The Instruments of Professional Risk Management: From Value at Risk and Expected Shortfall to Monte Carlo Simulation and Performance Attribution

Roland Rupprechter[1]

July 2026

Abstract

This paper systematises the quantitative instruments for measuring and managing financial risk that are available to the professionals in charge of wealth management — whether the wealth of private investors is pooled in a dedicated special fund, transferred to a foundation, or steered by a family office. Starting from the distributional foundations — moments, skewness, kurtosis, quantiles — it develops the central risk measures on a formula basis: Value at Risk (VaR) as a loss quantile in parametric form with the Cornish-Fisher correction, the Expected Shortfall (Conditional Value at Risk) as a coherent expectation measure of the exceedance losses (Artzner, Delbaen, Eber, and Heath, 1999; Rockafellar and Uryasev, 2000), the cash-flow-oriented quantities Cash-Flow at Risk and Earnings at Risk, the risk score as the product of probability of occurrence and severity, and the Lower Partial Moments as target-return-related downside measures (Bawa, 1975; Fishburn, 1977; Harlow, 1991).

Three methods of risk quantification are demonstrated on a running example portfolio of 10 million euros: the parametric calculation, historical simulation with bootstrap confidence intervals (Efron, 1979), and stochastic Monte Carlo simulation (Boyle, 1977). A dedicated chapter quantifies the default risk of an interest rate swap via simulated exposure profiles and the credit valuation adjustment (CVA); another decomposes active returns following Brinson, Hood, and Beebower (1986) and in the factor-based Barra framework (Rosenberg, 1974; Grinold and Kahn, 2000). A further chapter compares the three carrier structures — special fund, foundation, and family office mandate — by benefit, cost, and governance; the total-wealth perspective connects the instruments via target returns, withdrawal plans, and the supervisory reference measures of Basel (Expected Shortfall 97.5 percent) and Solvency II (VaR 99.5 percent) as quality benchmarks for private mandates.

[1]*R&B Wealth Management; Member, Austrian Association for Financial Analysis and Asset Management

The finding: the measures are complements, not competitors — the VaR standardises communication, the Expected Shortfall captures the tails, the Lower Partial Moments anchor the investor's goals, simulation carries the valuation of nonlinear positions, and attribution closes the cycle of risk-taking and outcome control. The paper is intended as a didactic synthesis for private wealth practice, not as an investment recommendation.

Keywords: Value at Risk · Expected Shortfall · Lower Partial Moments · skewness and kurtosis · Monte Carlo simulation · bootstrap · counterparty risk · performance attribution · special fund · foundation · family office.

JEL codes: G11, G17, G22, G23, G32.

Non-Technical Summary

Wealthy private individuals who organise their investments professionally — in their own special fund, through a foundation, or with a family office — must be able to answer one question at any time: how much can be lost, and with what probability? This paper first presents the three carrier structures with their advantages and disadvantages and then develops the methodological apparatus with which financial mathematics answers this question — each method demonstrated on a fully worked 10 million euro example portfolio.[2]

The best-known tool is Value at Risk: the loss that, with a chosen probability — say 95 or 99 percent — will not be exceeded within a given horizon. In the example, the portfolio loses no more than about 19 million euros over one year with 95 percent confidence. Value at Risk is silent, however, about how bad things get in the remaining five percent of cases — that is answered by the Expected Shortfall, the average of the losses beyond the threshold: here about 23 million euros. Because financial returns are more skewed and heavier-tailed than the bell curve, the paper also shows how skewness and kurtosis are measured and built into the risk calculation.

Alongside these stand tools for special questions: Cash-Flow at Risk and Earnings at Risk transfer the logic to cash flows and corporate results; the Lower Partial Moments measure only the downward deviations from a self-chosen target return — closer to what investors really perceive as risk; the sensitivity measures (Greeks) capture the reaction of derivative positions to market moves. For the default risk of swap transactions, the paper simulates thousands of interest rate paths with the Monte Carlo method and calculates from them what the possible default of the counterparty costs. And performance attribution finally decomposes the achieved return into its sources: what came from the division across asset classes, what from stock selection, what from style factors such as value or momentum?

Private wealth also has a second side: planned withdrawals of the family, promised distributions of a foundation, margin obligations from loans and derivatives. What is steered is therefore not the portfolio alone but the surplus over these obligations — with the same tools, applied to the difference between the two sides. The rules of the supervised world serve as the quality benchmark: banks calculate under Basel with the Expected Shortfall, insurers under Solvency II with a Value at Risk of 99.5 percent — standards against which a private mandate can also measure itself.

The most important message of the paper is a warning against the single number: every metric has blind spots, and the great losses of financial history regularly hit houses that confused a metric

[2] None of the metrics presented measures risk completely; professional risk management always rests on the interplay of several measures and on the understanding of their respective blind spots.

with reality. Professional risk management means using several tools side by side — and knowing what each one does not see.

Contents

1 Introduction

Risk management is the second half of investing. While the return side of investing has always stood in the limelight, the risk side decides survival: where wealthy private investors organise their wealth through a special fund, a foundation, or a family office mandate, the professionals in charge of wealth management — the specialists of banks, investment firms, and fund companies — must quantify potential losses before they occur, steer risk positions before they tip, and put numbers on probabilities of occurrence before the investment committee, foundation board, or family council asks. Financial mathematics provides for this a methodological repertoire grown over many decades — from the moments of the return distribution through quantile measures such as Value at Risk to stochastic simulation methods and factor-based performance decomposition — whose components this paper develops systematically, works through numerically, and probes for their limits.[3]

The scientific development of these instruments can be told along a few milestones. Markowitz (1952) made the variance the first operational risk measure and Roy (1952) simultaneously formulated the safety-first principle from which the downside measures emerged; Bawa (1975), Fishburn (1977), and Harlow (1991) generalised them into the Lower Partial Moments. Mandelbrot (1963) and Fama (1965) showed early that financial market returns are more skewed and heavier-tailed than the normal distribution — findings that Cont (2001) consolidated into the stylised facts. Value at Risk became the market standard with RiskMetrics (J.P. Morgan, 1996) and was given its scientific place by Jorion (2007) and Duffie and Pan (1997); its conceptual weaknesses moved Artzner, Delbaen, Eber, and Heath (1999) to the axiomatics of coherent risk measures, whose practical representative — the Expected Shortfall — was operationalised by Rockafellar and Uryasev (2000) as well as Acerbi and Tasche (2002) and today forms the Basel market risk standard. The simulation methods go back to Boyle (1977) for Monte Carlo and Efron (1979) for the bootstrap; attribution analysis to Brinson, Hood, and Beebower (1986) and factor decomposition to Rosenberg (1974), whose work founded the Barra models.

The contribution of this paper lies in the integrated, application-oriented synthesis of this apparatus for wealthy private investors and their carrier structures — special fund, foundation, and family office: all metrics — Value at Risk, Expected Shortfall, Cash-Flow at Risk and Earnings at Risk, the risk score, sensitivity measures, Lower Partial Moments, skewness and kurtosis — are developed on the same 10 million euro example portfolio and placed in relation to one another; the method families (parametric, historical with bootstrap, Monte Carlo) are

[3] All worked examples in this paper are hypothetical and serve to illustrate the methods; they constitute no investment, legal, or tax advice. Regulatory statements reflect the state as of mid-2026 in simplified form.

brought together on the default risk of an interest rate swap, and the Brinson and Barra attribution closes the control loop from risk-taking to outcome control. All examples are hypothetical and serve as illustration. In addition, the paper makes an empirical contribution of its own: the econometric case study of Section 13 applies the apparatus developed here to six decades of real S&P 500 data and tests whether the New Economy and the AI era have measurably changed the return properties of the market.

The paper is organised as follows. Section 1 (the present section) introduces the question and situates the literature. Section 2 compares the three carrier structures of private wealth — special fund, foundation, and family office mandate — by benefit, cost, and governance. Section 3 lays the distributional foundations: moments, skewness, kurtosis, and the concept of the quantile. Section 4 develops Value at Risk in parametric form with the Cornish-Fisher correction, Section 5 the Expected Shortfall as its coherent complement. Section 6 contrasts the method families — historical simulation, bootstrap, and Monte Carlo — and Section 7 transfers the logic to cash flows (Cash-Flow at Risk, Earnings at Risk) and the risk score. Section 8 treats the Lower Partial Moments, Section 9 the sensitivity measures. Section 10 quantifies the default risk of interest rate swaps via simulated exposure profiles and the credit valuation adjustment, Section 11 decomposes active returns following Brinson and in the Barra factor framework. Section 12 connects the instruments for steering total wealth including withdrawals, obligations, and supervisory reference measures, Section 13 subjects the S&P 500 monthly returns of the years 1966 to 2025 to a stand-alone econometric case study across periods, and Section 14 closes with a conclusion and outlook.

2 Three Routes to Professional Wealth Management: Special Fund, Foundation, Family Office

2.1 The Special Fund: a Tailor-Made Investment Vehicle

The special fund is the institutional answer to a private need: a dedicated, legally separate pool of assets whose investment guidelines the investor determines personally. A management company (KVG) administers the fund, an independent depositary controls the holdings, and the investor — frequently as sole subscriber — receives a vehicle that is supervised like a mutual fund but steered like a segregated mandate: own investment limits, own benchmark, own reporting, an investment committee with a voice of its own. In practice, viability begins — owing to the fixed costs of administration, custody, and audit — at volumes of roughly ten to fifty million euros; below that, mutual funds and discretionary mandates render the same service at lower cost.[4]

The advantages lie in control and pooling: all of a family's investments can be consolidated in one wrapper, reallocations occur within the fund without ongoing tax realisation at the investor level, reporting is uniform, and institutional terms (share classes, negotiating power vis-à-vis managers) become accessible. The disadvantages are the mirror image: running fixed costs regardless of success, committee and coordination effort with the management company and depositary, regulatory limits for illiquid investments, and a certain inertia, since changing the vehicle means changing a contractual construct rather than a brokerage account.

2.2 The Foundation: Wealth with Purpose and a Claim to Permanence

The foundation is not an administrative but an ownership construct: the wealth is transferred to a legal person of its own and thereby permanently separated from the founder's private wealth and from the family cycle. The founder lays down purpose, beneficiaries, and governance in the charter — board, advisory council, control rights of the family — and the foundation outlasts successions, divorces, and creditor claims against the person. For the investment of the capital, a particular clarity arises: the investment objective is typically real capital preservation plus the distributions required by the charter — a precise target return on which, as this paper will show, the entire risk measurement can be aligned.[5]

The advantages are continuity and protection: orderly succession across generations, the binding of wealth to a purpose, a professional governance structure, protection against fragmentation. The disadvantages demand honesty before establishment: the transfer is in essence irreversible — the founder gives up ownership for good —, running costs and (depending on the design) state supervision arise, the charter binds future generations to today's decisions, and the tax consequences are multi-layered. A foundation solves succession and protection questions — it is neither a return instrument nor a substitute for investment discipline.

2.3 The Family Office Mandate: Steering instead of a Legal Shell

The family office is — unlike special fund and foundation — not a legal shell but a service: the central steering instance over the total wealth, in whatever vehicles, accounts, and

[4] In Germany, special AIFs are open to professional and semi-professional investors; wealthy private individuals attain semi-professional investor status from a subscription of 200,000 euros with an examination of experience and expertise. The tax treatment is presented in strongly simplified form and belongs in professional advice.

[5] Meant are the German (family) foundation and the Austrian private foundation; the tax treatment — such as the interim tax in Austria or the substitute inheritance tax in Germany — is complex and deliberately bracketed out here.

countries it sits. Its core includes the consolidated wealth overview across all banks, the strategic asset allocation, selection and monitoring of the managers, fee negotiations, and the family's reporting and meeting routine; in the wider sense, succession coordination and family governance are added. Compensation is typically via fixed fees or basis points on assets under supervision — independent houses without products of their own avoid the conflicts of interest of commission-driven distribution.[6]

The advantage is the total view: only the consolidated level makes concentration risks visible across banks, only it can assign risk budgets and measure managers against uniform standards — the family office is the natural user of almost all the metrics of this paper. The disadvantages: the service costs money permanently and pays off only from sufficient wealth size and complexity; quality disperses considerably, and the selection creates a principal-agent problem of its own — delegating control requires the ability to control the controller. Precisely for this, the wealth owner needs the language of the following sections.

2.4 Choosing the Structure — and Why It Does Not Replace Risk Management

Table 1 juxtaposes the three routes. The choice follows less the return than the question of which problem is to be solved: the special fund pools and professionalises the investments, the foundation orders ownership and succession, the family office creates overview and steering. Frequently the answer is a combination — the foundation as owner, the special fund as investment wrapper, the family office as helmsman.[7]

[6] Single family offices serve one family and, in experience, pay off only from wealth in the hundreds of millions; multi family offices pool several families and open the service from considerably smaller fortunes.

[7] The structures are not mutually exclusive: in practice, a foundation frequently holds units of a special fund, and a family office steers both.

Table 1: Special fund, foundation, and family office mandate compared

Feature	Special fund	Foundation	Family office mandate
Legal nature	segregated pool with management company and depositary	legal person of its own, binding of wealth	service contract, no legal shell
Core benefit	tailor-made, pooled investment	succession, protection, purpose binding	total view, steering, manager selection
Typical minimum size	ca. €10–50m	sensible from low double-digit millions	MFO from ca. €5–20m, SFO far above
Cost character	fixed costs of management company, depositary, audit	establishment, governing bodies, running administration	fee/basis points on supervised assets
Governance	investment committee, investment guidelines	charter, board, advisory council	mandate contract, reporting duties
Main disadvantage	fixed costs and committee effort	irreversibility, complexity	quality dispersion, delegation risk

Source: Own compilation.

Decisive for this paper is what all three routes have in common: they delegate the administration but not the responsibility. Whether the investment committee of the special fund, the foundation board, or the family council — in the end, the wealth owner sits in front of a report and must judge how much risk sits in the wealth, whether it was paid for, and whether the manager delivers what the mandate promises. Exactly these three questions are answered by the risk analysis instruments of the following sections — the distribution measures, simulation methods, and attribution calculations are the common language in which special fund, foundation, and family office render account.

3 Distributional Foundations: Moments, Skewness, Kurtosis, and Quantiles

3.1 Returns as Random Variables and the First Two Moments

Every quantitative risk calculation begins with the same abstraction: the future return of a portfolio is a random variable, and risk is a property of its probability distribution. The first two moments of this distribution — the expected value and the variance, or its root, the volatility — have formed the foundation of portfolio theory since Markowitz (1952): the expected value measures the location of the distribution, the volatility its width. For the running example portfolio of this paper — 10 million euros, invested 60 percent in equities (volatility 18 percent) and 40 percent in bonds (volatility 5 percent) with a correlation of 0.2

— portfolio theory delivers a total volatility of 11.37 percent per year; diversification pushes the risk below the weighted average of the individual volatilities of 12.8 percent.[8]

Two moments suffice, however, only if the distribution is fully described by them — as with the normal distribution. Precisely that is not the case in financial markets, and the deviations are no footnote but the core of risk management: they concern the tails of the distribution, that is, exactly those rare, large losses that risk measurement is about.

3.2 The Third and Fourth Moments: Skewness and Kurtosis

Equations (1) and (2) define the standardised higher moments:[9]

$$S = E[(r - \mu)^3] \,/\, \sigma^3 \tag{1}$$

Equation (1): The skewness S measures the asymmetry of the distribution. Negative skewness — the rule for equity returns, with typical values between –0.3 and –0.8 — means: the large moves lie predominantly on the loss side; small gains are frequent, large slumps rare but severe. For the risk manager, negative skewness is a warning signal that volatility alone does not show.

$$K = E[(r - \mu)^4] \,/\, \sigma^4 - 3 \tag{2}$$

Equation (2): The excess kurtosis K measures the heaviness of the tails. Daily equity index returns show values from 3 to over 10 — extreme moves are many times more likely than the normal distribution admits (Mandelbrot, 1963; Fama, 1965; Cont, 2001). Figure 1 contrasts a skewed, heavy-tailed distribution with the normal distribution.

---

[8] All metrics of this paper can be referred to either returns or losses; the loss quantity L is the negative value change of the portfolio over the observation horizon.

[9] Customarily stated is the excess kurtosis, i.e. the kurtosis minus the normal distribution value of 3; positive values mean heavier tails than under the normal distribution.

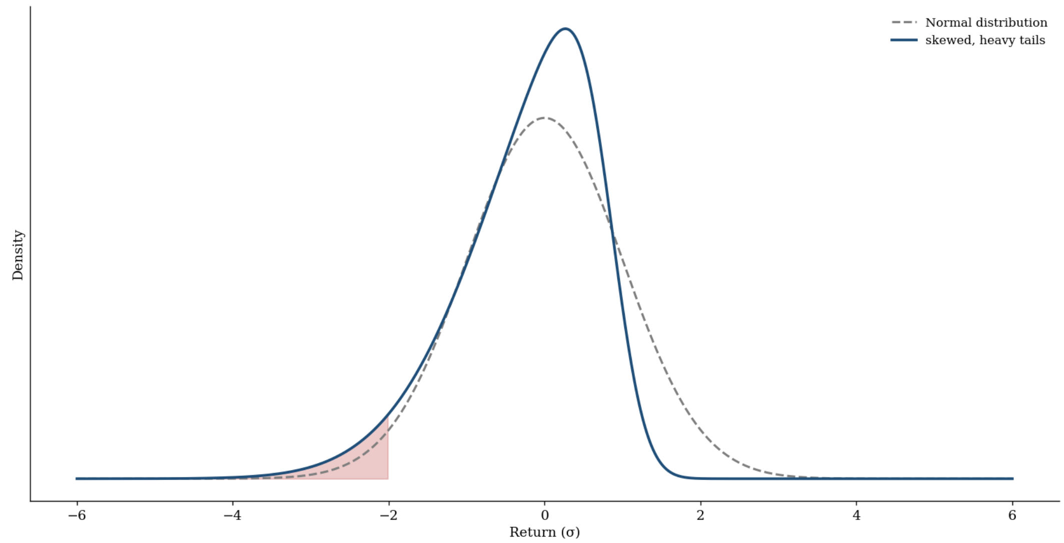

Figure 1: Normal distribution and a left-skewed distribution with heavy tails compared; the marked area shows the elevated probability of large losses observed in financial market returns (schematic).
Source: Own presentation and calculation.

Figure 1 shows why the higher moments are no academic subtlety for wealth owners: two portfolios with identical volatility can differ several-fold in the risk of large losses when skewness and kurtosis diverge — for instance because one of the two systematically sells options or holds credit risks whose return profile trades many small premiums against rare large losses. Sections 4.3 and 8 show two ways of bringing this information into the risk calculation: the Cornish-Fisher correction of Value at Risk and the Lower Partial Moments.

3.3 Quantiles, Quartiles, and Percentiles: the Language of Risk Measures

The third foundation is the concept of the quantile. The p-quantile of a distribution is the value that the random variable does not exceed with probability p: the median is the 50 percent quantile, the quartiles mark 25, 50, and 75 percent, and the percentile ranks of fund comparison place each manager within the distribution of their peer group — the first quartile gathers the best quarter, the fourth the weakest. This familiar statistic of performance comparison is at the same time the exact language of risk measurement: the Value at Risk of Section 4 is nothing other than a high quantile of the loss distribution — the 95 or 99 percent quantile — and the ability to read quartiles is the ability to read Value at Risk. Table 2 summarises the concepts.[10]

Table 2: Distribution metrics at a glance

Metric	Definition	Meaning	Typical use
Expected value μ	first moment	location of the distribution	return forecast, plan figure
Volatility σ	root of the second central moment	width of the dispersion	portfolio theory, Sharpe ratio
Skewness S	third standardised moment (Equation 1)	asymmetry: loss vs. gain moves	warning signal for premium-selling strategies
Excess kurtosis K	fourth standardised moment (Equation 2)	heaviness of the tails	tail risk, stress tests
Quantile q_p	value not exceeded with probability p	threshold of the distribution	VaR, percentile ranks
Quartiles Q1–Q3	25/50/75 percent quantiles	quartering of the distribution	manager and fund comparison

Source: Own compilation based on the data of Cont (2001).

With these foundations, the risk analysis instruments are ordered: Sections 4 to 8 develop risk measures that each single out a different property of the loss distribution — its quantile (Value at Risk), the expectation of its tail (Expected Shortfall), its transfer to cash flows (Cash-Flow at Risk), and its deviation from an investor's target (Lower Partial Moments). None of these measures contains more information than the distribution itself; their value lies in the condensation — and their risk in confusing the condensation with reality.

3.4 Stylised Facts: How Financial Markets Actually Behave

Three empirical regularities shape risk measurement beyond skewness and kurtosis. First, volatility clustering: calm and stormy phases alternate in blocks — a large move is followed by the next with elevated probability. For practice this means: risk is forecastable even where returns are not, and models with time-varying volatility beat static windows. Second, correlation asymmetry: correlations between risky assets rise in downturns — the diversification of Equation (4) is weakest when it is needed most urgently. Third, aggregational normality: as the time horizon grows, return distributions approach the normal distribution — daily data are more volatile than annual data, which is why confidence level, horizon, and data frequency must be chosen together.[11]

[10] Quartiles divide the distribution into four, quintiles into five, deciles into ten, and percentiles into a hundred equally likely sections; all are special cases of the quantile concept.

[11] Cont (2001) catalogues the stylised facts across markets and asset classes; they count as the empirical minimum requirements for every risk model.

These facts are the reality benchmark for all that follows: a risk model that knows no clusters, no crisis correlations, and no heavy tails is not badly calibrated but misspecified. The methods of Section 6 differ precisely in how much of this reality they capture — and at what price in complexity.

3.5 Data and Conventions: the Fine Print of Risk Calculation

Before the first metric is calculated, three inconspicuous decisions are made that shape every result. The return definition: discrete or continuous (logarithmic) — the former adds across portfolio components, the latter across time; risk systems must know which they are mixing. The frequency: daily, weekly, or monthly data carry different amounts of noise and differently heavy tails — the kurtosis of Section 3.2 shrinks with aggregation. And the annualisation: the square-root-of-time rule of Equation (5) holds exactly only without autocorrelation; smoothed valuations — real estate, private equity, thinly traded bonds — show artificially low short-term volatility whose extrapolation drastically understates the annual risk.[12]

The last point deserves particular attention: smoothing is the most frequent source of seemingly superior risk-adjusted returns of illiquid investments. Professional houses de-smooth such time series before they enter allocation or risk models — otherwise the illusion of calm migrates into the portfolio as an overweight of illiquid investments and takes its toll in the first liquidity crisis of Section 7.3.

4 Value at Risk: the Loss Quantile as Industry Standard

4.1 Definition and Interpretation

Value at Risk is the most widely used risk measure in finance, and its definition is deliberately plain: it is the loss that, over a fixed period, is not exceeded with a specified probability — the confidence level c. Equation (3) casts this as a quantile condition:[13]

$$P(L > VaR_c) = 1 - c \tag{3}$$

Equation (3): The loss L exceeds the VaR at confidence level c only with the residual probability $1 - c$. The VaR is thus exactly the c-quantile of the loss distribution of Section 3.3: a 99 percent VaR of 530,000 euros says that in 99 of 100 comparable ten-day periods the loss stays below it — and says nothing about what happens in the hundredth.

---

[12] For small returns, discrete and continuous calculation are practically identical; for large moves and long horizons the differences become material.

[13] The combination of confidence level and horizon is convention: banks internally often calculate with 99 percent and one to ten days, insurers under Solvency II with 99.5 percent and one year, asset managers frequently with 95 percent.

The strength of the VaR lies in its communicability: one number, in monetary units, with a clear probability statement — intelligible to board, supervisory board, and supervisor alike. With RiskMetrics (J.P. Morgan, 1996) it became the market standard, and Jorion (2007) as well as Duffie and Pan (1997) systematised it scientifically. Its weakness is the flip side of the same plainness: it is blind to everything beyond the threshold — a gap that Section 5 closes.

4.2 The Parametric VaR on the Example Portfolio

Under the normality assumption, the VaR can be calculated in closed form. Equation (4) delivers the portfolio volatility, Equation (5) the parametric VaR:[14]

$$\sigma_P = \sqrt{(w'\Sigma w)} = \sqrt{(w_1^2\sigma_1^2 + w_2^2\sigma_2^2 + 2w_1w_2\rho\sigma_1\sigma_2)} \tag{4}$$

Equation (4): The portfolio volatility σ_P follows from weights w, individual volatilities σ, and correlation ρ. For the example portfolio (60 percent equities at 18 percent, 40 percent bonds at 5 percent, $\rho = 0.2$), the result is 11.37 percent per year — the diversification effect relative to the weighted average amounts to almost one and a half percentage points.

$$VaR_c = z_c \cdot \sigma_P \cdot \sqrt{t} \cdot V \tag{5}$$

Equation (5): The parametric VaR scales the portfolio volatility with the normal quantile z_c (1.645 for 95 percent, 2.326 for 99 percent), the square root of the time horizon t, and the portfolio value V. For the 10 million euros of the example: over ten trading days 529,000 euros at the 99 percent level, over one year 1.87 million euros at the 95 percent level. Figure 2 shows both values in the loss distribution.

[14] The square-root-of-time scaling assumes independent, identically distributed returns; under volatility clustering it understates the risk of longer horizons.

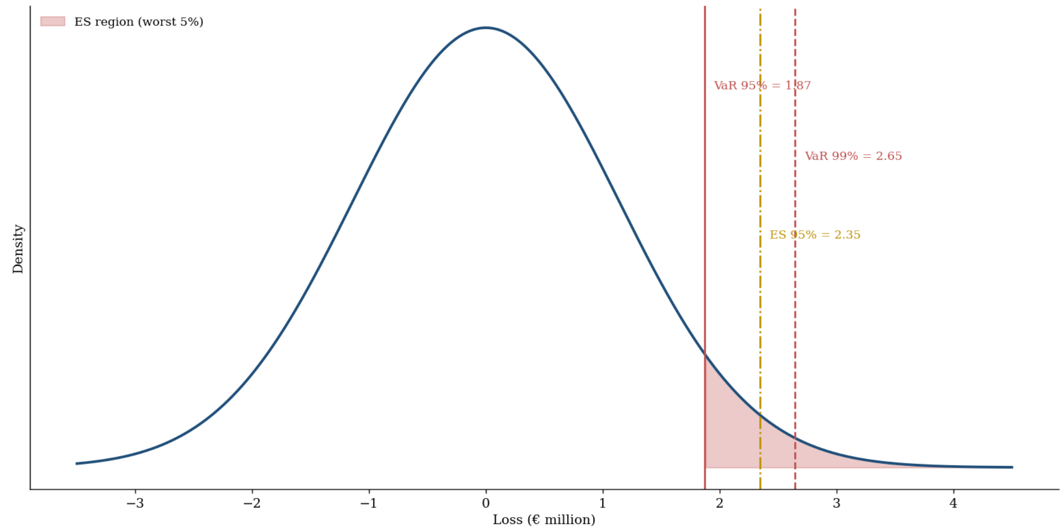

Figure 2: Loss distribution of the example portfolio over one year with Value at Risk at 95 and 99 percent and Expected Shortfall at 95 percent; the marked area shows the worst five percent of cases, whose average is the Expected Shortfall.
Source: Own presentation and calculation.

Figure 2 makes the anatomy of risk measurement visible: the VaR is a threshold in the distribution, not a maximum loss. The phrase "maximum loss under normal market conditions" is therefore to be taken literally — the VaR describes the normal conditions and is silent about the abnormal ones. Berkowitz and O'Brien (2002) show, on the trading books of large US banks, that internal VaR models are rather conservative in everyday conditions but are breached frequently and in clusters in stress phases — exactly the pattern the higher moments of Section 3.2 lead one to expect.

4.3 Skewness and Kurtosis in the VaR: the Cornish-Fisher Correction

The normality assumption of Equation (5) contradicts the stylised facts of Section 3.2. The Cornish-Fisher expansion remedies this without giving up the closed form: it corrects the quantile z_c with skewness and kurtosis terms according to Equation (6):[15]

$$z_{CF} = z + (z^2 - 1) \cdot S/6 + (z^3 - 3z) \cdot K/24 - (2z^3 - 5z) \cdot S^2/36 \tag{6}$$

Equation (6): The corrected quantile z_{CF} systematically widens the VaR under negative skewness S and positive excess kurtosis K (Cornish and Fisher, 1938; Favre and Galeano, 2002). In the example: with $S = -0.6$ and $K = 3$, the 99 percent quantile rises from 2.326 to 2.451 — the ten-day VaR grows from 529,000 to 557,000 euros. A good five percent more risk

[15] The approximation of Cornish and Fisher (1938) is reliable for moderate skewness and kurtosis; at extreme values, switching to simulation or extreme value methods is recommended (McNeil, Frey, and Embrechts, 2015).

capital, solely because the distribution is described more realistically: the price of the heavy tails in a single formula.

Table 3: Value at Risk of the example portfolio (10 million euros, volatility 11.37 percent per year)

Metric	Level / horizon	Value	Method
VaR	95 % / 1 year	€1.87m	parametric, Equation (5)
VaR	99 % / 10 days	€529,000	parametric, Equation (5)
VaR (Cornish-Fisher)	99 % / 10 days	€557,000	Equation (6), S = −0.6, K = 3
Expected Shortfall	95 % / 1 year	€2.35m	Equation (8)
Expected Shortfall	99 % / 10 days	€606,000	Equation (8)
VaR (historical)	99 % / 1 day	€154,000	250 trading days, Section 6.1

Source: Own calculations according to Equations (5), (6), and (8).

Table 3 gathers the risk calculation of the example portfolio and already shows the interplay of the measures that the following sections develop. Two remarks are essential for the interpretation. First, the values scale linearly neither with the confidence level nor with the time horizon; a comparison of metrics therefore presupposes knowledge of both conventions. Second, the difference between Value at Risk and Expected Shortfall constitutes a measure of the weight of the distribution tails — and thus itself contains risk-relevant information.

4.4 Backtesting: the Metric under Supervision

A risk metric is only as good as its subsequent verification. In backtesting, one counts how often the actual losses exceeded the forecast VaR: a correctly calibrated 99 percent VaR should be breached on around two and a half of 250 trading days — significantly more exceptions betray an overly optimistic model, significantly fewer an overly expensive one. Kupiec (1995) formalised the test for the correct hit rate, Christoffersen (1998) added the test for independence of the exceptions: breaches that cluster are a warning sign even when the overall count is right, for they show that the model misses volatility phases. Banking supervision has institutionalised this logic in the traffic-light approach and ties the number of exceptions directly to capital add-ons.[16]

For private mandates beyond the bank balance sheet, backtesting is equally obligatory — not in regulation but in method: a VaR without a documented hit rate is an assertion, not a metric. Practice tests quarterly for level and clustering and escalates model deviations to the investment committee — the small feedback loop that turns a computed figure into a steering instrument.

[16] The Basel traffic-light approach counts the breaches of the 99 percent VaR over 250 trading days: up to four the model counts as green, from ten as red, with capital add-ons and model review.

4.5 Limits and Misunderstandings

Three misinterpretations have accompanied the VaR since its invention — and all three have, in the past, contributed to the collapse of financial institutions. The first is semantic: the VaR is no maximum loss but a threshold that is exceeded by design — a 99 percent VaR is breached two to three times in a normal year, and about the size of the breach it says nothing. The second is statistical: the VaR is an estimate with estimation error, whose magnitude the bootstrap of Section 6.2 reveals; without error bars the number suggests a precision it does not have. The third is systemic: when many financial houses react to the same signals with the same models, VaR limits become procyclical — rising volatility forces position reductions everywhere at once, which drives volatility further. Danielsson, Embrechts, Goodhart, Keating, Muennich, Renault, and Shin (2001) raised this endogeneity of risk early against a mechanical model-based regulation.[17]

None of these arguments devalues the metric — they devalue its naive use. The VaR remains the best available communication measure for market risk; but it needs the accompaniment this paper develops: the Expected Shortfall for the tails, backtesting for truthfulness, the bootstrap for precision, and the stress test of Section 6.4 for everything that stands in no sample.

4.6 From Measuring to Managing: Component VaR and Risk Budgets

A total VaR quantifies the risk — steering happens through its decomposition. The marginal VaR measures by how much the portfolio VaR rises when a position is increased slightly; the component VaR weights these marginal contributions with the position sizes and decomposes the total VaR additively across positions, mandates, or risk factors. The decomposition is regularly more instructive than the total: in the example portfolio, the equity position accounts for 60 percent of the capital but — owing to higher volatility and correlation effects — for over 90 percent of the risk contribution; viewed in risk units, the seemingly balanced portfolio is an equity portfolio with a bond admixture.[18]

On this decomposition rests the risk budgeting of private total wealth: the investment committee assigns risk contributions instead of capital quotas to mandates and strategies, and utilisation is measured continuously against the budget — capital follows risk, not the other way round. Thought through to its end, the idea leads to the equalisation of risk contributions

[17] The collapse of Long-Term Capital Management in 1998 counts as the principal documented case: the fund's models priced scenarios as practically impossible that occurred within weeks — model risk beats market risk.

[18] Formally, the decomposition rests on Euler's theorem for homogeneous functions: the total VaR is the sum of the position-weighted marginal VaR contributions.

(risk parity) — for this paper the insight suffices: reading allocations only in capital weights leaves the portfolio incompletely understood.

5 Expected Shortfall: the View behind the Threshold

5.1 From the Quantile to the Expectation of the Tail

The VaR answers how high the loss is at most with probability c — and is silent about how high it becomes when the threshold falls. Exactly this question is answered by the Expected Shortfall: the conditional expectation of the losses in the worst 1 − c cases according to Equation (7):[19]

$$ES_c = E[L \mid L > VaR_c] \tag{7}$$

Equation (7): The Expected Shortfall ES_c is the average of the losses that exceed the VaR — in the image of Figure 2, the centre of gravity of the marked tail area. It answers the board's follow-up question to every VaR report: and if it gets worse — how bad on average?

$$ES_c = \sigma_P \cdot \sqrt{t} \cdot V \cdot \varphi(z_c) \;/\; (1 - c) \tag{8}$$

Equation (8): Under normality, the Expected Shortfall also possesses a closed form; φ denotes the density of the standard normal distribution. For the example portfolio: 2.35 million euros at the 95 percent level over one year — 1.25 times the corresponding VaR — and 606,000 euros at 99 percent over ten days. The heavier the tails, the wider this gap opens: the ratio of ES to VaR is a built-in tail indicator.

5.2 Coherence: Why the Expected Shortfall Satisfies the Better Axiom

The deeper reason for the rise of the Expected Shortfall is conceptual. Artzner, Delbaen, Eber, and Heath (1999) asked axiomatically which properties a reasonable risk measure must have, and identified subadditivity as the critical one: the risk of a combined portfolio must not exceed the sum of the individual risks — diversification must never be penalised computationally. The VaR violates precisely this axiom in relevant cases: for strongly asymmetric positions — say portfolios of a few default-prone bonds or sold options — the VaR of two merged books can lie above the sum of the individual VaRs and thus set false incentives against spreading. The Expected Shortfall, by contrast, is provably coherent (Acerbi and Tasche, 2002) and moreover well-behaved for portfolio optimisation: Rockafellar and Uryasev

---

[19] Common synonyms are Conditional Value at Risk (CVaR), Average Value at Risk, and Tail Loss; meant is always the expectation of the losses beyond the VaR.

(2000) show that CVaR minimisation can be solved as a linear programme — risk steering and optimisation speak the same language.[20]

Supervision has drawn the consequence: with the fundamental review of the trading book (FRTB), the Basel Committee replaced the 99 percent VaR with the Expected Shortfall at the 97.5 percent level as the standard measure of market risk (Basel Committee on Banking Supervision, 2019). For private mandates the implication is practical: current risk reports should present both quantities — the VaR for comparability and communication, the Expected Shortfall for the tails and for steering.

5.3 Expected Shortfall in Practice: Estimation and Review

The progress of the Expected Shortfall has a price that practice must know: it is harder to estimate and harder to test than the VaR. Harder to estimate, because it averages the rarest observations — in a window of 250 days, the 97.5 percent ES rests on a good six values, and every single extreme observation moves the metric palpably; the bootstrap confidence intervals come out correspondingly wider than for the VaR. Harder to test, because the ES lacks the property of elicitability (Gneiting, 2011): there is no hit counter as simple as in the VaR backtesting of Section 4.4 — the customary procedures test quantile and tail average jointly.[21]

For everyday reporting, a proven division of labour follows: the VaR remains the limit and backtesting quantity, the Expected Shortfall the steering and capital quantity — this is also how banking regulation holds it, using the ES for capital measurement while continuing to hang backtesting on the VaR quantile. The simultaneous disclosure of both quantities including their estimation uncertainty combines the strengths of both measures without concealing their respective weaknesses.

6 Method Families: Parametric, Historical, Bootstrap, Monte Carlo

6.1 Historical Simulation

Historical simulation is the most widespread nonparametric method and captivates through its openness: it assumes no distribution but values today's portfolio with the actual market moves of the past. From 250 trading days arise 250 hypothetical daily results; their empirical

---

[20] A risk measure is called coherent if it satisfies monotonicity, translation invariance, positive homogeneity, and subadditivity (Artzner, Delbaen, Eber, and Heath, 1999); the VaR violates subadditivity in general.

[21] Gneiting (2011) showed that the Expected Shortfall is not elicitable, i.e. permits no point-wise forecast verification of its own; practicable backtests therefore test it jointly with the underlying VaR quantile.

99 percent quantile — in practice the second- to third-worst day — is the historical VaR. For the example portfolio, a simulated trading year with realistically heavy tails delivers a one-day VaR of 154,000 euros against 167,000 euros from the normal-distribution calculation; the worst days of the window (178,000, 159,000, and 157,000 euros) show at the same time how thin the data base becomes at the tail — the 99 percent quantile hangs on a handful of observations. Figure 3 shows the loss distribution of the window.[22]

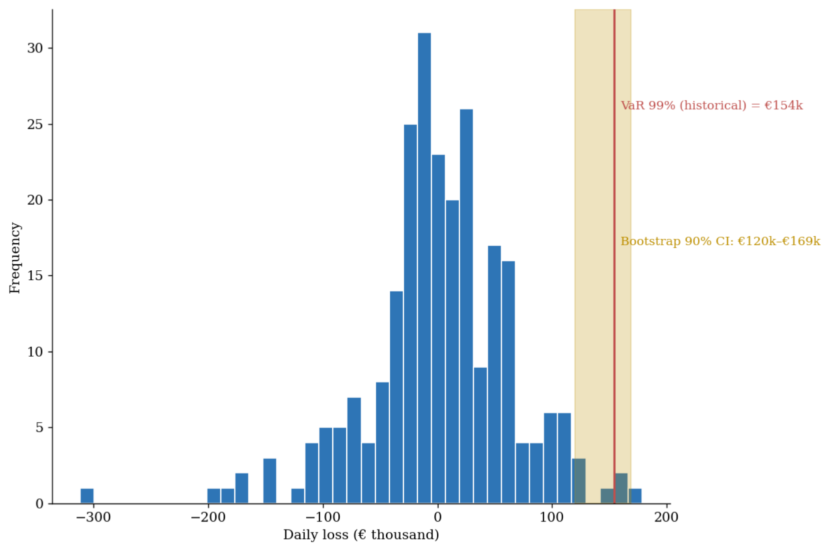

Figure 3: Historical simulation: distribution of 250 hypothetical daily losses of the example portfolio with historical 99 percent VaR and bootstrap confidence interval.
Source: Own presentation and calculation.

The strengths of the method are its freedom from models — skewness, kurtosis, and correlation breaks of the sample enter automatically — and its explainability: every VaR value can be traced back to concrete historical days. Its weakness is the rear-view mirror: what did not occur in the past does not exist for the model, and after calm years the historical VaR is systematically too low — the procyclical basic problem of every data-driven risk measurement.

6.2 Bootstrap: the Precision of the Estimate

How reliable is a VaR estimated at the tail of a sample of 250 days? The bootstrap answers this question with a simple device: from the observed returns, thousands of artificial samples of equal length are generated by random drawing with replacement, and in each the metric is

---

[22] The choice of the observation window is the central control variable: short windows react quickly, long ones smooth; practice frequently calculates with 250 to 500 trading days and adds a stressed window from a crisis period.

recalculated — the dispersion of these resampling results estimates the uncertainty of the original estimate (Efron, 1979). For the historical VaR of the example portfolio, 2,000 bootstrap samples yield a 90 percent confidence interval of 120,000 to 169,000 euros — the seemingly precise point figure of 154,000 carries a range of almost a third of its value.[23]

This openness is the real benefit of the bootstrap in risk management: it makes the estimation uncertainty of the risk figures themselves visible and prevents the most dangerous illusion of reporting — the confusion of decimal places with accuracy. Professional risk reports therefore increasingly state quantile estimates with confidence intervals, and the backtesting of Section 4.4 verifies whether the ranges hold.

6.3 Monte Carlo Simulation and the Method Comparison

Monte Carlo simulation inverts the historical logic: instead of repeating the past, it specifies a stochastic model of the risk factors — say geometric Brownian motions for equities and mean-reversion processes for interest rates — and generates from it tens of thousands of future paths on which the portfolio is revalued. From the loss distribution thus generated, VaR, Expected Shortfall, and every other metric can be read off. Its strengths are flexibility and orientation to the future: it values nonlinear positions (options, structured products) correctly, can contain scenarios that history does not know, and is the only method that delivers path-dependent quantities such as the exposure profiles of Section 10. Its price is model risk — the results are only as good as the assumed processes and correlations — and computational effort. Table 4 juxtaposes the method families.[24]

[23] The method goes back to Efron (1979); for time series with volatility clusters, the block bootstrap is used, which draws contiguous sections instead of single days.

[24] The method's name stems from the Manhattan Project; it was introduced into finance by Boyle (1977) for the valuation of options; the overview of methods and variance reduction is given by Glasserman (2004).

Table 4: Methods of risk quantification compared

Criterion	Parametric	Historical simulation	Bootstrap	Monte Carlo
Distributional assumption	explicit (mostly normal)	none (empirical)	none (resampling)	explicit (processes)
Heavy tails	only via corrections (Equation 6)	from the data	from the data	modellable
Nonlinear positions	approximated only	fully	fully	fully
Future scenarios	via parameters	no	limited	yes
Estimation uncertainty	analytical	high at the tail	measured itself	simulation error controllable
Effort	minimal	low	medium	high
Typical use	linear portfolios, limits	trading books, reporting	confidence intervals	derivatives, exposure profiles, total wealth

Source: Own compilation based on the data of Jorion (2007) and Glasserman (2004).

Table 4 makes clear that the methods are not competitors but instruments for different purposes: the parametric approach suits limit systems and approximate estimates, historical simulation the daily risk reporting, the bootstrap the quantification of estimation uncertainty, and Monte Carlo simulation the valuation of nonlinear or path-dependent positions. Large financial institutions therefore run all four methods in parallel — and interpret divergences between the methods not as a nuisance but as a source of information about model risk.

6.4 Beyond Distributions: Stress Tests and Scenario Analysis

All the metrics developed so far share a silent assumption: the future is a draw from an estimable distribution. Stress tests deliberately abandon this assumption and ask deterministically: what does a specific scenario cost — October 1987, autumn 2008, the rate shock of 2022, a hypothetical combination of equity slump, spread widening, and currency dislocation? The answer is not a probability statement but an impact statement — and precisely herein lies its analytical value: it retains its validity where distribution-based methods reach their limits, namely at structural breaks, newly introduced instruments, and the breakdown of historical correlation patterns. In supervisory law, stress tests are firmly anchored; methodologically, they form the necessary corrective to the quantile measures' implicit trust in the statistical extendability of the past.[25]

Good stress programmes combine three classes: historical scenarios (real crises applied to today's portfolio — the supreme discipline of historical simulation), hypothetical scenarios (combinations designed by the risk committee that history does not know), and sensitivity

[25] The reverse stress test turns the question around: instead of calculating the effect of a scenario, it seeks the scenario that would hit the institution existentially — and examines how plausible it is.

shocks (single-factor moves such as plus 200 basis points in rates — the bridge to the Greeks of Section 9). The methodological challenge lies in scenario selection: a stress test that considers only moderate scenarios performs a reassuring, not an insight-generating function. A complete programme therefore includes the reverse stress test — the systematic identification of the scenario that exceeds the bearing capacity of one's own wealth.

6.5 Time-Varying Volatility: EWMA and the GARCH Family

The volatility clusters documented in Section 3.4 require models that represent risk as a time-varying quantity. The simplest approach is exponentially weighted smoothing (EWMA): more recent observations receive the highest weight, while older ones decay geometrically. The volatility estimate thereby adapts to a regime change within days, where an equally weighted 250-day window needs months. The GARCH family (Engle, 1982; Bollerslev, 1986) formalises the same idea as a process with reversion to the long-run level and additionally delivers multi-day forecasts that replace the naive square-root-of-time rule. For VaR practice the effect is considerable: the same historical sample delivers, with time-varying volatility, a risk measure that signals crises earlier and does not burden calm phases with crisis capital.[26]

The price is the procyclicality of Section 4.5 in sharpened form: the more reactive the model, the more strongly limits and capital requirements fluctuate with the market. Regulation counters this with stressed calibrations — the Expected Shortfall of the trading book rules is calculated over a historical stress period, not over the friendly recent window — and well-run financial houses combine reactive steering quantities with stable capital quantities.

6.6 Extreme Value Theory: Statistics for the Outermost Tails

For confidence levels beyond 99 percent — the world of solvency calculation at 99.5 percent — every empirical sample becomes thin: the sought quantile lies outside the observations. Extreme value theory closes this gap with a remarkable result: the distribution of exceedances over a high threshold converges — almost independently of the initial distribution — to a known parametric family whose shape parameter measures the heaviness of the tail. The outermost quantiles can thus be extrapolated from the middle tail observations instead of being guessed from two or three extreme days (McNeil, Frey, and Embrechts, 2015).[27]

Within the apparatus of this paper, extreme value theory is the specialist method for solvency levels, reinsurance, and catastrophe questions — superior where the methods of

[26] RiskMetrics (J.P. Morgan, 1996) uses exponential smoothing with a decay factor of 0.94 for daily data; the GARCH family goes back to Engle (1982) and Bollerslev (1986).

[27] The peaks-over-threshold approach models all exceedances of a high threshold with the generalised Pareto distribution; the standard reference for the financial application is McNeil, Frey, and Embrechts (2015).

Section 6 reach their data limits, and oversized for daily reporting. Its central practical insight is qualitative in nature: the estimated shape parameter of real financial series regularly lies in the heavy-tail region — the stylised facts thus find their confirmation with the instruments of mathematical statistics.

6.7 Craft Questions of Simulation: Path Count, Convergence, Variance Reduction

Monte Carlo results carry a simulation error that falls with the square root of the path count: one additional decimal place of accuracy requires a hundredfold of computational work. For expectations such as the exposure profile of Section 10, ten thousand paths usually suffice comfortably; for high quantiles — the 99.5 percent tail of the solvency calculation — only a small fraction of the paths is decisive, and naive simulation wastes almost all its work on the irrelevant middle of the distribution. The variance reduction techniques of the simulation literature (Glasserman, 2004) start exactly here, first and foremost importance sampling, which shifts the sampling density into the tail and reweights the results afterwards.[28]

For the application, two methodological principles are decisive. First, the simulation error is an integral part of the result: a simulated metric without its standard error does not constitute a completed calculation; convergence is to be checked by doubling the number of simulation paths and monitoring the stability of the target quantity. Second, reproducibility takes precedence over elegance: only fixed random seeds and documented model versions make a simulation an auditable body of calculation; without them, every result remains a non-replicable one-off.

7 From Market Value to Cash Flow: CFaR, EaR, and the Risk Score

7.1 Cash-Flow at Risk and Earnings at Risk

The VaR measures value changes of market positions — the right approach for the portfolio view, frequently the wrong one for family entrepreneurs and distributing foundations: a family-owned company or a foundation with promised payment obligations does not fail on a bad valuation day but on missing cash flows. Cash-Flow at Risk (CFaR) and Earnings at Risk (EaR) therefore transfer the quantile logic to operating quantities: they measure the maximum

[28] Common variance reduction techniques are antithetic paths, control variates, and importance sampling, which steers the draws deliberately into the risk-relevant tail regions (Glasserman, 2004).

deviation of the planned cash flow or result from the plan figure that is not exceeded within the planning horizon with a specified probability — driven by sales volumes, commodity prices, exchange rates, and interest rates. Equation (9) shows the parametric form:[29]

$$CFaR_c = z_c \cdot \sigma_{CF} \tag{9}$$

Equation (9): The Cash-Flow at Risk at level c is z_c times the dispersion σ_{CF} of the operating cash flow around the plan. An example: with a planned operating cash flow of 50 million euros and a dispersion of 8 million euros aggregated from the volatilities and sensitivities of the drivers, the CFaR at the 95 percent level amounts to 13.2 million euros — with 95 percent confidence, at least 36.8 million euros will thus flow. This lower bound, not the plan figure, is the reliable quantity for investment and distribution decisions.

In practice, the distribution of the cash flow is rarely generated parametrically but by Monte Carlo simulation of the drivers — Section 6.3 supplies the methodology, including the correlations between volumes, prices, and currencies. Earnings at Risk follows the same logic at the level of the reported result and is, for listed companies, the bridge to capital market communication: it quantifies the probability with which the earnings guidance holds — and thereby connects risk management with expectations management.

7.2 The Risk Score: Probability times Severity

Not every risk carries a market price distribution. Operational risks, legal risks, project risks — for them risk management works with the simplest of all condensations, the risk score according to Equation (10):[30]

$$\mathit{Risk\ score} = p \cdot \mathit{severity} \tag{10}$$

Equation (10): The risk score aggregates probability of occurrence p and severity into the expected value of the loss — an event with a five percent annual probability and ten million euros of damage carries a risk score of 500,000 euros per year. The quantity is suitable for ranking and for budgeting countermeasures; as an expected value, however, it conceals precisely the dispersion — two risks with the same risk score can differ by orders of magnitude in the extreme loss.

The risk score is thus the link between quantitative and enterprise-wide risk management: ordered in the risk matrix by probability and impact, it makes heterogeneous risks comparable and prioritisable. For the large, rare losses, however, the same lesson applies as with the VaR:

[29] For industrial companies, the term CorporateMetrics framework has become established; the horizons, at one to five years, are considerably longer than in market risk, the distributions correspondingly wider.

[30] The risk matrix orders risks by probability of occurrence and severity in a grid; it is the standard instrument of enterprise-wide risk management under the relevant governance standards.

the expected value is the beginning of the analysis, not its end; steering by risk scores alone addresses the middle of the distribution and overlooks its tails.

7.3 The Cash Flow as Bottleneck: Liquidity Risk

The cash flow perspective of the CFaR and EaR quantities leads to the most underestimated risk of large private fortunes: liquidity risk. It occurs in two forms. Market liquidity concerns the question of the discount at which positions can be sold when selling is unavoidable — in stress phases, bid-ask spreads widen exactly when everyone wants through the same door. Funding liquidity concerns the question of whether payments falling due — margin calls on derivatives, capital calls, redemptions — can be met at all times. The two reinforce each other in a crisis: meeting margin calls through forced sales depresses prices, which triggers further margin calls.[31]

The response to this is cash-flow-based: a liquidity maturity ladder across all due dates and contingent obligations, stress scenarios for waves of margin calls and redemptions, and a liquidity reserve sized on the worst plausible month — not on the average one. The Cash-Flow at Risk of Equation (9) supplies the distributional basis; the total-wealth application of Section 12 shows what happens when this calculation is missing.

7.4 From Measuring to Acting: CFaR as the Basis of Hedging Programmes

The real purpose of the CFaR calculation is the sizing of hedging programmes. The logic: from the driver decomposition of the cash flow risk — how much of the dispersion stems from currencies, how much from commodities, rates, volumes? — follows which hedge lowers the CFaR most strongly per euro deployed; the programme is then dimensioned not by intuition but to the target floor of the cash flow. An example in the logic of Section 7.1: if hedging the dollar position lowers the cash flow dispersion from 8 to 5 million euros, the 95 percent floor rises by almost five million — a measurable risk relief that can be set against the hedging costs.[32]

This sizing logic disciplines in both directions: it prevents the under-hedging of existential risks as well as the popular over-hedging — the complete hedging away of risks whose bearing would be rewarded and which diversify in the portfolio context. What threatens solvency is

[31] The lesson of the recent crises — from 2008 through March 2020 to the margin crisis of the family office Archegos in 2021 — is uniform: institutions fail on the cash flow before they fail on the wealth.

[32] The hedging instruments themselves — forwards, options, swaps — are developed by Rupprechter (2026h); here the sizing logic stands in the foreground.

hedged; what the business model rewards is carried — the CFaR calculation makes this boundary visible and negotiable.

8 Lower Partial Moments: Risk as Falling Short of a Target

8.1 From the Safety-First Principle to the LPM Family

Volatility and VaR treat fluctuation symmetrically or in purely distributional terms — but investors perceive risk asymmetrically: as the danger of missing a target. This intuition has a long theoretical lineage: Roy (1952) formulated, with the safety-first principle, the maximisation of safety against falling below a catastrophe threshold, and Markowitz (1959) himself considered the semivariance the conceptually more appropriate measure. Bawa (1975) and Fishburn (1977) generalised both into the family of Lower Partial Moments (LPM), and Harlow (1991) introduced them into portfolio practice. Equation (11) defines the family:[33]

$$LPM_n(\tau) = E[\max(\tau - r, 0)^n] \tag{11}$$

Equation (11): The Lower Partial Moment of order n measures the deviations of the return r below the target return τ — and only these. The order n steers the risk aversion: LPM_0 is the probability of missing the target (shortfall probability), LPM_1 the expected shortfall amount, and LPM_2 the semivariance generalisation, whose root is reported as the downside deviation. Gains above τ do not — unlike with volatility — enter the risk measure.

$$Sortino = (\mu - \tau) \,/\, \sqrt{LPM_2(\tau)} \tag{12}$$

Equation (12): The Sortino ratio sets the excess return over the target in relation to the downside deviation (Sortino and van der Meer, 1991) — the LPM counterpart to the Sharpe ratio (Sharpe, 1966), which fairly compares strategies with asymmetric return profiles: upside fluctuation is not counted as risk; downside fluctuation is not rewarded as a source of return.

8.2 Worked Example and Interpretation

A compact example shows the calculation. Six equally likely annual return scenarios of a portfolio — −18, −6, +2, +7, +12, and +21 percent — are measured against a target return of $\tau = 2$ percent; Figure 4 shows the scenarios. Two scenarios miss the target: LPM_0 amounts to 0.333 — in a third of the cases τ is missed. The expected shortfall amount LPM_1 is 4.67

[33] The target return τ is freely selectable: zero (nominal capital preservation), the inflation rate (real capital preservation), the distribution rate of a foundation, or a benchmark return — the LPM family adapts to the investor's goal, not the other way round.

percentage points, and from LPM_2 follows a downside deviation of 8.8 percent — well below the ordinary volatility of the scenarios of 12.5 percent, because the upside moves do not count. With an average return of 3 percent, Equation (12) yields a Sortino ratio of 0.11 — the portfolio earns its excess over the target with considerable downside risk.[34]

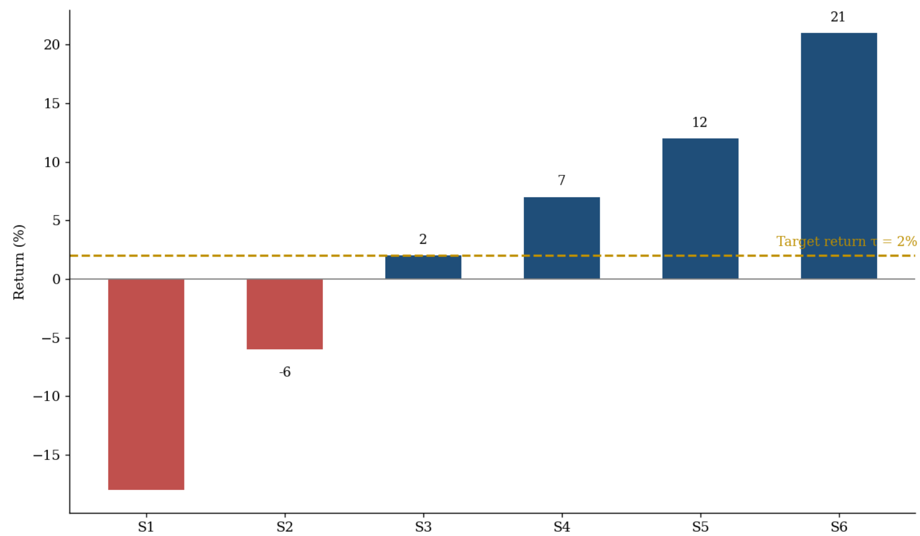

Figure 4: Six annual return scenarios of the LPM worked example against the target return of two percent; only the shortfalls enter the Lower Partial Moments.
Source: Own presentation and calculation.

Table 5: Lower Partial Moments of the worked example (target return $\tau = 2$ percent, six equally likely scenarios)

Metric	Value	Interpretation
LPM_0 (shortfall probability)	33.3 %	target is missed in a third of the scenarios
LPM_1 (expected shortfall amount)	4.67 %-points	average shortfall across all scenarios
$\sqrt{LPM_2}$ (downside deviation)	8.79 %	dispersion of the shortfalls only; volatility for comparison: 12.54 %
Sortino ratio (Equation 12)	0.11	excess return per unit of downside risk

Source: Own calculation according to Equations (11) and (12).

Table 5 shows the practical appeal of the LPM family for wealthy private investors and their structures: the target return τ makes the risk measure mandate-ready. A foundation sets τ to distribution rate plus inflation and thereby measures exactly the risk that threatens its purpose; a drawdown portfolio sets τ to the planned withdrawal rate, a special fund mandate to its target. The LPMs thus span the arc between the distribution statistics of Sections 3 to

[34] The six-scenario example is deliberately kept small; in practice the LPMs are estimated from historical time series or simulated distributions with thousands of observations.

5 and the obligation logic of Section 12 — they are the downside measure that answers the investor's question, not the statistician's.

8.3 Lower Partial Moments in Asset Allocation

The LPM family is not only a reporting but an optimisation quantity: Harlow (1991) showed how portfolios can be optimised for minimal LPM_2 relative to the target return instead of minimal variance. The difference is conceptually significant: variance optimisation punishes investments with welcome upside dispersion — convertible bonds, trend-following strategies, investments with option character — downside optimisation does not. Mirror-wise, it exposes the strategy whose calm volatility conceals an asymmetric loss profile: premium-selling strategies look attractive in the variance world and expensive in the LPM world — the skewness of Equation (1) here becomes allocation information.[35]

The price of the method is statistical: LPM estimates rest only on the observations below the target return and are correspondingly noisier than variance estimates; the bootstrap intervals of Section 6.2 come out wide, and short histories carry no reliable downside optimisation. Practice counters this with simulated distributions from the Monte Carlo repertoire of Section 6.3 — the downside objective remains, the data base is broadened with model support.

8.4 Quartiles in Practice: the Peer Comparison and Its Pitfalls

The quantile language of Section 3.3 meets the private investor most frequently in the context of peer group comparisons: fund universes are ordered by return or a risk-adjusted metric and divided into quartiles, with the first quartile counting as a mark of quality and the fourth as grounds for terminating the mandate. Three methodological pitfalls, however, counsel restraint in the interpretation. First, survivorship bias: a ranking restricted to surviving funds rests on a distorted population. Second, the group delineation: quartile ranks are only as meaningful as the homogeneity of the peer group — a defensive mixed fund in an aggressive universe harvests structural fourth-quartile ranks without any statement about management. Third, the lack of persistence: quartile ranks migrate strongly between periods, in line with the chance findings of Fama and French (2010) — yesterday's first quartile is no predictor for tomorrow.[36]

Used correctly, the quartile comparison remains useful nonetheless: over long windows, against cleanly delineated groups, on risk-adjusted quantities such as the Sortino ratio of

---

[35] Harlow (1991) shows the effect on historical allocations: LPM-optimised portfolios shift weights in favour of investments with asymmetrically favourable profiles without giving up the target return.

[36] Survivorship bias arises because closed and merged funds disappear from the databases; since predominantly weak products close, the remaining ranks overstate the average quality.

Equation (12) instead of raw returns — and always as the starting point for the attribution analysis of Section 11, never as its substitute. The rank says that something happened; the attribution says what.

9 Sensitivity Measures: the Greeks in the Portfolio Context

9.1 Delta, Gamma, Vega, and the Taylor Logic

For portfolios with derivatives, distribution metrics do not suffice, for their risk changes with the market level itself. The sensitivity measures — the Greeks — therefore measure the reaction of the portfolio value to marginal changes in individual market parameters: the delta to the price of the underlying, the gamma, as the rate of change of the delta, the curvature of the value profile, the vega to implied volatility; theta (time decay) and rho (interest rate) complete the set. Equation (13) bundles them in the Taylor approximation of the value change:[37]

$$dV \approx \Delta \cdot dS + \tfrac{1}{2} \cdot \Gamma \cdot (dS)^2 + Vega \cdot d\sigma \tag{13}$$

Equation (13): The value change dV of a derivatives-bearing portfolio decomposes into the linear delta term, the quadratic gamma term, and the volatility term. The equation is the hinge between the sensitivity and the distribution world: in the parametric VaR calculation of Section 4 it replaces the revaluation (delta-gamma VaR), and in the Monte Carlo simulation of Section 6.3 it accelerates the portfolio revaluation across tens of thousands of paths.

In the risk report of sophisticated private mandates, the Greeks have two roles. First, as limit quantities: delta, gamma, and vega limits bound the risk-taking of overlay and derivatives mandates where distribution measures would be too sluggish — sensitivities react immediately, quantiles only with the statistics. Second, as explanatory quantities: the daily result decomposition according to Equation (13) (P&L attribution to risk factors) reveals whether a result stems from the intended position or from an unintended side risk — the microscopic sister of the performance attribution of Section 11. A portfolio whose vega position is unknown to the wealth management has no volatility opinion — it has a volatility bet.

9.2 Sensitivities beyond Derivatives: Duration, Spread Duration, Beta

The logic of the Greeks carries far beyond options — every asset class has its sensitivity measures, and good risk reports speak this language throughout. For bonds, the duration is

[37] The detailed development of the option Greeks including delta hedging of a dealer book is given by Rupprechter (2026h); here the risk reporting perspective stands in the foreground.

the delta with respect to the interest rate and the convexity the gamma; the spread duration measures the sensitivity to credit spreads and thus separates rate risk from credit risk. For equity portfolios, the beta takes the delta role with respect to the market, supplemented by the factor loadings of Section 11.2. The benefit of the uniform language is aggregability: sensitivities can be condensed across asset classes into total positions per risk factor — the house then knows not only what each book holds but what the institution as a whole has positioned against rates, spreads, the equity market, and volatility.[38]

This factor view is at the same time the gateway to the parametric risk calculation: describing a portfolio as a vector of sensitivities to a few dozen risk factors permits calculating VaR and Expected Shortfall from the covariance matrix of these factors — the path that RiskMetrics (J.P. Morgan, 1996) industrialised and that the commercial risk models follow to this day. Sensitivities are thus not the counterpart to the distribution world of Sections 3 to 6, but its vocabulary.

9.3 Sensitivities in Execution: the Overlay Practice

In institutional practice, the aggregated sensitivities of Section 9.2 are steered by overlay programmes: a rates overlay keeps the duration of the total wealth within the target band of the investment guidelines, a currency overlay the open foreign currency quotas, an equity overlay the beta — each with the liquid futures instruments whose mechanics Rupprechter (2026h) develops. The appeal of the architecture lies in the division of labour: the underlying mandates optimise selection in their markets, the overlay steers the factor totals centrally — and the sensitivity language connects both levels without loss.[39]

The risk calculation of the overlay itself follows the rules of this paper: futures positions create the margin obligations of Section 7.3 and belong in the liquidity maturity ladder; their effect on VaR and Expected Shortfall is measured via the factor decomposition, and their success control runs through the attribution logic of Section 11 — against the steering mandate, not against the market. An overlay without this embedding is not a steering but a position risk.

[38] The key rate durations refine the picture by disclosing the sensitivity per maturity point of the yield curve — the instrument against curve twists that a single duration number does not capture.

[39] Overlay mandates steer defined risk factors of the total portfolio with derivatives without moving the underlying holdings; their success is measured against the steering mandate, not against a market benchmark.

10 Default Risk in Swap Transactions: Exposure Simulation and CVA

10.1 Why Swaps Carry a Special Credit Risk

The default risk of a loan is simple: the outstanding amount is fixed. The default risk of a swap is not — its replacement value is zero today and tomorrow a random variable that fluctuates with the markets and can change sign. Losses arise only when two things coincide: the counterparty defaults, and the contract happens to have positive value from the surviving partner's point of view. Risk measurement therefore needs the distribution of future contract values over the entire term — a path-dependent quantity that only the stochastic simulation of Section 6.3 can deliver (Duffie and Singleton, 1999; Zhu and Pykhtin, 2007). The core quantities are defined by Equations (14) and (15):[40]

$$EE(t) = E[max(V_t,\ 0)] \tag{14}$$

Equation (14): The expected exposure EE(t) is the expectation of the positive part of the contract value V_t at the future time t — only positive values are lost in default. In addition, the Potential Future Exposure (PFE) measures the high quantile of the same distribution — the VaR idea of Section 4 applied to future replacement values.

$$CVA = LGD \cdot \Sigma\ EE(t_i) \cdot \Delta PD(t_i) \cdot DF(t_i) \tag{15}$$

Equation (15): The credit valuation adjustment CVA is the market value of the expected default loss: loss given default LGD times expected exposure times marginal default probability, discounted over the term and summed. It makes counterparty risk priceable — as a deduction from the contract value and as a cost component of every uncollateralised trade.

10.2 Worked Example: Monte Carlo CVA of an Interest Rate Swap

The worked example: a family's special fund holds a payer interest rate swap of 100 million euros notional with a ten-year term against an uncollateralised counterparty with a CDS spread of 120 basis points. The Monte Carlo simulation generates ten thousand interest rate paths, revalues the swap on every path at every grid point, and averages the positive values into the exposure profile of Figure 5: the expected exposure rises with interest rate uncertainty, falls with the melting residual term, and reaches its peak of around 3.5 million euros — three

[40] After the reforms since 2009, standardised swaps are predominantly centrally cleared and bilateral trades collateralised; the exposure thinking remains fundamental for residual risks, collateral gaps, and pricing.

and a half percent of the notional — after about three and a half years; the 95 percent PFE reaches 14.5 percent at its peak. According to Equation (15), a CVA of around 0.24 million euros or 24 basis points of the notional results — the fair price of the default risk of this one trade.[41]

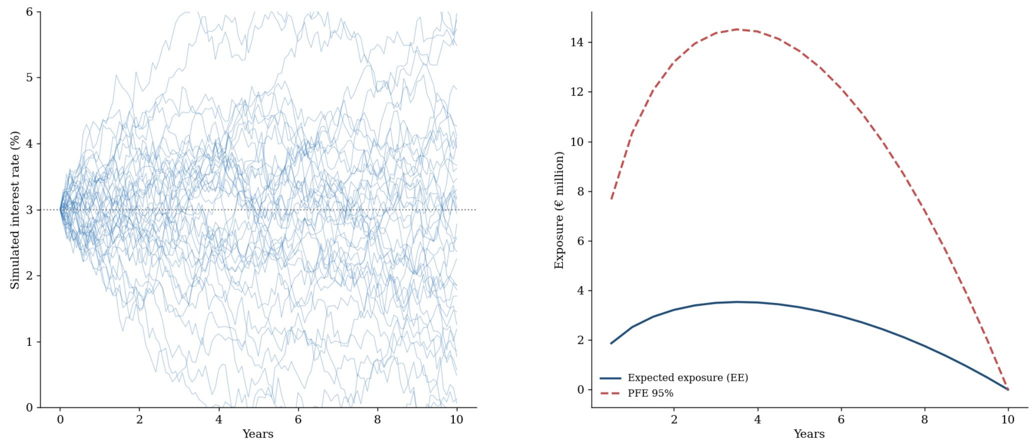

Figure 5: Monte Carlo simulation of the interest rate swap example: simulated interest rate paths (left) and the profile of expected exposure and potential future exposure obtained from them (right).
Source: Own presentations and calculations.

Table 6: Counterparty risk of the interest rate swap example (notional 100 million euros, term ten years, uncollateralised)

Quantity	Value	Explanation
Interest rate volatility	0.80 % p.a.	dispersion of the simulated swap rate
EE peak (Equation 14)	€3.5m (3.5 %)	after around 3.5 years; thereafter the residual term dominates
Average EE	€2.5m	average over the term
PFE 95 %	€14.5m (14.5 %)	high quantile of the future replacement value
Counterparty CDS spread	120 bp	implies around 2 % annual default probability at LGD 60 %
CVA (Equation 15)	€0.24m (24 bp)	market value of the expected default loss

Source: Own calculation; Monte Carlo simulation with 10,000 paths according to Equations (14) and (15).

Table 6 and Figure 5 show what the simulation calculation contributes to risk management: the humped exposure profile — first growing uncertainty, then a melting residual term — cannot be captured by any static metric, yet it determines collateral needs, limits,

[41] The example uses a mean-reversion interest rate model in the spirit of Vasicek (1977) with 80 basis points of rate volatility; the default probabilities are derived from a CDS spread of 120 basis points at a 60 percent loss given default.

and the price of the trade. The same architecture carries the alternatives: historical simulation replaces the model paths with historical rate changes and answers what the portfolio would have cost in the actual crises; the bootstrap of Section 6.2 places confidence intervals around the EE profile and the CVA and shows how strongly the price quantity hangs on the data base. Only the triad of methods turns the single CVA number into a reliable basis for decision.

10.3 Collateral, Wrong-Way Risk, and the xVA Family

Two refinements complete the counterparty calculation. The first is collateralisation: with daily collateral exchange under the market standard of bilateral annexes, the expected exposure of Equation (14) shrinks to the movement of a few days between the last exchange and liquidation — the CVA of the worked example would fall from 24 basis points to a fraction. This explains the architecture of post-crisis regulation: collateralisation and central clearing replace credit risk with liquidity risk — the margin calls of Section 7.3 are the price of the low CVA.[42]

The second is wrong-way risk: the unfavourable dependence between the size of the exposure and the default probability of the counterparty. A swap held with a bank whose value rises exactly in those scenarios in which the bank deteriorates has an exposure profile that the independence assumption of Equation (15) systematically understates — the Monte Carlo architecture permits building in this dependence via the joint simulation of market factors and credit quality (Zhu and Pykhtin, 2007). The insight is simpler than the calculation: avoid counterparties whose credit quality correlates with one's own claim — the principle that the financial crisis taught exemplarily in the case of the monoline insurers.

10.4 From Single Trade to Portfolio: Netting and Credit Portfolio Models

Counterparty risk is steered not per trade but per counterparty: within a netting master agreement, positive and negative contract values offset into one balance, and the expected exposure of Equation (14) is simulated over the entire netting set — a new swap can lower the exposure of an existing relationship if it works in the opposite direction. Only this portfolio view makes limits and pricing consistent: the same trade is risk-increasing with one counterparty and risk-reducing with another.[43]

Above the counterparty lies the credit portfolio as a whole. CreditMetrics (Gupton, Finger, and Bhatia, 1997) showed how the distribution logic of this paper can be transferred to default

---

[42] The CVA has been joined by further valuation adjustments — for one's own default probability (DVA), funding costs (FVA), and capital commitment (KVA) — bundled under the collective term xVA.

[43] CreditMetrics (Gupton, Finger, and Bhatia, 1997) transferred the VaR methodology to credit portfolios and became the blueprint of bank-internal credit risk modelling.

and migration risks: rating transitions are simulated, portfolio losses aggregated, and from the loss distribution follow credit VaR and economic capital. The load-bearing quantity is here, too, the correlation — joint defaults determine the tail of the distribution, as the tranche logic of Rupprechter (2026h) illustrates. With this, the methodological circle closes: market, credit, and counterparty risk speak the same distribution language and differ in the risk factors, not in the method.

10.5 Central Counterparties from the Investor's Perspective

For cleared trades, the counterparty calculation shifts fundamentally: the central counterparty with daily collateralisation takes the place of the bilateral exposure profile — the CVA of Section 10.2 shrinks to residual quantities, and in its place come new, liquidity-shaped metrics: the initial margin tied up, the fluctuation of the daily variation margins, and the concentration risk vis-à-vis the clearing house itself. The risk report of a derivatives-intensive house therefore today discloses margin metrics where CVA figures used to stand — the same shift from credit to liquidity risk that runs through the entire post-crisis architecture.[44]

The stress calculation follows: simulated is no longer only the market value but the margin requirement under extreme moves — including the procyclical margin increases of the clearing houses in crises. The Archegos episode of Section 12.1 is the case study of this calculation; its absence was not a model gap but a programme gap.

11 Performance Attribution: Brinson Decomposition and the Barra Factor Framework

11.1 The Brinson Decomposition: Allocation, Selection, Interaction

Risk-taking without systematic outcome control remains without insight — performance attribution closes the steering loop by decomposing the active return of a portfolio relative to its benchmark into the contributions of the individual investment decisions. The classic framework of Brinson, Hood, and Beebower (1986) and Brinson and Fachler (1985) distinguishes, per segment, the weighting decision from the selection decision according to Equations (16) and (17):[45]

[44] The margin logic of central counterparties — initial and variation margin, default fund — is treated in detail by Rupprechter (2026h); here the effect on the risk metrics counts.

[45] Brinson, Hood, and Beebower (1986) showed moreover that investment policy — the strategic weighting of the asset classes — explains the overwhelming share of the variation in results of institutional portfolios.

$$A = \Sigma_i\ (w_{P,i} - w_{B,i}) \cdot (r_{B,i} - R_B) \tag{16}$$

Equation (16): The allocation contribution A rewards overweights of segments that performed better than the overall benchmark R_B. In the worked example — portfolio 60/40 in equities/bonds against a 50/50 benchmark, segment returns 10 and 3 percent — the overweight of the better-performing equities contributes 0.70 percentage points.

$$S = \Sigma_i\ w_{B,i} \cdot (r_{P,i} - r_{B,i}) \tag{17}$$

Equation (17): The selection contribution S measures the stock selection within the segments at benchmark weights. In the example, the equity selection delivers two percentage points of excess return in the segment and thus a contribution of 1.00 percentage points; the slightly weaker bond selection costs 0.10 points — in net terms 0.90 points of selection across both segments. With the interaction term of 0.22 points, the contributions sum exactly to the active return of 1.82 percentage points — the decomposition adds up. Figure 6 shows the bridge.

11.2 The Barra Framework: Attribution through Factors

The Brinson decomposition thinks in segments — modern portfolios think in factors. The framework founded by Rosenberg (1974) and operationalised in the Barra models explains every individual return through common factors according to Equation (18):[46]

$$r = \Sigma_k\ \beta_k \cdot f_k + \varepsilon \tag{18}$$

Equation (18): The return r decomposes into factor loadings β_k times factor returns f_k — industries and style factors such as value, size, momentum, and quality (Fama and French, 1993; Carhart, 1997) — plus a specific residual ε. At the portfolio level, the active return thus becomes the sum of active factor positions times factor returns plus stock-specific selection: the same decomposition as with Brinson, only along the dimensions in which risk premiums are actually earned.

$$IR = \alpha \,/\, TE \tag{19}$$

Equation (19): The information ratio condenses the attribution into a steering quantity — active return α per unit of active risk (tracking error TE). Grinold and Kahn (2000) make it the central measure of active management; a value of 0.4 as in the example (1.2 percent alpha at 3 percent tracking error) counts as good in practice.

---

[46] The commercial Barra models go back to the work of Rosenberg (1974); their factors cover industries and styles (value, size, momentum, quality, volatility) and are estimated in the cross-section.

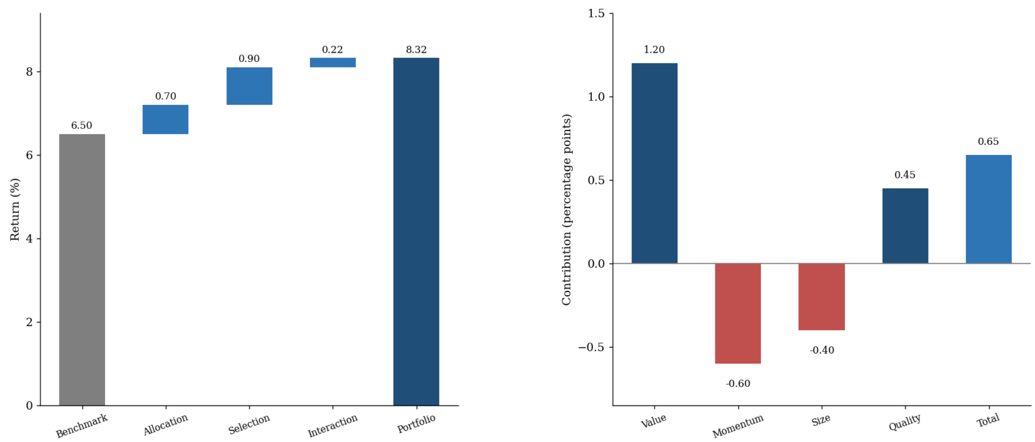

Figure 6: Attribution of the worked example: Brinson bridge from the benchmark to the portfolio return (left) and decomposition of the active return into style factor contributions (right).
Source: Own presentation and calculation.

Table 7: Attribution of the worked example: Brinson decomposition and factor contributions (active return 1.82 percentage points)

Level	Contribution	Value	Interpretation
Brinson	Allocation (Equation 16)	+0.70 %-pts	overweight of the better-performing equities
Brinson	Selection (Equation 17)	+0.90 %-pts	stock selection, predominantly in the equity segment
Brinson	Interaction	+0.22 %-pts	interplay of overweight and outperformance
Factors	Value	+1.20 %-pts	active value loading 0.3 at 4 % factor return
Factors	Momentum	−0.60 %-pts	underloading of a well-performing factor
Factors	Size	−0.40 %-pts	small-cap overweight at negative factor return
Factors	Quality	+0.45 %-pts	quality loading 0.15 at 3 % factor return

Source: Own calculation according to Equations (16) to (18); factor sum 0.65 percentage points, remainder stock-specific.

Table 7 shows the insight gained from the complementary perspectives: the Brinson view attests successful allocation and selection — the factor view reveals that a substantial part of the "selection" was in truth a value position that contributed 1.2 percentage points in the period, partly consumed by momentum and size positions. For mandate steering this difference is decisive: factor premiums can be replicated cheaply, genuine stock-specific selection cannot — active fees should be paid in the knowledge of which of the two is being bought (Grinold and Kahn, 2000). Attribution is thus risk management of manager selection: the same decomposition logic, one level higher.

11.3 Limits of Attribution: Chance, Benchmark, and Data Quality

Attribution answers where a result came from — not whether it was skill. The distinction is statistically demanding: at an information ratio of 0.4 as in the worked example, decades of data are needed to separate skill from luck at customary significance levels, and Fama and French (2010) show for the US fund universe that the distribution of fund alphas after costs is barely distinguishable from pure chance dispersion. For the practice of manager selection, two things follow: restraint towards short winning streaks, and the precedence of process analysis over outcome analysis. The attribution demonstrates whether a manager performs the activity for which they were mandated; about its future earning power it makes no statement.

Added to this are two methodological pitfalls. For one, the benchmark choice influences the result decisively: against an unsuitable comparison measure, a mere style tilt appears as selection skill — the factor perspective of Section 11.2 serves here as the corrective. For another, data quality sets limits to every analysis: attribution on a daily basis requires valuation- and transaction-synchronous data; every gap settles as a residual in the selection component. A large unexplained residual is therefore to be read not as a blemish but as an indicator of insufficient data quality.

11.4 Attribution for Bond Portfolios

For bond mandates, the Brinson logic falls short, for their result drivers are not segments but curve movements. Fixed income attribution therefore decomposes the active return along the risk factors of the asset class: the carry contribution from yield and roll-down; the rates contribution from parallel shift and twist of the curve — measured via the active key rate durations —; the spread contribution from the development of credit spreads, weighted with the active spread duration; the currency contribution; and a selection residual at the security level. It is thus the consistent application of the sensitivity language of Section 9.2 to the result calculation — the same metrics that bounded the risk beforehand explain the result afterwards.[47]

Precisely this mirror-image quality is the mark of quality of every attribution architecture: risk model and attribution model should use the same factors, so that every position taken possesses a measurable result line. If an institution steers its risk by key rate durations but explains its results at the segment level, a translation gap arises between the two levels of presentation in which unintended positions can persist undetected.

[47] The decomposition uses the sensitivity measures of Section 9.2: duration and key rate durations carry the rates contribution, the spread duration the credit contribution.

11.5 Attribution and Incentives: the Compensation Question

Attribution unfolds its strongest effect where it is coupled to compensation — and there it is most dangerous. Rewarding raw excess return also rewards the risk with which it was obtained: a manager who lies behind the benchmark shortly before year-end has, under raw-return incentives, every reason to double the active risk — the option logic of the compensation works like the sold premium strategies of Section 8.3: asymmetrically at the expense of the capital provider. Risk-adjusted assessment quantities — the information ratio according to Equation (19), the Sortino ratio according to Equation (12) — and multi-year periods take the edge off this incentive.[48]

The attribution supplies the control of the source: what is to be remunerated is what was commissioned — stock-specific selection in the active mandate, factor fidelity in the style mandate, hedging quality in the overlay. A compensation system that ignores the attribution lines regularly pays factor premiums at the price of active selection — the most expensive purchasing error of the asset management industry (Grinold and Kahn, 2000).

12 The Total-Wealth Perspective: Withdrawals, Obligations, and Supervisory Benchmarks

12.1 From Portfolio to Surplus: Duration Gap and Immunisation

A large private fortune, too, is a balance sheet: the assets A face obligations L — the promised distributions of a foundation, the family's planned withdrawals, loans and commitments — and both sides react to the same interest rates. What is steered is therefore not the portfolio value but the surplus over these obligations, with all the instruments of this paper applied to the difference. The classic first instrument is the duration gap according to Equation (20):[49]

$$\Delta Surplus \approx -(D_A \cdot A - D_L \cdot L) \cdot \Delta y \tag{20}$$

Equation (20): The surplus change under a rate shift Δy hangs on the duration gap — the difference of the market-value-weighted durations of assets and obligations (Macaulay, 1938; Redington, 1952). An example: a foundation with assets of 10 million euros (duration 8), whose promised future distributions correspond to a present value of 9 million euros with a duration of 15, loses on balance around 550,000 euros of surplus when rates fall by one percentage point

[48] Proven practice are multi-year assessment periods, hurdles on risk-adjusted quantities, and clawbacks — the compensation should mirror the investment horizon of the mandate, not the calendar year.
[49] The immunisation idea goes back to Redington (1952), the duration concept to Macaulay (1938); Leibowitz and Henriksson (1988) as well as Sharpe and Tint (1990) transferred portfolio theory into the surplus framework.

— although its assets gain 800,000 euros in value. Measuring only the asset side shows a gain where the balance sheet carries a risk; extending duration via the interest rate derivatives from Rupprechter (2026h) is the standard tool against this.

On the duration gap build the distribution-based surplus quantities: the surplus at risk transfers the VaR logic of Section 4 to the surplus, the shortfall probability — LPM_0 from Section 8 with the funding ratio as target — measures the danger of underfunding, and the stochastic simulation of Section 6.3 generates the joint paths of capital market and obligations on which both are calculated. Leibowitz and Henriksson (1988) and Sharpe and Tint (1990) formalised this surplus view: the same investment can be risky in the asset-only frame and risk-reducing in the surplus frame — long bonds are the prime example.

How existential the interplay of surplus steering and liquidity risk is, private wealth practice regularly shows through leverage: securities loans and futures positions create margin obligations at exactly the moment when prices fall, and forced selling into falling markets accelerates the loss. The collapse of the family office Archegos (2021), whose leverage built via total return swaps triggered forced sales of more than ten billion US dollars within days, is the extreme case study. The insight is that of Section 7.3: every levered or derivatives-based steering needs a liquidity calculation for its margin calls, stressed on market moves that the model considers nearly impossible.

12.2 The Regulatory Anchor: Basel and Solvency II

Wealthy private investors are subject to no solvency supervision — yet the instruments of this paper are at the same time the language of the supervision of banks and insurers, and their regimes set the standards against which a good private mandate can measure itself: in the regulated sector, risk measurement is not a choice but an own-funds requirement. Table 8 orders the core quantities: banks have calculated their market risk since the trading book reform as an Expected Shortfall at the 97.5 percent level over stressed periods (Basel Committee on Banking Supervision, 2019), insurers their solvency capital under Solvency II as a Value at Risk of basic own funds at the 99.5 percent level over one year — methodologically exactly the quantities of Sections 4 and 5, applied to the total balance sheet. The backtesting of Section 4.4, the counterparty risk quantities of Section 10, and the scenario calculations of the simulation methods are each anchored in regulation.[50]

[50] Added to these are the liquidity metrics of banking regulation and the own-funds requirements for operational risk; the table confines itself to the market and balance-sheet-risk-related core quantities.

Table 8: Regulatory risk metrics at a glance (simplified)

Rulebook	Addressee	Metric	Calibration	Reference in this paper
Basel III / FRTB	banks	Expected Shortfall	97.5 % / stressed periods	Sections 5, 6.5
Basel III	banks	VaR backtesting (traffic light)	99 % / 250 days	Section 4.4
Basel III (SA-CCR/CVA)	banks	counterparty exposure, CVA capital	supervisory factors	Section 10
Solvency II	insurers	solvency capital requirement (SCR)	VaR 99.5 % / 1 year	Sections 4, 12.1
Solvency II (ORSA)	insurers	own risk and solvency assessment	scenario- and simulation-based	Sections 6, 12.1
IORPs/pension funds	occupational pension institutions	funding ratio, stress tests	differs nationally	Sections 8, 12.1

Source: Own compilation based on the data of the Basel Committee on Banking Supervision (2019) and the Solvency II framework directive (2009); as of mid-2026, simplified.

The table moreover explains why the choice of metric is no purely methodological question: confidence level and horizon are regulated, and the capital requirement follows directly from the measured number. This creates the familiar incentive to calibrate models to the requirement rather than to reality — the model risk argument that Danielsson, Embrechts, Goodhart, Keating, Muennich, Renault, and Shin (2001) raised early against a mechanical model faith of regulation. The answers of practice are model validation, stress tests beyond the models, and the diversity of methods of Section 6.

12.3 The Integrated Apparatus: a User's Guide

In the result, the apparatus of risk analysis assembles into a complementary structure with a clear assignment of functions. The distribution moments of Section 3 characterise the empirical data base; the Value at Risk defines the communicable loss threshold, the Expected Shortfall quantifies the losses lying beyond this threshold; the Cornish-Fisher expansion and the simulation methods capture heavy distribution tails and nonlinearities; CFaR and EaR translate the results into the language of financial planning, the risk score into that of enterprise-wide risk management; the Lower Partial Moments anchor the investor's goal, the sensitivity measures the daily steering, the exposure simulation the counterparty risk, and the attribution verifies whether the risks taken were rewarded. For foundations and drawdown fortunes, the duration gap and the surplus quantities are added as the connecting bracket over the asset and liability sides of the wealth.[51]

[51] The assignment is a heuristic of practice, not a norm; decisive is that every metric is reported together with its blind spot.

From this follows the same methodological guiding principle that runs through all sections: no metric carries a decision on its own. The Value at Risk without the Expected Shortfall leaves the distribution tails unlit; the Expected Shortfall without backtesting remains an unsubstantiated assertion; the simulation without the bootstrap conceals its estimation uncertainty; and every statistical metric without a complementary stress test implicitly assumes that the future constitutes a draw from the distribution of the past. Professional risk steering can in this respect be characterised as institutionalised scepticism towards one's preferred single metric.

12.4 The Organisational Anchoring: Limits, Committees, Lines of Defence

Metrics unfold no steering effect of their own accord — steering is done by people, and the organisational framework decides whether the instruments presented in this paper become effective. Three elements have proved load-bearing. First, the limit architecture: a hierarchy that starts from the risk budget of the total wealth and is broken down into VaR, sensitivity, and stress limits per mandate, with defined escalation paths upon approaching and exceeding thresholds — limits without escalation consequences remain ineffective. Second, organisational independence: the risk-measuring unit must not report to the risk-taking one; to the second line of defence fall model responsibility, data sovereignty, and the right of veto. Third, the regularity of governance: an investment committee that does not merely receive the reports but discusses them substantively — with backtesting results, model deviations, and stress test outcomes as standing agenda items.[52]

The loss cases of financial history have almost never been measurement errors but organisational errors: metrics that no one escalated, limits that were quietly raised, models whose warnings counted as technical malfunctions. The risk analysis apparatus of this paper unfolds its value where its results are taken seriously — the last and most important precondition of professional risk management.

12.5 Reporting: from Number Work to Decision Document

The best risk analysis apparatus fails on a poor report. The proven form is the one page that carries everything that must be decided: the core quantities — VaR and Expected Shortfall with confidence intervals, limit utilisations, the largest risk contributions according to Section 4.6 — alongside the exceptions of the backtesting, the results of the stress scenarios, and three

[52] The model of the three lines of defence — market and portfolio responsibility, independent risk controlling, internal audit — is standard in the relevant governance frameworks.

lines of plain language on what has materially changed since the last report. Everything else is appendix. Decisive is the stability of the format: a risk report unfolds its information content essentially in comparison with the preceding periods; every redesign of the metric system destroys precisely this comparability.

The second rule concerns the honesty of the presentation: every estimate is to be disclosed with its estimation uncertainty, model changes are to be disclosed analogously to changes in accounting policy, and all metrics are to be presented together with their methodological limits — the Value at Risk with the reference to the distribution tails, the simulation with its assumptions. A report that suggests certainty where uncertainty prevails constitutes no information but itself a risk.

13 An Econometric Case Study: Six Decades of the S&P 500 — from the Old Economy to the AI Era

13.1 Research Question, Data, and Study Design

The preceding sections developed the instruments; this section applies them. The research question is: have the New Economy — the market carried by technology and momentum stocks since the turn of the millennium — and the AI era since 2016 measurably changed the statistical properties of S&P 500 returns, and do the more recent market phases produce more extreme values in the econometric tests as in the classic return and risk metrics than the six-decade baseline?

The data base are the monthly data of Robert J. Shiller for the S&P 500 from January 1966 to December 2025 — 720 monthly returns, in logarithms, based on the index levels; in addition, the yield of ten-year US government bonds (GS10) from the same source (Shiller, 2026). Two properties of the data are essential for the interpretation: first, Shiller's index levels are monthly averages of daily prices, which mechanically raises the first-order autocorrelation; second, the returns are pure price returns without dividends.

The study design compares three telescoping — that is, nested — observation windows: the full period January 1966 to December 2025 ($n = 720$) as the baseline, the New Economy period from January 2000 ($n = 312$), and the AI era from January 2016 ($n = 120$). This telescoping arrangement is methodologically intentional: it answers the practically relevant question of whether an investor calibrating today with the long history under- or overestimates the properties of the more recent market phases. To each period the same six-part econometric procedure is applied.

13.2 Econometric Procedures and Interpretive Framework

(1) Unit root test (Dickey and Fuller, 1979): It tests whether price shocks act permanently (unit root, random walk) or whether the price level reverts to a trend. Normal for equity indices is the non-rejection of the unit root — the t-statistic lies above the critical value of -2.86. Conspicuous on the downside would be mean reversion of the level (prices would be forecastable from their own history); conspicuous on the upside an explosive path with ρ well above one.[53]

(2) Autocorrelation test (Ljung and Box, 1978): It tests whether returns depend on their preceding months. Normal are values near zero (weak-form market efficiency); markedly positive coefficients would mean tradable momentum, markedly negative ones systematic overreaction. With time-averaged data, the first lag must be adjusted for the averaging artefact of around $+0.25$.

(3) Heteroskedasticity test (ARCH-LM, Engle, 1982): It tests whether the variance clusters over time. For financial returns a significant result is the normal case, not an anomaly; the absence of ARCH effects would be the truly conspicuous finding and points to strong aggregation or low test power.

(4) t-test and distribution tests (Jarque and Bera, 1980): The t-test examines whether the average return differs from zero — the existence question of the equity premium. Skewness and excess kurtosis measure the deviation from the normal distribution: values of zero are "normal"; negative skewness means that crashes are more violent than upswings, positive excess kurtosis that extremes occur more often than the bell curve permits (fat tails; cf. Mandelbrot, 1963; Fama, 1965).

(5) Interest rate dependence: Tested is whether equity returns differ at high interest rate levels (median split, F- and t-test) and how they correlate with the rate change ΔGS10. Theoretically expected, via the discount rate channel, is a negative reaction to rate rises; for the rate level itself no stable relationship is to be expected. A sign change of the correlation between periods indicates a regime change in the equity-bond nexus.

(6) Classic metrics: Return p.a., volatility p.a., their quotient, best and worst month, and the maximum drawdown anchor the test results in the quantities of investment practice.

13.3 Results across Periods

Table 9 juxtaposes all metrics of the three periods; p-values are in brackets. All calculations are reproducible on a formula basis in an accompanying Excel workbook.

[53] The t-statistic of the Dickey-Fuller test follows no t-distribution under the null hypothesis; the critical values used (-3.43/-2.86/-2.57 for the model with constant) stem from the response surface estimates of MacKinnon (1996).

Table 9: Econometric metrics and classic return-risk measures of the S&P 500 — full period, New Economy, and AI era

Metric	1966–2025	2000–2025	2016–2025
n (months)	720	312	120
Return p.a. (logarithmic)	7.2 %	6.0 %	12.0 %
Volatility p.a.	12.6 %	13.2 %	12.6 %
Return/volatility	0.57	0.46	0.96
Best month	+11.4 %	+11.4 %	+7.9 %
Worst month	−22.8 %	−22.8 %	−21.2 %
Maximum drawdown	−50.8 %	−50.8 %	−20.3 %
Dickey-Fuller t (level)	+0.98	+1.46	−0.06
Autocorrelation $\rho(1)$	+0.22	+0.20	+0.09
Ljung-Box Q(12) [p]	48.6 [<0.001]	22.8 [0.030]	9.6 [0.649]
ARCH-LM(4) [p]	10.9 [0.028]	3.7 [0.448]	0.5 [0.972]
t-statistic mean = 0 [p]	4.41 [<0.001]	2.33 [0.021]	3.04 [0.003]
Skewness	−1.20	−1.82	−2.31
Excess kurtosis	4.98	7.65	11.11
Jarque-Bera [p]	916 [<0.001]	933 [<0.001]	723 [<0.001]
Share $\|z\| > 3$	1.39 %	1.60 %	0.83 %
Correlation r ~ ΔGS10 [p]	−0.10 [0.005]	+0.20 [<0.001]	+0.03 [0.767]

Source: Own calculation based on the data of Robert J. Shiller (2026). Logarithmic monthly returns of the index levels (monthly averages, without dividends); Dickey-Fuller model with constant; ARCH-LM with four lags.

13.4 Findings and Interpretation

First, the random walk finding is regime-independent: in none of the three periods can the unit root be rejected — the Dickey-Fuller statistics lie far above the critical value. The price level remains unforecastable in the Old Economy as in the AI era; what changes between the regimes are the properties of the return distribution, not the fundamental unpredictability of the path.

Second, the measured autocorrelation falls from +0.22 through +0.20 to +0.09. Since the averaging artefact of the data construction contributes around +0.25 equally in all periods, the decline is remarkable: in the AI era the measured autocorrelation lies well below the artefact level, which points to a genuinely negative correlation component — rapid corrections of exaggerations within the month. One cannot speak of “more momentum” at the monthly level in the more recent periods; rather, the speed of price adjustment has increased.[54]

Third, the ARCH signals lose their significance in the subperiods ($p = 0.45$ and 0.97). This finding is to be interpreted with restraint: with $n = 312$ and all the more with $n = 120$, the test power falls considerably, so that the absence of a detection must not be read as evidence

[54] Working (1960) shows that forming period averages from a random walk generates an artificial autocorrelation of the average changes of around +0.25. Since Shiller’s index levels are monthly averages of daily prices, the measured autocorrelation of the full period is to be read almost entirely as an averaging artefact.

for the end of volatility clustering — the more so as the phases 2020 and 2022 were evidently marked by clustering. The methodologically appropriate reading is: at the monthly level, the clustering is moderately pronounced and, in short observation windows, no longer statistically detectable.[55]

Fourth — the core finding — the distribution tails become monotonically more extreme across the regimes: the skewness falls from −1.20 through −1.82 to −2.31, the excess kurtosis rises from 4.98 through 7.65 to 11.11. The AI era has the highest annual return (12.0 % against 7.2 % of the baseline) at nearly unchanged volatility (12.6 %) — and at the same time the most strongly left-skewed, fattest-tailed distribution of all three periods. The premium of the momentum and technology stocks was thus paid for not with more running fluctuation but with an intensification of the tail risk: rarer, but harder. For the instruments of this paper, that is the central message — variance-based measures do not capture this shift; Expected Shortfall, Cornish-Fisher correction, and Lower Partial Moments become more important with every regime.

Fifth, the equity-rate correlation reverses its sign: in the full period the discount rate channel dominates (−0.10, $p = 0.005$), since 2000 the correlation is significantly positive (+0.20, $p < 0.001$) — rising yields on ten-year government bonds went along with rising equity prices, as a growth-driven regime with an expansionary central bank reaction to phases of weakness would lead one to expect. In the AI era, the relationship is statistically no longer distinguishable from zero. The practical consequence reaches into portfolio construction: the diversification property of bonds relative to equities is regime-dependent and must not be treated as a constant.

Sixth, the classic metrics confirm the picture of asymmetric extremisation: the best month of the AI era (+7.9 %) stays behind that of the baseline, the worst (−21.2 %, March 2020) almost reaches the full-period minimum (−22.8 %, October 2008) — the upside is smoothed, the downside is not. That the maximum drawdown of the AI era (−20.3 %) nonetheless lies well below that of the longer periods (−50.8 %) puts this into perspective only seemingly: ten years are simply too short to contain the full cycle of bull market and structural bear market.

13.5 Answer to the Research Question, and Limitations

The answer to the research question is differentiated. "More extreme values in all respects" the New Economy and the AI era have not produced: random walk character, volatility, and autocorrelation prove stable or declining. What has changed is, rather, specifically the shape

[55] The test power falls with the sample length: that the Ljung-Box and ARCH tests no longer register in the short subperiods does not demonstrate the disappearance of the effects but, in the first instance, only their non-detectability at $n = 312$ and $n = 120$ respectively.

of the distribution tails — the skewness falls monotonically, the kurtosis rises monotonically — combined with a markedly higher average return of the AI era and a sign change of the interest rate correlation. The risk of the momentum markets manifests itself accordingly not in the running range of fluctuation but in the intensity of the rare setbacks; a risk measurement resting solely on the standard deviation thus misses exactly the dimension that has changed.

The limitations of the case study are stated: the observation windows are nested (telescoping design) and of different length, which is why significance comparisons between the periods must take the test power into account; the monthly-average construction of the data lifts the autocorrelation; price returns without dividends understate the investor return; and the period boundaries 2000 and 2016 are economically motivated but not endogenously dated. Formal structural break tests (such as Chow or Bai-Perron), daily data, and the extension to GARCH estimates mark the natural continuation of this case study.

14 Conclusion and Outlook

Which instruments are available to the professionals in charge of wealth management — whether in the special fund, in the foundation, or in the family office mandate? The answer of this paper is an ordered risk analysis apparatus with a common foundation: all metrics are condensations of the same loss distribution, and they differ in which section of it they seize — the quantile (Value at Risk), the tail behind it (Expected Shortfall), the higher moments (skewness, kurtosis, Cornish-Fisher), the shortfall from a target (Lower Partial Moments), the cash flow (Cash-Flow at Risk and Earnings at Risk), the sensitivity (Greeks), the future counterparty exposure (EE, PFE, CVA), or the sources of the result (Brinson and Barra attribution). Understanding the distribution as the common denominator also clarifies the relationship of the measures: they do not compete, they complement one another — and only their interplay, carried by the three method families of the parametric calculation, historical simulation with bootstrap, and Monte Carlo simulation, yields a reliable picture.[56]

For practice, three guidelines can be derived. First, the choice of metric follows the question: limit steering requires quickly available parametric quantities, reporting the historical simulation, derivatives and counterparty books the Monte Carlo calculation, investor mandates the target-related Lower Partial Moments, and balance sheet steering surplus quantities; the reverse practice of applying one preferred metric to all questions leads in the majority of cases to inaccurate answers. Second, every metric is to be furnished with a statement of its estimation uncertainty and a retrospective validation: bootstrap intervals and backtesting are obligatory components of every report, since a risk metric without a statement

[56] This paper serves exclusively informational and educational purposes; it constitutes no investment recommendation.

of precision and a hit-rate control possesses no empirical evidential value. Third, decisive importance falls to the distribution tails: skewness, kurtosis, and the difference between Expected Shortfall and Value at Risk are no subordinate quantities but the core of the task — investment strategies fail not at the middle of the distribution but at its tails.

The econometric case study of Section 13 underpins these guidelines empirically: over six decades of the S&P 500, the random walk character of prices persists regime-independently, yet the distribution tails become monotonically more extreme from the Old Economy through the New Economy to the AI era — the skewness falls to -2.3, the excess kurtosis rises to over 11, while volatility remains nearly unchanged and the equity-rate correlation changes its sign after 2000. The results confirm the architecture of this paper: measuring risk by the standard deviation alone misses the dimension in which modern markets have actually changed; Expected Shortfall, Cornish-Fisher correction, and Lower Partial Moments are no academic refinements but the measures appropriate to the data.

The limits of this paper mark at the same time the open fields. The worked examples are stylised, the model world stationary; real markets change regimes, and extreme value theory, time-varying volatility and correlation models, and liquidity risk remained largely bracketed here. The outlook is shaped by two developments: growing computing power, which brings full simulations of entire balance sheets within reach of medium-sized financial institutions as well, and regulation, which with the transition to the Expected Shortfall has already completed the orientation towards the distribution tails. Both developments underline the central message of this paper: the methodological apparatus is mature and available — the limiting factor is the consistency to deploy it in combination, transparently, and in awareness of its methodological limits.

References

Acerbi, C., and D. Tasche (2002): On the Coherence of Expected Shortfall, Journal of Banking and Finance 26 (7), 1487–1503.

Artzner, P., F. Delbaen, J.-M. Eber, and D. Heath (1999): Coherent Measures of Risk, Mathematical Finance 9 (3), 203–228.

Basel Committee on Banking Supervision (2019): Minimum Capital Requirements for Market Risk, Bank for International Settlements, Basel.

Bawa, V. S. (1975): Optimal Rules for Ordering Uncertain Prospects, Journal of Financial Economics 2 (1), 95–121.

Berkowitz, J., and J. O'Brien (2002): How Accurate Are Value-at-Risk Models at Commercial Banks?, Journal of Finance 57 (3), 1093–1111.

Bollerslev, T. (1986): Generalized Autoregressive Conditional Heteroskedasticity, Journal of Econometrics 31 (3), 307–327.

Boyle, P. P. (1977): Options: A Monte Carlo Approach, Journal of Financial Economics 4 (3), 323–338.

Brinson, G. P., and N. Fachler (1985): Measuring Non-US Equity Portfolio Performance, Journal of Portfolio Management 11 (3), 73–76.

Brinson, G. P., L. R. Hood, and G. L. Beebower (1986): Determinants of Portfolio Performance, Financial Analysts Journal 42 (4), 39–44.

Carhart, M. M. (1997): On Persistence in Mutual Fund Performance, Journal of Finance 52 (1), 57–82.

Christoffersen, P. F. (1998): Evaluating Interval Forecasts, International Economic Review 39 (4), 841–862.

Cont, R. (2001): Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues, Quantitative Finance 1 (2), 223–236.

Cornish, E. A., and R. A. Fisher (1938): Moments and Cumulants in the Specification of Distributions, Revue de l'Institut International de Statistique 5 (4), 307–320.

Danielsson, J., P. Embrechts, C. Goodhart, C. Keating, F. Muennich, O. Renault, and H. S. Shin (2001): An Academic Response to Basel II, Special Paper 130, LSE Financial Markets Group, London.

Dickey, D. A., and W. A. Fuller (1979): Distribution of the Estimators for Autoregressive Time Series with a Unit Root, Journal of the American Statistical Association 74 (366), 427–431.

Duffie, D., and J. Pan (1997): An Overview of Value at Risk, Journal of Derivatives 4 (3), 7–49.

Duffie, D., and K. J. Singleton (1999): Modeling Term Structures of Defaultable Bonds, Review of Financial Studies 12 (4), 687–720.

Efron, B. (1979): Bootstrap Methods: Another Look at the Jackknife, Annals of Statistics 7 (1), 1–26.

Engle, R. F. (1982): Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation, Econometrica 50 (4), 987–1007.

Fama, E. F. (1965): The Behavior of Stock-Market Prices, Journal of Business 38 (1), 34–105.

Fama, E. F., and K. R. French (1993): Common Risk Factors in the Returns on Stocks and Bonds, Journal of Financial Economics 33 (1), 3–56.

Fama, E. F., and K. R. French (2010): Luck versus Skill in the Cross-Section of Mutual Fund Returns, Journal of Finance 65 (5), 1915–1947.

Favre, L., and J.-A. Galeano (2002): Mean-Modified Value-at-Risk Optimization with Hedge Funds, Journal of Alternative Investments 5 (2), 21–25.

Fishburn, P. C. (1977): Mean-Risk Analysis with Risk Associated with Below-Target Returns, American Economic Review 67 (2), 116–126.

Glasserman, P. (2004): Monte Carlo Methods in Financial Engineering, Springer, New York.

Gneiting, T. (2011): Making and Evaluating Point Forecasts, Journal of the American Statistical Association 106 (494), 746–762.

Grinold, R. C., and R. N. Kahn (2000): Active Portfolio Management, 2nd edition, McGraw-Hill, New York.

Gupton, G. M., C. C. Finger, and M. Bhatia (1997): CreditMetrics — Technical Document, J.P. Morgan, New York.

Harlow, W. V. (1991): Asset Allocation in a Downside-Risk Framework, Financial Analysts Journal 47 (5), 28–40.

Jarque, C. M., and A. K. Bera (1980): Efficient Tests for Normality, Homoscedasticity and Serial Independence of Regression Residuals, Economics Letters 6 (3), 255–259.

Jorion, P. (2007): Value at Risk: The New Benchmark for Managing Financial Risk, 3rd edition, McGraw-Hill, New York.

J.P. Morgan (1996): RiskMetrics — Technical Document, 4th edition, New York.

Kupiec, P. H. (1995): Techniques for Verifying the Accuracy of Risk Measurement Models, Journal of Derivatives 3 (2), 73–84.

Leibowitz, M. L., and R. D. Henriksson (1988): Portfolio Optimization within a Surplus Framework, Financial Analysts Journal 44 (2), 43–51.

Ljung, G. M., and G. E. P. Box (1978): On a Measure of Lack of Fit in Time Series Models, Biometrika 65 (2), 297–303.

Macaulay, F. R. (1938): Some Theoretical Problems Suggested by the Movements of Interest Rates, Bond Yields and Stock Prices in the United States since 1856, NBER, New York.

MacKinnon, J. G. (1996): Numerical Distribution Functions for Unit Root and Cointegration Tests, Journal of Applied Econometrics 11 (6), 601–618.

Mandelbrot, B. (1963): The Variation of Certain Speculative Prices, Journal of Business 36 (4), 394–419.

Markowitz, H. (1952): Portfolio Selection, Journal of Finance 7 (1), 77–91.

Markowitz, H. (1959): Portfolio Selection: Efficient Diversification of Investments, Wiley, New York.

McNeil, A. J., R. Frey, and P. Embrechts (2015): Quantitative Risk Management: Concepts, Techniques and Tools, 2nd edition, Princeton University Press, Princeton.

Redington, F. M. (1952): Review of the Principles of Life-Office Valuations, Journal of the Institute of Actuaries 78 (3), 286–340.

Rockafellar, R. T., and S. Uryasev (2000): Optimization of Conditional Value-at-Risk, Journal of Risk 2 (3), 21–41.

Rosenberg, B. (1974): Extra-Market Components of Covariance in Security Returns, Journal of Financial and Quantitative Analysis 9 (2), 263–274.

Roy, A. D. (1952): Safety First and the Holding of Assets, Econometrica 20 (3), 431–449.

Rupprechter, R. (2026h): The Benefits of Derivatives for the Investor: Forms, Pricing, Hedging, and Leverage — from Options to Credit Default Swaps, Working Paper, R&B Wealth Management, Dornbirn, www.rbvm.at/science.html.

Sharpe, W. F. (1966): Mutual Fund Performance, Journal of Business 39 (1), 119–138.

Sharpe, W. F., and L. G. Tint (1990): Liabilities — A New Approach, Journal of Portfolio Management 16 (2), 5–10.

Shiller, R. J. (2026): U.S. Stock Markets 1871–Present and CAPE Ratio, online data set, Yale University, www.econ.yale.edu/~shiller/data.htm (retrieved July 2026).

Sortino, F. A., and R. van der Meer (1991): Downside Risk, Journal of Portfolio Management 17 (4), 27–31.

Vasicek, O. (1977): An Equilibrium Characterization of the Term Structure, Journal of Financial Economics 5 (2), 177–188.

Working, H. (1960): Note on the Correlation of First Differences of Averages in a Random Chain, Econometrica 28 (4), 916–918.

Zhu, S. H., and M. Pykhtin (2007): A Guide to Modeling Counterparty Credit Risk, GARP Risk Review 37, 16–22.