Lesson 4.1: Measuring Market Risk

Module 4: The Risk Problem

Prof. Dr. Joerg Osterrieder

Digital Finance — BSc Course

How Much Can We Lose Tomorrow?

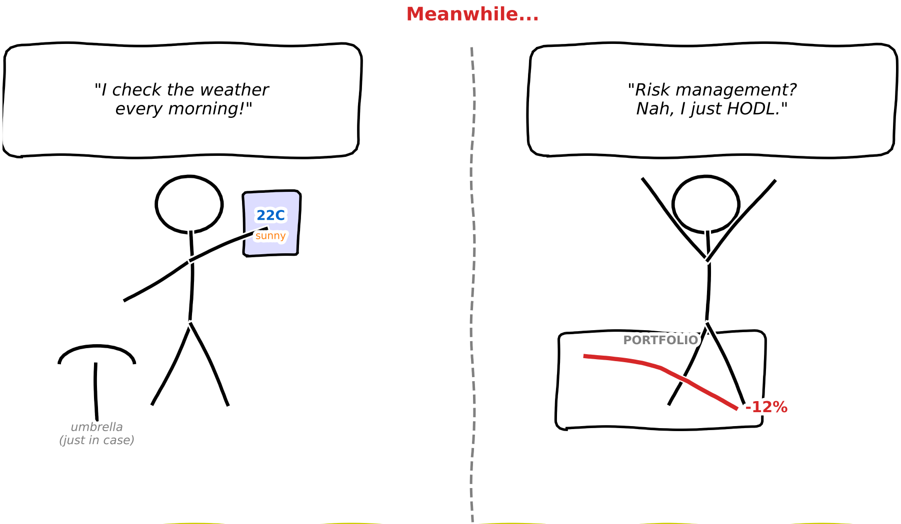

We prepare for a 30% chance of rain but not a 5% chance of losing everything.

Learning Objectives

After completing this lesson, you will be able to:

1. **Explain** what Value-at-Risk (VaR) measures using the "sorted returns" intuition [Understand]
2. **Compare** the three VaR methods: historical simulation, variance-covariance, Monte Carlo [Understand]
3. **Calculate** a simple VaR from a list of sorted returns [Apply]
4. **Distinguish** VaR from Expected Shortfall (ES / CVaR) [Understand]
5. **Explain** why financial returns have fat tails and what that means for risk [Analyze]
6. **Illustrate** how diversification reduces portfolio risk using correlation [Analyze]

Bloom's levels covered: Understand, Apply, Analyze

Objectives follow Bloom's taxonomy: Understand → Apply → Analyze.

Module 4: The Risk Problem

Where we have been (teaching order):

- M1 Cost: why payments are expensive
- M2 Access: who gets served by finance, who does not
- M6 Infrastructure: the rails that clear and settle money
- M3 Trust: cryptographic foundations, consensus, smart contracts, DeFi stablecoins

Module 4 asks a different question:

- **"What could go wrong?"**
- Financial systems can lose money, sometimes *a lot* of money
- Risk management is the discipline of **measuring** and **controlling** potential losses
- Every bank, fund, and fintech *must* answer: "How much can we lose tomorrow?"

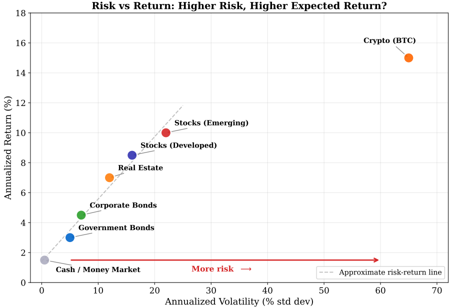

More return usually means more risk. Can we measure that risk?

Teaching sequence: M3 Trust → M4 Risk (here) → M5 Automation next. We built financial systems; now we ask what could go wrong.

Risk Is Like Weather Forecasting

Weather forecast:

- "There is a 30% chance of rain tomorrow"
- You decide: bring an umbrella or not
- The forecast does **not** tell you *exactly* what will happen
- It tells you the **range of possibilities** and their **probabilities**

Risk forecast:

- "There is a 5% chance we lose more than CHF 50,000 tomorrow"
- You decide: hedge, reduce position, or accept the risk
- Risk management does **not** predict the future
- It quantifies the **range of possible losses**

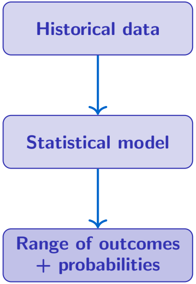

Risk management is not about predicting the future — it is about preparing for a range of possible outcomes.

Value-at-Risk (VaR): The Core Idea

Definition: Value-at-Risk

Value-at-Risk (VaR) answers: "What is the **maximum loss** I can expect on **X% of days**?" Equivalently: "On 95% of days, my loss will be **no worse** than this number."

The sorted-returns intuition:

1. Collect 1,000 days of past daily returns
2. Sort them from worst to best
3. The **50th worst return** (out of 1,000) is your **95% VaR**
4. Interpretation: on 95% of days, you did better than this; on 5%, you did worse

Analogy: If you ranked your 1,000 worst-to-best exam scores, your 95% VaR is the 50th lowest score. "On 95% of exams, I did better than this."

VaR is a single number that summarizes downside risk.

VaR = sort your returns, count from the worst. The 5th percentile is your 95% VaR.

VaR on a Histogram: The Shaded Tail

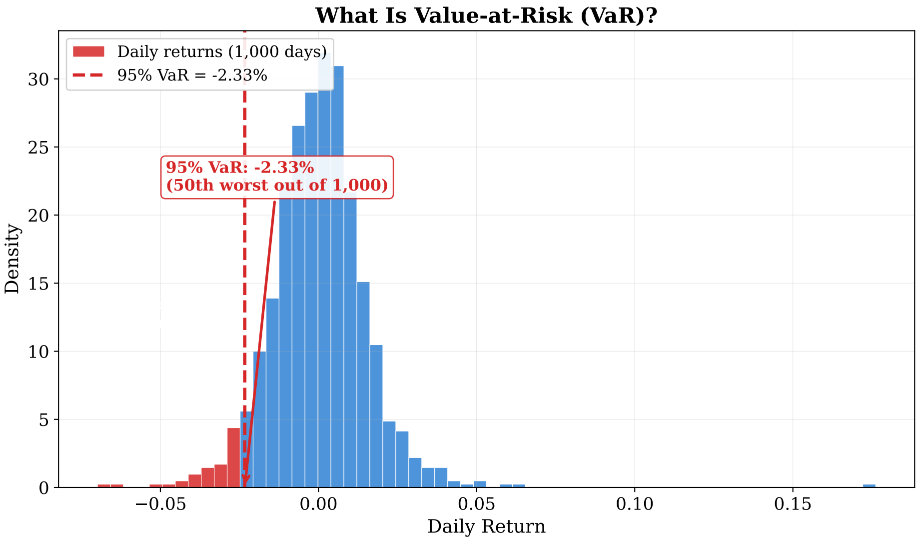

VaR Worked Example

Scenario: You invest CHF 100,000 in a stock fund. Over the past 1,000 trading days, you recorded daily returns.

Step 1: Sort all 1,000 returns from worst to best.

Rank	Meaning	Return
1st worst	Biggest loss ever	-4.2%
2nd worst	Second-biggest loss	-3.8%
⋮	⋮	⋮
50th worst	**95% VaR cutoff**	-2.1%
⋮	⋮	⋮
1,000th (best)	Best day	$+3.5\%$

Step 2: 95% daily VaR $= -2.1\% \times$ CHF 100,000 $=$ **CHF 2,100**.

Interpretation: "On 95% of trading days, I expect to lose **no more than CHF 2,100**. On 5% of days, I could lose more."

VaR does NOT say how much you lose on the worst 5% of days — only where the boundary is.

Three Ways to Estimate VaR

	Historical Simulation	Variance-Covariance	Monte Carlo
Idea	Sort *actual* past returns	Assume returns are normal, use formula	Simulate thousands of random scenarios
Analogy	"Look at what *actually* happened"	"Assume a bell curve"	"Roll the dice 10,000 times"
Pros	Simple, no assumptions	Fast, elegant	Flexible, handles complex portfolios
Cons	Assumes the past repeats	Returns are *not* normally distributed	Computationally expensive

Which is "best"? None is perfect. Each method has strengths. Real risk teams often use **all three** and compare.

Historical simulation is the most intuitive. Variance-covariance is the fastest. Monte Carlo is the most flexible.

VaR Method Comparison

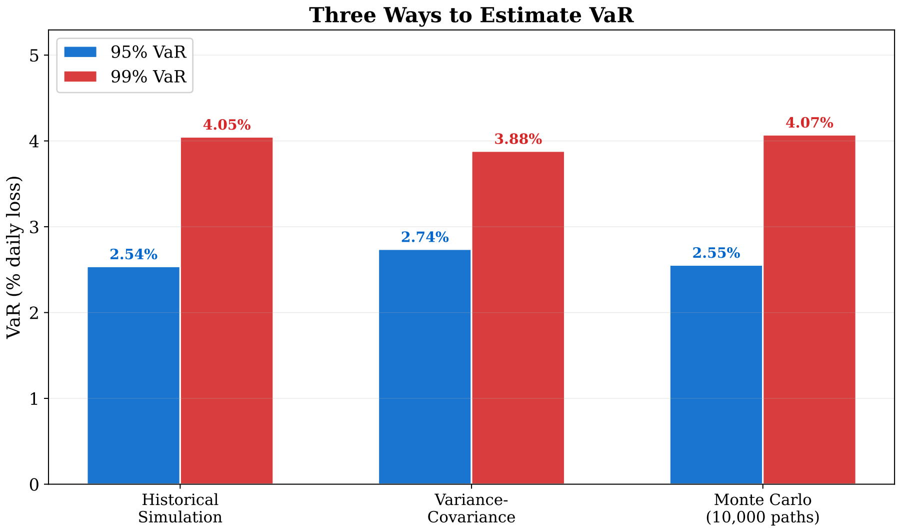

- **What you see:** Three methods produce slightly different VaR estimates for the same data

Monte Carlo Simulation: Start with a Coin Flip

Imagine a simple game:

- You flip a coin. Heads: you win CHF 110. Tails: you lose CHF 100.
- Positive expected value: $0.5 \times 110 + 0.5 \times (-100) = +5$ per flip.
- But what is the **risk** of playing 20 rounds?

Monte Carlo approach:

1. Simulate 10,000 players, each flipping 20 times
2. Record each player's total profit/loss
3. Sort all 10,000 outcomes from worst to best
4. The 500th worst outcome is the 95% VaR of this game

Result: Even though the game has positive expected value, about 5% of players will lose more than CHF 400 over 20 rounds. That is the 95% VaR.

Monte Carlo = repeat a random experiment many times and count the outcomes.

This exact logic scales from coin flips to portfolios with thousands of assets.

Monte Carlo: 500 Possible Futures for Your Portfolio

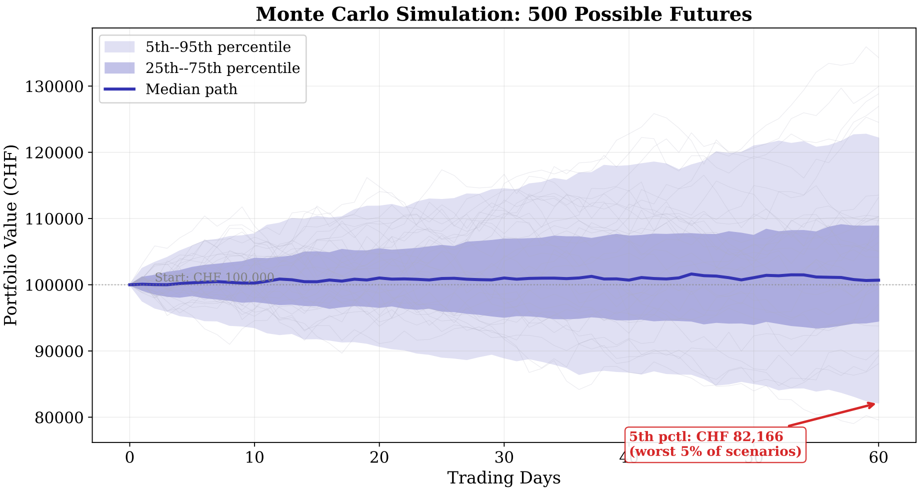

Each line is one possible future. The fan shows the range of outcomes. VaR is the bottom edge of the fan.

The Problem with VaR: It Stops at the Threshold

VaR tells you the boundary of the worst 5%. But what happens *beyond* that boundary?

Two portfolios, same 95% VaR of −2%:

	Portfolio A	**Portfolio B**
95% VaR	−2.0%	−2.0%
Worst 5% outcomes:		
Average loss	−2.3%	−8.5%
Maximum loss	−3.0%	−25.0%

VaR says these portfolios have the same risk. They clearly do not.

Portfolio B has a **catastrophic tail** — the losses beyond VaR are enormous. VaR misses this because it only looks at the *threshold*, not what happens beyond it.

VaR is blind to the severity of tail losses. This is its most important limitation.

Expected Shortfall (ES): Looking Beyond VaR

Definition: Expected Shortfall (ES / CVaR)

Expected Shortfall (also called CVaR or Conditional VaR) is the **average loss** in the worst X% of scenarios. If VaR asks "how bad is the boundary?", ES asks "**how bad is it on average** when things go wrong?"

The sorted-returns intuition:

1. Sort 1,000 daily returns from worst to best
2. The 50th worst is the 95% VaR (the boundary)
3. The **average of the 50 worst** is the 95% Expected Shortfall

Analogy: VaR is like knowing "the speed limit is 120 km/h." ES is like knowing "when people speed, they average 145 km/h." ES tells you **how bad** the bad days really are.

ES $\geq$ VaR always. ES captures tail severity. Regulators now prefer ES over VaR.

Basel III/IV moved from VaR to Expected Shortfall as the primary risk measure for banks.

VaR vs Expected Shortfall: Visual Comparison

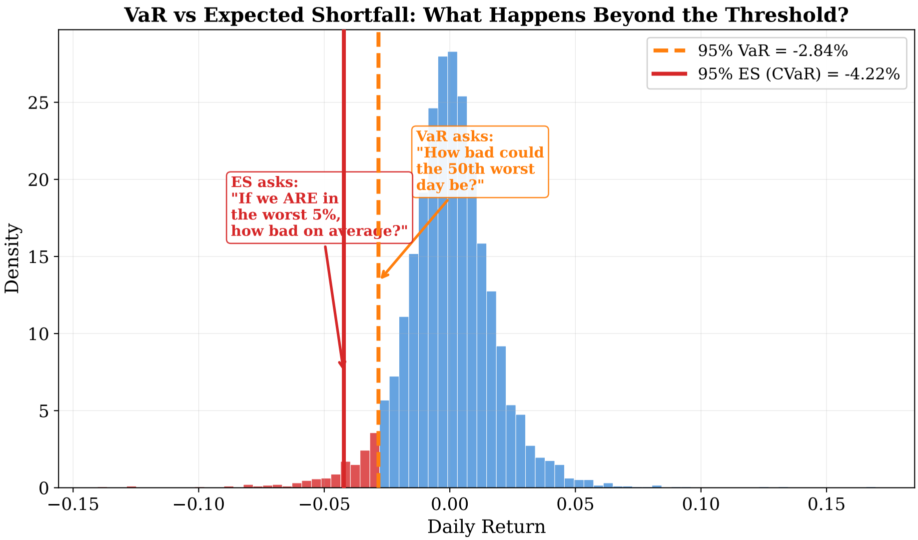

VaR is the threshold; ES is the average loss beyond the threshold. ES is always worse (larger loss).

Volatility: How Much Does the Price Wiggle?

Definition: Volatility

Volatility is the **standard deviation of returns** — a measure of how much an asset's price fluctuates. Higher volatility = larger daily swings = more risk.

Intuition:

- A government bond might move $\pm 0.2\%$ per day (low volatility)
- A stock might move $\pm 1.5\%$ per day (moderate volatility)
- Bitcoin might move $\pm 5\%$ per day (high volatility)

Key fact: Volatility is the **building block** of VaR.

- Under the variance-covariance method: $\text{VaR}_{95\%} \approx 1.65 \times \sigma_{\text{daily}}$
- Higher volatility $\rightarrow$ higher VaR $\rightarrow$ more risk

Example: For a \$10M equity portfolio with daily volatility of 1.5%, 99% VaR = \$10M × 2.326 × 0.015 = \$348,900. There is a 1% chance of losing more than \$348,900 in one day.

Annualized volatility $= \sigma_{\text{daily}} \times \sqrt{252}$ (252 trading days per year)

Example: If daily volatility = 1.5%, then annualized $= 1.5\% \times \sqrt{252} \approx 1.5\% \times 15.87 = 23.8\%$.

Volatility is the single most important input to any risk model.

Volatility Clustering: Big Moves Follow Big Moves

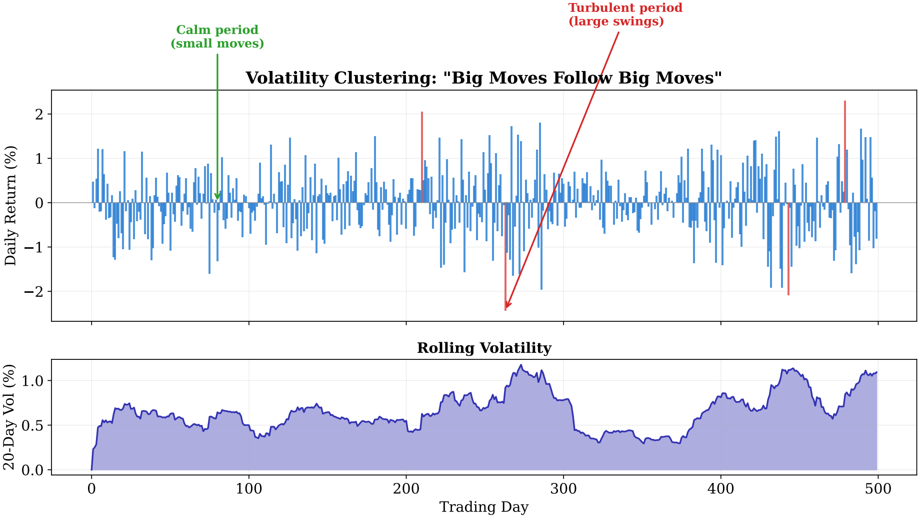

Key observation: Volatility is **not constant**. Calm periods and turbulent periods **cluster together**. Today's large move

Estimating Volatility: Rolling Window vs. EWMA

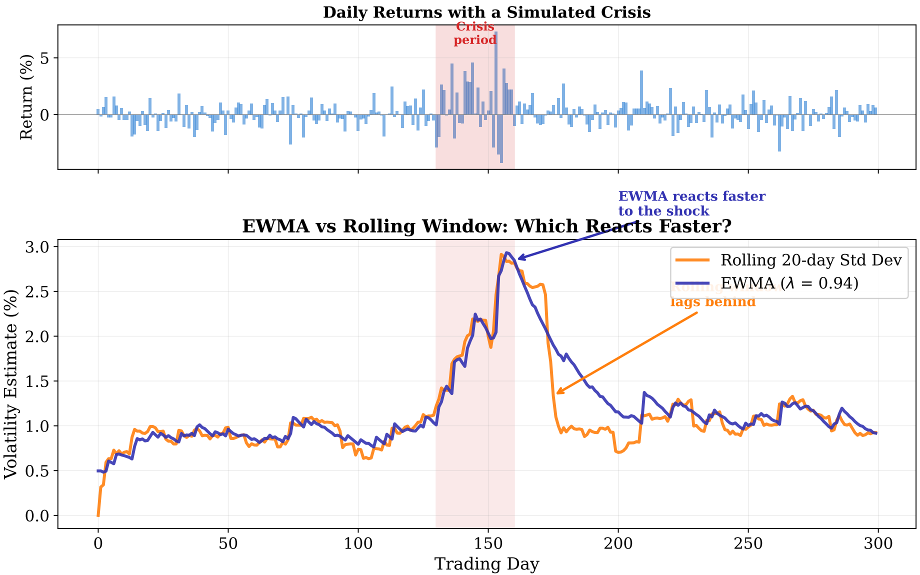

GARCH: A Name You Should Know (Conceptual Only)

The idea behind GARCH:

- "Big moves follow big moves" — we saw this in volatility clustering
- GARCH (Generalized Autoregressive Conditional Heteroskedasticity) is a **formal statistical model** that captures this pattern
- It says: tomorrow's volatility depends on **today's volatility** and **today's return**

Why it matters:

- GARCH models are used by banks and hedge funds to forecast volatility
- Tim Bollerslev (1986) introduced GARCH; Robert Engle won the 2003 Nobel Prize for the related ARCH model
- You do **not** need to know the equation — just the concept

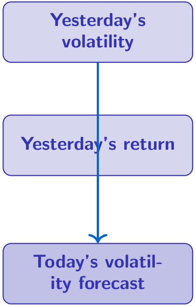

GARCH combines yesterday's volatility with yesterday's surprise to forecast today's risk.

GARCH = "big moves follow big moves" formalized. Know the concept; the equation is graduate-level.

Fat Tails: The Bell Curve Underestimates Extreme Events

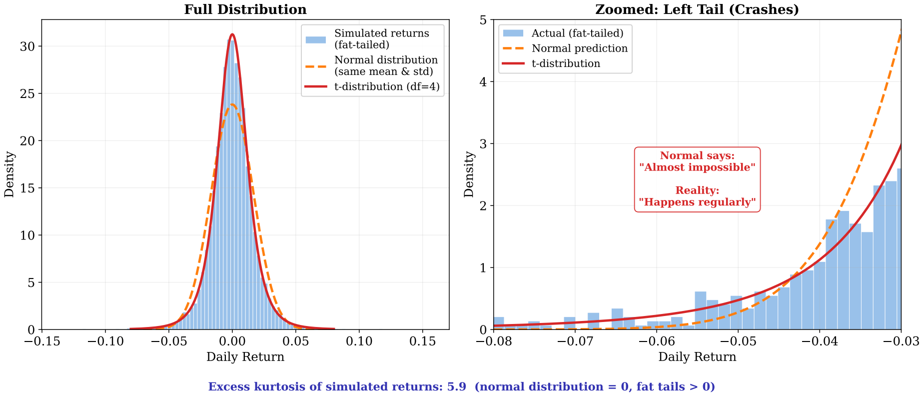

Real financial returns have "fat tails": extreme events (crashes, spikes) occur far more often than a normal distribution predicts.

Kurtosis: How “Fat” Are the Tails?

Definition: Kurtosis

Kurtosis measures how heavy the tails of a distribution are compared to a normal distribution. A normal distribution has excess kurtosis $= 0$. Financial returns typically have excess kurtosis of 3–10$+$, meaning extreme events are **far more common** than the bell curve predicts.

Why this matters for risk:

- The variance-covariance VaR method assumes returns are normal
- If returns have fat tails (high kurtosis), this method **underestimates** the true VaR
- Historical simulation and Monte Carlo can capture fat tails naturally

Real-world example:

- Under a normal distribution, a -5σ daily move should occur roughly once every 14,000 trading days ($\approx$55 years), or once every 3.5M calendar days if we do not condition on market days
- In reality, it happens roughly once every 10–20 years (e.g., 2008, 2020) — fatter tails than the model allows

---

“Fat tails” means extreme losses happen far more often than textbook models predict.

Correlation: Do Assets Move Together?

Definition: Correlation

Correlation measures how two assets move relative to each other. It ranges from -1 (perfect opposites) through 0 (unrelated) to $+1$ (perfect lockstep).

Examples:

Pair	**Typical Correlation**	**Interpretation**
EU stocks vs. US stocks	$+0.85$	Move together (high positive)
Stocks vs. gov. bonds	-0.15	Tend to move oppositely
Gold vs. stocks	$+0.05$	Nearly independent

Key insight: Low or negative correlation is the **engine of diversification**. If assets do not move together, combining them reduces total portfolio risk.

Correlation is the most important number for portfolio construction: it determines how much diversification you get.

Correlation Heatmap: Which Assets Diversify?

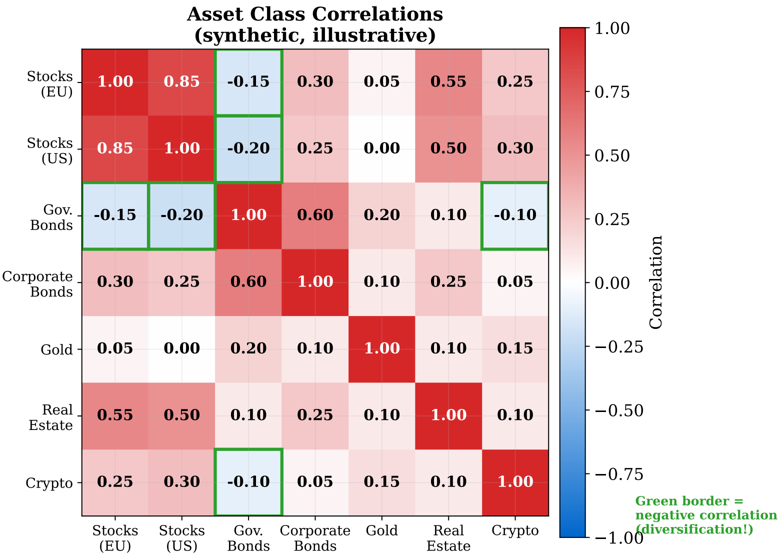

Diversification: "The Only Free Lunch in Finance"

Why diversification works:

- When one stock falls, another may rise (if correlation < 1)
- Gains and losses **partially cancel each other out**
- Portfolio volatility is **lower** than the average stock's volatility

The portfolio variance formula (two assets):

$$\sigma_p^2 = w_A^2\sigma_A^2 + w_B^2\sigma_B^2 + 2w_A w_B \rho \sigma_A \sigma_B$$

- If $\rho < 1$: portfolio risk $<$ weighted average of individual risks
- If $\rho = -1$: risk can be **completely eliminated** (in theory)
- If $\rho = +1$: no diversification benefit at all

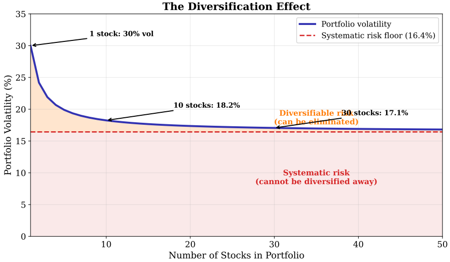

Adding stocks rapidly reduces portfolio volatility — but a floor remains (systematic risk).

Harry Markowitz called diversification "the only free lunch in finance" — you reduce risk without reducing expected return.

How Many Stocks Do You Need?

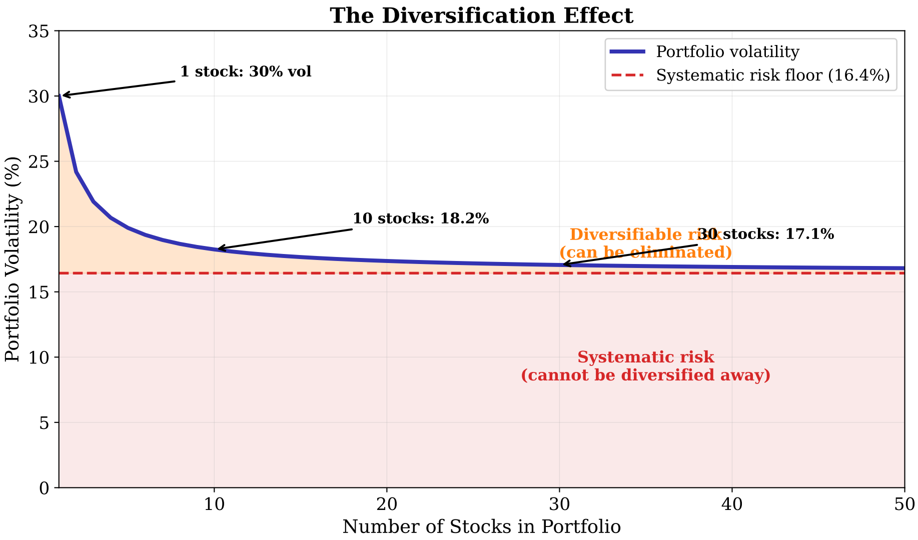

Warning: Correlations Change in a Crisis

The diversification trap:

- In calm markets, correlations between stocks are moderate (≈ 0.3–0.5)
- In a crisis (2008, 2020), correlations **spike toward 1.0**
- Everything falls together — diversification fails precisely when you need it most

Real examples:

Asset Pair	Normal Times	March 2020 Crash
Stocks vs. corporate bonds	$+0.30$	$+0.70$
US stocks vs. EM stocks	$+0.65$	$+0.90$
Stocks vs. gold	$+0.05$	-0.20 (gold held up)

Lesson: Correlations estimated in calm times **understate** the risk in a crisis. This is one reason VaR models failed in 2008.

"Diversification works until you need it" — correlations spike during market stress.

The Risk-Return Trade-off

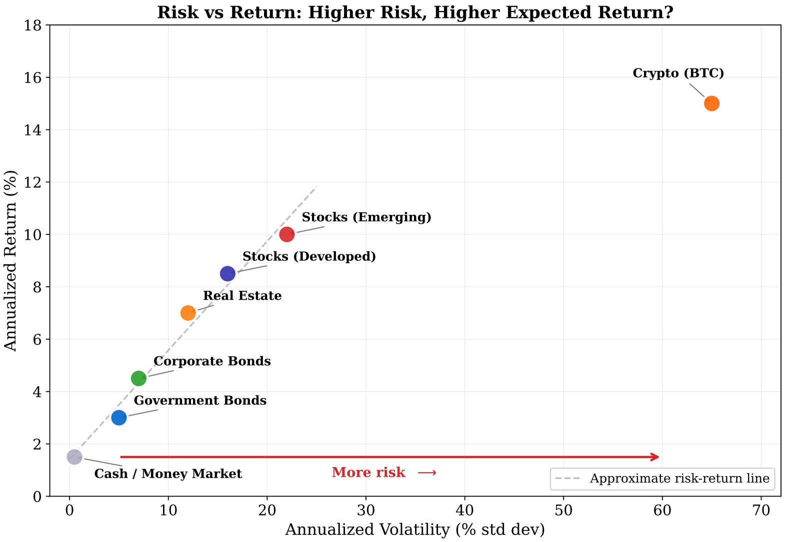

One Last Thought. . .

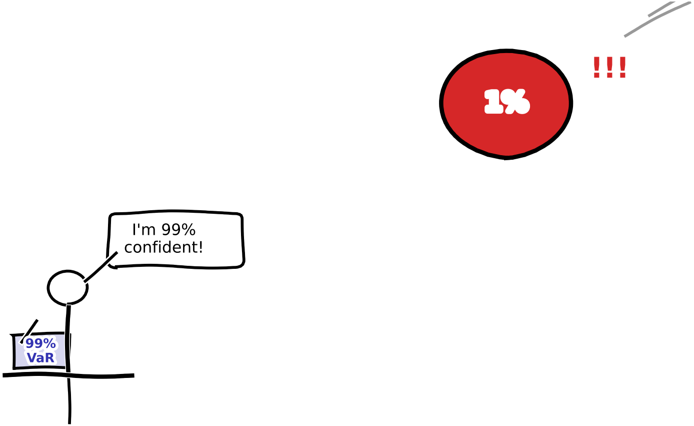

VaR tells you the boundary. It doesn't promise you won't cross it.

Sometimes the best way to remember a concept is to laugh about it.

Key Takeaways

1. **VaR** answers "how much can I lose?" — sort 1,000 returns; the 50th worst is the 95% VaR
2. **Three VaR methods:** historical simulation (sort actual data), variance-covariance (assume normal), Monte Carlo (simulate thousands of scenarios)
3. **Expected Shortfall (ES)** asks "how bad is it *on average* beyond VaR?" — it captures tail severity
4. **Volatility** is the standard deviation of returns and the key input to all risk models
5. **Volatility clusters:** big moves follow big moves (GARCH captures this)
6. **Fat tails:** extreme events occur far more often than the normal distribution predicts (kurtosis > 0)
7. **Diversification** reduces portfolio risk — but cannot eliminate systematic (market-wide) risk
8. **Correlations spike in crises** — diversification weakens when you need it most

VaR, ES, volatility, fat tails, correlation, diversification — these are the building blocks of risk management.

Summary / Next Lesson Preview

This lesson: We learned to measure market risk using VaR and ES, explored why the normal distribution underestimates extreme events, and saw how diversification reduces — but does not eliminate — portfolio risk.

Key vocabulary:

- Value-at-Risk (VaR)
- Expected Shortfall (ES / CVaR)
- Historical simulation
- Monte Carlo simulation
- Variance-covariance method
- Volatility (standard deviation)
- EWMA / GARCH (concepts)
- Fat tails / kurtosis
- Correlation
- Diversification
- Systematic vs. diversifiable risk
- Risk-return trade-off

Next lesson (M4L2): *Credit Risk and Counterparty Risk* — What happens when the other side of a trade cannot pay? We will explore credit scoring, probability of default, and how the 2008 financial crisis was fundamentally a credit risk failure.

Review: Can you explain VaR using the "sort 1,000 returns" intuition? Can you distinguish VaR from ES?

Common Misconceptions About Measuring Market Risk

Misconception	Reality
"VaR tells you the worst case"	VaR is a *quantile* (e.g., 99%), not a maximum. It says nothing about the 1% of days worse than the threshold — that is precisely where banks blow up. Use ES for tail.
"Normal-distribution VaR is fine for 95%"	Equity returns are leptokurtic (fat-tailed). Normal VaR underestimates 99%+ losses by 3–10x. See M4L1 slide on kurtosis.
"Stress tests prove the bank is safe"	Stress tests reveal vulnerability to *specified* scenarios. The 2008 crisis featured correlations $\to 1$ — a scenario none of the pre-2008 stress tests modelled.
"Diversification eliminates risk"	Diversification eliminates *idiosyncratic* risk. Systemic risk (2008, March 2020 COVID sell-off) moves everything together.
"Volatility is risk"	Volatility measures dispersion around the mean. Risk is downside. Sortino and drawdown capture this more honestly than sigma alone.

Taleb's aphorism: "Never cross a river 4 feet deep on average." VaR can be a 4-foot river.

Self-Check: Measuring Market Risk

Attempt these before turning the page.

1. [Understand] State the difference between parametric VaR, historical VaR, and Monte Carlo VaR in one sentence each.
2. [Apply] Portfolio value \$10M. Daily return $\mu = 0.05\%$, $\sigma = 1.5\%$, assumed normal. Compute 1-day 99% parametric VaR. Then compute the 10-day VaR using the $\sqrt{T}$ scaling rule.
3. [Analyze] Why does Expected Shortfall (ES) not suffer from VaR's tail-blindness? Name two regulatory regimes that replaced VaR with ES.

Solutions hidden unless `\solutionstrue` is set before compiling.