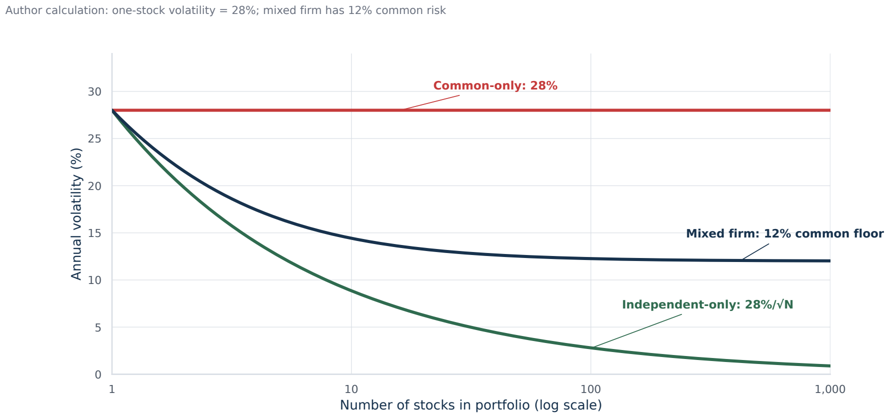

The curve falls rapidly at first because idiosyncratic shocks are being averaged. It approaches, but does not cross, the common-risk floor.

[1]

Ten holdings reveal which component can be diversified

Suppose each individual stock has 28% volatility.

COMMON RISK ONLY

$$\sigma_{p,10} = 28\%.$$

All ten stocks move together. Combining them does not reduce the common component.

INDEPENDENT RISK ONLY

$$\sigma_{p,10} = \frac{28\%}{\sqrt{10}} = 8.9\%.$$

Independent firm shocks offset inside the portfolio at the square-root rate.

BOTH COMPONENTS

$$\sigma_{\text{id}} = \sqrt{(28\%)^2 - (12\%)^2} = 25.3\%, \qquad \sigma_{p,10} = \sqrt{(12\%)^2 + \frac{(25.3\%)^2}{10}} = 14.4\%.$$

Portfolio volatility falls from 28% to 14.4%, but it cannot fall below the 12% common-risk floor.

$\sigma_{p,N}$: volatility of an equal-weight portfolio of N securities. The mixed case assumes zero covariance between common and idiosyncratic components, and independent idiosyncratic shocks across holdings.

Diversifiable risk does not earn a separate premium

Suppose a broad portfolio retained an expected premium after its firm-specific risk had been almost eliminated.

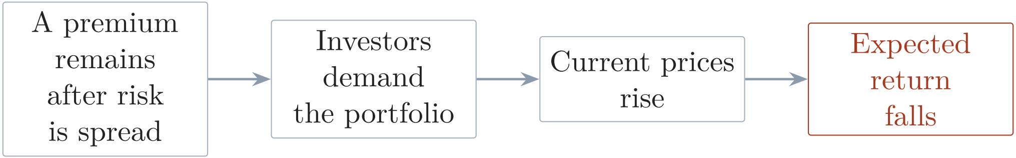

- Competition removes a separate premium for risk that investors can eliminate by combining holdings.
- Systematic risk remains even in a broad portfolio, so bearing it requires compensation.

This is a competitive-pricing argument, not a claim that a finite portfolio is literally risk-free.

Diversification lowers volatility by averaging firm-specific movements

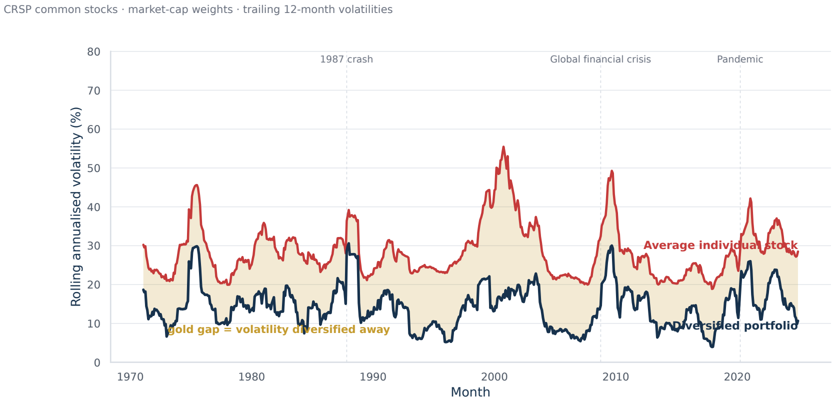

The gold gap is the volatility removed by diversification. The portfolio remains volatile because common market movements survive.

[5]

Diversification benefits can weaken during market crises

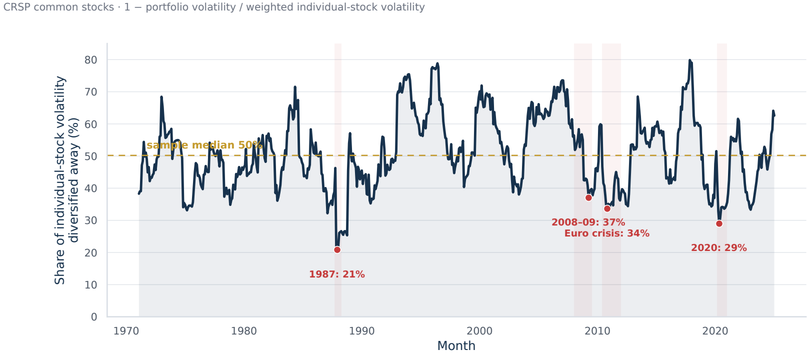

In this sample, the share diversified away fell from 74% at end-2006 to 43% at end-2008.

Diversification still helped, but co-movement made the market-wide component a larger part of total risk.

Diversified-away share $= 1 - \sigma_{\text{portfolio}}/\bar{\sigma}_{\text{individual}}$. A lower value means less of average individual-stock volatility is removed by forming the portfolio.

[5]

More years is not the same as more independent holdings

Can waiting longer diversify risk in the same way as holding more independent assets?

MORE HOLDINGS: AVERAGE

$$\bar{R}_t = \frac{1}{N}\sum_{i=1}^{N} R_{i,t}$$

- Different assets experience different shocks at the same date.
- Independent gains and losses can offset in the portfolio average.

MORE YEARS: COMPOUND

$$\pounds 100 \xrightarrow{-40\%} \pounds 60 \xrightarrow{+40\%} \pounds 84$$

- The second return is earned on the £60 that remains.
- A later $+40\%$ return does not cancel an earlier -40% return.

More holdings can average independent shocks; more years compound one realised sequence of shocks.

N: number of holdings; $R_{i,t}$: return on asset i in period t; $\bar{R}_t$: equal-weighted portfolio return in period t.

Why time compounds one realised path

1. Update wealth for one period

$$W_t = W_{t-1}(1 + R_t).$$

2. Substitute repeatedly

$$W_2 = W_1(1 + R_2) = \big[W_0(1 + R_1)\big](1 + R_2).$$

$$\boxed{W_T = W_0 \prod_{t=1}^{T}(1 + R_t)}.$$

More periods extend one compounded path; they do not diversify it.

W_t: end-of-period wealth; R_t: period-t return; T: number of holding periods.

Systematic risk and required return

How can the risk that survives diversification be measured and translated into a return benchmark?

Define the market benchmark and beta, then scale the expected market risk premium.

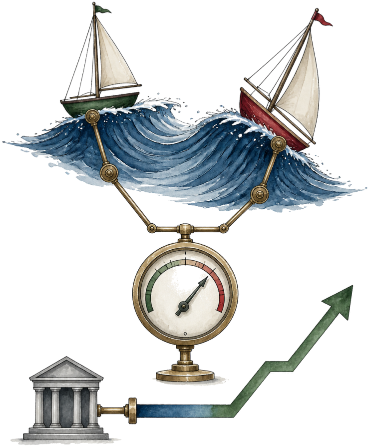

A market-wide benchmark is needed for common risk

We need a portfolio whose movements represent broad shocks rather than firm-specific surprises.

ALL TRADED RISKY CLAIMS
Stocks, bonds and other risky assets

Market portfolio
Broadly diversified conceptual benchmark

OBSERVED PROXY
A broad equity index in practical measurement

- An **efficient portfolio** cannot reduce risk further without lowering expected return.
- The **market portfolio** is the portfolio of all traded risky securities and is treated here as the common-risk benchmark.
- A stock index is observable and useful, but it is only a proxy for that broader portfolio.

Beta measures sensitivity to the market benchmark

Beta is the least-squares slope of security i's return on the market return:

$$\boxed{\beta_i = \frac{\mathrm{Cov}(R_i, R_M)}{\mathrm{Var}(R_M)}}.$$

Numerator: co-movement
$\mathrm{Cov}(R_i, R_M)$ records whether the security and market tend to be above or below their means together.

Denominator: market scale
$\mathrm{Var}(R_M)$ normalises that co-movement by the size of market fluctuations.

> If the market return is 1 percentage point above its mean, the security's conditional expected return is approximately β_i percentage points above its mean.

$\beta_i < 0$: opposite direction $\beta_i = 1$: one-for-one $\beta_i > 1$: amplified

R_i: return on security i; R_M: market-benchmark return. Covariance and variance use the same scenarios or sample periods. Beta is dimensionless and describes an average linear relationship, not a period-by-period promise.

[1]

A two-state calculation makes beta visible

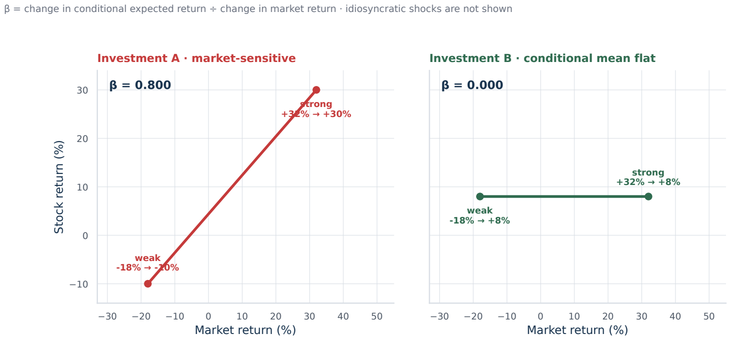

$$\beta_A = \frac{30\% - (-10\%)}{32\% - (-18\%)} = \frac{40}{50} = 0.80.$$

$$\beta_B = \frac{8\% - 8\%}{32\% - (-18\%)} = 0.$$

Investment B may still have firm-specific uncertainty; beta zero only says its conditional mean is unchanged across these two market states.

[1]

Beta and volatility answer different questions

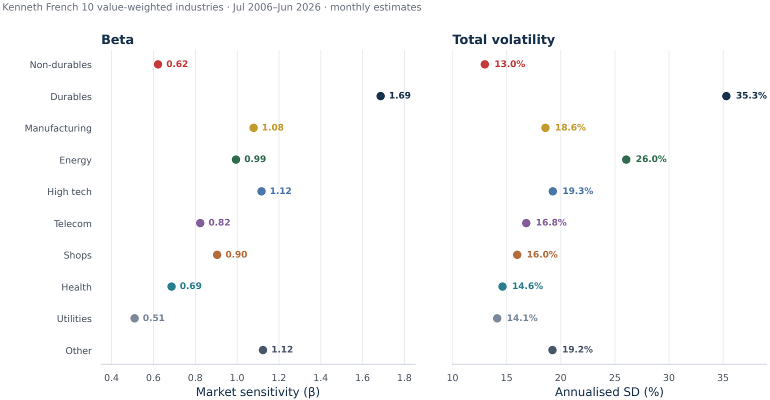

Volatility measures total standalone dispersion.

Beta measures co-movement with the market benchmark.

[3]

The market risk premium is an expected excess return

The return benchmark must compensate for time and for exposure to market-wide risk.

$$\text{Market risk premium} = \mathbb{E}[R_M] - r_f.$$

RISK-FREE RATE r_f

Compensation for postponing consumption over the stated horizon, in the stated currency.

MARKET RISK PREMIUM

Additional expected return for bearing one unit of market exposure.

R_M: market-portfolio return; r_f: matched one-period risk-free rate. $\mathbb{E}[R_M] - r_f$ is forward-looking and is not automatically equal to a historical average excess return.

Beta scales the required premium

What expected return would compensate investors for time and for the market risk that cannot be diversified away?

$$\boxed{\underbrace{r_i}_{\text{required expected return}} = \underbrace{r_f}_{\text{time compensation}} + \underbrace{\beta_i\big(\mathbb{E}[R_M] - r_f\big)}_{\text{systematic-risk compensation}}}.$$

TIME COMPENSATION

r_f

The matched risk-free rate compensates for postponing consumption over the stated horizon and currency.

SYSTEMATIC-RISK COMPENSATION

$$\underbrace{\beta_i}_{\text{risk quantity}} \times \underbrace{\big(\mathbb{E}[R_M] - r_f\big)}_{\text{premium per unit}}$$

Beta scales the market risk premium. It does not scale the risk-free rate.

r_i: benchmark expected return required for security i, not a promised realised return. All terms must use the same period and currency.

[1]

Each term in the return benchmark has one job

r_i	**Required for what?**	Benchmark expected return; not a promised realised return.
r_f	**What is time worth?**	Matched horizon and currency, without market exposure.
β_i	**How much market risk?**	Quantity of exposure; beta 0.8 means 0.8 units.
$\mathbb{E}[R_M] - r_f$	**Premium per unit?**	Expected excess return for one unit of market risk.

$$\underbrace{4\%}_{r_f} + \underbrace{0.8}_{\beta_i} \times \underbrace{5\%}_{\mathbb{E}[R_M]-r_f} = \underbrace{8\%}_{r_i}.$$

The systematic-risk premium is $0.8 \times 5\% = 4\%$; adding the risk-free rate gives $r_i = 8\%$.

All rates share a horizon and currency; $\mathbb{E}[R_M]$ is forward-looking.

Higher systematic sensitivity raises the return benchmark

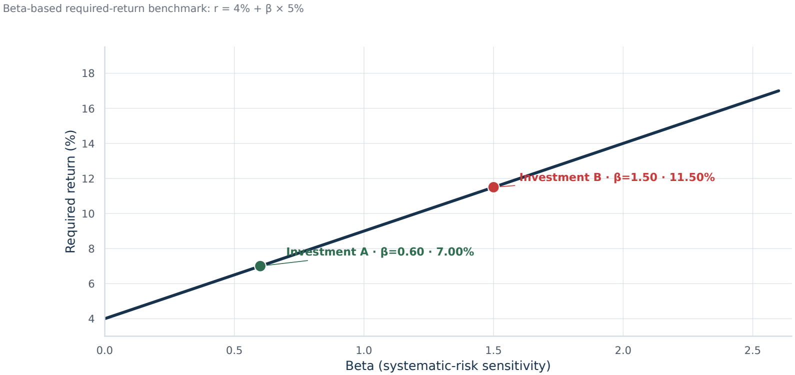

At $r_f = 4\%$ and a 5% market risk premium, beta 0.60 implies $r_A = 7.0\%$.

Beta 1.50 implies $r_B = 11.5\%$ because the market exposure is larger.

If $\beta_i < 0$, the benchmark premium is negative: the payoff tends to help when the market is weak, so investors may accept an expected return below r_f.

[1]

Recap: from realised returns to portfolio-relevant risk

01	**Read market history**	Realised returns reveal wealth paths and drawdowns, but not the exact return investors expected beforehand.
02	**Measure uncertainty**	Expected return describes the centre of a distribution; volatility describes its dispersion.
03	**Separate the shocks**	Independent firm-specific risk can be diversified; common market risk remains.
04	**Measure systematic risk**	Beta records the expected response of a security's return to a market movement.
05	**Set a return benchmark**	The risk-free rate compensates for time; beta scales the expected market risk premium.

The central distinction is not risky versus safe. It is diversifiable risk versus the systematic risk that a broad portfolio still bears.

Technical appendix

- **A1** Holding-period marker
- **A2** Expected return
- **A3** Variance and volatility
- **A4** Realised-return derivation
- **A5** Compounding identity
- **A6** Why use $T-1$?
- **A7** Confidence interval audit
- **A8** Arithmetic mean and CAGR
- **A9** Volatility-drag approximation
- **A10** Variance of an average

Use these pages for the algebra, substitutions and assumptions behind the corresponding lecture result.

← **Back to lecture outline**

A1. One five-year marker from the annual returns

For U.S. large stocks beginning in 1980, the five annual returns are

$$31.7\%,\ -4.7\%,\ 20.4\%,\ 22.3\%,\ 6.1\%.$$

Using the source values at calculation precision,

$$\begin{aligned} W_{1980,5} &= £100(1.317352)(0.952976)(1.204191)(1.223372)(1.061461) \\ &= £196.31. \end{aligned}$$

The chart therefore plots the marker (1980, £196.31).

Moving the start to 1981 removes the 1980 return and adds the 1985 return of 31.2%:

$$W_{1981,5} = £100(0.952976)(1.204191)(1.223372)(1.061461)(1.312351) = £195.56.$$

The two windows share four annual returns, but each marker represents a separate £100 investment with its own starting year. The return list is rounded for reading; the substitution shows calculation precision.

← **Return to holding-period marker**

Appendix map

A2. Expected return: full weighted calculation

State j	Probability p_j	Return R_j	Contribution $p_j R_j$
Weak	0.25	-20%	-5.0%
Middle	0.50	10%	5.0%
Strong	0.25	50%	12.5%
Total	1.00		12.5%

$$\mathbb{E}[R] = 0.25(-0.20) + 0.50(0.10) + 0.25(0.50) = 0.125 = 12.5\%.$$

- Probabilities sum to one, so the contributions form a weighted average.
- The result is the distribution's centre, not a promised outcome.

← **Return to expected return**

Appendix map

A3. Variance and volatility: full calculation

State	p_j	R_j	$d_j = R_j - 0.125$	$p_j d_j^2$
Weak	0.25	−0.20	−0.325	0.026406
Middle	0.50	0.10	−0.025	0.000313
Strong	0.25	0.50	0.375	0.035156
Total	1.00			0.061875

$$\mathrm{Var}(R) = \sum_j p_j d_j^2 = 0.061875, \qquad \sigma = \sqrt{0.061875} = 0.2487 \approx 24.9\%.$$

- Subtract the expected return to measure each state's distance from the centre.
- Square before weighting so positive and negative deviations cannot cancel.
- Take the square root so volatility is again measured in return units.

d_j: deviation in state j; $\mathrm{Var}(R)$: variance in squared-return units; σ: population standard deviation, or volatility.

← **Return to volatility**

Appendix map

A4. Realised return from beginning and ending wealth

An investor pays P_0 at date 0. At date 1, ending wealth is the sale price plus the dividend: $W_1 = P_1 + D_1$. Hence

$$R_1 = \frac{W_1 - P_0}{P_0} = \frac{P_1 + D_1 - P_0}{P_0} = \underbrace{\frac{D_1}{P_0}}_{\text{dividend yield}} + \underbrace{\frac{P_1 - P_0}{P_0}}_{\text{capital-gain rate}}.$$

With $P_0 = £100$, $D_1 = £2$, and $P_1 = £148$,

$$R_1 = \frac{2}{100} + \frac{148 - 100}{100} = 2\% + 48\% = 50\%.$$

For a bond, the coupon replaces the dividend. Ignoring the cash payment understates total return.

← **Return to realised return**

Appendix map

A5. Why multi-period returns contain a cross-product

Two periods of wealth accumulation give

$$W_2 = W_0(1 + R_1)(1 + R_2).$$

Therefore

$$\begin{aligned} 1 + R_{0\to 2} &= (1 + R_1)(1 + R_2), \\ R_{0\to 2} &= R_1 + R_2 + R_1 R_2. \end{aligned}$$

For $R_1 = 18\%$ and $R_2 = -18\%$,

$$R_{0\to 2} = 0.18 - 0.18 + (0.18)(-0.18) = -0.0324 = -3.24\%.$$

Why addition fails

The second return applies to W_1, not to W_0.

What the cross-product records

$R_1 R_2$ captures the changed wealth base carried into period 2.

← **Return to multi-period returns**

Appendix map

A6. Why sample variance divides by $T - 1$

Define each sample deviation as $d_t = R_t - \bar{R}$. Because $\bar{R}$ is estimated from the same data,

$$\sum_{t=1}^{T} d_t = 0 \quad \Longrightarrow \quad d_T = -\sum_{t=1}^{T-1} d_t.$$

Once $T - 1$ deviations are known, the last is fixed: only $T - 1$ deviations are free.

$$s^2 = \frac{1}{T-1} \sum_{t=1}^{T} (R_t - \bar{R})^2.$$

If the true mean were known

Deviations from that fixed mean could be averaged using T.

When the mean is estimated

Dividing by $T - 1$ corrects the downward bias caused by fitting $\bar{R}$ to the sample.

This Bessel correction is justified for the usual independent-sampling variance estimator.

← **Return to sample estimates**

Appendix map

A7. Every step in the approximate 95% interval

For large stocks, $s = 19.40\%$, $T = 98$, and $\bar{R} = 11.85\%$.

$$\mathrm{SE}(\bar{R}) = \frac{s}{\sqrt{T}} = \frac{19.40\%}{\sqrt{98}} = 1.9597 \text{ percentage points.}$$

The lecture uses the rounded two-standard-error rule:

$$\begin{aligned}
\text{margin of error} &= 2(1.9597) = 3.9194 \text{ pp},\\
\text{lower endpoint} &= 11.85 - 3.9194 = 7.9306\% \longrightarrow 7.9\%,\\
\text{upper endpoint} &= 11.85 + 3.9194 = 15.7694\% \longrightarrow 15.8\%.
\end{aligned}$$

Using the exact standard-normal critical value instead,

$$11.85\% \pm 1.96(1.9597 \text{ pp}) = [8.0\%, 15.7\%] \text{ after rounding.}$$

The value 1.96 cuts off 2.5% in each tail of a standard normal distribution. The calculation assumes independent observations from a stable distribution; dependence or structural change can widen uncertainty.

← **Return to estimation precision**

Appendix map

A8. Arithmetic mean and CAGR answer different tasks

For returns of $+25\%$ and -15%, the arithmetic mean is

$$\bar{R} = \frac{25\% - 15\%}{2} = 5\%.$$

It gives equal weight to the two annual observations.

The compound annual growth rate is the constant g that reproduces terminal wealth:

$$(1+g)^2 = (1.25)(0.85) = 1.0625 \quad \Longrightarrow \quad g = \sqrt{1.0625} - 1 = 3.08\%.$$

Expected one-period return
Under a stable sampling model, estimate it with the arithmetic mean.

Realised multi-period performance
Report the constant rate that reproduces the compounded wealth path.

← **Return to choice of average**

Appendix map

A9. Why volatility lowers compound growth

In the two-period example, deviations from the 5% arithmetic mean are +20% and −20%. Their population variance is

$$\sigma^2 = \frac{(0.20)^2 + (-0.20)^2}{2} = 0.04, \qquad \sigma = 20\%.$$

For moderate returns, the log approximation $\log(1+R) \approx R - R^2/2$ implies

$$g \approx \bar{R} - \frac{1}{2}\sigma^2 = 5\% - \frac{1}{2}(0.20)^2 = 3.0\%.$$

The exact CAGR is 3.08%, so the approximation captures the direction and approximate size of the gap.

Greater dispersion lowers compound growth because losses reduce the wealth base available for subsequent gains.

The approximation is not the definition of CAGR. Exact compound growth always comes from the product of gross returns.

← **Return to the two wealth paths**

Appendix map

A10. Covariance determines whether averaging reduces risk

For $\bar{X} = N^{-1} \sum_{k=1}^{N} X_k$,

$$\operatorname{Var}(\bar{X}) = \frac{1}{N^2} \left[\sum_{k=1}^{N} \operatorname{Var}(X_k) + 2 \sum_{i<j} \operatorname{Cov}(X_i, X_j) \right].$$

INDEPENDENT DEFAULTS

$$\operatorname{Cov}(X_i, X_j) = 0$$

$$\operatorname{Var}(\bar{X}) = \frac{p(1-p)}{N}.$$

Hence $\operatorname{SD}(\bar{X}) = \sqrt{p(1-p)}/\sqrt{N}$.

PERFECTLY COMMON SHOCK

$$X_1 = X_2 = \cdots = X_N$$

$$\operatorname{Var}(\bar{X}) = p(1-p).$$

Averaging adds no protection because every indicator moves together.

Covariance measures co-movement. The square-root rule is therefore an independence result, not a mechanical consequence of holding many positions.

← **Return to the square-root rule**

Appendix map

References

[1] Jonathan Berk and Peter DeMarzo (2020), *Corporate Finance*, 5th ed., Pearson, Chapter 10.

[2] Aswath Damodaran, *Historical Returns on Stocks, Bonds and Bills*, NYU Stern.

[3] Kenneth R. French, Data Library: U.S. size, industry and research-factor portfolios.

[4] Federal Reserve Bank of St. Louis, FRED series CPIAUCNS and USREC.

[5] Center for Research in Security Prices, CRSP U.S. Stock Database via WRDS.