Market Risk Measures

Theory, Statistical Estimation and Applications

Noufel Frikha

Université Paris 1 Panthéon-Sorbonne

6 November 2026



Two loss distributions

A deliberately simple example: two positions over the same horizon, with losses expressed in the same monetary unit.

	Position A	Position B
Loss with probability 99%	0	0
Loss with probability 1%	100	$1\,000$

Discussion: Would you assign the same risk to these two positions?

- Which features of these distributions would you compare?
- What information might a single risk number leave out?

From portfolio value to loss

Let V_t denote the portfolio value at time t and $h > 0$ the risk horizon.

$$\mathrm{P\&L}_{t,t+h} = V_{t+h} - V_t, \qquad L_{t,t+h} = -\mathrm{P\&L}_{t,t+h}.$$

- V_t is known at time t; V_{t+h} is uncertain. Risk is assessed using the information available at time t.
- A positive loss is an adverse outcome; a negative loss is a gain.
- We neglect discounting and assume no intermediate cash flows or external cash injections in this introductory framework.

Convention: Unless stated otherwise, risk measures will be applied to losses.

What should a risk measure capture?

Different summaries retain different information about a loss L.

- **Dispersion:** how variable are the outcomes?
  Variance and standard deviation.
- **Loss threshold:** which level is exceeded with a small probability?
  Value-at-Risk (VaR).
- **Tail severity:** how large are losses in the upper tail?
  Expected Shortfall (ES).

In the opening example, loss probabilities coincide, but the adverse outcome is ten times larger for position B.

Learning objectives and prerequisites

By the end of this course, you should be able to:

- Define and compare variance, VaR and ES.
- Analyse their properties and construct counterexamples.
- Compute risk measures in simple models and estimate them from observations or simulations.
- Explain their use in capital allocation and risk budgeting.

Prerequisites: Basic tools of Probability

Probability distributions, expectation, variance and covariance; laws of large numbers and the central limit theorem; linear algebra and multivariable differentiation. Key notions will be recalled when needed.

Course organisation

Lectures: 18 hours, 9:00 am–12:00 noon

- Fridays 6, 13 and 20 November 2026.
- Mondays 23 and 30 November, and 7 December 2026.

Tutorials

- 30 November and 7 December, 12:30–3:30 pm.

Assessment and materials

- Final examination: 14 December 2026 (time to be confirmed).
- Course questions and one or two exercises/problems.
- Slides, exercises and past examination papers on the EPI.
- No RETAKE

Why measure financial risk?

The choice of a measure depends on the decision it supports.

- **Compare positions** and their loss profiles.
- **Set limits** on acceptable risk.
- **Assess capital** to absorb adverse outcomes.
- **Choose portfolios** balancing return and risk.



Chicago Board of Trade, 1993.

Photograph: Jeremy Kemp, public domain.

Source: Wikimedia Commons.

Risk measurement supports decisions; it does not eliminate uncertainty.

Market risk and related risks

- **Market risk:** losses from changes in prices, interest rates, exchange rates, credit spreads and volatility.
- **Credit and counterparty risk:** losses associated with default or deterioration in credit quality.
- **Liquidity risk:** difficulty trading without substantial price impact, or meeting payment obligations.
- **Operational risk:** losses from failures of processes, people or systems, or from external events.
- **Model risk:** adverse consequences of incorrect models or inappropriate model use.

These risks interact. Climate-related shocks can affect several categories; systemic risk concerns propagation across the financial system.

From a loss distribution to a risk estimate

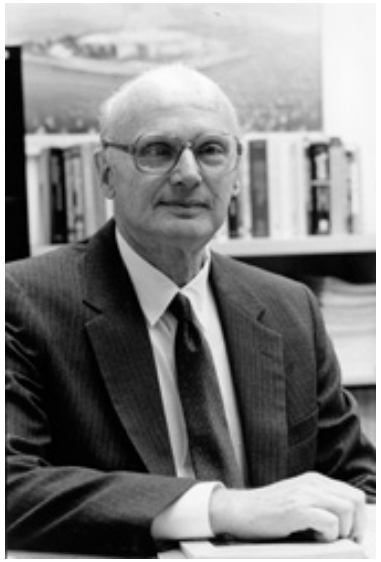
$$\begin{aligned} \text{Loss model} &\longrightarrow \rho(L) \quad \text{(target)}, \\ \text{Loss data } L_1, \ldots, L_n &\longrightarrow \widehat{\rho}_n \quad \text{(estimate)}. \end{aligned}$$

The data may be observed or simulated. The target depends on the chosen measure and loss model.

- **Interpretation:** what feature of the loss distribution does ρ capture?
- **Properties:** how should it respond to larger losses, diversification or additional capital?
- **Estimation:** how close is $\widehat{\rho}_n$ to $\rho(L)$, and at what computational cost?

Risk measurement before widespread VaR adoption

- **Exposures and limits:** position size and concentration.
- **Sensitivities:** response to market moves. Bond duration: Macaulay (1938).
- **Variance and covariance:** dispersion and diversification. Markowitz (1952).
- **Downside criteria:** frequency and severity of shortfalls. Roy (1952); Markowitz (1959).



Harry Markowitz

Photograph: Henry Foster, undated.

These approaches address different questions and continue to coexist.

Sources: Macaulay (1938); Markowitz (1952, 1959); Roy (1952).
Portrait: University of Chicago Photographic Archive, apf1-13352,
Hanna Holborn Gray Special Collections Research Center, University of Chicago Library.

Archive record. Educational use; commercial reuse requires permission.

Downside risk: losses below a target

Here risk is expressed in returns: for $V_0 > 0$ and no intermediate cash flows, $R = (V_T - V_0)/V_0$ and $L = -V_0 R$. Let $r_\star$ be a target return.

- **Safety first:** control the probability of falling below a critical return (Roy, 1952):

$$\mathbb{P}(R < r_\star).$$

- **Downside dispersion:** penalise only adverse deviations:

$$\mathbb{E}\big[(r_\star - R)_+^2\big], \qquad x_+ = \max(x, 0).$$

  This is a lower partial moment of order two. At $r_\star = \mathbb{E}[R]$, it is a common definition of semi-variance (Markowitz, 1959).

Variance treats deviations above and below the mean symmetrically.

Sources: Roy (1952), Econometrica 20, 431–449; Markowitz (1959), Portfolio Selection; Markowitz (1990), Nobel lecture.

Maximum drawdown: risk along a portfolio path

For strictly positive portfolio values $V_0, \dots, V_N$, define

$$M_k = \max_{0 \le j \le k} V_j, \qquad D_k = 1 - \frac{V_k}{M_k}, \qquad \mathrm{MDD}_N = \max_{0 \le k \le N} D_k.$$

Path	Final return	Maximum drawdown
$100 \to 120 \to 90 \to 110$	$+10\%$	25%
$100 \to 90 \to 120 \to 110$	$+10\%$	10%

- These are realised paths. Under a stochastic model, MDD is itself random and depends on the entire path.
- Historical MDD depends on the observation period and frequency; it does not by itself specify a future loss probability.

Further reading: Magdon-Ismail et al. (2004), On the Maximum Drawdown of a Brownian Motion, J. Appl. Prob. 41, 147–161.

RiskMetrics: a milestone in the diffusion of VaR

- **October 1994:** J.P. Morgan released its internal risk methodology and free volatility and correlation data.
- **Adoption:** public documentation and data helped RiskMetrics become a common reference for market risk measurement.
- **2001 retrospective:** Christopher Finger stressed that a risk measure, its model and its numerical implementation must be distinguished.

The 1994 release helped popularise VaR; it was not the invention of VaR.

Source: C. Finger (2001), Return to RiskMetrics: the Evolution of a Standard, MSCI.

VaR, a model and a computational method

- **Risk measure:** for a loss L and confidence level $\alpha \in (0,1)$,

$$\mathrm{VaR}_\alpha(L) = \inf\{\ell : \mathbb{P}(L \leq \ell) \geq \alpha\}.$$

  It is a loss threshold, not a maximum possible loss.

- **Model:** specify risk factors and their distribution. Classic RiskMetrics uses Gaussian risk-factor returns and exponentially weighted volatility estimates.

- **Computation:** use an analytical approximation or simulation to obtain the chosen measure under the model.

VaR is not tied to a single model. Stress scenarios and other measures provide complementary information.

Source: C. Finger (2001), Return to RiskMetrics: the Evolution of a Standard, MSCI.

What does VaR leave out?

Recall the two loss distributions from the opening example:

	L_A	L_B
Loss with probability 99%	0	0
Loss with probability 1%	100	$1\,000$

With the lower-quantile definition given above,

$$\mathrm{VaR}_{0.95}(L_A) = \mathrm{VaR}_{0.95}(L_B) = 0.$$

- Both positions have the same 95% loss threshold.
- Their losses beyond that threshold differ substantially.

Equal VaR does not imply equal risk: a quantile does not describe the severity of losses beyond its threshold.

Diversification and the choice of a risk measure

- **A portfolio is more than a list of positions.** Its risk depends on how losses occur together, as well as on the loss distribution of each position.
- **Diversification raises a consistency question.** How should a risk measure assess the combination of several positions? Combining positions can produce a VaR larger than the sum of their separate VaRs.
- **An axiomatic approach emerged in the late 1990s.** Artzner, Delbaen, Eber and Heath proposed explicit requirements for risk measures, leading to the notion of *coherent risk measures*.

We will examine these requirements and their financial interpretation in Chapter 2.

Reference: Artzner, Delbaen, Eber and Heath (1999), Coherent Measures of Risk, Mathematical Finance 9, 203–228.

Expected Shortfall: looking beyond the quantile

For $\alpha \in (0,1)$, Expected Shortfall averages losses in the worst $1-\alpha$ fraction of outcomes. In the opening example, the worst 5% of outcomes comprise:

- the 1% of outcomes with positive losses;
- a further 4% of outcomes with zero loss.

Hence

$$\mathrm{ES}_{0.95}(L_A) = \frac{0.01 \times 100 + 0.04 \times 0}{0.05} = 20, \qquad \mathrm{ES}_{0.95}(L_B) = 200.$$

- ES distinguishes tail severity, but does not describe all risk.
- For integrable losses, ES satisfies the coherence requirements studied in Chapter 2, where we will give its formal definition.

Reading guide

- **Main reference: modelling and statistical methods**
  A. J. McNeil, R. Frey and P. Embrechts, *Quantitative Risk Management: Concepts, Techniques and Tools*, 2nd ed., Princeton University Press, 2015.

- **Mathematical foundations and risk measures**
  H. Föllmer and A. Schied, *Stochastic Finance: An Introduction in Discrete Time*, 5th ed., De Gruyter, 2025.

- **Applications and financial risk management**
  T. Roncalli, *Handbook of Financial Risk Management*, Chapman & Hall/CRC, 2020.

- **Complementary reading in French**
  T. Roncalli, *La gestion des risques financiers*, Economica, 2009.

Questions guiding the course

1. When does variance provide a useful description of risk?
2. What information do VaR and ES retain or leave out?
3. How can we estimate risk measures with controlled error?
4. How can we use them to allocate capital and construct portfolios?

We begin with variance, covariance and Markowitz portfolio theory.