


Market Risk Management

From **VaR** to **drawdown**: a practitioner's guide to short-, medium- and long-term risk for investment funds, mapped to **UCITS**, **AIFMD** and **MiFID II**.



THREE HORIZONS · ONE FRAMEWORK

01


SHORT TERM Monitoring
1 day – 1 month



VaR · ES · EWMA · Backtesting

02


MEDIUM TERM Testing
weeks – 1 year



Stress · Reverse stress · Liquidity

03

LONG TERM Testing
1 year +



Drawdown · Scenarios · PRIIPs

18

chapters in 3 parts

16

worked examples

11

figures and charts

Python

reproducible toolkit

UCITS

AIFMD · Annex IV

MiFID II

Basel · FRTB

EMIR · CSA

PRIIPs

Yini Gao

English edition · Version 2.1 · September 2026

Regulatory content as of 30 September 2026. Every numerical example independently recomputed. Study notes only – not investment, legal or compliance advice.

How to Use This Guide & Notation

The handbook is organised by risk horizon. **Part I, short-term monitoring** (1 day to 1 month): probability language, VaR and ES, volatility models, historical simulation, non-Normal corrections and backtesting, with reproducible S&P 500 case studies. **Part II, medium-term testing** (weeks to a year): stress testing, including sensitivity, historical, hypothetical, reverse and liquidity stress tests. **Part III, long-term testing** (beyond a year): drawdown analysis and multi-period scenario analysis. Each part sets out the corresponding AIFMD, UCITS and MiFID II requirements.

Five coloured boxes recur throughout:

- **Deep dive:** theory beginners can compute but struggle to explain;
- **In practice:** details the source material leaves open but a real bank or fund must decide;
- **Pitfall:** frequent computational or conceptual errors, including typos in the sources;
- **Worked example:** fully numerical examples;
- **Regulatory mapping:** related AIFMD, UCITS and MiFID II requirements.

Cut-off date • Regulatory content as of 30 September 2026 / Day • 2026 9 30 Day

Regulatory content reflects rules and public information as of 30 September 2026. The following are still moving; check the latest texts before use: (1) AIFMD II Level 2 measures and the new supervisory reporting template (expected after October 2027); (2) the Central Bank of Ireland's revised UCITS rules and AIF Rulebook; (3) EU FRTB application from 1 January 2027 and transitional relief to end-2029; (4) ESMA guidelines and technical standards on liquidity management tools and the integrated reporting framework.

Contents

Part I Short-term risk monitoring

- 1 Foundations: returns, volatility, time scaling
- 2 VaR as a percentile
- 3 Expected Shortfall and coherence
- 4 Volatility: rolling window, EWMA, GARCH
- 5 Historical and filtered historical simulation
- 6 VaR methods: non-Normal corrections, simulation, decomposition, curves
- 7 Backtesting: procedure, tests, traffic lights
- 8 Regulation of short-term risk (incl. commitment approach and counterparty risk)
- 9 Reproducible case studies

Part II Medium-term risk testing

- 10 Stress testing: theory
- 11 Stress testing: practice and worked examples
- 12 Regulation of medium-term risk testing
- Part III Long-term risk testing**
- 13 Drawdown: theory and practice
- 14 Scenario analysis: theory and practice
- 15 Regulation of long-term risk testing
- 16 Integrated framework, checklist and errata
- 17 AIFMD Annex IV reporting, step by step
- 18 Risk management for derivative-heavy funds (incl. CSA)
- Appendices: Python, Excel, glossary

The three-horizon framework

	Short-term monitoring	Medium-term testing	Long-term testing
Horizon	1 day – 1 month	Weeks – 1 year	1 – 10 years
Question	How much can we lose in normal markets?	Can we survive a severe but plausible shock?	How deep and long can losses run?
Tools	VaR, ES, volatility, backtesting	Sensitivity, historical, hypothetical, reverse, liquidity stress tests	Drawdown, multi-period scenarios, Monte Carlo
Basis	Distributions, quantiles	Conditional scenarios, sensitivities	Path dependence, narrative, simulation
Frequency	Daily; monthly backtest	Monthly to quarterly	Quarterly to annually
Key rules	UCITS global exposure; AIFMD limits and leverage; PRIIPs MRM; MiFID 10% loss report	UCITS/AIFMD stress tests; ESMA LST guidelines; IFD internal capital	MiFID product scenario analysis; PRIIPs performance scenarios; HWM; sustainability risk

Notation

Symbol	Meaning
X	Portfolio P&L or return, a random variable; losses negative
$L = -X$	Loss expressed as a positive number
$c, \alpha = 1 - c$	Confidence level (e.g. 99%) and tail probability (e.g. 1%)
$F_X(x) = \Pr(X \leq x)$	Cumulative distribution function
Φ, φ, Φ^{-1}	Standard Normal CDF, PDF and inverse CDF (quantile function)
μ, σ	Mean and standard deviation of returns (daily unless stated)
u_t, r_t	Day- t log-return $\ln(S_t/S_{t-1})$
h (or T)	Holding period / risk horizon in days, e.g. 1D, 10D
n	Estimation window length (observations), unrelated to h
S, K	Cornish–Fisher skewness and excess-kurtosis parameters
γ_1, γ_2	Sample skewness and sample excess kurtosis
T, x (backtest)	Number of eligible comparisons and number of breaches

Pitfall • Sign conventions / •

Following the course, formulas write VaR as a **negative return percentile** (e.g. -2.33σ); a breach occurs when the realised return is below VaR. Reports and regulators usually quote VaR and ES as **positive loss amounts**. The two differ only by a sign, yet the comparison operator in Excel (< vs >) is a classic source of errors: use one convention per sheet and state it in the header.

Short-Term Risk Monitoring · 1 day to 1 month

Short-term monitoring answers "how much can we lose in normal markets?". It rests on distributions and quantiles, runs daily with monthly backtests, and underpins UCITS global exposure and bank market-risk capital.

1 Foundations: Returns, Volatility and Time Scaling

1.1 Why log-returns

Log-returns $u_t = \ln(S_t/S_{t-1})$ are additive over time: the 10-day return is the sum of ten daily returns, $r_{10D,t} = \ln(S_{t+10}/S_t) = \sum_{i=1}^{10} u_{t+i}$. That is exactly what makes variances add and volatility scale with $\sqrt{h}$. For daily data log and simple returns are close, but the gap becomes material at 10-day horizons and in large sell-offs.

1.2 Variance is the average squared return

$$\hat{\sigma}_{1D}^2 = \frac{1}{n-1} \sum_{i=1}^n (u_i - \bar{u})^2 \xrightarrow{\bar{u} \approx 0} \hat{\sigma}_{1D}^2 \approx \frac{1}{n} \sum_{i=1}^n u_i^2$$

The daily mean ($\approx 0.03\%$) is more than an order of magnitude smaller than the daily standard deviation ($\approx 1\%$), and its estimation error is huge (s.e. = $\sigma/\sqrt{n}$). Risk managers therefore set $\mu \rightarrow 0$, and variance reduces to the plain average of squared returns. This simplification is the key to EWMA, which merely replaces the simple average with an exponentially weighted one.

Whether the window holds 21, 42 or 66 observations, the resulting σ is always a **daily** quantity. The window length n controls responsiveness and noise; the horizon h controls scaling. They are different concepts.

1.3 Square-root-of-time rule and annualisation

$$\sigma_{hD} = \sqrt{\sigma_{1D}^2 + \dots + \sigma_{1D}^2} = \sigma_{1D} \sqrt{h}, \quad \mu_{hD} = h \mu_{1D}, \quad \sigma_{1Y} = \sigma_{1D} \sqrt{252}$$

• $\sqrt{h}$ In Result / Deep dive • Is $\sqrt{h}$ a consequence of the log?

Not entirely. Two layers do different jobs:

- **The log makes returns additive:** the multi-day log-return is the sum of daily ones, $r_{hD} = u_1 + \dots + u_h$. It only turns the problem into a sum.
- **$\sqrt{h}$ is statistical:** how the variance of that sum grows depends on how the terms relate to each other.

Whether the multiplier is $\sqrt{h}$ or something else is decided by the second layer. With simple returns the result holds approximately for small daily moves; the log merely makes the summation exact.

$$\text{Var} \left(\sum_{i=1}^h u_i \right) = \sum_{i=1}^h \text{Var}(u_i) + 2 \sum_{i < j} \text{Cov}(u_i, u_j) \xrightarrow{\text{Cov}=0, \text{Var}(u_i)=\sigma^2} h\sigma^2 \Rightarrow \sigma_{hD} = \sigma \sqrt{h}$$

So $\sqrt{h}$ needs two conditions: (1) daily returns are **uncorrelated** (zero covariance terms; this is the key, and Normality is **not** required); (2) daily variance is **constant** (otherwise $\sqrt{\sigma_1^2 + \dots + \sigma_h^2}$). In one line: **the log makes returns additive; "uncorrelated, equal variance" makes variances additive; $\sqrt{h}$ follows from the latter.**

In practice • When $\sqrt{h}$ breaks down / • $\sqrt{h}$

1. **Autocorrelation:** momentum (positive correlation) makes true 10-day volatility larger than $\sqrt{10}$ times daily; mean reversion makes it smaller. Test with the variance ratio $VR(h) = \text{Var}(r_{hD})/(h \text{Var}(r_{1D}))$ against 1 (Lo–MacKinlay).
2. **Volatility clustering:** high σ today means high σ over the next days. Today's $\sigma \times \sqrt{10}$ understates risk while volatility rises and overstates it while it falls. EWMA implies a flat term structure; GARCH reverts to the long-run mean.
3. **Fat tails:** $\sqrt{h}$ scales σ , not quantiles directly. Summed fat-tailed daily returns move towards Normal (central limit theorem), so the 10-day tail is relatively thinner and 1-day VaR $\times \sqrt{10}$ **overstates** extreme quantiles.
4. **Mean:** the mean grows with h , σ with $\sqrt{h}$. Negligible at short horizons, the mean matters more as h grows: one reason VaR is a short-horizon measure.
5. **Constant positions:** $\sqrt{10}$ assumes no trading or hedging for ten days. In reality books are cut or hedged, which is why FRTB uses liquidity horizons by risk factor (10–120 days).
6. **Regulatory convenience:** Basel allows $\sqrt{10}$ to convert 1-day VaR to 10-day mainly because estimating 10-day returns directly is data-hungry: non-overlapping samples are one tenth the size, overlapping ones introduce serial correlation (§7.2). $\sqrt{10}$ is an accepted approximation, not an exact result.

What to check in practice is whether "uncorrelated, equal variance" holds for your data and horizon.

1.4 Portfolio volatility

$$\sigma_{\Pi}^2 = \mathbf{w}'\Sigma\mathbf{w}, \quad \text{two assets: } \sigma_{\Pi}^2 = \sigma_X^2 + 2\rho_{XY}\sigma_X\sigma_Y + \sigma_Y^2 \quad (\text{in money terms})$$

Worked examples 1–2 • Analytical VaR / 1–2 • VaR

Ex. 1: a £5,000,000 position with 1% daily volatility. Annualised volatility $0.01\sqrt{252} \approx 15.9\%$; 10-day s.d. $0.01\sqrt{10} = 3.16\%$. 99%/10-day VaR = $2.326 \times 0.0316 \times 5,000,000 \approx \text{£}367,800$ (£368,400 with the rounded factor 2.33).

Ex. 2: two assets, £100,000 each, both 1% daily volatility (£1,000), correlation 0.3. $\sigma_{\Pi}^2 = 2.6 \times 10^6$, $\sigma_{\Pi} = \text{£}1,612.45$. 99%/5-day VaR = $2.33 \times 1612.45 \times \sqrt{5} \approx \text{£}8,401$. With correlation 1 the VaR would be £10,420; the difference is the diversification benefit.

In practice • Do not annualise VaR / • VaR

VaR and ES are **short-horizon** measures (1-day, 10-day). Over six months or a year the portfolio is rebalanced and distributional assumptions drift, so 6M or 1Y VaR carries little meaning. Use stress tests for the medium term (Part II) and drawdown and scenario analysis for the long term (Part III); they have no fixed distributional horizon but reflect what investors actually experience.

2 Value at Risk as a Percentile

2.1 Definition and probability statement

VaR answers: "Over a given horizon, what is the loss we are c -confident will not be exceeded?" In terms of P&L X , the 99% VaR is the 1st percentile of the distribution of X :

$$\Pr(X \leq \text{VaR}_c(X)) = 1 - c$$

$$\text{VaR}_c(X) = \sup\{x \in \mathbb{R} : F_X(x) \leq 1 - c\} \quad (\text{P\&L, negative number})$$

$$\text{VaR}_c(L) = \inf\{x \in \mathbb{R} : F_L(x) \geq c\} \quad (\text{loss } L = -X, \text{ positive number})$$

Deep dive • Why "VaR is a percentile" / • "VaR "

The definition $\Pr(X \leq \text{VaR}) = 1 - c$ is literally the definition of a percentile: rank all possible P&L outcomes from worst to best, and VaR is the point that cuts off the worst $1 - c$. Three angles:

1. **Ranking:** sort 500 daily returns; the 99% VaR is about the 5th worst ($500 \times 1\%$). No distributional assumption.
2. **Inverse CDF:** the CDF maps a value to a cumulative probability; VaR goes the other way, $\text{VaR} = F^{-1}(1 - c)$.
3. **Normal formula:** -2.326 is the standard Normal 1st percentile; $\text{VaR} = \mu - 2.326\sigma$ just shifts and scales it to your portfolio.

Being a percentile is also VaR's limitation: it marks where the tail **starts**, not how **deep** it goes.

Deep dive • CDF as a ranking table / • " " CDF CDF

The CDF $F(x) = \Pr(X \leq x)$ gives the share of outcomes at or below x . Sort ten daily returns from smallest to largest (table below):

- **Forward (CDF):** value in, cumulative share out. $F(-2\%) = 20\%$.
- **Inverse (inverse CDF):** share in, value out. "Where do the worst 10% end?" $\rightarrow F^{-1}(10\%) = -3\%$, the 90% VaR.

The standard Normal Φ is the continuous version of the same table: $\Phi(-2.326) = 1\%$ is the forward lookup (Excel `NORM.S.DIST(-2.326, TRUE)`) and $\Phi^{-1}(1\%) = -2.326$ the inverse (`NORM.S.INV(0.01)`). **The CDF answers "how common is this loss?"; the inverse CDF answers "how large is a loss this common?"; VaR is the latter.**

Return x	-3%	-2%	-1.5%	-1%	-0.5%	0%	...	2%
F(x)	10%	20%	30%	40%	50%	60%	...	100%




Figure 1. Left: the PDF is a height; probability is area under the curve. Right: the CDF is the accumulated area; VaR is read off it in reverse.

Deep dive • Why probability comes from the CDF, not the PDF / • Probability CDF PDF

For a continuous distribution, probability is **area** under the PDF; the CDF is the area accumulated up to x . The PDF itself is a **height**, and a height is not a probability.

1. **A single point has probability 0:** a return of exactly $-2.32635\ldots\%$ has probability zero; only intervals carry probability: $\Pr(a \leq X \leq b) = \int_a^b f(x) dx = F(b) - F(a)$.
2. **The PDF is a density:** probability per unit of x , like population per square kilometre; multiply by a width to get a probability. A density can exceed 1: with $\sigma = 1\%$, the Normal peak on a return axis is about 39.9. Anything that can exceed 1 cannot be a probability.
3. **Applied to VaR:** $\Pr(X \leq \text{VaR}) = 1\%$ asks for the area of the tail left of VaR, i.e. $F(\text{VaR}) = 1\%$, so the CDF (and its inverse) is used.
4. **The ϕ in the ES formula:** $\phi(-2.326) = 0.02665$ is **not** a 2.665% probability. It comes from the integral: tail mean $= \int x f(x) dx \div \text{tail area } (1 - c)$; since the antiderivative of $z\phi(z)$ is $-\phi(z)$, ϕ appears as a height in the result. The denominator $1 - c$ is the probability, and it comes from the CDF.
5. **Exception: discrete distributions:** for discrete variables such as breach counts, each value has positive probability and the PMF is a probability. For $T = 250$, exactly 2 breaches has probability 25.7% (`BINOM.DIST(2, 250, 0.01, FALSE)`); traffic lights use the cumulative version (`TRUE`).

In one line: the PDF shows where probability is dense; the CDF tells you how much has accumulated. Asking for a probability means asking for an area.

inf / Deep dive • Why supremum and infimum / • sup

If the distribution is continuous and strictly increasing, the x solving $F(x) = 1 - c$ is unique and $F^{-1}(1 - c)$ suffices. Real P&L is often **discrete or mixed**: historical simulation has finitely many points, bond default is a 0/1 event, option books jump around strikes. The CDF is then a staircase: there may be **no** x with $F(x) = 1 - c$ exactly, or a **whole interval** of them.

The infimum says: "among all loss levels whose cumulative probability reaches at least c , take the smallest". It guarantees a unique answer for any distribution. The supremum on P&L is the mirror image of the same idea.

Lower-case x is a specific value (dummy variable); upper-case X is the random variable. " $x < \text{VaR}(X) < 0$ " reads: a specific loss worse than VaR.

2.2 The fundamental formula: standardise, then invert

Assume $X \sim N(\mu, \sigma^2)$. Standardising both sides of the event leaves the probability unchanged; applying Φ^{-1} then removes the probability, because $\Phi(z) = \Pr(Z \leq z)$ is the definition of the CDF:

$$\Pr\left(Z \leq \frac{\text{VaR} - \mu}{\sigma}\right) = 1 - c \implies \Phi\left(\frac{\text{VaR} - \mu}{\sigma}\right) = 1 - c \implies \boxed{\text{VaR}_c = \mu + \Phi^{-1}(1 - c)\sigma}$$

$\Phi^{-1}(1 - c)$ is the **Factor**, the standardised percentile. At 99%, $\Phi^{-1}(0.01) = -2.32635$ (Excel: `NORM.S.INV(0.01)`). Every parametric refinement of VaR swaps either the factor (Student-t, Cornish-Fisher, EVT) or the volatility (rolling window, EWMA, GARCH). **VaR = Factor × Volatility** is the backbone of this guide.

Confidence c	Tail $1-c$	Normal factor $\Phi^{-1}(1-c)$	Normal ES factor $\phi(z)/(1-c)$	ES/VaR	Use Typical use
90%	10%	-1.2816	1.7550	1.369	· internal limits
95%	5%	-1.6449	2.0627	1.254	/ UCITS · fund reporting
97.5%	2.5%	-1.9600	2.3378	1.193	FRTB ES PRIIPs VaR
99%	1%	-2.3263	2.6652	1.146	Basel · Basel IMA, backtesting
99.9%	0.1%	-3.0902	3.3671	1.090	· economic capital

In practice • Keep or drop μ ? / • μ

A positive μ shrinks VaR (more optimistic). Because μ is so unstable, most banks, and the course's backtest demo, **drop μ** , which is conservative. If kept, use a long, robust estimate scaled by $\mu_{hD} = h\mu_{1D}$. The $-\frac{1}{2}\sigma^2 T$ term in the PRIIPs formula is not a mean but a log-normal drift correction (Itô), present even when $\mu = 0$.

Pitfall • What VaR does not say / • VaR

A 99% VaR says "1% of outcomes are worse" but **not how much worse**. Two portfolios can share a VaR while one has a thin tail and the other a very fat one. VaR is also not subadditive (bond example, §3.4). These two facts drove the regulatory move from VaR to ES.

Regulatory mapping • VaR definition and parameters / • VaR

Framework	English
UCITS	VaR mandatory for complex strategies; absolute VaR $\leq 20\%$ NAV or relative $\leq 2\times$ reference; 99%, 20 days, ≥ 1 year of data, daily (§8.1).
AIFMD	No prescribed VaR; requires "appropriate" risk measurement documented in policy; if used, VaR and method are reported in Annex IV.
MiFID II	No VaR limit; PRIIPs uses a 97.5% VaR to set the MRM class, which feeds target-market and suitability decisions.

3 Expected Shortfall and Coherent Risk Measures

3.1 Two equivalent definitions

$$\text{ES}_c(X) = \mathbb{E}[X \mid X \leq \text{VaR}_c(X)] \quad (\text{conditional tail expectation})$$

$$\text{ES}_c(X) = \frac{1}{1-c} \int_0^{1-c} \text{VaR}_u(X) du \quad (\text{average of tail percentiles})$$

First definition: **given** that we are in the tail, what is the average loss? Second: average every VaR inside the tail (99%, 99.1%, 99.2%, ... 99.99%, ...). For continuous distributions they coincide; for discrete ones the second is more robust (always coherent). Both are **model-free**: the working formula depends on the distribution of X .

In short: **VaR marks the door to the tail; ES tells you what is inside**.

3.2 Deriving Normal ES, step by step

Route A: from the average of tail percentiles. Rewrite the integral over the upper tail, substitute $\text{VaR}_u = \mu + \sigma\Phi^{-1}(1-u)$ and pull out the mean:

$$\text{ES}_c = \frac{1}{1-c} \int_c^1 (\mu + \sigma\Phi^{-1}(1-u)) du = \mu + \frac{\sigma}{1-c} \int_c^1 \Phi^{-1}(1-u) du$$

Substitute $u = \Phi(z)$, $du = \varphi(z)dz$, limits $[c, 1] \mapsto [\Phi^{-1}(c), \infty)$, and use symmetry $\Phi^{-1}(1-\Phi(z)) = -z$:

$$= \mu - \frac{\sigma}{1-c} \int_{\Phi^{-1}(c)}^{\infty} z \varphi(z) dz$$

Key trick: the antiderivative of $z\varphi(z)$ is $-\varphi(z)$, since $\frac{d}{dz}e^{-z^2/2} = -ze^{-z^2/2}$. At the upper limit $\varphi(\infty) = 0$, hence:

$$\boxed{\text{ES}_c = \mu - \sigma \frac{\varphi(\Phi^{-1}(c))}{1-c}} \quad c = 99\% : \frac{\varphi(-2.32635)}{0.01} = \frac{0.026652}{0.01} = 2.6652 \Rightarrow \text{ES}_{99\%} = \mu - 2.6652\sigma$$

Route B: first tail moment ÷ tail probability. Tail mean $= \int_{-\infty}^{\mu-n\sigma} xf(x)dx / \int_{-\infty}^{\mu-n\sigma} f(x)dx$. The denominator is the tail mass $1-c$, the numerator the first tail moment. The substitution $x = \mu - z\sigma$ yields the same result. Either route is fine for an interview.

The "inverse CDF inside the PDF" ($\varphi(\Phi^{-1}(c))$) looks odd until you plug in numbers: find the percentile -2.326 , then the density height there, 0.02665 .




Figure 2. 99% VaR (solid) and 99% ES (dashed) under the standard Normal. ES is the centre of mass of the shaded tail.

Worked example 3 • The tail beyond 2σ / $3 \cdot 2\sigma$

The share of the standard Normal left tail beyond -2σ is $\Phi(-2) = 2.275\%$, i.e. confidence 97.725%. The mean of that tail is $-\varphi(2)/\Phi(-2) = -0.05399/0.02275 = -2.373$. In a Normal world the average loss beyond 2σ is only about 19% worse than 2σ itself: Normal tails are thin.

3.3 The ES/VaR diagnostic ratio and tail heaviness

$$\rho_c = \frac{\text{ES}_c}{\text{VaR}_c} \quad \text{Normal}(\mu = 0) : \rho_c = \frac{\varphi(z_c)}{(1-c)|z_c|} \xrightarrow{c \rightarrow 1} 1; \quad \text{GPD tail} : \rho \xrightarrow{\text{VaR} \gg u} \frac{1}{1-\xi}$$

ρ is unit-free: how far the mass behind the door sits from the door. Under the Normal, ρ is moderate (1.146 at 99%) and **drifts slowly towards 1** as $c \rightarrow 1$ (table in §2.2): the deeper you go, the thinner a Normal tail looks. For power-law (fat) tails ρ converges to the constant $1/(1-\xi) > 1$: the tail never thins.

The numerator $\int x f(x) dx$ can explode for fat tails while the denominator (tail mass) stays finite; that is why ES can be undefined ($\xi \geq 1$) although VaR always exists. Hull's range $\xi \in [0.1, 0.4]$ keeps ρ between 1.1 and 1.7.

Note: the course slide describes Gaussian ρ as "fixed at every c "; strictly it is moderate and slowly decreasing towards 1, which is the contrast with the constant limit of fat tails.

ξ	Tail type	$\rho \approx$
0.10	Hull	1.11
0.25	Moderate	1.33
0.40	Hull	1.67
0.50	ES = 2×VaR	2.00
$\rightarrow 1$	Cauchy ES	$\rightarrow \infty$

3.4 Coherence and subadditivity

Artzner et al. (1999) define a coherent risk measure $\rho(\cdot)$ (on positive losses) by four axioms:

- Monotonicity:** a portfolio with always-larger losses is not less risky;
- Subadditivity:** $\rho(X + Y) \leq \rho(X) + \rho(Y)$; merging does not add risk, i.e. diversification;
- Positive homogeneity:** double the position, double the risk;
- Translation invariance:** adding cash m reduces risk by m .

ES satisfies all four; VaR generally fails axiom 2 (except for Normal or elliptical distributions). Subadditivity lets ES capital be decomposed and aggregated cleanly across desks.

Worked example 4 • Discrete defaults / 4 • VaR ES

Three independent bonds A, B, C, face value £1,000 each, default probability 0.5% each, loss given default = face value. The portfolio loss distribution is tabulated below.

VaR: a single bond defaults with $0.5\% < 1\%$, so $\text{VaR}_{99}(A) = £0$. For the portfolio the cumulative probability at £0 is only $98.51\% < 99\%$, so $\text{VaR}_{99}(ABC) = £1,000 > 0+0+0$. **VaR violates subadditivity:** diversification "increases" VaR.

ES: for one bond, half of the 1% tail is a £1,000 loss and half is £0, so $\text{ES}(A) = 0.005 \times 1000 / 0.01 = £500$. For the portfolio, 2 and 3 defaults occupy 0.0075% and 0.0000125% of probability; the remaining conditional mass 0.992525 is one default: $\text{ES}(ABC) = £1,007.49 \leq £1,500$. **ES is subadditive.**

Note how the calculation treats the tail as exactly 1% of probability, splitting the £1,000 atom where it straddles the 99th percentile. This is the "average of tail percentiles" definition at work, and it is why that definition is preferred for discrete P&L.

Defaults	Loss	Probability	Cumulative
0	£0	98.5075%	98.5075%
1	£1,000	1.4850%	99.9925%
2	£2,000	0.0075%	≈100%
3	£3,000	0.0000125%	100%

3.5 Why FRTB chose 97.5% ES

Under the Normal the 97.5% ES factor 2.338 almost equals the 99% VaR factor 2.326: for well-behaved distributions capital is roughly unchanged, but for fat-tailed books 97.5% ES is materially higher than 99% VaR. It penalises exactly the tail shape. That is the calibration logic behind the Basel move from 99% VaR to 97.5% ES.

The price is harder backtesting: ES is a conditional mean rather than a binary event and is not *elicitable* (Gneiting 2011), so there is no test as clean as counting breaches. FRTB therefore still backtests **VaR** at 97.5% and 99% while capital is measured with ES.

Regulatory mapping • Expected Shortfall / • ES

Framework	English
UCITS	ES not required; CESR encourages supplementing VaR for tail-heavy strategies (arbitrage, structured), with ES a common choice.
AIFMD	Not required; managers may use ES as an internal limit metric, documented in the risk policy.
MiFID II	Through IFR/CRR for large investment firms' trading books: the FRTB internal model uses 97.5% ES (EU from 2027).

4 Volatility: Rolling Window, EWMA and GARCH

4.1 Rolling-window standard deviation

The simplest approach: each day compute the sample s.d. of the latest n daily returns (Excel STDEV.S), shifting the window by one observation. The course demo uses $n = 21$ (about a month) and a backtest period of 200+ trading days.

- n too short (10–15): noisy estimates, jumpy VaR;
- n too long (60–100): insensitive to recent volatility, clustered breaches at the start of a crisis;
- the Normal factor presumes at least ~25 observations.

Weakness: equal weights inside the window, zero outside. A shock lifts volatility when it enters and drops it **abruptly** when it exits n days later (the "ghost effect").

4.2 EWMA (RiskMetrics)

$$\sigma_t^2 = \lambda \sigma_{t-1}^2 + (1 - \lambda) u_{t-1}^2 \iff \sigma_t^2 = (1 - \lambda) \sum_{i \geq 1} \lambda^{i-1} u_{t-i}^2, \quad \tau_{1/2} = \frac{\ln 0.5}{\ln \lambda}$$

EWMA stands for Exponentially Weighted Moving Average. It is a "probability filter": it weights a **prior** (yesterday's variance σ_{t-1}^2) against **news** (yesterday's squared return u_{t-1}^2) using λ .

- Low λ : more weight on news, very responsive, but more procyclical and noisy;
- High λ : more weight on the prior, slow to react, close to a long window.

RiskMetrics (J.P. Morgan, 1994) uses $\lambda = 0.94$ for daily data, a half-life of about 11.2 days; $\lambda = 0.97$ gives about 22.8 days, and observations older than 60–80 days barely matter. EWMA has no fixed window and hence no ghost effect.




Figure 3. Weight on past squared returns. EWMA decays smoothly; the 21-day window drops to zero on day 22.

In practice • Three EWMA implementation details / • EWMA

- 1. Initialisation:** seed σ_0^2 with the average squared return over a long history (e.g. the first 250 days). With a long enough burn-in the seed is immaterial ($\approx 150+$ days for $\lambda=0.97$).
- 2. Indexing:** $\sigma_{t|t-1}$ must use returns up to $t-1$ only to be a genuine forecast; including u_t gives an ex-post estimate and look-ahead bias in backtests. The Bank of England paper shows the two conventions lead to different FHS outcomes.
- 3. Choosing λ :** calibrate so breach counts match theory, or so devolatilised residuals look i.i.d., but beware over-fitting; the course workbook's RiskMetrics sheet uses $\lambda=0.88$ as a sensitivity illustration.

In practice • Setting exponential weights with λ / • λ

1. Two equivalent computations: the recursion $\sigma_t^2 = \lambda \sigma_{t-1}^2 + (1-\lambda)u_{t-1}^2$ needs only yesterday's variance and return, one Excel formula filled down; the explicit form $\sigma_t^2 = \sum_i w_i u_{t-i}^2$ with $w_i = (1-\lambda)\lambda^{i-1}$. With finite data (n observations) the weights sum to $1-\lambda^n$, so **normalise:** $w_i = (1-\lambda)\lambda^{i-1}/(1-\lambda^n)$. At $\lambda=0.94$ truncation at 250 days is negligible; at $\lambda=0.99$ it is not.

2. Choosing λ :

Method	How
Industry convention	RiskMetrics: 0.94 daily, 0.97 monthly
From half-life	$\lambda = 0.5^{1/au_{1/2}}$; a 20-day half-life gives $\lambda \approx 0.966$
From mean look-back	mean weight lag $1/(1-\lambda)$: about 17 days at 0.94, 33 at 0.97
Statistical fit	minimise the squared error of σ^2 forecasts against realised u^2 , or maximum likelihood
Backtest calibration	breach rate close to theory, without clustering

Covariance matrices need **one common λ** across assets, otherwise the matrix may not be positive semi-definite and portfolio variance can even turn negative; yet the Bank of England study finds different "optimal λ " by asset (WTI 1.00, Gasoil 0.958, NatGas 0.973), a trade-off to manage.

3. Initialisation: seed σ_0^2 with the average squared return over the first 1–2 years; allow a burn-in of about 5 half-lives (≈ 60 days at $\lambda=0.94$, ≈ 115 at 0.97).

4. In historical simulation: give scenario i days ago a normalised weight, sort P&L from worst to best, and the scenario where cumulative weight first reaches 1% is the 99% VaR (deep dive after §5.2).

5. Other parameterisations: exponential weights can equally be set by half-life or effective window length, all interchangeable; GARCH adds a long-run term ω on top of the λ structure (§4.3).

Method	How
Industry convention	RiskMetrics Day 0.94 0.97
From half-life	$\lambda = 0.5^{1/au_{1/2}}$ 20 $\rightarrow \lambda \approx 0.966$
From mean look-back	$1/(1-\lambda)$ 0.94 17 0.97 33
Statistical fit	$\sigma^2 u^2$ or
Backtest calibration	Breach rate close to theory, without clustering

4.3 GARCH(1,1) versus EWMA

$$\sigma_{t+1}^2 = \omega + \alpha r_t^2 + \beta \sigma_t^2, \quad \bar{\sigma}^2 = \frac{\omega}{1 - \alpha - \beta}, \quad \text{EWMA} = \text{GARCH with } \omega = 0, \alpha = 1 - \lambda, \beta = \lambda$$

GARCH stands for Generalized Autoregressive Conditional Heteroskedasticity: volatility varies over time (conditional heteroskedasticity) and depends on past squared returns and past variance (autoregressive); Bollerslev (1986) generalised Engle's (1982) ARCH model by adding the lagged variance term.

	EWMA	GARCH(1,1)
Parameters	One: λ	3 ω, α, β MLE
Mean reversion	None; forecast is flat across horizons	Yes, towards $\bar{\sigma}^2$ at speed $\alpha + \beta$
Term structure	$\sigma_{hD} = \sigma \sqrt{h}$	$\sum_k [\bar{\sigma}^2 + (\alpha + \beta)^{k-1} (\sigma_{t+1}^2 - \bar{\sigma}^2)]$
Practitioner view	Simple, transparent, uniform across assets; more procyclical	More parameters, unstable estimates; forecast pulled towards the long-run mean, may under-react in a crisis

Deep dive • Why practitioners prefer EWMA / • EWMA

GARCH's mean reversion is academically more "correct", but risk management cares about volatility over the **next week or two**. When a crisis erupts, GARCH forecasts are pulled back towards the long-run mean and look optimistic exactly when conservatism is needed; three parameters also need periodic re-estimation across hundreds of risk factors. EWMA has one parameter, no estimation error and is uniform across factors; these engineering virtues usually outweigh the theoretical shortcoming.

Regulatory mapping • Volatility estimation / •

Framework	English
UCITS	History ≥ 1 year (250 days), shorter only if volatility rises significantly; data refreshed at least quarterly. Weighting schemes such as EWMA should be justified in the RMP, including the effective observation period.
AIIFMD	Free choice of method, subject to independent validation and periodic review (Level 2 Art. 45); document the rationale for λ and window length.
MiFID II	PRIPs Category 2 MRM estimates σ , skewness and kurtosis from 5 years of returns (at least 2 years if daily).

5 Historical and Filtered Historical Simulation

5.1 Historical simulation, step by step

1. Identify the risk factors driving portfolio value (equity prices, FX, curve nodes, implied vols, ...) and collect their historical changes over the lookback (e.g. 250 or 500 days);
2. Assume "tomorrow replays historical day i ": **revalue** today's positions under day i 's factor moves, giving N scenario P&Ls;
3. Sort the N P&Ls and take the $1 - c$ percentile as VaR; average the scenarios worse than VaR for ES.

Single-index shortcut: sort historical returns and take PERCENTILE(., 0.01).

Pitfall • A tail of 2.5 observations / • 2.5

A 1% tail of a 250-day window holds 2.5 observations; 300 days gives 3. VaR is set by one or two worst days and ES is an average of two or three numbers, so sampling error is large. Conventions for "which observation" also differ: Excel PERCENTILE.INC interpolates between neighbours, PERCENTILE.EXC uses position $(N + 1)p$, and many banks take the $\lceil Np \rceil$ -th worst. In small samples they can differ by tens of basis points; **document the choice**.

5.2 Plateaux and insensitivity

Equal-weighted HS VaR moves in **plateaux**: it changes only when the worst observation leaves the window or a worse one enters. Two consequences:

- after a shock, VaR stays elevated for a full window length (250 days) even if markets have calmed;
- when markets turn from calm to stressed, VaR reacts slowly and breaches cluster.

Moreover, no unfiltered HS model can ever produce a VaR **larger than the worst loss in its window**; it cannot react to stress beyond what it has "seen".

Regulation and practice: Basel III / FRTB require at least a 250-day lookback, and FRTB adds a stressed calibration window. Front-office internal VaR often uses 60–150 days for responsiveness (more procyclical); more commonly banks keep a long window and apply EWMA decay or volatility scaling to historical returns: the "best of both worlds".

Deep dive • Plateau effect and exponential weighting: problem and remedy / •

What the plateau effect is: historical simulation or an equal-weighted rolling window gives every observation inside the window the same weight and zero outside, so VaR moves only when one of the worst observations enters or leaves, producing step-like plateaux. Three consequences:

- **Slow reaction:** when calm turns to stress the window still holds calm data, VaR is too low and breaches cluster (as in Feb–Mar 2025 in §9.3);
- **Ghost effect:** an extreme shock keeps VaR high for a full window length, then VaR drops abruptly on day $n + 1$ when it leaves;
- **Ceiling:** unfiltered HS VaR can never exceed the worst loss in the window.

It appears in regulatory 250-day HS VaR (Basel, UCITS), equal-weighted rolling standard deviations (e.g. a 21-day window), and CCP and bilateral margin models built on HS.

How exponential weighting fixes it: the observation i days ago gets weight $(1 - \lambda)\lambda^{i-1}$, so old data fade smoothly instead of dropping off a cliff. The half-life is about 11 days at $\lambda = 0.94$ and about 23 days at 0.97. Main applications:

1. **EWMA volatility and covariance:** the volatility input to parametric VaR (§4.2);
2. **Filtered historical simulation:** devolatilise and revolatilise with EWMA volatility (§5.3);
3. **Age-weighted historical simulation** (Boudoukh–Richardson–Whitelaw, 1998): weight historical scenarios by λ^{i-1} and take a weighted quantile;
4. **Special case of GARCH:** EWMA is GARCH(1,1) without the long-run mean term (§4.3).

The price: exponential weighting is smoother and more responsive but more **procyclical**: VaR is low in calm periods and jumps in a crisis, and if it drives margin it concentrates margin calls in the crisis. In the Bank of England study, peak-to-trough at $\lambda = 0.95$ is about 1.8 times that of unfiltered HS (2.97 vs 1.69, §5.4). A common compromise keeps a 250+ day window as a floor with exponential weighting on top.

5.3 Filtered HS: devolatilise, then revolatilise

$$x_j^F(i) = x_j(i) \cdot \frac{\sigma_{Q,j}^\mu(t)}{\sigma_{P,j}^\lambda(i)} \quad \text{devol by the vol on day } i, \text{ revol by the vol on VaR day } t$$

The FHS idea (Boudoukh et al. 1998; Hull & White 1998): historical returns carry valuable information about tail shape, but they occurred under different volatility regimes. Standardise each historical return by its own-day volatility (devol, giving residuals), multiply by today's volatility (revol), and obtain the scenario set "as if history replayed at today's volatility"; then run ordinary HS.

The Bank of England paper compares two typical designs:

- **LTV** (long-term devol): devol with a 500-day equal-weighted volatility, revol with EWMA(0.97). All historical returns get the **same** factor, so skewness, kurtosis and correlations are preserved; the distribution is simply rescaled.
- **STV** (short-term devol, Hull–White type): each historical day is devolatilised by its own EWMA(0.97) vol. Factors differ by observation, so skewness, kurtosis, autocorrelation and cross-correlations **change**.

In the paper, WTI sample excess kurtosis rises from 3.56 to 7.35 under STV; most of UK NatGas autocorrelation disappears after scaling.

Deep dive • Why scale a 99th percentile by volatility? / • 99%

Volatility is driven mostly by the **centre** of the distribution, yet we care about the **tail**. We still scale by it because volatility needs far less data to estimate than a 99th percentile: for the same accuracy, vol scaling makes the model much more reactive. The implicit assumption is that the **ratio of volatility to the 99th percentile is stable**. If the tail lengthens or shrinks independently of the centre (e.g. jump risk rises while day-to-day volatility does not), it fails. The paper also finds devolatilised residuals are usually still non-Normal and not fully independent (Jarque–Bera statistics far above the 9.2 critical value): FHS often works partly because a "ghost" of non-normality survives devolatilisation.

5.4 Procyclicality

FHS gives low VaR in calm periods and raises it fast in turmoil. When VaR sets initial margin (CCPs or bilateral OTC), this timeliness produces concentrated margin calls during a crisis and amplifies liquidity stress. The paper measures it by:

- **Peak-to-trough:** max / min margin for a constant portfolio over the observation period;
- **n-day procyclicality:** the largest margin increase over any n days, as % of notional.

For a WTI position (historical data): unfiltered HS peak-to-trough 1.69; LTV($\lambda=0.95$) 2.97, LTV(0.97) 2.53, LTV(0.99) 2.02, STV 2.82. The smaller λ , the more reactive and the more procyclical. Largest 5-day increase: unfiltered 0.39%, LTV(0.95) 3.18%.

Mitigants: EMIR requires CCPs to adopt one of three tools, one being a **floor** at a 10-year equal-weighted VaR; one can also use a partial scaling factor $\frac{K+f}{K+1}$ ($K \geq 0$), between full scaling f and none.

Pitfall • The ratchet model and over-fitting / • " "

The paper builds a cynical "ratchet" VaR: 500-day HS that moves to a worse observation if there were ≥ 3 breaches in two years and to a better one if ≤ 1 . It has no model of the return process yet often passes Kupiec and even reacts faster than HS. Lesson: **passing a backtest $\neq$ a correct model**. Simple tests have low power; validation should also examine residuals, stress scenarios and parameter stability, and monitor calibration drift over time.

Framework	English
UCITS	Parametric, HS, FHS and Monte Carlo are all acceptable if they capture all material risks (including option non-linearity); initial and periodic validation required.
AIFMD	No method restriction; FHS procyclicality should be considered in liquidity and leverage management, especially for funds posting derivative margin.
MiFID II	PRIIPs Category 3 (non-linear) products use bootstrap simulation (10,000+ resamples of historical returns) for the 97.5% VaR, essentially an HS variant.

6 VaR Methods: Non-Normal Corrections, Simulation, Decomposition and Yield Curves

6.1 The Cornish–Fisher expansion

$$z_{CF} = z + \frac{S}{6}(z^2 - 1) + \frac{K}{24}(z^3 - 3z) - \frac{S^2}{36}(2z^3 - 5z), \quad z = \Phi^{-1}(p), \quad \text{VaR} = -\frac{1}{2}\sigma^2 T + z_{CF} \sigma \sqrt{T}$$

Cornish–Fisher uses the first four moments to transform the standard Normal quantile z into one matching the data's shape: negative skewness S and positive excess kurtosis K both push the left-tail quantile further out. It replaces only the factor in the VaR formula; the volatility part is unchanged.

Substituting z gives the coefficients at each level (table below).

On the course S&P 500 sample (2013–2017), $\gamma_1 = -0.437$ and $\gamma_2 = 2.870$; plugged in directly they give a 99% CF factor of -3.247 , about 40% deeper than the Normal -2.326 .

p	z	S coefficient	K coefficient	S ² coefficient
2.5%	-1.960	+0.474	-0.0687	+0.146
1%	-2.326	+0.735	-0.2338	+0.376




Figure 4. The 99% Cornish–Fisher factor as a function of K and S . Dashed: Normal factor.

Deep dive • S and K are parameters, not sample moments / • S K

After the CF transform, the **actual** skewness γ_1 and excess kurtosis γ_2 are complex non-linear functions of S, K , not S, K themselves. The rigorous route: compute sample γ_1, γ_2 , then **solve numerically** for S, K (e.g. `scipy.optimize.fsolve`); the solution is not unique and depends on starting values. Reference values (RiskConcile) are tabulated below.

For small $|\gamma_1|, \gamma_2$ the two nearly coincide, so direct substitution is an acceptable quick approximation.

Validity domain: the transform must be monotone so quantile order is preserved. This requires $|S|$ below about 2.49 and K within an S -dependent band; with $S = 0, 0 \leq K \leq 8$, and K never above 11.55. Implied limits on actual moments: $|\gamma_1| \leq 4.36, 0 \leq \gamma_2 \leq 43.28$. In practice $|\gamma_1| > 2$ is rare, corresponding to $|S|$ below 0.5–0.7.

γ_1	γ_2	S	K
0.1	0.2	0.0958	0.1872
-0.2	0.5	-0.1821	0.4317
1.0	2.0	0.8833	3.5875

In practice • Estimation window / •

Estimating skewness and kurtosis from 21–42 observations is meaningless: one or two extreme days swing them wildly (the course notebook's 62-day rolling exploration confirms it). In practice estimate S, K over 200–500+ days and hold them as **periodically recalibrated parameters**; use the short window only for σ . RiskConcile's 90-day rolling analysis of the FXI ETF shows short-window sample moments periodically leaving the validity domain, distorting VaR.

6.2 Cornish–Fisher VaR in PRIIPs

$$\text{VaR}_{\text{return}} = \sigma \sqrt{N} \left(-1.96 + 0.474 \frac{\mu_1}{\sqrt{N}} - 0.0687 \frac{\mu_2}{N} + 0.146 \frac{\mu_1^2}{N} \right) - \frac{1}{2} \sigma^2 N$$

The European Supervisory Authorities (ESMA, EBA, EIOPA), in the PRIIPs RTS (final draft March 2016), use a 97.5% Cornish–Fisher VaR for Category 2 products, converted to a "VaR-equivalent volatility" (VEV) and mapped to a 1–7 market-risk class. Here μ_1, μ_2 are the skewness and excess kurtosis of **daily** returns and N the number of trading days in the recommended holding period; the $1/\sqrt{N}$ and $1/N$ scaling reflects how skewness and kurtosis decay under aggregation of i.i.d. returns.

Pitfall • The sign of the S^2 term / • S^2

The course notes and the RiskConcile note write the term as " $-0.146S^2$ ". Substituting $z = -1.96$ into the expansion: $-\frac{1}{36}(2z^3 - 5z) = -\frac{1}{36}(-15.06 + 9.80) = +0.146$. The PRIIPs text also has **+0.146**. Trust the derivation; copying the minus sign into a spreadsheet understates VaR for strongly skewed products.

6.3 Extreme value theory: peaks over threshold

$$\text{VaR}_q = u + \frac{\beta}{\xi} \left\{ \left[\frac{n}{n_u} (1 - q) \right]^{-\xi} - 1 \right\}, \quad \text{ES}_q = \frac{\text{VaR}_q + \beta - \xi u}{1 - \xi}$$

Gnedenko–Pickands–Balkema–de Haan: for a high enough threshold u , the excess loss $L - u$ converges to a generalised Pareto distribution (GPD) with scale β and shape ξ . n is the total number of observations and n_u the number above u .

In practice: set u near the empirical 95th percentile (course workbook: 500 scenarios, $u = 160, n_u = 22$); fit β, ξ by maximum likelihood: compute $\ln g_i$ for each exceedance and maximise the sum over β, ξ (Excel Solver). Hull suggests $\xi \in [0.1, 0.4]$. EVT's value is **extrapolation**: it can estimate quantiles beyond the worst observed loss, which HS cannot.

6.4 Method overview

Method	Factor	Volatility	Strengths	Weaknesses
Normal analytical	$\Phi^{-1}(1-c)$	/ EWMA	Simple, decomposable, closed-form ES	Understates fat tails
Student-t	$t^{-1} \cdot \sqrt{(v-2)/v}$	/ EWMA	One parameter for fat tails; suits single stocks	Symmetric; v unstable
Cornish–Fisher	$z_{\text{CF}}(S, K)$	/ EWMA	Skew and kurtosis; used by regulators	Validity domain; needs long window for S, K
Historical simulation (HS)	empirical	implicit	Model-free; full revaluation	Plateaux, slow, few tail points
HS / FHS	residual empirical	EWMA revol	Long history + current vol	Alters moments and correlations; procyclical
EVT (GPD)	GPD	—	Extrapolates the far tail	Threshold-sensitive; data hungry

6.5 Three VaR frameworks: parametric, historical and Monte Carlo

All three compute the same quantity, the $1 - c$ quantile of portfolio P&L. They differ only in **where the distribution comes from**.

- **Parametric** (variance–covariance, delta-normal): assume a distribution, estimate $\sigma_P = \sqrt{\mathbf{w}' \Sigma \mathbf{w}}$, $\text{VaR} = \mu + \text{factor} \times \sigma$. Variants use Normal, Student-t or Cornish–Fisher factors, and delta-gamma approximations for options (Chs 2–4, §6.1).
- **Historical simulation**: assume "tomorrow replays a past day", apply each of the last N days' factor moves to today's positions with full revaluation, sort and take the quantile. Variants: equal-weighted, age-weighted, filtered HS (Ch. 5).
- **Monte Carlo**: specify a stochastic model for risk factors (multivariate Normal or t, GARCH, jumps, copulas), simulate 10,000–100,000 scenarios, fully revalue each and take the quantile.

	Parametric	Historical simulation	Historical	Monte Carlo
Source of distribution	Assumed (Normal, t, CF)	Actual historical moves		Specified stochastic model
Fat tails, skew	Needs a corrected factor	Automatic within the window		Depends on the model
Non-linearity	Poor; delta-gamma approximation	Good (full revaluation)		Good, including path dependence
Correlation	Covariance matrix, may be low in calm periods	Historical co-movement		Model-set, can be stressed
Beyond history	Yes (extrapolates)	No; capped at the worst day		Yes
Data and compute	Least, fastest	Needs complete history		Heaviest (N revaluations)
Transparency	High	High, easy to explain		Low, most model risk
Typical weakness	Understates tails	Plateaux, few tail points		Misspecification, sampling error

Worked example 5 • One dataset, three methods / 5 • Method

The last 500 S&P 500 daily returns (Jan 2016–Jan 2018, daily volatility 0.64%), 1-day 99%: see the table below.

The Normal parametric result is the most optimistic; historical simulation captures the sample's actual tail. Monte Carlo depends on the model: an unconstrained maximum-likelihood t fit gives about 1.9 degrees of freedom, implying infinite variance and a meaningless ES, a vivid example of model risk.

Method	VaR	ES
Parametric (Normal)	1.49%	1.71%
Historical simulation	1.87%	2.55%
Monte Carlo (t, 4 d.f., scaled to sample σ)	1.68%	2.31%

In practice • Risk-factor mapping and Monte Carlo sampling error / •

- **Equities:** map single stocks to an index via β , but keep single-name factors for concentrated holdings or idiosyncratic risk is understated.
- **Bonds and rates:** map cash flows to key tenors (e.g. 3M, 1Y, 2Y, 5Y, 10Y, 30Y) with PV01 per tenor as the exposure; RiskMetrics mapping splits each cash flow between adjacent vertices preserving present value and variance. Add spread factors (CS01) for credit.
- **FX:** map net exposure by currency; non-base-currency assets carry both an asset and an FX factor (the cross term in §11.2).
- **Options:** delta-gamma-vega in parametric VaR, or delta-gamma-Cornish-Fisher using the first four moments of the quadratic P&L in the CF expansion. With heavy option books, switch to full-revaluation HS or Monte Carlo.
- **Monte Carlo sampling error:** the quantile's standard error is about $\sqrt{p(1-p)/N}/f(q)$, roughly ± 0.13 percentage points at $N = 10,000$ in Example 5. Use enough paths or variance reduction (antithetic variates, importance sampling), and fix the random seed for reproducibility.

Choosing a method and regulatory use

Framework	Method	Key points
UCITS CESR/10-788	All three allowed	Must capture all material risks including option gamma/vega and basis (e.g. swap–bond spread); independent validation after development; monthly backtesting. Limits apply whatever the method.
AIFMD	Manager's choice	Documented in the risk policy and independently validated; Annex IV reports the VaR and its method (Ch. 17).
MiFID II / PRIIPs	Set by product category	Category 2 (linear) uses parametric Cornish–Fisher VaR; Category 3 (non-linear) uses bootstrap simulation (10,000+ resamples of historical returns), between HS and Monte Carlo.
Banks (and CRR investment firms)	Not prescribed	Mostly HS, some Monte Carlo; FRTB PLA effectively requires full revaluation of non-linear positions, which pure parametric models struggle to pass.
EMIR (CCP margin)	Not prescribed	Mainly HS and FHS, with procyclicality controls (§5.4).

In practice: linear books (equities, bonds, spot FX) can use parametric VaR with a t or CF factor; options and structured products need full-revaluation HS or FHS; path-dependent products (barriers, autocallables) need Monte Carlo. Most banks and alternative UCITS run HS/FHS as the main engine, use parametric VaR for fast decomposition (§6.6), and Monte Carlo for complex products and long horizons.

6.6 VaR decomposition: marginal, component and incremental VaR

$$MVaR_i = \frac{\partial VaR}{\partial w_i} = z \frac{(\Sigma \mathbf{w})_i}{\sigma_P}, \quad CVaR_i = w_i MVaR_i, \quad \sum_i CVaR_i = VaR, \quad IVaR = VaR(\mathbf{w} + \Delta \mathbf{w}) - VaR(\mathbf{w}) \approx \Delta \mathbf{w}' MVaR$$

VaR is positively homogeneous, so by Euler's theorem the **component VaRs** (position $\times$ marginal VaR) add up exactly to portfolio VaR. This answers three daily questions:

- **Marginal VaR:** how much does VaR change per extra €1 of an asset? Guides which trade reduces risk most;

- **Component VaR:** where does portfolio risk come from? Used for risk budgets and limit allocation;
- **Incremental VaR:** how much would a new trade change VaR? Used in pre-trade approval; recompute directly for large trades rather than using the marginal approximation.

For HS or Monte Carlo, take each position's average P&L in scenarios where portfolio P&L is near VaR (or decompose ES the same way, which is more stable).

Worked example 6 • Decomposing VaR for three assets / 6 • VaR

€100m portfolio: equities €50m (16% annual volatility), government bonds €30m (5%), credit €15m (6%); correlations equity–govt −0.2, equity–credit 0.5, govt–credit 0.4. 1-day 99% parametric VaR = **€1.23m**; decomposition in the table below.

Stand-alone VaRs sum to €1.52m against €1.23m for the portfolio; the €0.29m gap (about 19%) is the diversification benefit. Equities are half the notional but 92% of the risk; government bonds, negatively correlated with equities, contribute almost nothing. Incremental VaR: another €10m of equities adds about €0.23m of VaR; another €10m of government bonds adds only about €0.004m.

Position	VaR		Share
Equities	€1.17m	€1.14m	92.4%
Government bonds	€0.22m	€0.007m	0.6%
Credit	€0.13m	€0.086m	7.0%
Total	€1.52m	€1.23m	100%

6.7 Illiquid assets and data gaps

AIFs often hold private equity, real estate and private credit without daily prices. VaR computed on appraisal series badly understates risk because appraisals are smoothed. Common treatments:

- **Proxies:** map to listed comparables (private equity → sector small-cap index with a leverage-adjusted β; private credit → high-yield or leveraged-loan index);
- **Unsmoothing:** apply the Geltner adjustment $r_t = (r_t^* - \phi r_{t-1}^*) / (1 - \phi)$ to appraisal returns r_t^* , with ϕ the first-order autocorrelation, to restore suppressed volatility;
- **Look-through:** funds of funds and multi-layer structures should look through to underlying holdings; if impossible, use fund-level return proxies with a conservative buffer;
- **Missing data:** for new issues or mismatched holidays, backfill, proxy or interpolate, and record the scope and effect of each treatment;
- **Complementary measures:** for such assets, pair VaR with stress tests, liquidity buckets and drawdown analysis (Parts II and III).

6.8 Multi-factor models and multi-factor VaR

$$r_i = \alpha_i + \sum_k \beta_{ik} f_k + \varepsilon_i, \quad \Sigma = B \Omega B' + D, \quad \text{VaR} = z \sqrt{\mathbf{b}' \Omega \mathbf{b} + \mathbf{w}' D \mathbf{w}}, \quad \mathbf{b} = B' \mathbf{w}$$

A multi-factor model explains many asset returns with a few common factors (market, rates, spreads, industries, styles). Its core job is **dimension reduction**: the risk of thousands of assets becomes the risk of dozens of factors plus each asset's specific risk. B holds factor exposures, Ω the factor covariance ($K \times K$, $K \ll N$), D the diagonal specific-risk matrix (residuals assumed uncorrelated).

Type	Factor source	Examples
Macroeconomic	Observable economic and market variables	GDP, inflation, rates, oil, USD
Fundamental / style	Asset characteristics	Industry, country, value, size, momentum, quality (Barra, Axioma-type models)
Statistical	Extracted from return data	Principal components, e.g. level, slope and curvature of the yield curve

Type	Factor source	Examples
Macroeconomic	Observable economic and market variables	GDP Rates
Fundamental / style	Asset characteristics	Barra Axioma
Statistical	Extracted from return data	PCA Rates

Multi-factor models in regulatory frameworks

Framework	How factors are used
FRTB standardised approach	A regulator-specified multi-factor model: prescribed risk factors (e.g. 10 GIRR tenors, credit spreads, equity, commodity, FX), delta/vega/curvature sensitivities times prescribed risk weights, aggregated with prescribed correlations. EU from 2027; also used for K-NPR by CRR investment firms.
FRTB internal models	Each risk factor must pass the eligibility test (RFET); failures are NMRFs capitalised by stress scenarios; stressed ES is calibrated on a reduced factor set explaining at least 75% of full-set ES.
UCITS CESR/10-788	The VaR model must cover every risk factor with more than negligible influence; in practice, map positions to factors (§6.5).
AIFMD Annex IV	Net equity delta, bucketed DV01, CS01 and vega are factor exposures in all but name (Ch. 17).
EMIR uncleared margin	UCITS and AIFs typically compute initial margin with ISDA SIMM, a standardised multi-factor sensitivity model similar in structure to SBM.
Solvency II (insurance)	The standard formula charges interest, equity, property, spread, currency and concentration modules and aggregates with correlation matrices; internal models often use factor-based economic scenario generators.
EBA / ECB stress tests	Macro factor paths mapped through satellite models to credit losses and market shocks (§10.4, §14.2).
US SEC Rule 18f-4 (reference)	Fund derivatives VaR limits in the UCITS mould: relative VaR $\leq 200\%$ of a designated reference portfolio, or absolute VaR $\leq 20\%$ of net assets; 99%, 20 days, at least 3 years of data.

Implementing multi-factor VaR

1. **Choose factors:** market, industry and style for equities; tenors and spreads for bonds; plus FX.
2. **Estimate exposures B :** time-series regression or characteristics; DV01, CS01 and delta directly for bonds and derivatives.
3. **Estimate factor covariance Ω :** typically EWMA with shrinkage or positive-definiteness repair; stress key correlations where needed.
4. **Estimate specific risk D :** residual variances; do not ignore it for concentrated holdings.
5. **Compute:** parametric with the formula above; for HS or Monte Carlo, replay or simulate factor returns, bootstrap or simulate residuals separately, and fully revalue.

Why factors: 1,000 stocks need about 500,000 covariances, 40 factors only about 800, and more stable; new or data-poor securities get risk through their exposures (§6.7); VaR splits into systematic and specific parts and by factor (§6.6); stress tests and VaR share one factor set.

Worked example 12 • Multi-factor VaR / 12 • VaR

€100m equity portfolio of 40 stocks. Exposures: market $\beta = 1.1$, value 0.3, size -0.2 ; annual factor volatilities: market 16%, value 6%, size 5%; each stock's specific volatility 25%, about 4.0% at portfolio level after equal-weight diversification. Results in the table below.

The portfolio is essentially "levered market beta": the most effective risk reduction is hedging the market factor with index futures, not swapping stocks.

Result	Value
Portfolio annual volatility	17.8%
of which systematic / specific	17.3% / 4.0%
1 Day 99% VaR	≈ €2.60m
Variance share: market factor	≈ 96%
Variance share: specific risk	≈ 5%
Variance share: style factors	≈ 0 (size -0.9%)

Pitfall • Limits of multi-factor models / •

- **Omitted factors:** when new common drivers emerge in a crisis (liquidity, crowded trades) the model understates risk;
- **Unstable exposures and correlations:** correlations rise in crises and the "uncorrelated specific risk" assumption breaks;
- **Fat tails:** the parametric version still needs a t or Cornish–Fisher correction;
- **Validation:** check the bias statistic (standard deviation of realised return $\div$ forecast volatility, ideally about 1) and backtest regularly.

6.9 Yield curves: construction and use in risk management

$$DF(t) = e^{-z(t)t}, \quad f(t_1, t_2) = \frac{z(t_2)t_2 - z(t_1)t_1}{t_2 - t_1}, \quad DV01_k = -\frac{\partial V}{\partial r_k} \times 0.0001$$

Curves underpin fixed-income and derivative valuation and interest-rate risk: **zero rates** $z(t)$, **discount factors** $DF(t)$ and **forward rates** $f(t_1, t_2)$ are three views of one curve. In risk management each tenor point is a risk factor.

Multi-curve framework: since 2008 one currency needs several curves. The **discount curve** follows the collateral rate: cash-collateralised derivatives under a CSA are discounted at the overnight rate (€STR, SOFR, SONIA), i.e. OIS discounting; **projection curves** are built per floating index (e.g. 3M and 6M EURIBOR). USD and GBP LIBOR have ceased and contracts moved to risk-free rates. The CSA collateral currency sets the discount curve; where several currencies are eligible, the poster holds a cheapest-to-deliver option, and cross-currency collateral brings in FX basis.

Curve construction

- 1. **Choose inputs:** deposits or OIS at the short end, futures or FRAs in the middle, swaps at the long end; liquid and snapped at the same time.
- 2. **Bootstrap:** start at the shortest tenor and use known discount factors to solve for the next, so every input reprices exactly.
- 3. **Interpolate:** log-linear discount factors (simple, piecewise-flat forwards), monotone convex (smooth, no negative-forward spikes), splines (smoother but can oscillate locally). The choice directly affects forwards and key-rate sensitivities.
- 4. **Check:** zero repricing error on inputs; no spikes in forwards; floors in negative-rate regions; long-end extrapolation (insurers use Smith–Wilson converging to an ultimate forward rate).

Other curves	Use Use
Credit (CDS spreads → hazard rates, survival)	CS01, CVA (\$8.7), credit bond valuation
Repo and dividend curves	Equity forwards, futures basis, equity options
Inflation curve	Inflation-linked bonds and swaps; inflation scenarios (\$14)
Commodity forward curve	Contango/backwardation, roll yield; map futures VaR by maturity point
Volatility term structure and surface	Option valuation and bucketed vega (\$10.3.2, \$18.4)

Using curves in risk management

- **Key-rate DV01:** bump one tenor by 1bp at a time, rebuild the curve and reprice to get per-tenor sensitivities. State whether you bump **par quotes** or **zero rates**; results differ.
- **Map to regulatory buckets:** the 10 FRTB GIRR tenors (0.25, 0.5, 1, 2, 3, 5, 10, 15, 20, 30 years); the 4 UCITS duration-netting buckets (\$8.6); bucketed DV01 in Annex IV (Ch. 17).
- **Principal components (PCA):** curve moves are typically explained by three components: **level** (parallel, about 80–90% of variance), **slope** (steepening/flattening) and **curvature** (butterfly). Use them to reduce dimensions, design scenarios and set limits.
- **Curve scenarios:** parallel, steepener, flattener, twist, butterfly. Basel IRRBB prescribes six standard shocks: parallel up/down, steepener, flattener, short up/down. Short shocks decay as $e^{-t/4}$, long shocks build as $1 - e^{-t/4}$; steepener = $-0.65|\Delta R_{\text{short}}| + 0.9|\Delta R_{\text{long}}|$, flattener = $+0.8|\Delta R_{\text{short}}| - 0.6|\Delta R_{\text{long}}|$. After the Basel Committee's 2024 recalibration, EUR sizes from 1 January 2026 are parallel 225bp, short 350bp, long 200bp (implemented in the EU through EBA standards; check the latest text). The same shapes work well as interest-rate stress scenarios for funds.

Worked example 14 • Curve risk hidden by parallel DV01 / 14 • DV01

Portfolio: receive-fixed 2-year swap €200m (loses €39k per +1bp), pay-fixed 10-year swap €50m (gains €44k per +1bp), 30-year government bond €20m (duration 20, loses €40k per +1bp). Net parallel DV01: loses €35k per +1bp. Results under IRRBB shapes with 2026 EUR parameters are in the table below.

The three positions largely offset in parallel DV01, so rate risk looks small; yet a short-end rise loses almost as much as a parallel rise, a flattener loses €3.83m while a steepener makes money. Parallel DV01 alone badly understates curve risk here, so limits and stress tests should be set by key rate or principal component.

Scenario	2Y	10Y	30Y	P&L
Parallel up	+225	+225	+225	−€7.88m
Parallel down	−225	−225	−225	+€7.88m
Steepener	−67	+147	+180	+€1.88m
Flattener	+123	−87	−120	−€3.83m
Short rates up	+212	+29	0	−€7.02m
Short rates down	−212	−29	0	+€7.02m

Regulatory mapping • Non-Normal corrections / •

Framework	English
UCITS	Retail share classes' PRIIPs KID uses Cornish–Fisher VaR (\$6.2); the pre-2023 KIID used volatility bands (SRRI).
AIFMD	Retail-facing AIFs are also in PRIIPs scope; professional AIFs choose freely, and EVT is common in tail-risk reporting.
MiFID II	Distributors use the SRI for target markets: $VEV = [\sqrt{(3.842 - 2 \cdot VaR)} - 1.96]/\sqrt{T}$; MRM classes 1–7 break at 0.5%, 5%, 12%, 20%, 30%, 80%.

7 Backtesting: Procedure, Tests and Traffic Lights

7.1 The logic of a backtest

Picture this: each morning you compute a 99%/10-day VaR using **only data available then** and "seal" it; ten trading days later you compare it with the realised return over that period and check whether the worst-loss prediction was breached. If the model is right, breaches occur about 1% of the time and **independently**. A backtest tests exactly these two properties: **frequency** (unconditional coverage) and **independence**.

7.2 Step by step (course Excel / Python demo)

$$\text{VaR}_t = \Phi^{-1}(0.01) \times \sigma_t \sqrt{10} (+10\mu), \quad r_{10D,t} = \ln(S_{t+10}/S_t), \quad \text{breach}_t = 1\{r_{10D,t} < \text{VaR}_t\}$$

1. **Returns:** the day-1 price is only used to compute the first log-return.
2. **Rolling σ :** 21 returns are needed for the first (daily) s.d.; the window moves by one observation (returns 1–21, 2–22, ...).
3. **VaR:** from the day of the 21st return, compute VaR_t daily.
4. **Forward return:** $r_{10D,t}$ uses prices 10 days after t ; the last 10 days have none.
5. **Eligible comparisons:** $T = N_{\text{prices}} - 21 - 10$. Course data: 1,250 prices $\rightarrow$ **1,219** comparisons.
6. **Statistics:** (a) count and share of breaches; (b) conditional breach probability $\Pr(\text{breach}_t \mid \text{breach}_{t-1})$; (c) comparison with the binomial.
7. **Plot:** VaR_t and $r_{10D,t}$ together with breaches marked; look for clusters.

Pitfall • Overlapping 10-day returns / • 10 Day

Consecutive 10-day forward returns share 9 days of data, so breaches cluster mechanically: one large drop breaches several consecutive comparisons. The binomial "independent breaches" assumption fails and independence tests reject by construction. Regulatory backtesting therefore compares **1-day** VaR with **1-day** P&L (MAR32). To backtest a 10-day VaR, use non-overlapping periods (one tenth of the sample) or overlap-adjusted tests.

7.3 The binomial distribution and Basel traffic lights

$$X \sim \text{Binomial}(T, p = 1 - c), \quad \mathbb{E}[X] = Tp, \quad \Pr(X \leq k) = \sum_{j=0}^k \binom{T}{j} p^j (1-p)^{T-j} \quad (\text{Excel: BINOM.DIST}(k, T, p, \text{TRUE}))$$




Figure 5. Distribution of breach counts for $T=250$, $p=1\%$; numbers above bars are cumulative probabilities.

Basel (1996) rule: with k breaches, the model is **green** if $\Pr(X \leq k) < 95\%$, **yellow** if $95\% \leq \Pr(X \leq k) < 99.99\%$, **red** if $\geq 99.99\%$. For $T = 250$: green 0–4, yellow 5–9, red ≥ 10 .

The zone drives a capital multiplier applied to the VaR-based charge (Basel 2.5: 3 + plus factor; FRTB MAR32 bank-wide: 1.5 rising to 2.0). In the red zone a supervisor may also withdraw model approval.

Breaches	Zone	Basel 2.5	FRTB
0–4	Green	3.00	1.50
5	Yellow/amber	3.40	1.70
6	Yellow/amber	3.50	1.76
7	Yellow/amber	3.65	1.83
8	Yellow/amber	3.75	1.88
9	Yellow/amber	3.85	1.92
≥10	Red	4.00	2.00

Pitfall • Off-by-one zone boundaries / • Zone

Using `scipy.stats.binom.ppf(0.95, 250, 0.01) = 5` and writing "green ≤5" is off by one: `ppf` returns the smallest k whose cumulative probability **reaches** 95%, and that k already belongs to yellow. Correct: green if $k < \text{ppf}(0.95)$, red if $k \geq \text{ppf}(0.9999)$. Also, one course slide gives the yellow threshold as 0.9599; it should be **0.9999**. The slide's breach-probability table is in fact the **probability mass** for $T=252$, $p=5\%$ (`BINOMDIST(..., FALSE)`), not cumulative probabilities for a 99% VaR.

7.4 Kupiec proportion-of-failures test

$$\text{LR}_{\text{POF}} = -2 \ln \frac{(1-p)^{T-x} p^x}{(1-\hat{p})^{T-x} \hat{p}^x}, \quad \hat{p} = \frac{x}{T}, \quad \text{LR}_{\text{POF}} \sim \chi_1^2 \quad (5\% \text{ critical value } 3.841)$$

Null: the true breach probability equals p . The likelihood ratio compares how likely the observed breaches are under the theoretical p versus the observed frequency $\hat{p}$. The test is **two-sided**: too many breaches (aggressive model) and too few (over-conservative, wasted capital) are both rejected.

Example: $T = 249$, $x = 11$, $\hat{p} = 4.42\%$. $\text{LR} = 15.96$, p-value 6.5×10^{-5} : a strong rejection.

7.5 Christoffersen independence and conditional coverage

$$\pi_0 = \frac{n_{01}}{n_{00} + n_{01}}, \quad \pi_1 = \frac{n_{11}}{n_{10} + n_{11}}, \quad \text{LR}_{\text{ind}} = -2 \ln \frac{(1-\pi)^{n_{00}+n_{10}} \pi^{n_{01}+n_{11}}}{(1-\pi_0)^{n_{00}} \pi_0^{n_{01}} (1-\pi_1)^{n_{10}} \pi_1^{n_{11}}} \sim \chi_1^2, \quad \text{LR}_{\text{cc}} = \text{LR}_{\text{POF}} + \text{LR}_{\text{ind}} \sim \chi_2^2$$

n_{ij} counts days with state i yesterday and j today ($1 = \text{breach}$). The course's "probability of a breach given a breach in the previous period" is π_1 . For a good model $\pi_1 \approx \pi_0 \approx p$; $\pi_1 \gg p$ signals clustered breaches and a model that reacts too slowly to volatility changes. The conditional coverage test LR_{cc} (5% critical value 5.991) tests frequency and independence jointly.

In practice • Power and sample size / •

With 250 observations a year and 2.5 expected breaches, every test has low power: a model whose true breach rate is 2% has a sizeable chance of staying green. With few breaches the Christoffersen statistic is also unstable (n_{11} is usually 0 or 1). So: (1) use as long a history as possible (years, many portfolios); (2) backtest several levels (97.5% and 99%); (3) combine with P&L attribution, residual diagnostics and qualitative review rather than a single p-value.

7.6 Three ways to improve VaR

- Mean:** has μ been dropped? A positive μ biases VaR towards optimism.
- Volatility:** responsive enough? Too procyclical? GARCH adds parameters and pulls forecasts towards the long-run mean; practitioners use EWMA but must manage procyclicality.
- Factor:** the Normal percentile is "good enough" but inadequate for empirical P&L; Cornish–Fisher or t factors improve it.

Regulatory mapping • Backtesting / •

Framework	English
UCITS	Monthly backtest of 1-day VaR at least; more than 4 overshootings at 99% in 250 days must be reported to senior management and the regulator, with a model review (consistent with the Basel 0–4 green zone).
AIFMD	Level 2 Art. 45 requires periodic backtests "where appropriate", with no numeric threshold; practice often borrows the UCITS 4-overshooting standard.
MiFID II	Investment firms using CRR internal models follow bank backtesting rules (1-day 99% VaR vs actual and hypothetical P&L, traffic-light multiplier).

8 Regulation of Short-Term Risk Measurement

The three EU regimes differ sharply in how hard they bite on short-term measurement: **UCITS** has binding VaR limits and a backtesting threshold; **AIFMD** prescribes process, leverage calculation and reporting but no VaR cap; **MiFID II** itself sets no VaR rules; its requirements arrive through the PRIIPs KID, client loss reporting and the investment-firm prudential regime (IFR/IFD). This chapter covers each, then the bank framework (Basel/FRTB) as a reference point.

8.1 UCITS: global exposure

Item	Requirement CESR/10-788 Directive 2010/43/EU
Method	Default commitment approach : global exposure $\leq 100\%$ of NAV (conversion, netting and hedging rules in §8.6). VaR is mandatory for complex strategies, non-negligible exotic derivatives, or where commitment does not capture the risk.
VaR Limits	Relative VaR $\leq 2 \times$ reference-portfolio VaR; absolute VaR $\leq 20\%$ of NAV.
Parameters	99% one-tailed; 20 business days; ≥ 1 -year history (shorter only if volatility has risen significantly); data updated at least quarterly; calculated at least daily.
Alternatives	Confidence $\geq 95\%$ and horizon ≤ 1 month allowed, with the 20% limit rescaled by Normal quantiles and $\sqrt{t}$: 95%/5-day $\approx 7.07\%$.
Backtesting	At least monthly, 1-day VaR vs actual daily change; more than 4 overshootings at 99% in 250 days must be reported to senior management and the competent authority, with a model review.
Governance	The model must capture all material risks (including non-linearity), with initial and periodic validation; the method is documented in the risk management process filed with the regulator.
Disclosure	VaR funds disclose expected leverage (sum of notionals) in the prospectus; retail share classes provide a PRIIPs KID (replacing the KIID since 2023).

In practice • Running UCITS VaR day to day / • UCITS VaR Day

- Reference portfolio for relative VaR**: unleveraged and, in principle, free of derivatives (except to express shorts in long/short strategies or for currency-hedged indices); its risk profile must match the fund's objectives, policies and limits; if the risk/return profile changes often or no reference portfolio can be defined, relative VaR must **not** be used; switch to absolute VaR. Selection and maintenance belong in the risk management process and the reference portfolio is disclosed in the prospectus.
- Holding-period conversion**: with a 1-day VaR, rescale the 20% limit by $\text{VaR}_{x \text{ days}} \approx \text{VaR}_{1 \text{ day}} \sqrt{x/t}$ (assuming i.i.d. factor returns), or equivalently multiply the 1-day VaR by $\sqrt{20}$ and compare with 20%. Relative VaR needs no adjustment since both sides scale.
- Limit-usage monitoring**: compute "VaR ÷ limit" daily. Example: the \$6.6 portfolio has a 1-day 99% VaR of €1.23m, $\times \sqrt{20} \approx €5.50\text{m} = 5.5\%$ of NAV, i.e. 27.5% of the absolute VaR cap. Set internal warning (e.g. 80%) and hard (e.g. 90%) lines to keep a buffer.
- Annual report**: disclose at least the **lowest, highest and average** VaR-limit usage over the financial year and the model and inputs (model, confidence level, holding period, data length); disclose leverage employed (sum of notionals). The prospectus discloses expected leverage and the possibility of higher levels.
- Risk management process (RMP)**: method, parameters, reference portfolio, backtesting and stress-testing procedures and escalation are written into the RMP filed with the regulator; material changes require an RMP update.

Deep dive • Absolute versus relative VaR / • VaR VaR

Absolute VaR: the fund's own VaR as a percentage of NAV, compared with no benchmark; $\leq 20\%$ of NAV at 99%/20 days. Suited to absolute-return, multi-strategy and alternative UCITS without a fitting benchmark. Example: a 20-day VaR of 12% of NAV uses 60% of the cap.

Relative VaR: the **ratio** of fund VaR to reference-portfolio VaR, ≤ 2 . Suited to benchmark-driven funds using derivatives for overlays or leverage. Example: fund VaR 12%, reference portfolio (unleveraged MSCI World) VaR 8%, ratio 150%.

	Absolute VaR	Relative VaR
Yardstick	Fixed 20% of NAV	Reference VaR $\times 2$
Cap moves with markets?	No	Yes
In a crisis	Cap reached sooner; forced deleveraging possible	Reference VaR rises too, allowed VaR expands
In calm markets	Room for heavy leverage within 20%	Bounded by $2\times$
Horizon conversion	1-day VaR $\times \sqrt{20}$	Both sides scale; none needed
Prerequisite	None	An eligible reference portfolio

Two common mistakes: (1) relative VaR is **not** tracking-error VaR: the former is a ratio of two stand-alone VaRs, the latter the VaR of the active return "fund minus benchmark", measuring deviation from the benchmark; (2) the method is fixed in the RMP and prospectus and cannot be switched opportunistically to make compliance easier.

	VaR	VaR
Yardstick	20% NAV	VaR $\times 2$
Cap moves with markets?	No	Yes
In a crisis	Cap reached sooner; forced deleveraging possible	VaR VaR
In calm markets	Room for heavy leverage within 20%	Bounded by $2\times$
Horizon conversion	1 Day VaR $\times \sqrt{20}$	Both sides scale; none needed
Prerequisite	None	An eligible reference portfolio

Pitfall • Absolute VaR and leverage / • VaR

The absolute VaR cap limits "loss in normal markets", not leverage itself. ESMA's 2025 analysis found alternative UCITS on absolute VaR ran average gross leverage of about 800% of NAV, above hedge-fund AIFs at 448%. When volatility is low VaR shrinks and funds can lever up within the 20% cap; when volatility returns they may be forced into procyclical deleveraging. VaR limits must therefore be paired with stress testing (Part II) and leverage monitoring.

8.2 AIFMD: process, leverage and reporting

- **Risk function** (AIFMD Art. 15; Level 2 Arts 38–45): functionally separate from portfolio management; limits for market, credit, liquidity, counterparty and operational risk; "where appropriate", periodic backtests of risk measurement and stress tests and scenario analyses. Whether to use VaR is the manager's choice, recorded in the risk management policy.
- **Leverage** (Level 2 Arts 7–8): compute under both the **gross** and **commitment** methods; set and disclose a maximum per AIF; commitment leverage above $3\times$ NAV is "substantial" and triggers extra reporting. Regulators may cap leverage under Art. 25 (and ESMA's 2020 guidelines).
- **Annex IV reporting:** market risk profile (net equity delta, DV01, CS01, vega, VaR and its method, etc.), liquidity profile, stress-test results and monthly returns. See Chapter 17 for the step-by-step preparation process.
- **AIFMD II** (2024/927, Level 1 applicable from 16 April 2026; Ireland S.I. 181/2026): open-ended AIFs select at least two liquidity management tools; loan-originating AIFs capped at 175% (open-ended) and 300% (closed-ended) leverage. The new reporting template follows Level 2 measures after October 2027.
- If sold to retail investors, an AIF also needs a PRIIPs KID (see §6.2).

8.3 MiFID II: clients, products and firm capital

- **Clients:** suitability and appropriateness must reflect risk tolerance and capacity to bear losses; discretionary or leveraged accounts must be told by end of business day when value falls **10%** (and each further 10%) since the last report (Delegated Regulation 2017/565 Art. 62).
- **Products:** product governance uses the SRI (1–7) to express target-market risk tolerance; the SRI combines the PRIIPs MRM (97.5% Cornish–Fisher VaR $\rightarrow$ VEV, §6.2) with the credit risk class.
- **Firm capital** (IFR/IFD): investment firms with trading books compute K-NPR using CRR market-risk methods. EU FRTB applies from **1 January 2027**, with transitional relief to end-2029 under the June 2026 delegated act (e.g. PLA as monitoring only, weekly ES).

8.4 Reference: bank and margin frameworks

Framework	Measure	Key points
Basel 1996 / 2.5	99% 10 Day VaR + VaR, 2009	Capital = multiplier (3 + plus) × VaR; 1-day 99% backtest with traffic lights; ≥250 days of data.
FRTB Basel III MAR30–33	97.5% ES, liquidity horizons 10–120 days, stressed calibration	Desk-by-desk IMA approval: backtesting + PLA; bank-wide 1-day 99% VaR vs actual and hypothetical P&L over 250 days.
PRIIPs RTS	97.5% Cornish–Fisher VaR → VEV → 1–7	Standardised risk disclosure for retail packaged products.
EMIR CCP	99%+ VaR/ES FHS	Must limit procyclical changes: buffer, stressed weighting, or a 10-year lookback floor.

Deep dive • Bank market-risk capital formulas / •

Basel 2.5 (from 2009): capital = $\max\{\text{VaR}_{t-1}, m_c \overline{\text{VaR}}_{60}\} + \max\{s\text{VaR}_{t-1}, m_s s\overline{\text{VaR}}_{60}\}$, with VaR and stressed VaR at 99%/10 days, $\bar{\tau}_{60}$ the 60-business-day average, and multipliers $m_c, m_s \geq 3$ plus the backtesting add-on (§7.3). Stressed VaR is calibrated to a 12-month period of significant stress so capital does not collapse in calm markets.

FRTB (EU from 2027): 97.5% ES cascaded by liquidity horizon $LH_j \in \{10, 20, 40, 60, 120\}$ days using the formula below, where $\text{ES}_T(P, j)$ shocks only factors with liquidity horizon $\geq LH_j$. ES is calibrated on a reduced factor set over a stress period and scaled by the current full/reduced ratio, then blended with the sum of asset-class ES at $ho = 0.5$ (partial diversification), plus stressed capital for non-modellable risk factors (NMRF) and the default risk charge (DRC).

$$\text{ES} = \sqrt{\text{ES}_T(P)^2 + \sum_{j \geq 2} \left(\text{ES}_T(P, j) \sqrt{\frac{LH_j - LH_{j-1}}{T}} \right)^2}, T = 10$$

8.5 Three P&L definitions under FRTB

- **Actual P&L (APL)**: all realised P&L including intraday trading and daily valuation adjustments (excluding fees and commissions);
- **Hypothetical P&L (HPL)**: revalue yesterday's closing positions with today's market data, excluding intraday trades and fees; it measures the effect of market moves on yesterday's book;
- **Risk-theoretical P&L (RTPL)**: P&L produced by the risk model's valuation engine using only modelled risk factors; it measures whether the model explains P&L.

Backtesting: VaR vs APL and HPL. At desk level, more than 12 exceptions at 99% or 30 at 97.5% within 12 months removes IMA eligibility.

P&L attribution (PLA): RTPL vs HPL using Spearman rank correlation (green >0.80, red <0.70) and the Kolmogorov–Smirnov statistic (green <0.09, red >0.12). Red desks fall back to the standardised approach; amber desks attract a capital surcharge.

In practice • The hard part is the P&L itself / • P&L

Textbooks backtest a single index; real books hold equities, options, commodity futures, TRS and rates products. Producing clean, timely HPL and RTPL per desk, consistent with the risk system (snapshot timing, valuation-model consistency, market-data sources, stripping fees and valuation adjustments), often takes longer than the VaR formula itself. Aligning valuation and risk teams on "same timestamp, same curves, same models" is a prerequisite for passing PLA.

8.6 UCITS commitment approach: conversion, netting and hedging

UCITS not using VaR measure global exposure with the **commitment approach**: convert each derivative into the equivalent market value of its underlying, apply permitted netting and hedging, sum, and keep the total $\leq 100\%$ of NAV. Exposure from efficient portfolio management (repo, securities lending) where cash collateral is reinvested above the risk-free return also counts towards the same cap (CESR/10-788).

Instrument	ConversionMethod Conversion
Equity index / stock futures	Contracts × multiplier × index level or price
Bond futures	Contracts × notional size × market price of cheapest-to-deliver bond
Plain vanilla options	Contracts × multiplier × underlying price × delta (maximum delta for barriers)
Interest rate swaps	Market value of underlying (or fixed-leg notional)
CDS (protection seller)	Higher of reference-asset market value and notional
CDS (protection buyer)	Market value of reference asset
Total return swaps	Underlying market value; both legs for non-basic TRS
FX forwards / currency swaps	Notional of the currency leg(s)
Convertibles and embedded derivatives	Referenced shares × share price × delta
Variance swaps	Variance notional × current variance

- **Excluded**: (1) performance swaps that fully exchange the return of held assets for another reference return without extra leverage or optionality; (2) derivatives which, combined with cash held, equal a direct holding of the underlying (e.g. index futures plus equivalent cash), creating no incremental exposure.

- **Netting** (same underlying): derivatives on the same underlying, or a derivative and its underlying asset, may be netted even with different maturities, provided the aim is to eliminate risk.
- **Hedging** (not necessarily the same underlying) must meet all five criteria: aim to offset risk, not generate return; verifiable risk reduction at fund level; offsets both general and specific risk; same asset class; effective in stressed markets. Recognised: duration hedging of bonds with an IRS; beta hedging a diversified equity portfolio with index futures. **Not recognised**: hedging a share with a correlated different share, alpha-preserving beta hedges, merger arbitrage, cross-asset hedges (e.g. credit protection against equities).
- **Duration netting** (interest-rate derivatives): allocate to four buckets (0–2y, 2–7y, 7–15y, 15y+), convert to equivalent positions at the target duration (derivative MtM ÷ target duration × duration). Net fully within a bucket; count 40% of amounts netted between adjacent buckets, 75% one bucket apart, 100% for the most remote buckets and for unnetted positions. Not allowed for arbitrage strategies or hedging arrangements.

Worked example 7 • Commitment calculation / 7 •

NAV €100m: physical equities €40m, cash €60m; derivatives and their conversion as in the table below. Total commitment **€53m = 53% of NAV**, within the 100% cap.

If the €20m net long futures are fully cash-backed as an index-replication strategy, they can be excluded as cash-equivalent, cutting commitment to 33%. Physical equities do not count (they are not derivatives), but a short index future used to beta-hedge them can be offset if it meets the five hedging criteria.

Position	Conversion	Counted
Euro Stoxx 50 12	€30m	€20m (netted, same underlying)
Euro Stoxx 50 3	–€10m	
Long EUR/USD forward (not a hedge)	€15m	€15m
€2,000 Delta 0.4	€8m	€8m
CDS €10m MtM €950	€10m	€10m
Total		€5,300 = 53% NAV

8.7 Counterparty risk and CVA

$$CE = \max \left(\sum_{k \in \text{netting set}} \text{MtM}_k - C, 0 \right), \quad CVA \approx (1 - R) \sum_i EE(t_i) [PD(t_{i-1}, t_i)] DF(t_i), \quad \lambda \approx \frac{s}{1 - R}$$

Market risk measures losses from price moves; **counterparty risk** measures losses when a counterparty defaults while a contract is in the money. Key concepts:

- **Current exposure (CE)**: sum of MtM within a netting set (ISDA master agreement) less collateral C , floored at zero;
- **Potential future exposure (PFE)**: a high quantile (e.g. 95%) of future exposure; **expected exposure (EE)** is its mean and EPE its time average;
- **CVA**: the valuation adjustment for possible counterparty default, with recovery R and hazard rate λ approximated from the CDS spread s ;
- **Wrong-way risk**: default probability rising with exposure (e.g. buying sovereign CDS protection from a bank in that country) badly understates CVA;
- **Mitigants**: netting, collateral agreements (CSA), daily variation margin (VM), initial margin (IM), central clearing.

Worked example 8 • Estimating CVA / 8 • CVA

A 5-year interest rate swap with constant expected exposure $EE = €1\text{m}$ (uncollateralised); counterparty 5-year CDS spread 120bp, recovery 40%, risk-free rate 3%. Hazard rate $\lambda = 0.012/0.6 = 2\%$, 5-year cumulative default probability 9.5%. Summing quarterly gives $CVA \approx €53\text{k}$, about 5.3% of the exposure; the shortcut " $EE \times \text{spread} \times \text{annuity}$ " gives €55k. With a daily-VM CSA, EE falls to the price move over the margin period of risk (tens of thousands of euros) and CVA shrinks to a small fraction.

Framework	Counterparty-risk requirements
UCITS	Exposure to any single OTC derivative counterparty $\leq 5\%$ of assets, or 10% for credit institutions (UCITS Directive Art. 52); measured as positive MtM net of eligible collateral, including margin not protected by client-money rules. Collateral must be liquid, valued daily, of good credit quality, weakly correlated with the counterparty, diversified (ESMA 2014/937: generally $\leq 20\%$ of NAV per issuer), held by a third-party custodian and enforceable at any time; non-cash collateral may not be sold, reinvested or pledged, cash collateral only placed in prescribed low-risk assets; a haircut policy is required.
AIFMD	No numeric cap; due diligence and ongoing monitoring of counterparties with internal limits; Annex IV reports top-5 counterparty exposures in both directions, collateral and clearing shares (Ch. 17); collateral reuse counts in leverage.
EMIR (UCITS and AIFs are financial counterparties)	Standardised OTC derivatives above the clearing threshold must be centrally cleared; uncleared derivatives require daily variation margin and, above the size threshold, initial margin (usually ISDA SIMM); trades are reported to a trade repository.
MiFID II / IFR	Investment firms dealing on own account hold K-TCD (trading counterparty default) capital under the IFR: exposure $\times$ risk factor $\times$ CVA multiplier; large firms under CRR follow bank rules.
Banks	Default exposure via SA-CCR ($EAD = 1.4 \times (\text{replacement cost} + \text{PFE})$) or internal models; separate CVA risk capital (BA-CVA or SA-CVA).

9 Reproducible Case Studies

9.1 Replicating the course backtest: S&P 500, 99%/10-day, 21-obs rolling σ




Figure 6. 99%/10-day VaR (red) and 10-day forward return (blue), Jan 2013–Jan 2018; × marks breaches.

Recomputing independently from the 1,250 closes in the course workbook, following §7.2, reproduces the course notebook exactly: 1,219 comparisons, 25 breaches, 2.05%. Interpretation:

- the breach rate is about twice the theoretical level, so the 21-obs Normal VaR is aggressive;
- breaches cluster (Jan 2014, Sep–Dec 2014, the **August 2015** renminbi-devaluation shock). The 56% conditional probability, far above 1%, mixes genuine volatility clustering with the mechanical effect of overlapping 10-day windows;
- before August 2015 σ was low and the model could not "see" the coming shock; afterwards VaR deepened sharply and then snapped back 21 days later: the classic window effect.

9.2 Comparing 1-day VaR models (Jan 2014–Jan 2018, $T = 999$)

Model	Breaches x	x/T	LR _{POF} (p)	LR _{ind} (p)	Last 250 days	Zone Zone	VaR
Rolling σ (21), Normal	23	2.30%	12.51 (0.0004)	2.57 (0.11)	5	Yellow	−1.62%
EWMA 0.94, Normal	21	2.10%	9.31 (0.002)	3.17 (0.08)	4	Green Green	−1.65%
EWMA 0.94 + Cornish–Fisher*	11	1.10%	0.10 (0.75)	8.15 (0.004)	3	Green Green	−2.31%
Historical simulation 250	15	1.50%	2.20 (0.14)	5.63 (0.02)	3	Green Green	−2.14%
Filtered HS (EWMA 0.94)	14	1.40%	1.45 (0.23)	6.17 (0.01)	3	Green Green	−2.17%

* The CF factor uses full-sample skewness −0.437 and excess kurtosis 2.870 (not inverted to S,K) and so contains look-ahead bias; illustration only. All models use data up to $t - 1$ for day t ; 1-day returns do not overlap. Expected breaches: 9.99.




Figure 7. Three 1-day 99% VaRs: HS 250 plateaux, fast-reacting EWMA, and FHS combining the empirical tail with current volatility.

Takeaways: (1) With a Normal factor, breaches run at about twice the theoretical rate whatever the volatility model: the problem is the factor, not the volatility. (2) Switching to a CF factor or to empirical quantiles (HS, FHS) passes the frequency test. (3) Yet the independence test rejects for CF, HS and FHS: with only 11–15 breaches, two consecutive pairs inflate π_1 ; in small samples LR_{ind} is unstable and should be read alongside the chart. (4) HS VaR stays on a high plateau for about a year after August 2015, while FHS eases back as markets calm: more accurate, but more procyclical.

9.3 Reading an S&P 500 HS VaR backtest (Jan 2025–Jan 2026)




Figure 8. 99% 1-day HS VaR backtest: price and breach days (top), 1-day returns vs VaR (middle), cumulative breaches vs the binomial 95% band (bottom).

Number check: 249 comparisons, 11 breaches (4.42%), 2.49 expected. Independently recomputed Kupiec LR = 15.96, $p = 6.5 \times 10^{-5}$ (the chart's 0.0001 is rounded); $\Pr(X \geq 11) = 5.2 \times 10^{-5}$. The **red** verdict is correct.

Zone boundaries: the chart states "Green ≤5, Yellow ≤10"; under the Basel rule it should read **green 0–4, yellow 5–9, red ≥10** (the off-by-one pitfall of §7.3). The verdict here is unaffected, but counts of 5 or 10 would be misclassified.

What the chart shows:

- **A short lookback:** after the April 2025 tariff shock VaR sits near -6% for about seven weeks and is back near -2% by late May. With a 250-day equal-weighted HS the extreme April days would stay in VaR until April 2026, which suggests a window of roughly one to two months (to be confirmed against the code). With only 30–60 observations the 1% quantile is essentially "the worst recent day" (or an interpolation between the two worst), so **every new worst day is necessarily a breach**: the structural reason for 11 breaches.
- **Clustered breaches:** five breaches in Feb–Mar fall within about five weeks, as markets turned from calm to falling while the window held only January's low-volatility data. A classic independence failure; add a Christoffersen test.
- **Plateaux and steps:** the flat segments and steps in the middle panel are the signature of HS.

Suggestions: extend the window to ≥250 days (the regulatory minimum) or move to FHS (EWMA 0.94–0.97 revol); document the percentile interpolation convention; report 97.5% ES and a stressed VaR alongside; for a portfolio, backtest against HPL.

Medium-Term Risk Testing · weeks to one year

Medium-term testing asks "can the portfolio survive a severe but plausible shock?". It does not rest on a probability distribution: it fixes a shock, computes the loss, liquidity gap and limit breaches, and pre-plans hedging, deleveraging or liquidity management tools.

10 Stress Testing: Theory

10.1 How stress testing complements VaR

VaR is a **probability statement**: from past distributions, "on 99% of normal days we will not lose more than this". A stress test is a **conditional statement**: "if X happens, we lose this much", with no probability attached. They answer different questions:

- VaR marks the door to the tail; a stress test walks through it and examines losses beyond VaR;
- VaR depends on its window and cannot see crises absent from it (§5.2, "never beyond the worst day in the window"); stress tests can impose unprecedented shocks;
- VaR assumes stable correlation; a stress test can impose correlation breakdown explicitly (bonds and equities falling together);
- VaR usually ignores liquidity and redemptions; a stress test can shock assets and liabilities together.

10.2 Types of stress test

Type	How	Strength	Limitation
Sensitivity	One factor moved by a fixed amount	Transparent, additive, easy to limit	Ignores co-movement
Historical	Replay factor moves from real crises	Credible, internally consistent	The next crisis will differ
Hypothetical	Narrative-driven consistent shocks	Forward-looking, targets vulnerabilities	Subjective, hard to calibrate
Reverse	Fix an unacceptable loss, solve for the most plausible scenario	Reveals hidden concentrations	Many solutions; plausibility is model-based
Stressed VaR	VaR calibrated on a stress period	Same framework as VaR	Still a single-period quantile
Liquidity (LST)	Asset liquidation time/cost + redemptions and margin calls	Core fund regulatory requirement	Market-depth data are scarce

10.3 From shocks to P&L: sensitivities and full revaluation

$$\Delta P \approx \sum_i \delta_i \Delta x_i + \frac{1}{2} \sum_{i,j} \Gamma_{ij} \Delta x_i \Delta x_j, \quad \text{bond: } \frac{\Delta P}{P} \approx -D \Delta y + \frac{1}{2} C (\Delta y)^2, \quad \text{option: } \Delta V \approx \Delta \Delta S + \frac{1}{2} \Gamma (\Delta S)^2 + \nu \Delta \sigma$$

There are two ways to map shocks to P&L. **Sensitivities** use a Taylor expansion: fast, suited to daily monitoring and many scenarios. **Full revaluation** reprices every position under shocked market data: exact but slow.

Deep dive · Why Taylor fails for large shocks / ·

A Taylor expansion is accurate only for small moves. Stress shocks are large: at +250bp the bond convexity term matters; at -30% equity a short out-of-the-money put becomes deep in the money, delta moves from -0.2 towards -1, and a gamma approximation badly understates the loss. Barrier options and autocallables can even jump. Rule of thumb: **sensitivities are fine for linear products; positions with options or material convexity need full revaluation in stress tests.**

10.3.1 The Taylor expansion explained

$$f(x + \Delta x) \approx f(x) + f'(x) \Delta x + \frac{1}{2} f''(x) (\Delta x)^2 + \frac{1}{6} f'''(x) (\Delta x)^3 + \dots$$

A complicated function can be approximated near a point by a polynomial whose coefficients are its derivatives there. The **zeroth** term is the current value; the **first** is the tangent (direction and speed); the **second** corrects for curvature (how the tangent leaves the curve); higher terms refine further but fade quickly when Δx is small. Think of driving: first order is "where will I be in a minute at this speed", second order adds "am I accelerating or braking". In risk management Δx is the market shock: the smaller it is, the more first order suffices; the larger, the more higher terms matter, until the expansion itself breaks down.

Term	Bonds	Options	Risk measure
First order	Duration, DV01	Delta Vega Theta	Delta-normal VaR, sensitivity stress
Second order	Convexity	Gamma, cross-gamma	Delta-gamma VaR, CF correction

Worked example 9 • When the Taylor approximation is good enough / 9 •

Bond: 10-year zero-coupon at 3% (modified duration 9.71, convexity 103.7); results in the first table below.

Short put: spot 100, strike 90, volatility 20%, 3 months (delta -0.135, gamma 0.0217); seller's P&L in the second table.

The bond's convexity correction still works for large shocks. The option, at -30%, moves from out of the money to deep in the money, delta runs from -0.135 towards -1, the delta approximation understates the loss by almost 80%, and even delta-gamma understates it by about 28%.

Log returns: $\ln(1+x) \approx x - x^2/2$: 9.53% vs 9.50% at +10%; -35.7% vs -34.5% at -30%. This is also where the $-\frac{1}{2}\sigma^2$ term in log-normal models comes from.

Rates	Exact	First order	First + second order
+100bp	-9.21%	-9.71%	-9.19%
+300bp	-24.96%	-29.13%	-24.46%

Spot	Exact	Delta	Delta-Gamma
-5	-1.00	-0.67	-0.95
-10	-2.88	-1.35	-2.43
-30	-19.30	-4.05	-13.81

Rule of thumb: small shocks and linear products → first order; convexity or gamma → at least second order; large shocks with options or path dependence → full revaluation. Periodically compare the Taylor approximation with full revaluation and escalate when the gap exceeds a threshold (e.g. 5%), which is in effect what the FRTB P&L attribution test checks.

10.3.2 Full revaluation under large shocks: options and path-dependent products

Full revaluation skips derivative approximations: feed the full shocked market data into the same pricing models used for daily valuation, reprice every position, and take P&L = shocked value – base value. Steps:

1. **Base snapshot:** freeze all valuation-date market data: spot, yield curves, credit spreads, implied volatility surfaces, dividends and borrow costs, correlations, FX.
2. **Build shocked market data:**
 - spot: move by relative or absolute shocks;
 - volatility surfaces: implied vols usually rise and skew steepens in sell-offs, so shock vols too (e.g. equity -30% with ATM vol +15 points), and choose how the surface moves with spot: sticky strike or sticky moneyness/delta;
 - curves: shift, steepen or twist, then re-bootstrap, checking rate floors;
 - shock credit spreads, correlations and dividends alongside;
 - check the shocked surface is arbitrage-free (no calendar or butterfly arbitrage), or models may return nonsense.
3. **Reprice** with the same models and parameters as daily valuation (AIFMD Art. 19 valuation, UCITS valuation policy), so risk and valuation stay consistent.
4. **Time:** instantaneous shocks ignore time decay; multi-day scenarios should include theta, carry and roll-down, and the **path** of the shock (see path-dependent products below).
5. **Explain P&L:** split full-revaluation P&L into delta, gamma, vega and an unexplained residual; a large residual flags an unreliable Taylor approximation or a model problem.

Worked example 10 • Vanilla option: joint spot and volatility shock / 10 • Spot

The short put of Example 9, fully revalued with Black–Scholes; seller's P&L in the table below.

The vol shock alone loses about three times the premium (0.71): vega peaks near the money and out of the money. After a large spot fall the option is deep in the money, vega is small, and the extra vol shock adds little. Vega must be shocked together with spot in scenarios, not read off as a stand-alone sensitivity.

Scenario	P&L
Vol only +15 points (20% → 35%)	-2.07
Spot -30%, vol unchanged	-19.30
Spot -30%, vol to 35%	-19.76

Path-dependent products (barriers, autocallables, Asians, lookbacks) depend on the **path** of the underlying, not just the end point, so they rarely have closed forms and are priced by Monte Carlo or PDE methods. Key points in full revaluation:

- **Barrier status:** first check whether a barrier has already been hit; knocked-in or knocked-out products are priced as their new structure;
- **Path of the shock:** for multi-day scenarios (e.g. a 30% fall over 10 days), simulate that path and check barriers at the contractual monitoring frequency (daily, closing, at expiry) rather than only the end point; jumping straight to the end point can miss a knock-in along the way;
- **Observation dates and features:** autocall dates, coupon barriers, memory coupons and knock-in barriers must each be modelled;
- **Common random numbers:** use the **same random draws** for base and shocked valuations, or Monte Carlo noise swamps the true P&L difference;

- **Enough paths:** payoffs are discontinuous near barriers, so use more paths or a Brownian-bridge correction for discrete monitoring;
- **Greeks near barriers:** delta and gamma change violently or jump near a barrier, which is exactly where Taylor approximations fail worst.

Worked example 11 • Full revaluation of a barrier option / 11 • Full revaluation

One-year down-and-in put: strike 100, barrier 70 (daily monitoring), volatility 20%, zero rates. Monte Carlo with 100,000 paths and common random numbers for base and shocks. Base value 2.50, delta -0.26, gamma 0.054 (finite differences). Value changes in the table below.

Three observations: (1) delta-gamma **overstates** the change here (by 17–25%), the opposite direction to Example 9, so Taylor errors cannot be assumed conservative in advance; (2) spot moves only 2 points from 71 to 69, yet crossing the barrier makes the value jump by about 2.9, a discontinuity no Taylor expansion can express; (3) higher volatility raises the knock-in probability and adds another 3.7. Investors in reverse convertibles or autocallables are effectively **short** such options and bear the opposite P&L.

Spot	Full revaluation	Delta-Gamma
85	+8.03	+10.01
75	+19.80	+23.42
71 (not knocked in)	+25.80	+30.30
69 (knocked in, now a vanilla put)	+28.66	—
75, vol to 30%	+23.46	—

In practice • Making full revaluation tractable / • Full revaluation

- **Spot-vol grids:** pre-compute full revaluations of each product on a spot × vol grid (e.g. spot -50% to +50%, vol -10 to +30 points) and interpolate per scenario. Standard in risk systems, and suited to the hundreds or thousands of HS or Monte Carlo scenarios.
- **Tiering:** sensitivities for linear positions, closed-form revaluation for vanillas, grids or Monte Carlo for path-dependent products.
- **Surrogate models:** for very complex products, regression or machine-learning pricing proxies, validated regularly against the full model.
- **Compute:** parallelise, cache intermediate results, run overnight batches; stress tests can use more paths than daily VaR.
- **Structured products held by funds:** do not rely only on issuer quotes; use an independent model or third-party valuation, documented in the valuation policy. PRIIPs Category 3 VaR and performance scenarios require simulation anyway.
- **Regulatory mapping:** UCITS VaR models must capture option gamma and vega (CESR/10-788); AIFMD requires independent model validation (Level 2 Art. 45) consistent with valuation policy (Art. 19); FRTB PLA effectively requires full revaluation of non-linear positions.

10.4 Conditional scenarios: propagating core shocks

$$\mathbb{E}[\mathbf{x}_s | \mathbf{x}_c] = \Sigma_{sc} \Sigma_{cc}^{-1} \mathbf{x}_c \quad (\text{jointly Normal factors, zero mean})$$

Real portfolios may carry thousands of risk factors, but a scenario usually specifies only a few **core factors** (equity index, 10-year rate, credit spread, USD). The remaining **satellite factors** (single stocks, other tenors, other currencies) are derived by conditional expectation from the covariance matrix: given the core move, the satellites move "on average" by this much. This keeps the scenario internally consistent and is standard for supervisory and bank internal scenarios.

Caution: correlations shift in crises (in 2022 equities and bonds fell together as their correlation turned positive). Propagating with a calm-period Σ understates joint declines, so a stressed covariance matrix, or explicitly set key correlations, is common.

10.5 Plausibility and reverse stress testing

How plausible is a scenario? The **Mahalanobis distance** measures how far it lies from "normal" in standard-deviation units, accounting for correlation. Reverse stress testing asks: among all scenarios producing loss L , which is most plausible (smallest Mahalanobis distance)? For a linear portfolio $\Delta P = \mathbf{w}'\mathbf{x}$, a Lagrange multiplier gives a closed form:

$$d(\mathbf{x}) = \sqrt{\mathbf{x}'\Sigma^{-1}\mathbf{x}}, \quad \min_{\mathbf{x}} \mathbf{x}'\Sigma^{-1}\mathbf{x} \text{ s.t. } \mathbf{w}'\mathbf{x} = -L \implies \mathbf{x}^* = -L \frac{\Sigma\mathbf{w}}{\mathbf{w}'\Sigma\mathbf{w}}, \quad d(\mathbf{x}^*) = \frac{L}{\sigma_P}$$

Deep dive • The intuition behind reverse stress testing / •

$\mathbf{x}^* \propto \Sigma\mathbf{w}$: the most plausible loss scenario moves each factor in the direction of its contribution to portfolio risk; volatile factors correlated with the portfolio take more of the shock. The distance needed equals L/σ_P , "how many portfolio standard deviations this loss is". This links reverse stress testing to VaR: under Normality, the loss at $d = 2.326$ is exactly the 99% VaR.

But a Normal Mahalanobis distance **understates** how plausible extreme scenarios are: fat tails and crisis correlation shifts make an "8-sigma" event far more common than the Normal suggests (§11.3). Use it to **rank** scenarios, not to convert them into probabilities.

10.6 Liquidity stress testing

$$\text{days to liquidate}_i = \frac{Q_i}{\theta \cdot \text{ADV}_i}, \quad \text{cost}_i \approx \frac{1}{2} \text{spread}_i + k \sigma_i \sqrt{\frac{Q_i}{\text{ADV}_i}}, \quad \text{coverage}(h) = \frac{\text{assets liquidable within } h}{\text{stressed redemptions within } h}$$

- **Asset side:** days to liquidate each position at a participation rate θ (typically 10–25% of average daily volume) and the cost incurred; the second term is the common square-root market-impact model. Under stress, widen spreads and cut ADV together.
- **Liability side:** redemption scenarios from the worst historical outflow, investor concentration (top 5 investors redeem), peer outflows in past crises, plus derivative margin calls and repo roll failure.
- **Coverage** ≥ 1 means assets liquidable within h cover stressed redemptions.
- **Vertical slicing vs waterfall:** selling the most liquid assets first (waterfall) is cheapest but leaves remaining investors with a less liquid book, a first-mover advantage that invites more redemptions. Pro-rata selling (slicing) preserves the profile at higher cost. Report both.
- **Liquidity management tools (LMTs):** gates, extended notice, redemption fees, swing pricing, anti-dilution levies, redemptions in kind, suspension, side pockets. Swing pricing and anti-dilution levies push trading costs onto redeemers and structurally weaken the first-mover advantage.

10.7 Limits of stress testing

Stress tests carry no probability and cannot "prove" safety; scenarios are human-designed and tend to test only the risks people already think of; severity calibration is subjective. Good practice: a scenario library mixing historical and hypothetical cases, periodic independent challenge, reverse stress tests to find blind spots, and a genuine "use test" where results drive decisions rather than just reports.

11 Stress Testing: Practice and Worked Examples

11.1 Design process

1. **Identify vulnerabilities:** concentration, leverage, non-linearity, liquidity mismatch, counterparties, basis.
2. **Map factors:** map every position to core and satellite risk factors (§10.4).
3. **Scenario library:** sensitivity grid + 3–5 historical + 2–3 hypothetical + reverse + liquidity scenarios; review the library annually.
4. **Compute P&L:** sensitivities for linear positions, full revaluation for options and convexity (§10.3).
5. **Aggregate and limit:** compare with stress-loss limits, VaR limits and liquidity coverage; set amber warnings (e.g. at 80% of a limit).
6. **Escalate and act:** results go to the risk committee, with pre-agreed actions on breach (hedge, deleverage, activate LMTs).
7. **Document and validate:** record assumptions, data, code, results and decisions; independent periodic validation.

11.2 Worked example: stress testing a multi-asset UCITS

Take a €100m multi-asset UCITS: EUR equities €30m, US equities €20m (USD unhedged), EUR government bonds €30m (modified duration 7, convexity 60), IG credit €15m (duration 5, convexity 30), cash €5m. Shocks are instantaneous and **illustrative**, sized to the order of magnitude of the relevant historical episodes.

Scenario	Equities	Rates	Spreads	USD	Equity P&L	Government bonds	Credit	Total €m / % NAV
Sensitivity: equity -10%	-10%	0	0	0	-5.00	0.00	0.00	-5.00
Sensitivity: rates +100bp	0	+100	0	0	0.00	-2.01	-0.73	-2.74
Sensitivity: spreads +100bp	0	0	+100	0	0.00	0.00	-0.73	-0.73
Sensitivity: USD -10%	0	0	0	-10%	-2.00	0.00	0.00	-2.00
Historical: 2008 GFC-type	-40%	-150	+300	+15%	-18.20	+3.35	-1.07	-15.92
Historical: 2020 COVID-type	-34%	-50	+150	+3%	-16.60	+1.07	-0.73	-16.26
Historical: 2022 inflation-type	-20%	+250	+100	+12%	-8.08	-4.69	-2.35	-15.12
Hypothetical: stagflation	-25%	+150	+200	0	-12.50	-2.95	-2.35	-17.80

Reading the results:

- Under the covariance assumed in §11.3, the portfolio's 1-month 99% VaR is about 5.4% of NAV, while the scenarios lose 15–18%, roughly three times VaR: the tail that VaR does not see.
- In the 2008 and 2020-type scenarios, government bond gains and USD appreciation partly offset equity losses; in the 2022-type and stagflation scenarios **bonds and equities fall together** and the hedge fails. Stagflation produces the largest loss and is absent from the historical library, which shows the value of hypothetical scenarios.
- Sensitivities add linearly, but historical-scenario losses are **not** the sum of sensitivities: they include convexity and the equity $\times$ FX cross term.

11.3 Worked example: reverse stress test

Assume 1-month factor volatilities: equity 5%, rates 25bp, spreads 15bp, USD 3%; correlations: equity–rates +0.3, equity–spreads –0.6, equity–USD –0.3, rates–spreads –0.2, spreads–USD +0.2. Portfolio 1-month s.d. $\sigma_P = €2.33\text{m}$ (2.33% of NAV).

Target loss $L = €20\text{m}$ (20% of NAV, the UCITS absolute VaR cap). From §10.5 the most plausible scenario is **equity –40.0%, rates –5.5bp, spreads +74bp, USD +1.9%**, with Mahalanobis distance $d = 20/2.33 = 8.6$. For $L = €10\text{m}$ the shocks halve and $d = 4.3$.

Lesson: under Normality an 8.6-sigma monthly event has probability around 5×10^{-18} , practically impossible; yet the 2008-type historical scenario in §11.2 already loses 15.9%. The gap comes from fat tails, crisis correlation shifts and multi-period accumulation. Use Mahalanobis distance to **compare** scenarios; the final plausibility call must draw on history and expert judgement.

11.4 Worked example: liquidity stress test

Bucket	NAV	Days (normal)	Days (stressed, ×2)	Cost
Cash	5%	1	2	0.0%
Government bonds Govt bonds	30%	1	2	0.1%
Large-cap equity	45%	3	6	0.2%
IG credit	15%	10	20	1.0%
Small-cap equity	5%	20	40	3.0%

RedemptionScenario Redemption scenario	Redemption	Horizon	Liquidable in horizon	Coverage
Worst historical weekly outflow	10%	5 days	80%	8.0×
Stress: 20% in a week, liquidation time doubled	20%	5 days	35%	1.75×
Stress: 40% in a month, liquidation time doubled	40%	20 days	95%	2.38×

All three coverage ratios exceed 1, so daily dealing is appropriate. But in the "20% in a week" case, a waterfall sale of cash and government bonds lifts the share of illiquid assets (credit + small caps) held by remaining investors from 20% to 25%. Reverse liquidity test: with doubled liquidation times, about 35% of NAV can be met within a week; beyond that, gates or swing pricing are needed.

In practice • An example stress-limit framework / • Framework

- **Stress-loss limit:** worst historical/hypothetical loss $\leq 25\%$ of NAV; amber at 20%, escalated to the risk committee.
- **Single-factor limits:** e.g. loss at rates +100bp $\leq 5\%$ of NAV, at equity –10% $\leq 7\%$ of NAV.
- **Liquidity coverage:** 5-day stressed coverage ≥ 1.2 ; below 1.0 requires a pre-agreed LMT plan.
- **Action list:** hedge or de-risk proposal within 5 business days of an amber flag; immediate de-risking and board reporting on breach, notifying the regulator where required.

12 Regulation of Medium-Term Risk Testing

12.1 UCITS

- **Stress tests for VaR funds** (CESR/10-788): at least monthly, more often on material portfolio or market change; must cover risks VaR misses (correlation breakdown, liquidity, jumps); must consider leverage scenarios producing negative NAV; results must feed investment and risk decisions.
- **General risk management** (Directive 2010/43/EU Arts 38–40): the risk management process includes stress tests and scenario analyses "where appropriate", and liquidity stress tests under exceptional conditions. This applies to commitment-approach funds too.
- **Liquidity stress testing** (ESMA guidelines ESMA34-39-897, applicable from 30 September 2020): at least annually, quarterly or more often recommended; assets and liabilities; historical and hypothetical scenarios plus reverse stress testing; documented policy and independent validation; material risks and actions notified to the regulator; the depositary checks the process exists.
- **Liquidity management tools** (UCITS amendments from 16 April 2026; Ireland S.I. 182/2026): select at least two LMTs from the statutory list and disclose activation conditions in the prospectus.

12.2 AIFMD

- **Art. 15(3)(b)**: risks of each position and their overall portfolio effect must be identified, measured, managed and monitored on an ongoing basis, **including through appropriate stress-testing procedures**.
- **Art. 16(1)**: regular stress tests under normal and exceptional liquidity conditions, with investment strategy, liquidity profile and redemption policy kept consistent.
- **Level 2 Art. 45**: periodic stress tests and scenario analyses "where appropriate" covering adverse changes in market conditions; Art. 48 details liquidity management.
- **Annex IV**: results of market and liquidity stress tests are reported.
- The **ESMA LST guidelines** apply to AIFs too; from AIFMD II open-ended AIFs select at least two LMTs; leverage restrictions may follow under Art. 25.

12.3 MiFID II and investment-firm prudential rules

- **Organisation** (MiFID II Art. 16; Delegated Regulation 2017/565 Art. 23): policies and procedures to identify, manage and monitor firm risks, with an independent risk function.
- **Internal capital and liquidity** (IFD Art. 24): ongoing assessment of internal capital and liquid assets; supervisory review expects stress testing and orderly wind-down planning.
- **Trading book**: firms on CRR internal models currently compute stressed VaR; from FRTB in 2027 this is replaced by a 97.5% ES calibrated to a stress period.
- **Product governance**: manufacturers run scenario analysis of poor client outcomes, including deteriorating markets (Delegated Directive 2017/593 Art. 9(10); see §15.3).

Requirement	UCITS	AIFMD	MiFID II / IFD
Market stress	Monthly for VaR funds; "where appropriate" for all	Art. 15(3)(b); L2 Art. 45	Firm-level risk management; IFD assessment
LST	ESMA: ≥ annual, quarterly advised	Art. 16(1) + ESMA	IFD liquid-asset assessment
Reverse	Required in LST guidelines	Required in LST guidelines	Supervisory expectation
LMTs	≥ 2 (from April 2026)	≥ 2 for open-ended	—
Reporting	RMP; notify material LST issues	Annex IV stress results	SREP / supervisory reporting

Long-Term Risk Testing · one year and beyond

Long-term testing asks "how deep can losses go, how long can they last, and how long will recovery take?". It is about the **path** rather than the distribution at one date: drawdown measures the worst loss investors actually live through, and scenario analysis turns macro narratives into multi-year portfolio paths.

13 Drawdown: Theory and Practice

13.1 Definitions and common metrics

$$M_t = \max_{s \leq t} V_s, \quad DD_t = 1 - \frac{V_t}{M_t}, \quad MDD_T = \max_{t \leq T} DD_t, \quad \text{Calmar} = \frac{\text{CAGR}}{\text{MDD}}, \quad \text{Ulcer} = \sqrt{\frac{1}{n} \sum_t DD_t^2}$$

Metric	Meaning
Drawdown DD_t	M_t · Distance of NAV below its running peak (high-water mark)
MDD	· Largest drawdown over the period
Decline / recovery / time under water	· Peak to trough, trough to new peak, and their sum
Calmar	÷ " " · Annual return per unit of maximum pain
Ulcer	· RMS of drawdowns, penalising depth and duration
CDaR $_{\alpha}$	$(1 - \alpha)$ ES Chekhlov 2005 · Average of the worst $(1 - \alpha)$ drawdowns, the drawdown analogue of ES

Deep dive · The asymmetry of drawdowns / ·

After a loss d , a gain of $d/(1 - d)$ is needed to break even: -10% needs +11.1%, -20% needs +25%, -33% needs +50%, -50% needs +100%. Required gain and recovery time grow non-linearly with depth. That is why long-term investors care more about drawdown than volatility: volatility treats ups and downs alike, drawdown records only the pain.

13.2 Theory: path dependence and horizon

Path versus end point. VaR looks at the distribution of the terminal return; drawdown looks at the lowest point along the way. For driftless Brownian motion the **reflection principle** gives:

$$\Pr \left(\min_{0 \leq s \leq T} X_s \leq -a \right) = 2 \Pr (X_T \leq -a), \quad \mathbb{E}[MDD_T] \approx \sqrt{\frac{\pi}{2}} \sigma \sqrt{T} \approx 1.25 \sigma \sqrt{T} \quad (\text{log scale, zero drift})$$

Meaning: the probability of **touching** the 99% one-year VaR level at some point during the year is about **twice** the probability of ending below it (2% rather than 1%). Simulation check: with daily discrete paths, 1.01% end below it and 1.80% touch it (discrete sampling sits slightly below the theoretical 2%).

Expected maximum drawdown is proportional to $\sigma\sqrt{T}$, the same root as VaR's $\sqrt{h}$ rule but with a larger constant (about 1.25, where VaR uses a quantile). With positive drift μ drawdown growth slows and eventually grows like $(\sigma^2/\mu) \ln T$ (Magdon-Ismail et al. 2004).




Figure 9. Monte Carlo: maximum drawdown (mean and 95th percentile) by horizon for an asset with 16% annual volatility.

Horizon	Theory, zero drift	Mean, 0% drift	95th pct, 0%	Mean, 6% drift	95th pct, 6%
1 year	18.2%	17.4%	29.9%	15.5%	27.0%
3 years	29.3%	29.0%	47.3%	24.1%	39.9%
5 years	36.1%	36.1%	56.7%	29.1%	47.1%
10Y	47.0%	47.3%	70.0%	36.2%	55.8%

In practice • "Is this drawdown unusual?" / • "

Drawdowns are often read as evidence that "something is wrong" with a manager. The table gives a baseline: a strategy with 16% volatility and 6% expected return has a median 3-year maximum drawdown around 23% and a 5% chance of exceeding 40%. A 25% drawdown alone does not prove the strategy has failed; compare the realised drawdown with the simulated distribution for the same volatility and horizon, then look at its causes.

13.3 Worked example: S&P 500 (2013–2018)




Figure 10. Top: index and high-water mark (dashed), with the maximum drawdown's time under water shaded; bottom: drawdown curve.

-14.2% · MDD Day	190 / 107 / Day · decline / recovery days	0.92 Calmar CAGR 13.1%	-8.9% monthly-sampled MDD ·
----------------------------	---	----------------------------------	---------------------------------------

The maximum drawdown starts at the 2015-05-21 peak, runs through the August 2015 renminbi shock and the early-2016 oil slump, bottoms on 2016-02-11 (-14.2%) and regains the peak on 2016-07-11, about 297 business days under water. **Sampling frequency matters:** sampled at month-ends the same data show an MDD of only -8.9%, because month-end prices miss intra-month lows. Always state the frequency and use total-return NAV net of fees.

13.4 Drawdown limits and risk control

- **Tiered triggers:** a common pattern is review at -8%, halve risk at -12%, stop adding risk and escalate to the board at -15%. Calibrate triggers against the simulated distribution in §13.2 so they do not fire on normal noise.
- **Drawdown-constrained allocation** (Grossman & Zhou 1993): if NAV must never fall below $(1 - \delta)$ of the high-water mark, the optimal risky exposure is proportional to the cushion $V_t - (1 - \delta)M_t$; less risk the closer to the floor. CPPI is a simplified version.
- **Costs:** stop-losses lock in losses, whipsaw in choppy markets, and sell together in broad declines, which is procyclical.
- **High-water mark and performance fees:** the HWM is the M_t in the drawdown definition; performance fees are only payable after past drawdowns are recovered (§15.1).

	VaR / ES	MDD
Perspective	· Single-period terminal distribution	· The whole path
Horizon	1 Day 10 Day 20 days · Fixed	Horizon · Grows with the observation period
Aggregation	ES · Decomposable (ES subadditive)	· Not simply decomposable
Estimation	· Few tail points	MDD · One MDD per path, noisier still
Use	Day · Daily limits, capital, regulatory caps	Framework · Investor experience, strategy review, stop-loss

14 Scenario Analysis: Theory and Practice

14.1 Scenario analysis versus stress testing

A stress test is usually an **instantaneous, single-period** shock: "what if X happens tomorrow?". Scenario analysis is a **multi-period, narrative-driven** path: "what if the world evolves like this over the next few years?". Growth, inflation, rates, equities, spreads and FX must be consistent year by year, with rebalancing, cash flows, fees and strategy responses included. The output is not one loss figure but a set of paths: terminal value, annualised return, maximum drawdown along the way, recovery time, and the chance of missing the investment objective.

14.2 Building a scenario

1. **Narrative:** one paragraph on cause, transmission and resolution (e.g. "a supply shock triggers stagflation, central banks hike, inflation fades after two years").
2. **Core variable paths:** year-by-year GDP, inflation, policy rate, 10-year yield, equity index, spreads, FX.
3. **Map to portfolio factors:** via regressions or the conditional expectation of §10.4.
4. **Project the portfolio:** period-by-period P&L with rebalancing, flows, fees and pre-agreed management actions.
5. **Outputs and decisions:** terminal value, annualised return, maximum drawdown, recovery time, probability of meeting objectives; discuss allocation or product-design changes.

14.3 Types of scenario

Type	Description
Deterministic narrative	3–5 Scenario · 3–5 distinct scenarios (base, adverse, severe, structural shift), easy to communicate
Historical analogue	1970 2000–2002 2008–2009 · Apply multi-year paths of past episodes to today's portfolio
Stochastic (Monte Carlo)	or 9 · GBM, bootstrap or regime-switching paths giving distributions of terminal value and drawdown
Climate	NGFS None " "Scenario Horizon 10–30Y · NGFS orderly, disorderly and hot-house scenarios, separating transition risk (carbon prices, policy) from physical risk; horizons of 10–30 years

14.4 Worked example: five-year scenarios for a 60/40 portfolio

Portfolio: 60% equities, 40% government bonds (duration 7, 3% carry), rebalanced annually. Four **illustrative** scenarios set yearly equity returns and yield changes; bond return $\approx 3\% - 7 \times \Delta y$.

Scenario	Equities Equity by year	Rates Yield change (bp)	5 years Terminal	p.a.	MDD	Recovered
Base	+6% × 5	0	126.4	4.8%	0.0%	—
Stagflation	−15, −5, +5, +8, +8	+150, +50, 0, −25, −25	101.1	0.2%	−14.8%	Year 5
Recession & recovery	−30, +20, +15, +8, +7	−150, +50, +25, 0, 0	119.5	3.6%	−12.6%	Year 3
None Disorderly transition	−10, −5, +4, +6, +6	+50, +25, 0, 0, 0	104.1	0.8%	−8.5%	Year 5




Figure 11. Five-year paths of a 60/40 portfolio under four scenarios (annual, illustrative).

Reading: the recession's first-year fall (equities -30%) is larger than stagflation's, but the bond rally cushions it and the portfolio recovers within three years; in stagflation bonds and equities fall together, the portfolio keeps sliding after year one and only breaks even in year five, at 0.2% a year. For a product targeting 4% a year, both stagflation and disorderly transition miss the objective: exactly the "poor outcomes" MiFID product governance asks firms to identify. Annual data also understate intra-year drawdowns (the sampling point in §13.3), so true maximum drawdowns would be deeper.

14.5 PRIIPs performance scenarios

A PRIIPs KID shows retail investors four performance scenarios, **stress, unfavourable, moderate and favourable**, at the recommended holding period (RHP) and at one year when the RHP is longer. Since 2023 (Delegated Regulation 2021/2268) the Category 2 method is:

- **Unfavourable / moderate / favourable:** over at least 10 years of history, take the outcomes of all rolling RHP-length windows and select the worst, median and best;
- **Stress:** compute rolling volatility over 21-day ($RHP \leq 1$ year) or 63-day ($RHP > 1$ year) sub-intervals, take the 99th or 90th percentile as stressed volatility, then a Cornish–Fisher 1% or 5% quantile.

This is "regulator-standardised scenario analysis", entirely historical, so it looks optimistic after long bull markets. Refer to the regulation for full computational detail.

14.6 Limits and governance

Scenarios are human-designed, prone to recency bias and hard to attach probabilities to; multi-period scenarios must also assume management actions and second-round effects (e.g. redemption-driven selling). Governance: review at least annually; build jointly across risk, investment and economics; record assumptions and sources; use results in allocation, product design and board discussion rather than an appendix.

15 Regulation of Long-Term Risk Testing

15.1 UCITS

- **No drawdown limit:** UCITS sets none, but the risk management process must keep the risk profile consistent with objectives and strategy.
- **Performance fees and HWM** (ESMA guidelines, from January 2021): fees only on positive performance over the reference period; where an HWM model's reference period is shorter than the fund's life it must be at least a rolling 5 years; past losses or underperformance must be recovered first; crystallisation generally no more than once a year.
- **PRIIPs KID:** performance scenarios (§14.5); past-performance chart (up to 10 years); previous performance scenarios updated monthly and published.
- **Sustainability risk** (Delegated Directive 2021/1270, from August 2022): integrated into the risk management process; climate scenario analysis is common practice.
- **Marketing** (ESMA marketing guidelines): past performance must be fair and clear, with risk warnings.

15.2 AIFMD

- **Consistent risk profile** (Art. 15(3)(c)): the risk profile must match the AIF's size, portfolio structure, strategy and objectives; long-horizon scenarios help evidence this.
- **Annex IV**: monthly gross and net returns (from which regulators can compute drawdowns), plus stress-test results.
- **Investor disclosure** (Art. 23): historical performance where available; PRIIPs performance scenarios if sold to retail.
- **Sustainability risk** (Delegated Regulation 2021/1255, from August 2022): integrated into risk management and due diligence.
- **AIFMD II loan-originating funds**: leverage caps and 5% retention address long-term credit risk; closed-ended structures should match loan maturities.
- The ESMA performance-fee guidelines also apply to certain retail-facing AIFs.

15.3 MiFID II

- **Product-governance scenario analysis** (Delegated Directive 2017/593 Art. 9(10)): manufacturers must run scenario analysis assessing the risk of **poor outcomes** for end clients and when they may arise, including deteriorating markets, distress at the manufacturer or a third party, the product ceasing to be commercially viable, and demand far above expectations. Art. 9(14) requires regular product review.
- **Suitability**: investment horizon and capacity for loss; periodic suitability reports; sustainability preferences (Delegated Regulation 2021/1253).
- **10% loss reporting** (2017/565 Art. 62): in effect a **drawdown trigger** measured from the last report.
- **Performance in marketing** (2017/565 Art. 44): past performance covers the preceding 5 years in complete 12-month periods; forward-looking figures must rest on reasonable, objectively supported assumptions and use **performance scenarios in different market conditions, both negative and positive**.
- **Firm level** (IFD Art. 24): the internal capital assessment typically spans a multi-year planning horizon, including orderly wind-down.

Requirement	UCITS	AIFMD	MiFID II
Drawdown-related	Performance-fee HWM (≥ 5-year reference)	Annex IV monthly returns; retail AIF fee rules	10% loss reporting
Scenario Scenarios	PRIIPs performance scenarios	Risk-profile consistency; PRIIPs if retail	Product governance Art. 9(10); marketing scenarios
/ Sustainability	2021/1270	2021/1255	2021/1253 (suitability)
Frequency	KID reviewed at least annually; monthly previous scenarios	Per risk policy	Regular product review

16 Integrated Framework, Checklist and Errata

16.1 Linking the three horizons

The three toolsets do not replace one another; they layer: **VaR/ES** covers normal days, **stress testing** severe shocks, **drawdown and scenario analysis** multi-year bad paths. A complete framework shares one risk-factor mapping and dataset across all three, with limits calibrated to each other:

- high VaR-limit usage usually comes with high stress losses; joint warnings mean more than single ones;
- the worst stress loss should sit inside the drawdown triggers, otherwise one shock trips the stop-loss directly;
- poor outcomes found in long-term scenarios should feed back into product design (target market, SRI, choice of LMTs) and asset allocation.

Frequency	Short-term	Medium-term	Long-term	Main rule
Day Daily	VaR/ES, leverage, limit usage	Sensitivity snapshot (optional)	Drawdown monitoring	UCITS daily VaR; MiFID 10% report
Monthly	Backtesting	Market stress tests (VaR funds)	—	CESR/10-788
Quarterly	Parameter and data refresh	Liquidity stress tests (advised)	Scenario review	ESMA LST; Annex IV (size-dependent)
Annually	Model validation	Scenario library and LMT review	Multi-year scenarios, product review, KID update	AIFMD L2 Art. 45; MiFID product governance; PRIIPs

16.2 Decision checklist

#	Decision	Common choices	Document
1	Return definition	Log (equity, commodity); absolute changes (rates, spreads, vols)	Shock type per factor class
2	Horizon Confidence and horizon	99%/1D, 99%/10D, 97.5% ES; $\sqrt{10}$ or direct 10-day	Test of the scaling assumption
3	Window	σ : 21–66 or EWMA; HS: 250–500 + stress; S,K: ≥ 250	Data span, gaps, holidays
4	Mean	0 · usually zero	If kept, how it is estimated
5	EWMA λ	0.94 (daily), 0.97; seed and burn-in	Indexing convention
6	Quantile convention	Interpolation / [Np]-th worst / smoothing	Difference vs Excel/Python defaults
7	Revaluation	Full revaluation; delta–gamma; sensitivities $\times$ shocks	Approximation error and PLA
8	Backtesting	1-day 99% and 97.5%; APL and HPL; Kupiec + Christoffersen + traffic light	Eligible comparisons; sign convention
9	Procyclicality	Long-run floor, buffers, partial scaling	Peak-to-trough, n-day increase
10	Ongoing monitoring	Recalibration frequency, residual diagnostics, change control	Triggers for model review
11	Scenario Stress library	Sensitivity + historical + hypothetical + reverse + liquidity; annual review	Sources, severity rationale, correlation assumptions
12	Drawdown	Total-return NAV, daily frequency, tiered triggers	Trigger calibration vs simulated distribution
13	Scenario Long-term scenarios	3–5 narratives + Monte Carlo; climate scenarios	Assumptions, management actions, comparison with objectives

16.3 Errata and clarifications on the source materials

Location	As written	/ Correction / note
Slides: traffic zones	$p_Y = 0.9599$	0.9999 Should be 0.9999 (the notes have it right).
Probability Slides: breach table	"Cum Binomial ... (T=252, 0.01)"	Value T=252 p=5% Probability Cumulative 99% VaR CumulativeProbability 5 Values are the PMF for T=252, p=5%, not cumulative. See Figure 5 for 99%.
PRIIPs Notes, whitepaper: PRIIPs	$-0.146 S^2$	$+0.146 S^2$ S K $1/\sqrt{N}$ $1/N$ §6.2 Should be $+0.146 S^2$, with S, K scaled by $1/\sqrt{N}$, $1/N$.
Notes: diagnostic ratio	Gaussian ρ "fixed at every c"	$\rho \rightarrow 1$ 1.254 $\rightarrow$ 1.146 $\rightarrow$ 1.090 Moderate but slowly decreasing to 1.
Solutions: two-asset	—	None 1 VaR £10,420 Correct; with $p=1$ VaR is £10,420, illustrating diversification.
8 Figure 8 chart	Green ≤ 5 , Yellow ≤ 10	Green 0–4 5–9 Red ≥ 10 Should be 0–4 / 5–9 / ≥ 10 .

17 AIFMD Annex IV Reporting: A Step-by-Step Process

Annex IV is the periodic supervisory report an AIFM files with its regulator under AIFMD Articles 3(3)(d) and 24, with the template set by Annex IV of Delegated Regulation 231/2013 plus ESMA's reporting guidelines and XML schema (XSD). It has an **AIFM report** (manager level) and an **AIF report** (one per fund, and per compartment of an umbrella). The AIF report pulls together almost every output in this handbook: VaR and sensitivities (Part I), stress tests and liquidity (Part II), monthly returns (Part III). It therefore closes the book as the regulatory landing point of the three-horizon framework.

17.1 Process overview and timeline

When	Step	Owner	Output
Annually / pre-period	Scope and frequency	· Compliance	Frequency map, NCA list, portal access
Pre-period	Governance and master data	· Compliance, operations	RACI, identifiers, strategy classification
T+0 – T+5	Positions and valuation	· Administrator	Period-end positions, NAV, trades, flows
T+5 – T+15	Exposures and risk metrics	· Risk	Market, counterparty, liquidity, leverage, stress
T+15 – T+22	Validation	· Risk, compliance	Checks, variance explanations
T+22 – T+28	XML, approval, submission	· Compliance	Signed-off filing
From T+30	Post-submission	· Compliance	Amendments, records, lessons

T = reference date (last business day of March, June, September, December); the deadline is one month later (last calendar day of the following month in Ireland), with 15 extra days for fund-of-funds AIF reports.

17.2 Step 1: Determine scope and frequency

(a) Compute regulatory AuM (Level 2 Arts 2–5): unlike NAV, convert derivatives to underlying equivalents under the commitment method, include assets acquired through leverage, aggregate all AIFs managed, and exclude UCITS the AIFM manages as a management company. Compute at least annually and monitor threshold crossings continuously.

(b) Set the frequency using the table below. For unleveraged AIFMs whose AIFs have 5-year lock-ups, the half-yearly band starts at €500m rather than €100m.

(c) Set the content layer: Art. 24(1) for all AIFs (main instruments, markets, principal exposures and concentrations); Art. 24(2) for authorised AIFMs (risk profile, liquidity, leverage, stress tests); Art. 24(4) for substantial leverage, commitment above 3× NAV (sources of leverage and more).

(d) Special cases: after authorisation, report from the next quarter; switch frequency on crossing a threshold; file a "nil return" (no-reporting flag set) even when there is nothing to report; non-EU AIFMs marketing under national private placement regimes file separately with each host regulator.

Condition	Frequency
Registered (sub-threshold) AIFM	Annual, Art. 24(1) content only
AIFM AuM €100m–€1bn (€500m–€1bn if unleveraged with 5-year lock-up)	· Half-yearly
AIFM AuM > €1bn	· Quarterly
Single AIF > €500m	AIF · Quarterly for that AIF
Unleveraged AIF acquiring control of non-listed companies	AIF · Annual for that AIF

17.3 Step 2: Governance and master data

- **RACI:** compliance owns the process and filing; the administrator provides positions, NAV and flows; risk computes metrics and stress tests; the investment team confirms strategy classification; the board or delegate approves. A written procedure sets cut-offs and escalation.
- **Identifiers:** AIFM and AIF national codes (the Central Bank C-code in Ireland), LEIs, share-class ISINs, ECB codes where relevant, master–feeder links, LEI/BIC of prime brokers and counterparties.
- **Static attributes:** base currency, inception date, domicile, marketing jurisdictions, predominant AIF type (hedge fund, private equity, real estate, fund of funds, other) and strategy breakdown as % of NAV (summing to 100%).
- **Portal readiness:** portal accounts, user rights and the "uploader ≠ signer" four-eyes rule; track the XSD version and validation-rule updates published by the regulator.

17.4 Step 3: Positions and valuation data

- Obtain final reference-date positions and NAV from the administrator and reconcile to the published NAV.
- Map each position to ESMA **sub-asset type codes**, market (MIC), region, currency and counterparty.
- Aggregate: top 10 principal exposures, top 5 portfolio concentrations, principal markets, turnover by sub-asset class over the period, long/short exposure by currency, holdings with dominant influence over non-listed companies (private equity).
- Collect the **monthly series** for the period: gross and net returns, NAV changes, subscriptions and redemptions, which are the inputs for the drawdown analysis in Part III.

17.5 Step 4: Risk metrics (Arts 24(2) and 24(4))

Annex IV Section	What is reported	In this handbook
Market risk	Net equity/FX/commodity delta; DV01 and CS01 by maturity bucket; vega; VaR and its method	Chs 2–6; §10.3 sensitivities
Counterparty	Top-5 counterparty exposures (both directions), collateral, share cleared via CCP vs bilateral	§8.2
Portfolio liquidity	% of NAV liquidable within 1 day, 2–7, 8–30, 31–90, 91–180, 181–365, >365 days; unencumbered cash	§10.6 days-to-liquidate; §11.4
Investor liquidity	Redeemable % in the same buckets; dealing frequency and notice; lock-ups; gates, side pockets; top-5 investor concentration	§10.6 liabilities; §12.1 LMTs
Financing liquidity	Borrowing by remaining term bucket	§10.6
Operational / historical	Open positions; monthly gross/net returns, NAV changes, flows	Ch. 13 drawdown
Result Stress tests	Results of market and liquidity stress tests (free text, Arts 15(3)(b) and 16(1))	Chs 10–11
Leverage	Gross and commitment leverage; cash and securities borrowing; derivative exposure; re-hypothecation; (24(4)) top-5 sources of leverage	§8.2

In practice • Writing the stress-test section / • Result

The stress-test field is free text, and regulators most often criticise boilerplate. Use a fixed structure: (1) scenario and shock sizes; (2) NAV impact (%); (3) liquidity coverage; (4) comparison with internal limits; (5) actions taken or planned. For example: "2022-type inflation scenario (equity –20%, rates +250bp, spreads +100bp, USD +12%): NAV –15.1%, within the 25% stress limit; 5-day stressed coverage 1.75×; no action required." Keep it consistent with the ESMA LST policy and the risk management policy.

17.6 Step 5: Validation

- Internal reasonableness:** reconcile NAV with the administrator; strategy shares and liquidity buckets each sum to 100%; monthly returns agree with NAV changes and flows; portfolio liquidity agrees with dealing terms and LST conclusions; leverage methods match the prospectus; explain large period-on-period moves (e.g. ±20%).
- ESMA XSD:** structure and format.
- ESMA business rules:** conditional fields, enumerations, cross-field consistency (separate rule sets for AIF and AIFM files).
- Regulator-specific rules:** e.g. Ireland cross-checks against its fund register, and an unreported AIF blocks completion.
- Four-eyes review and sign-off:** separate risk and compliance review, with records kept.

17.7 Step 6: Generate XML and submit (Ireland as example)

- Generate AIFM and AIF files under the ESMA XSD (Ireland allows several AIFs in one file). The AIFM report is in EUR at the reference-date ECB rate; the AIF report is in the fund's base currency.
- File naming: C999100_YYYYMMDD_AIFM.xml (AIFM) or C999100_YYYYMMDD_AIF.xml (AIF), where C999100 is the institution's C-code.
- Log in to the Central Bank of Ireland Portal: Returns → View/Edit Returns → Add Return → choose the reference date (and C-code for AIF returns).
- Load a File; the portal validates automatically and emails the result; fix errors and re-upload.
- A **different user** finalises and signs off. An **unsigned return is not filed**; once signed off it is locked.
- Use INIT for the first filing and AMND for corrections, resending the whole record rather than changed fields only.

17.8 Step 7: After submission

- Keep the acknowledgement, XML, source data, workpapers and approvals for regulatory queries and audit.
- On finding an error, ask the regulator to unlock and file a correction; late or incorrect returns may lead to supervisory action.
- If a fund of funds uses the 15-day extension, the AIF report follows the AIFM report; confirm consistency and amend the AIFM report if needed.
- Run a post-mortem each period: delays, validation failures, data-quality issues; update procedures and controls.
- Changes in manager–fund relationships (launch, liquidation, change of AIFM) must be notified separately as the regulator requires.

Pitfall • Frequent Annex IV errors / • Annex IV

- Using NAV instead of regulatory AuM, so the frequency is wrong;
- AIFM report not converted to EUR, or the wrong FX date;
- Enumerations outside ESMA's list, or confusing "registration status" with the "AIF reporting code";
- Conditional fields out of step: a flag set to "no" with sub-fields populated (or the reverse);
- Reusing INIT, or sending only changed fields on correction;
- Portfolio liquidity contradicting dealing terms or LST conclusions (e.g. a daily-dealing fund reporting most assets as needing 90+ days);
- A stress-test field that says only "stress tests performed, results acceptable";
- Missing nil returns, or omitting a compartment of an umbrella.

17.9 Outlook: the new reporting framework under AIFMD II

AIFMD II mandates ESMA to redesign supervisory reporting and brings UCITS into comparable periodic reporting for the first time. ESMA's direction: one **modular template** integrating AIFMD, UCITS, MMF and ECB statistical reporting; **ISO 20022 XML** with a common data dictionary (LEI, ISIN, CFI); **security-by-security** reporting where useful; a single collection point that "reports once, uses many times". Technical standards are still being drafted, and Level 2 measures are expected only after October 2027. Preparation now: map data overlaps across reports, assess the IT impact of security-level reporting, confirm the national collection authority, and build the risk data from all three parts of this handbook into a traceable data store.

18 Risk Management in Practice for Derivative-Heavy Funds

Alternative UCITS, macro and CTA strategies, volatility strategies, option overlays and hedge-fund AIFs rely mainly on derivatives: futures, interest rate swaps, CDS, total return swaps, options, variance swaps and FX forwards. Notionals can run to several times NAV, risk is highly non-linear, and the fund depends on daily margin and collateral liquidity. This chapter strings the earlier methods into an executable daily process, mapped item by item to UCITS, AIFMD and MiFID II requirements.

18.1 New-instrument approval: before the first trade

1. **Eligibility** (UCITS): the underlying must be eligible assets, financial indices, interest rates, FX rates or currencies (no direct commodities, only via eligible indices); OTC counterparties must be prudentially supervised institutions in categories approved by the regulator; OTC derivatives need reliable, verifiable daily valuation and must be closable at fair value at any time (UCITS Directive Art. 50(1)(g), Directive 2007/16/EC Art. 8).
2. **Risk management process (RMP)**: new derivative types must be written into the RMP; in Ireland a UCITS files its RMP with the Central Bank before using derivatives, and types not covered require an update first. AIFs update the risk policy and maximum-leverage disclosure.
3. **Legal documentation**: ISDA master and CSA, clearing agreements, prime brokerage agreements; confirm netting is enforceable in the counterparty's jurisdiction (legal opinions).
4. **Systems and models**: booking, pricing model, risk-factor mapping (§6.5), whether VaR captures gamma, vega and basis (§10.3), commitment conversion (§8.6), Annex IV coding (Ch. 17).
5. **Valuation source**: an independent price source or model, not counterparty quotes alone.
6. **Operations**: margin and collateral processes, settlement, EMIR trade reporting.
7. **Dry run and sign-off**: test the full chain with small or simulated trades; risk, compliance, operations and valuation sign off together.

18.2 The daily risk process

When	Step	Key points
T+0 morning	Position reconciliation	Reconcile fund system, prime brokers, counterparties and administrator; clear unconfirmed trades and breaks the same day
T+0	Market data and valuation	Closing prices, curves, vol surfaces; independent price verification; start CSA dispute resolution when valuation gaps exceed thresholds
T+0	Risk calculation	VaR/ES, Greeks ladders, leverage (sum of notionals, commitment, gross), counterparty exposure, margin forecast, liquidity coverage
T+0	Limit checks and escalation	Warning and hard limits; same-day escalation to PM and CRO, with actions and deadlines logged
By T+1	Margin and collateral	Issue and meet margin calls, move collateral, update cash forecast

18.3 Leverage on several bases

The same fund shows very different leverage numbers under different rules, so always state the basis (table below).

Example: a macro UCITS holds €3bn of interest-rate futures and swaps, €0.5bn of equity index futures and €0.5bn of FX forwards on €0.5bn NAV. Sum of notionals is 800% of NAV; because much of it is short-dated rates, duration-adjusted risk is modest and the 20-day 99% VaR may be only 8–10% of NAV. An 800% notional is not a breach, but it must match the expected leverage in the prospectus; persistently exceeding it calls for a disclosure review.

Basis	Rule	Use
Sum of notionals	UCITS VaR	Expected leverage in prospectus, actual leverage in annual report
Commitment (with netting and hedging)	UCITS AIFMD	$\leq 100\%$ NAV UCITS > 3 AIFMD
Gross (no netting or hedging)	AIFMD	Annex IV
Basis VaR	UCITS VaR	Absolute 20% / relative 2× limits

18.4 Non-linear risk and Greek limits

- **Greek limits**: set delta, gamma and vega limits by underlying and expiry bucket; bucket vega by tenor since short and long vols move differently.
- **Spot–vol ladder**: show full-revaluation P&L daily for spot $\pm 5\%$, $\pm 10\%$, $\pm 20\%$, $\pm 30\%$ combined with vol -5 to $+20$ points (§10.3.2); it exposes short-option tails far better than a single VaR.
- **Spotting negative gamma**: short options, autocallable hedges and short variance swaps earn steady returns with low VaR in calm markets but lose at an accelerating rate in large moves. Give negative-gamma books a dedicated stress-loss limit.
- **Expiry concentration and pin risk**: many options on one expiry with spot near the strike make delta jump before expiry and hedging costs unpredictable.
- **Jump and gap risk**: overnight or weekend gaps cannot be hedged dynamically; include gaps in scenarios.

18.5 Margin and collateral liquidity

Derivative-heavy funds most often fail not from losses as such but from **being unable to post margin in time**: variation margin (VM) is paid in cash daily, and initial margin (IM) rises with volatility at CCPs or under SIMM. In March 2020 and in the 2022 UK LDI crisis, many funds were solvent yet forced into fire sales because they could not raise cash fast enough.

Margin coverage = (cash + assets immediately usable as collateral) ÷ stressed VM and IM calls over 5 days. An internal floor of 1.2 is common; below 1, actions must be pre-arranged.

Worked example 13 • Margin stress test / 13 •

A €500m NAV fund: interest rate swaps with total DV01 of €400k per bp (receiving fixed), long equity index futures of €1bn notional, current initial margin €80m; cash and T-bills €180m. Stress: rates +100bp, equities –15%, IM up 30%. Cash needs in the table below.

Coverage = $1.8 \div 2.14 = 0.84 < 1$. The fund is not insolvent (VM losses cut NAV by about 38% but it stays positive), yet faces a cash gap of about €34m and must sell other assets under stress. Pre-emptive measures: a larger cash buffer, options instead of some futures to cap VM, several clearing brokers, pre-arranged repo lines, and a margin-coverage limit.

Item	Cash need
IRS VM: €0.4m × 100	€40m
Index futures VM: €1bn × 15%	€150m
IM increase: €80m × 30%	€24m
Total	€214m

18.6 Counterparties, collateral and legal terms

- **Limits**: UCITS ≤ 5% of assets per OTC counterparty (10% for credit institutions), net of eligible collateral (§8.7); AIFs set internal limits. Diversify prime brokers and counterparties so one default cannot trap positions.
- **CSA terms**: thresholds, minimum transfer amounts, eligible collateral and haircuts, rating triggers, dispute timelines; term differences directly change exposure and liquidity (see §18.9).
- **Counterparty-default scenario**: assume the largest counterparty defaults; compute replacement cost (re-establishing the same positions) plus losses on collateral liquidation.
- **Collateral reuse**: rehypothecation adds leverage and counterparty risk; include it in leverage and report under SFTR.

18.7 Stress-scenario checklist for derivative funds

Scenario	What it tests
Vol spike (+20 points) with spot fall	Short gamma and vega; IM increases
Overnight gap (–10%)	Failure of dynamic hedging
Correlation flip (bonds and equities fall together)	Relative-value and hedged books
Basis widening (futures–cash, CDS–bond, swap spreads)	Arbitrage and hedge assumptions
Margin spiral (§18.5)	Cash and collateral liquidity
Largest counterparty default	Replacement cost and collateral
Hedge-instrument illiquidity	Cost of being unable to close or rebalance
Reverse: what takes NAV down 20% or margin coverage below 1	Hidden concentrations (§10.5)

18.8 Regulatory mapping and control frequency

Control	Frequency	UCITS	AIFMD	MiFID II
VaR/ES and limit usage	Daily	Absolute 20% / relative 2× (§8.1)	Internal limits (Art. 15)	Mandate limits for discretionary accounts
Leverage bases	Daily	Commitment ≤ 100% or disclose notionals	Both methods + maximum; > 3× extra reporting	10% loss report for leveraged accounts
Counterparty and collateral	Daily	5% / 10%; ESMA collateral criteria	Due diligence, Annex IV top 5	Client asset safeguarding
Margin liquidity	Daily forecast, weekly stress	LST includes margin calls	Art. 16 liquidity	IFD liquid assets
Stress tests and Greek ladders	Daily ladder, monthly full	At least monthly for VaR funds	Art. 15(3)(b); Annex IV results	Product-governance scenarios
Backtesting	Monthly	> 4 must be reported (§7.3)	Level 2 45	—
Model and valuation validation	Annually and on material change	Independent validation; RMP update	Art. 19 valuation; L2 Art. 45	Appropriateness (derivatives are complex)
Disclosure and reporting	Annual / periodic	Prospectus expected leverage; annual VaR usage and leverage	Annex IV (Ch. 17)	EMIR reporting; cost disclosure

Pitfall • Derivative funds / •

- Watching VaR but not notionals and Greeks: a low VaR can hide the tail of heavy short optionality;
- Treating "not insolvent" as "safe": the binding constraint is often cash, not NAV;
- Relying on counterparty valuations: in a dispute you cannot independently judge exposure or margin;
- Hedge ratios failing under stress: basis and correlation shifts make "hedged" positions lose together;
- Trading new instruments before the process is ready: derivatives outside the RMP or not yet mapped leave VaR and leverage understated.

18.9 CSA requirements and collateral management

The **CSA** (Credit Support Annex) is an annex to the ISDA master agreement setting when, in what assets and at what valuation the parties exchange collateral. It directly drives counterparty exposure, margin cash needs and even derivative valuation (CSA discounting, §6.9).

Document	Features
1994 New York law CSA	Security interest: the poster keeps ownership, the taker holds a security interest
1995 English law CSA	Title transfer: ownership passes to the taker and is netted on default
2016 VM CSA	Built for regulatory variation margin: zero threshold, daily exchange
2018 IM documents (CSA / CSD) with custodian control agreements	Regulatory initial margin, segregated at a third-party custodian, no reuse

Key terms

Term	Meaning and risk impact
Threshold	Exposure below it is uncollateralised
MTA	Calls below it are not made, avoiding small transfers
Independent amount / IM	An extra buffer independent of MtM
Eligible collateral and haircuts	Which assets and at what haircut; sets usable liquidity
Valuation agent, dates, settlement	Who calculates, how often, settlement lag (usually T+1)
Interest on cash collateral	Usually the overnight rate (e.g. €STR) ± spread, with negative-rate handling; sets the discount curve
Dispute resolution	Timelines, recalculation and escalation when valuations differ
Substitution and reuse	Whether collateral can be swapped and whether the taker may rehypothecate
Additional termination events	Rating downgrades, NAV declines (common for funds, e.g. -15% in 1 month, -25% in 3 months, -35% in 12 months), key-person events triggering early termination

Regulatory requirements

Framework	Requirement
EMIR margin RTS	Uncleared OTC derivatives: VM calculated daily with same-day calls and zero VM threshold; combined VM and IM minimum transfer amount $\leq$ €500k; IM exchanged when both parties exceed €8bn average aggregate notional (final phase from September 2022), with a group IM threshold of up to €50m; IM segregated and not reused; eligible-collateral list, standard haircuts and concentration limits; physically settled FX forwards exempt from IM.
EMIR risk mitigation	Portfolio reconciliation: daily for 500+ outstanding trades with a counterparty, weekly for 51–499, quarterly for up to 50; disputes above €15m outstanding for over 15 business days reported monthly to the regulator.
UCITS	Collateral received must meet ESMA criteria (liquidity, daily valuation, credit quality, low correlation, diversification, third-party custody, enforceability, §8.7); non-cash collateral may not be sold, reinvested or pledged; unprotected collateral posted (e.g. non-segregated IM) counts as counterparty exposure; uncollateralised exposure within the CSA threshold counts towards the 5% / 10% limit.
AIFMD	Collateral reuse counts in leverage; the depositary oversees assets and cash flows; Annex IV reports collateral; reuse is reported under SFTR.
MiFID II	Investment firms may not conclude title-transfer collateral arrangements with retail clients (Art. 16(10)); client assets must be safeguarded and segregated.

Worked example 15 • CSA threshold and the UCITS counterparty limit / 15 • CSA UCITS

A €150m NAV UCITS has a legacy CSA with a **non-credit-institution** counterparty: threshold €10m, MTA €0.5m. Current MtM with that counterparty is + €12.3m and collateral held is €1.5m. Required collateral = $12.3 - 10 = €2.3\text{m}$, so a €0.8m call is made ($\geq$ MTA). Uncollateralised exposure afterwards = $12.3 - 2.3 = €10\text{m} = 6.7\%$ of NAV, above the 5% limit for non-credit institutions. Conclusion: UCITS CSAs should use zero thresholds (2016 VM CSA), otherwise the threshold alone can cause a breach.

In practice • The daily margin process / • Day

1. Morning: compute collateral receivable or payable per CSA from prior-day closing valuations;
2. Issue or receive calls and agree valuations; start disputes when gaps exceed internal thresholds;
3. Choose collateral to deliver, weighing haircuts, reuse rights and cheapest-to-deliver while keeping §18.5 margin coverage;
4. Settle and reconcile: confirm receipt, log pending and failed transfers;
5. Monitor: ageing of open disputes, uncollateralised exposure per counterparty as % of NAV, distance to NAV triggers.

18.9.1 Applying and calculating the threshold

$$CSA = \max(E + IA_{\text{cpty}} - IA_{\text{own}} - TH, 0), \quad \text{Value} = MV \times (1 - h), \quad \text{Delivery} = CSA - \text{Value} \text{ if } \geq \text{MTA}$$

The threshold (TH) is the amount of exposure the counterparty may leave uncollateralised, in effect an agreed unsecured credit line. Three steps:

1. **Credit Support Amount:** exposure E plus independent amounts owed by the counterparty, minus those owed by us, minus the counterparty's threshold, floored at zero; without independent amounts it is $\max(\text{MtM} - \text{threshold}, 0)$.
2. **Compare with collateral held:** collateral value = market value $\times$ (1 – haircut h). Call if the shortfall $\geq$ MTA; return if the excess $\geq$ MTA; otherwise do nothing.
3. **Round:** calls usually round up, returns down (e.g. to €100k); once above the MTA the **whole** difference moves, not just the excess.

Worked example 16 • Daily margining under a threshold / 16 • Day

The fund is in the money; counterparty threshold €5m, MTA €0.5m, rounding to €0.1m; the counterparty delivers government bonds with a 2% haircut. Figures in € millions in the table below.

Whenever MtM exceeds the threshold, uncollateralised exposure is pinned near the threshold (plus at most just under one MTA). The threshold is thus a standing unsecured exposure.

Day	MtM	Credit support amount	Difference	Action	Held	Uncollateralised
1	4.0	0	0	None (within threshold)	0	4.0
2	6.2	1.2	+1.2	Call 1.2 (bonds MV $\approx$ 1.224)	1.2	5.0
3	6.5	1.5	+0.3	None $<$ MTA	1.2	5.3
4	8.0	3.0	+1.8	Call 1.8	3.0	5.0
5	5.3	0.3	–2.7	Return 2.7	0.3	5.0

Effect	Explanation
Exposure ceiling	Worst exposure $\approx$ threshold + MTA + MtM move over the margin period of risk (usually 10 business days bilaterally)
UCITS counterparty limit	Uncollateralised exposure within the threshold counts towards 5% / 10%; keep thresholds well inside the headroom, e.g. $\leq 1\text{--}2\%$ of NAV (Example 15)
CVA	Expected exposure is roughly capped near threshold + MTA + MPOR move; higher thresholds mean higher CVA (§8.7)
Bank capital (SA-CCR)	Replacement cost for margined netting sets $RC = \max(V - C, TH + MTA - NICA, 0)$; the threshold enters capital directly, so banks price high-threshold trades higher

Context	Threshold rule
EMIR variation margin	Must be zero between in-scope parties (2016 VM CSA)
EMIR initial margin	Up to €50m at group level, allocated across counterparties; IM due = $\max(\text{SIMM IM} - \text{allocated threshold}, 0)$
Minimum transfer amount	Combined VM and IM $\leq$ €500k
Legacy CSAs / out-of-scope trades	Positive thresholds possible, but UCITS remain bound by 5% / 10%
Rating-linked thresholds	Step down or to zero on downgrade; mitigates deterioration but can trigger concentrated calls under stress

In practice • Managing thresholds / •

1. Monitor "uncollateralised exposure $\div$ NAV" per CSA daily against counterparty limits;
2. Negotiate: UCITS should seek zero thresholds; AIFs accepting positive thresholds include them in internal counterparty limits;
3. Check both directions: a threshold granted to the fund helps liquidity (less collateral posted), but counterparties may charge more or shorten its validity;
4. Reconcile systems: threshold, MTA, rounding and haircuts in the collateral system must match the CSA text; mis-keyed parameters are a common source of margin disputes.

18.9.2 How much of the exposure does collateral cover?

$$\text{Coverage} = \frac{\sum_j MV_j (1 - h_j)}{E}, \quad \text{with a threshold: Coverage} \approx \frac{\max(E - TH, 0)}{E}$$

There is no fixed ratio; it depends mainly on the CSA terms (table below):

- **With a threshold:** coverage $\approx (E - TH) \div E$, zero when $E \leq TH$; the gap is the threshold itself (day 4 of Example 16: $3.0 \div 8.0 = 37.5\%$).
- **Zero-threshold VM CSA:** about 90–100%; the gap comes from amounts below the MTA, the T+1 settlement lag and intraday MtM moves.
- **VM + IM:** above 100% (over-collateralised), as IM covers potential losses between default and close-out.
- **Haircuts:** coverage uses post-haircut value; €1.224m of bonds at a 2% haircut counts as €1.2m.
- **Value under stress:** collateral itself falls in a crisis; if it is correlated with the counterparty's credit (wrong-way risk, e.g. its home sovereign or its own securities), effective coverage is lower.
- **UCITS:** only collateral meeting ESMA criteria, at post-haircut value, reduces the 5% / 10% counterparty limit; collateral the fund **posts** without segregation (e.g. non-segregated IM) **adds** to exposure to that counterparty.
- **Monitoring:** report coverage per CSA daily and recompute it under stress (collateral value falls, haircuts rise).

Arrangement	Typical coverage	Source of the gap
With a threshold	$(E - TH) \div E$; 0 when $E \leq TH$	The threshold itself (Example 16, day 4: $3.0 \div 8.0 = 37.5\%$)
Zero-threshold VM CSA	$\approx 90\text{--}100\%$	Amounts below MTA, T+1 settlement lag, intraday MtM moves
VM + IM	$> 100\%$ (over-collateralised)	IM covers potential losses from default to close-out

Appendix A Python Reference Implementation

The code below runs as is (numpy, pandas, scipy) and reproduces every result of §9.2 ($T=998$ due to a one-day alignment offset; breach counts are identical), plus the reverse stress test of §11.3, the liquidity coverage of §11.4 and the drawdown statistics of §13. Design rules: the forecast for day t uses data up to $t - 1$ only; zone boundaries use "strictly below ppf".

```
import numpy as np, pandas as pd
from scipy.stats import norm, binom, chi2

def log_returns(px: pd.Series) -> pd.Series:
    return np.log(px).diff().dropna()

def ewma_sigma(r: pd.Series, lam=0.94, seed_n=250) -> pd.Series:
    """sigma_{t|t-1}: forecast for day t using returns up to t-1."""
    v = np.empty(len(r)); v[0] = np.mean(r.values[:seed_n] ** 2)
    for t in range(1, len(r)):
        v[t] = lam * v[t-1] + (1 - lam) * r.values[t-1] ** 2
    return pd.Series(np.sqrt(v), index=r.index)

def cf_factor(p, S, K):
    z = norm.ppf(p)
    return z + (z*z-1)*S/6 + (z**3-3*z)*K/24 - (2*z**3-5*z)*S*S/36

def normal_es_factor(c):
    return -norm.pdf(norm.ppf(1-c)) / (1-c) # -2.6652 at c = 0.99

def var_forecasts(r, c=0.99, window=250, lam=0.94):
    p, z = 1 - c, norm.ppf(1 - c)
    sig = ewma_sigma(r, lam)
    out = pd.DataFrame(index=r.index)
    out["roll21"] = z * r.rolling(21).std().shift(1)
    out["ewma"] = z * sig
    out["hs"] = r.rolling(window).quantile(p).shift(1)
    resid = r / sig # devolatilise
    out["fhs"] = resid.rolling(window).quantile(p).shift(1) * sig # revolatilise
    return out

def backtest(r, var, p=0.01):
    m = var.notna(); hit = (r[m] < var[m]).astype(int).values
    T, x = len(hit), hit.sum(); ph = x / T
    ll = lambda q, k0, k1: (k0*np.log(1-q) if k0 else 0) + (k1*np.log(q) if k1 else 0)
    pof = -2 * (ll(p, T-x, x) - ll(ph, T-x, x))
    a, b = hit[:-1], hit[1:]
    n00, n01 = ((a==0)&(b==0)).sum(), ((a==0)&(b==1)).sum()
    n10, n11 = ((a==1)&(b==0)).sum(), ((a==1)&(b==1)).sum()
    pi0, pi1 = n01/(n00+n01), (n11/(n10+n11) if n10+n11 else 0.0)
    pi = (n01+n11)/(n00+n01+n10+n11)
    ind = -2 * (ll(pi, n00+n10, n01+n11) - ll(pi0, n00, n01) - ll(pi1, n10, n11))
    k250 = hit[-250:].sum()
    g, y = binom.ppf(0.95, 250, p), binom.ppf(0.9999, 250, p) # 5, 10
    zone = "green" if k250 < g else ("yellow" if k250 < y else "red")
    return dict(T=T, x=int(x), rate=ph, LR_pof=pof, p_pof=chi2.sf(pof, 1),
                LR_ind=ind, p_ind=chi2.sf(ind, 1), cc=pof+ind, p_cc=chi2.sf(pof+ind, 2),
                pil=pil, last250=int(k250), zone=zone)

# ----- Part II: stress testing -----
def reverse_stress(w, cov, loss):
    """Most plausible scenario x* (min Mahalanobis) with w'x = -loss."""
    w, cov = np.asarray(w, float), np.asarray(cov, float)
    var_p = w @ cov @ w
    x = -loss * cov @ w / var_p
    return x, loss / np.sqrt(var_p) # scenario, Mahalanobis distance

def coverage(buckets, redemption, horizon, stress_mult=1.0):
    """buckets: list of (share_of_nav, days_to_liquidate)."""
    avail = sum(a for a, d in buckets if d * stress_mult <= horizon)
    return avail / redemption

# ----- Part III: drawdown -----
def drawdown_stats(nav: pd.Series) -> dict:
    peak = nav.cummax(); dd = nav / peak - 1
    trough = dd.idxmin(); start = nav.loc[:trough].idxmax()
    after = nav.loc[trough:]; rec = after[after >= nav.loc[start]]
    yrs = (nav.index[-1] - nav.index[0]).days / 365.25
    cagr = (nav.iloc[-1] / nav.iloc[0]) ** (1 / yrs) - 1
    return dict(mdd=dd.min(), peak=start, trough=trough,
                recovery=rec.index[0] if len(rec) else None,
                calmar=cagr / -dd.min(), ulcer=np.sqrt((dd ** 2).mean()))

def mc_mdd(sigma=0.16, mu=0.0, years=3, n_paths=10000, steps=252, seed=1):
    rng = np.random.default_rng(seed)
    lr = (mu - 0.5 * sigma**2) / steps + sigma / np.sqrt(steps) * rng.standard_normal((n_paths, years * steps))
    lv = np.concatenate([np.zeros((n_paths, 1)), np.cumsum(lr, axis=1)], axis=1)
    mdd = 1 - np.exp(lv - np.maximum.accumulate(lv, axis=1))
    return mdd.max(axis=1) # distribution of MDD

if __name__ == "__main__":
```

```

px = pd.read_csv("sp500.csv", index_col=0, parse_dates=True)["Px"]
r = log_returns(px)
f = var_forecasts(r).iloc[251:]
for col in f:
    print(col, {k: (round(v, 4) if isinstance(v, float) else v)
                for k, v in backtest(r.loc[f.index], f[col]).items()})
print(drawdown_stats(px))
m = mc_mdd(years=3, mu=0.06); print("3y MDD median/95%:", np.median(m).round(3), np.quantile(m, .95).round(3))

```

Appendix B Excel Formula Cheat Sheet

Use Purpose	Formula	Note
Log-return	=LN(B8/B7)	
Normal factor	=NORM.S.INV(0.01)	-2.32635
ES Normal ES factor	=-NORM.S.DIST(NORM.S.INV(0.01),FALSE)/0.01	-2.6652
σ 21 obs Rolling σ	=STDEV.S(C8:C28)	1 · shift one row
10 Day VaR	=STDEV.S(C8:C28)*SQRT(10)*Factor_VaR	Factor_VaR
10 Day Forward 10D return	=LN(B39/B29)	VaR · align with VaR row
Breach flag	=IF(E28<D28,1,0)	Basis · negative convention
EWMA EWMA variance	=E5*\$I\$2+D5*(1-\$I\$2)	Column D holds u^2 , I2 holds λ
HS VaR	=PERCENTILE.INC(C8:C257,0.01)	or SMALL(range,k)
Cumulative Binomial CDF	=BINOM.DIST(k,250,0.01,TRUE)	· traffic light
Kupiec p	=CHISQ.DIST.RT(LR,1)	
/ Skew / excess kurtosis	=SKEW(range),=KURT(range)	KURT · KURT is excess

Appendix C Glossary

English	/ One-line meaning
Value at Risk (VaR)	Loss percentile at a given confidence and horizon ·
Expected Shortfall (ES), CVaR	Average loss beyond VaR ·
Quantile / percentile	Point where cumulative probability reaches a level ·
Inverse CDF (quantile function)	Maps probability to value ·
Supremum / infimum	Least upper / greatest lower bound ·
Subadditivity	Merging does not add risk ·
Coherent risk measure	Satisfies the four axioms ·
EWMA	Exponentially decaying weights ·
Plateau effect / ghost effect	Step-like VaR from equal-weighted windows ·
Parametric / historical / Monte Carlo VaR	Distribution assumed, historical, or simulated ·
Marginal / component / incremental VaR	Derivative, contribution and change of VaR by position ·
Absolute / relative VaR (UCITS)	$VaR \leq 20\%$ of NAV / fund-to-reference VaR ratio ≤ 2 ·
Tracking-error VaR	VaR of the active return ·
Reference portfolio (UCITS)	Unleveraged benchmark for relative VaR ·
Stressed VaR (sVaR)	VaR calibrated on a stress period ·
Unsmoothing (Geltner)	Restores volatility hidden by appraisal smoothing ·
Netting / hedging arrangements	The two ways to offset exposure under the commitment approach ·
Duration netting	Partial netting of rate derivatives across maturity buckets ·
CE / PFE / EE	Three measures of counterparty exposure ·
Credit valuation adjustment (CVA)	Valuation charge for counterparty default ·
Wrong-way risk	Default probability rising with exposure ·
Variation / initial margin (VM / IM)	Daily MtM settlement / cover for close-out risk ·
Taylor expansion	Polynomial approximation built from derivatives ·
Full revaluation	Repricing with the valuation model under shocked data ·
Sticky strike / sticky moneyness	Two rules for moving the vol surface with spot ·
Barrier option / autocallable	Typical path-dependent products ·
Common random numbers	Same draws for base and shocked Monte Carlo ·
Multi-factor model / factor exposure	Explains returns with a few common factors ·
Specific (idiosyncratic) risk	Residual risk not explained by factors ·
Sensitivities-based method (SBM) / ISDA SIMM	Prescribed factor-sensitivity models ·
Risk factor eligibility test (RFET)	Whether a factor can enter the internal model ·
Sum of notionals	Leverage basis disclosed by UCITS VaR funds ·
Margin coverage ratio	Available cash and collateral ÷ stressed margin calls ·
Pin risk / gap risk	Spot near strike at expiry / discontinuous price jumps ·
Independent price verification (IPV)	Valuation check independent of trading ·
Credit Support Annex (CSA)	ISDA annex governing collateral exchange ·
Threshold / minimum transfer amount (MTA)	Uncollateralised allowance / smallest call ·
Additional termination event (ATE)	E.g. NAV-decline early termination ·
Discount / projection curve	The two curve roles in multi-curve pricing ·
Bootstrapping	Sequentially solving discount factors ·
Key-rate DV01	P&L for a 1bp bump at one tenor ·
Principal component analysis (PCA)	Level, slope, curvature ·

Credit Support Amount	$\max(\text{exposure} - \text{threshold} \pm \text{independent amounts}, 0)$ ·
Margin period of risk (MPOR)	Time from default to close-out ·
Collateral coverage	$\text{Post-haircut collateral} \div \text{exposure}$ ·
Half-life	Days for weight to halve ·
Historical simulation (HS)	Revalue under historical scenarios ·
Filtered historical simulation (FHS)	HS on devol/revol returns ·
Devolatilise / revolatilise	Divide by own-day σ / multiply by today's σ ·
Procyclicality	Risk estimates amplify the cycle ·
Skewness / excess kurtosis	3 · Third and fourth standardised moments (kurtosis – 3)
Generalised Pareto distribution (GPD)	Limit law of threshold exceedances ·
Backtesting / breach / exception	Realised loss exceeds VaR ·
Eligible comparisons	Days with both VaR and realised return ·
Unconditional / conditional coverage	Right frequency / frequency and independence ·
Traffic-light approach	Green/yellow/red by breach count ·
Fundamental Review of the Trading Book (FRTB)	Basel III market-risk framework ·
Internal models approach (IMA) / standardised approach (SA)	
HPL / APL / RTPL	see §8.5 ·
P&L attribution (PLA) test	Consistency of RTPL with HPL ·
VaR-equivalent volatility (VEV)	Input to the PRIIPs risk class ·
Initial margin (IM)	Collateral for close-out risk ·
Stress test / sensitivity analysis	P&L under a given shock ·
Reverse stress test	Solve for scenarios causing an unacceptable loss ·
Mahalanobis distance	Standard-deviation distance accounting for correlation ·
Liquidity stress test (LST)	Stress on liquidation and redemptions ·
Coverage ratio	$\text{Liquidable assets} \div \text{stressed redemptions}$ ·
Liquidity management tool (LMT)	Gates, swing pricing, anti-dilution levies, etc. ·
Swing pricing	NAV adjusted for net flows ·
Drawdown / maximum drawdown (MDD)	Fall from the running peak ·
Time under water	Time from peak to new peak ·
High-water mark (HWM)	Highest NAV to date ·
Calmar ratio / Ulcer index	Drawdown-adjusted return and pain measure ·
Reflection principle	Touching probability is twice the terminal one ·
Scenario analysis	Multi-period, narrative-driven path analysis ·
Performance scenarios (PRIIPs)	Stress, unfavourable, moderate, favourable ·
Recommended holding period (RHP)	Horizon for PRIIPs scenarios and SRI ·
Global exposure (UCITS)	Derivative risk under commitment or VaR ·
Gross / commitment method (AIFMD)	Two leverage calculations ·
Summary risk indicator (SRI)	PRIIPs 1–7 risk class ·