

Beyond Delta–Gamma: Value-at-Risk and Expected Shortfall for Intrinsically Quadratic Portfolios under Elliptical Risk

Jules Sadefo Kamdem¹

¹Université de Montpellier, Faculté d'Économie, Montpellier Recherche en Économie
(MRE UR 209), France

Revised version, 2026

Abstract

We revisit and substantially extend an early framework for Value-at-Risk (VaR) and Expected Shortfall (ES) of intrinsically quadratic portfolios of securities under multivariate elliptical and mixture-of-elliptical risk factors. The quadraticity considered here is conceptually different from the usual delta–gamma approximation of derivative portfolios: it arises directly from the second-order expansion of the exponential map linking log returns to security prices. This distinction yields a structured quadratic loss whose coefficients are determined by portfolio holdings rather than option Greeks.

The revised analysis makes four substantive improvements. First, it separates the exact security loss, its linear approximation, and its quadratic approximation, and derives an explicit third-order remainder bound. Second, for long-only portfolios it identifies a canonical noncentral quadratic form and shows that the quadratic loss is bounded above by one half of initial portfolio value, thereby clarifying the local nature of the approximation and the origin of the classical restriction that the quadratic VaR remain below that bound. Third, it gives a clean joint characterization of VaR and ES through the distribution and lower partial first moment of a positive quadratic form. This yields transparent extensions to finite mixtures and precise existence conditions. Fourth, it develops modern computational routes for Gaussian and Student- t factors, combining generalized noncentral chi-square methods, scale-mixture representations, and Monte Carlo validation.

A reproducible simulation study compares exact revaluation, linear approximation, and quadratic approximation under Gaussian, Student- t , and mixture risks. We then apply the framework to an equally weighted portfolio of eleven French equities observed weekly from September 2019 to September 2026. On the historical sample, the quadratic approximation is nearly indistinguishable from exact revaluation at the 95%, 97.5%, and 99% levels, while the linear approximation materially overstates tail risk. A fitted multivariate Student- t model yields an estimated degree of freedom close to 5.1, confirming substantial tail thickness. We further conduct an out-of-sample-style horizon validation using daily and directly observed weekly S&P 500, NASDAQ Composite, and CAC 40 returns from September

2016 to September 2026, together with an equally weighted three-index portfolio. Across all markets and confidence levels, the quadratic approximation remains extremely close to exact exponential revaluation; its relative error increases from daily to weekly horizons but stays below 0.4% even for weekly 99% ES, whereas the linear approximation can overstate weekly 99% ES by about 7%. We additionally compare directly observed weekly VaR and ES with the conventional $\sqrt{5}$ scaling of daily measures. The scaling rule is not uniformly accurate: it generally overstates weekly tail risk for the S&P 500, NASDAQ, and the equally weighted portfolio, while it can understate some CAC 40 weekly risks. The analysis therefore separates Taylor-approximation error from temporal-scaling error. The resulting framework preserves the original contribution while providing a more rigorous, transparent, and computationally modern treatment of quadratic portfolio tail risk.

Keywords: Value-at-Risk; Expected Shortfall; quadratic portfolio; elliptical distributions; Student- t ; mixture distributions; generalized chi-square; tail risk; market risk.

JEL classification: C46; C63; G11; G32.

1 Introduction

Value-at-Risk remains one of the most widely used quantile-based measures of market risk, while Expected Shortfall has become the canonical tail-average companion because it incorporates the severity of losses beyond the VaR threshold. The distinction between quantile risk and tail-average risk is by now standard; see, among others, [Artzner et al. \(1999\)](#), [Rockafellar and Uryasev \(2000\)](#), and [Rockafellar and Uryasev \(2002\)](#). For portfolios whose profit-and-loss is linear in approximately Gaussian risk factors, both measures are straightforward. The situation becomes materially more difficult when the portfolio mapping is nonlinear and the risk factors are non-Gaussian.

The classical literature on “quadratic VaR” developed largely around delta–gamma approximations of derivative portfolios, in which the second-order term is generated by the Hessian of a derivative-pricing function with respect to underlying risk factors. Early numerical approaches include Fourier methods and quadratic-form approximations, while related work studied principal-component and nonlinear portfolio approximations; see [Cardenas et al. \(1997\)](#), [Brummelhuis et al. \(2002\)](#), and the later numerical treatment in [Kamdem and Genz \(2008\)](#).

The object studied here is different. For a portfolio made directly of securities, let X_i denote the log return of security i . Since $S_i(t+h) = S_i(t)e^{X_i}$, even a portfolio that is linear in security prices becomes nonlinear in log returns. A second-order expansion of e^{X_i} produces an *intrinsically quadratic securities portfolio*. The quadratic coefficient is determined by the current monetary exposure in each security, not by option Greeks. This observation was the basis of [Kamdem \(2003a\)](#), where a quadratic VaR and ES methodology was developed under multivariate elliptical and mixture-of-elliptical risk factors.

The purpose of the present paper is not merely to reproduce that derivation. It is to rebuild it in a form that is mathematically cleaner, economically more transparent, and computationally more robust. Three issues are especially important. First, the exact, linear, and quadratic portfolio mappings must be distinguished from the outset. Second, the second-order Taylor map is local: for a long-only portfolio the resulting quadratic loss is bounded above by one

half of initial portfolio value, whereas the exact loss is bounded above by the entire portfolio value. Third, VaR and ES should be formulated using a single loss convention and the modern quantile/integrated-quantile definitions, avoiding sign ambiguities.

The paper makes the following contributions.

- (i) We formalize the distinction between an intrinsically quadratic securities portfolio and a delta–gamma derivative portfolio, and derive a deterministic third-order remainder bound for the approximation.
- (ii) For long-only portfolios we show that the quadratic loss is an affine decreasing transform of a positive quadratic form. VaR is therefore determined by a *lower* quantile of that quadratic form, and ES by its lower partial first moment. This yields particularly simple formulas and makes explicit the factor $1/2$ inherited from the exponential Taylor expansion.
- (iii) We derive mixture formulas that separate distributional heterogeneity from the portfolio geometry. Under Gaussian risk factors the required distribution is a generalized noncentral chi-square law; under multivariate Student- t factors a scale-mixture representation reduces the problem to one-dimensional mixing over Gaussian quadratic-form probabilities. We also derive a conditional one-step VaR–ES theorem for time-varying elliptical predictive distributions.
- (iv) We provide a local approximation theorem that bounds the difference between exact and quadratic losses on a high-probability return region. This gives a quantitative diagnostic for whether quadratic VaR/ES should be trusted at a given horizon and confidence level.
- (v) We validate the framework numerically under Gaussian, Student- t , and mixture risks and compare exact revaluation, linear risk, and quadratic risk. We then apply the framework to a real eleven-stock French equity portfolio using 367 weekly observations and report VaR and ES at 95%, 97.5%, and 99%. Finally, using ten years of daily and directly observed weekly S&P 500, NASDAQ Composite, and CAC 40 returns, we quantify how the linear and quadratic approximation errors change with the horizon and tail probability, and we test the conventional $\sqrt{5}$ daily-to-weekly scaling rule against directly observed weekly VaR and ES.

The historical positioning is deliberately cautious. Earlier work studied VaR under non-Gaussian and elliptical-type specifications, and [Landsman and Valdez \(2003\)](#) derived tail conditional expectations for elliptical distributions. The contribution of [Kamdem \(2003a\)](#) was to study VaR and ES jointly for a quadratic portfolio of securities, distinct from delta–gamma derivative portfolios, under elliptical and mixture-elliptical risk factors. The present paper sharpens that framework rather than claiming priority over all forms of quadratic VaR.

The rest of the paper is organized as follows. [Section 2](#) defines the exact, linear, and quadratic losses and establishes approximation properties. [Section 3](#) recalls elliptical and mixture representations. [Section 4](#) gives the canonical quadratic-form reduction. [Section 5](#) derives VaR and ES. [Section 6](#) treats mixtures. [Section 7](#) specializes the results to Gaussian and Student- t risks. [Section 8](#) presents numerical algorithms. [Section 9](#) reports a reproducible simulation study.

Section 10 applies the framework to an eleven-stock French equity portfolio. Section 11 validates the Taylor approximations at daily and weekly horizons for the S&P 500, NASDAQ Composite, CAC 40, and an equally weighted three-index portfolio. Section 12 discusses model risk and approximation limits. Section 14 concludes.

2 Exact, linear, and intrinsically quadratic portfolio losses

2.1 Exact portfolio mapping

Consider n traded securities with strictly positive prices $S_i(t)$ and holdings θ_i . Let

$$P_0 = \sum_{i=1}^n \theta_i S_i(0)$$

be the initial portfolio value and define the monetary exposure

$$a_i = \theta_i S_i(0), \quad a = (a_1, \dots, a_n)^\top.$$

For a horizon h , define log returns

$$X_i = \log\left(\frac{S_i(h)}{S_i(0)}\right), \quad X = (X_1, \dots, X_n)^\top.$$

Then

$$P_h = \sum_{i=1}^n a_i e^{X_i}, \quad \Delta P = P_h - P_0 = \sum_{i=1}^n a_i (e^{X_i} - 1).$$

We adopt the loss convention

$$L^E(X) := -\Delta P = \sum_{i=1}^n a_i (1 - e^{X_i}).$$

A positive number is therefore a loss.

2.2 Linear and quadratic approximations

The first- and second-order Taylor approximations of e^x at zero are

$$e^x = 1 + x + \frac{x^2}{2} + R_3(x).$$

This gives the linear loss

$$L^L(X) = -a^\top X, \tag{1}$$

and the intrinsically quadratic loss

$$L^Q(X) = -a^\top X - \frac{1}{2} X^\top W X, \quad W := \text{diag}(a_1, \dots, a_n). \tag{2}$$

This is not a delta–gamma approximation. In a derivative portfolio, one typically writes

$$L^{\Delta\Gamma}(Z) \approx c - \Delta^\top Z - \frac{1}{2} Z^\top \Gamma Z,$$

where Δ and Γ are first- and second-order sensitivities of derivative prices. In (2), by contrast, W contains current monetary holdings in the securities themselves.

Proposition 2.1 (Deterministic third-order remainder). *For every $x \in \mathbb{R}$, there exists ξ between 0 and x such that*

$$e^x - 1 - x - \frac{x^2}{2} = \frac{e^\xi}{6} x^3.$$

Consequently, if $a_i \geq 0$ for all i and $\|X\|_\infty \leq \delta$, then

$$|L^E(X) - L^Q(X)| \leq \frac{e^\delta}{6} \sum_{i=1}^n a_i |X_i|^3 \leq \frac{P_0 e^\delta \delta^3}{6}. \quad (3)$$

Proof. Taylor’s theorem with Lagrange remainder gives the scalar identity. Summing the absolute remainder terms with nonnegative a_i yields the first inequality; the second follows from $|X_i| \leq \delta$ and $\sum_i a_i = P_0$. $\square$

Remark 2.2. The bound (3) is deliberately deterministic. It remains valid for heavy-tailed distributions such as Student- t , for which exponential moments need not exist. This is preferable to an L^1 remainder argument when the log-return specification has polynomial tails.

2.3 A structural bound for long-only portfolios

The next result, implicit in the original 2003 restriction, is central to understanding the approximation.

Proposition 2.3 (Maximum quadratic loss). *Suppose $a_i \geq 0$ and $P_0 = \sum_i a_i > 0$. Then*

$$L^Q(X) = \frac{P_0}{2} - \frac{1}{2} (X + \mathbf{1})^\top W (X + \mathbf{1}), \quad (4)$$

and therefore

$$L^Q(X) \leq \frac{P_0}{2}. \quad (5)$$

Equality holds iff $X_i = -1$ for every i with $a_i > 0$.

Proof. For each coordinate,

$$-x_i - \frac{x_i^2}{2} = \frac{1}{2} - \frac{1}{2} (x_i + 1)^2.$$

Multiplying by a_i and summing yields (4). Since W is positive semidefinite, the quadratic term is nonnegative. $\square$

Remark 2.4 (Local nature of the approximation). The exact long-only loss satisfies $L^E(X) < P_0$ and approaches P_0 as all $X_i \rightarrow -\infty$. The quadratic approximation instead satisfies $L^Q \leq P_0/2$ and eventually decreases when returns become more negative than -1 . Hence the quadratic

mapping must not be interpreted as a global tail model. It is a local approximation whose adequacy must be checked relative to horizon, volatility, and confidence level.

Proposition 2.5 (CDF approximation on a high-probability region). *Let $A_\delta = \{\|X\|_\infty \leq \delta\}$ and*

$$\varepsilon_\delta = \frac{P_0 e^\delta \delta^3}{6}.$$

For every $z \in \mathbb{R}$,

$$F_Q(z - \varepsilon_\delta) - \mathbb{P}(A_\delta^c) \leq F_E(z) \tag{6}$$

$$\leq F_Q(z + \varepsilon_\delta) + \mathbb{P}(A_\delta^c), \tag{7}$$

where F_E and F_Q are the CDFs of L^E and L^Q .

Proof. On A_δ , Proposition 2.1 implies $|L^E - L^Q| \leq \varepsilon_\delta$. Thus $\{L^Q \leq z - \varepsilon_\delta\} \cap A_\delta \subseteq \{L^E \leq z\}$ and $\{L^E \leq z\} \cap A_\delta \subseteq \{L^Q \leq z + \varepsilon_\delta\}$. Taking probabilities gives the result. $\square$

3 Elliptical and mixture-of-elliptical risk factors

Definition 3.1 (Elliptical distribution). A random vector $X \in \mathbb{R}^n$ has an elliptical distribution with location $\mu \in \mathbb{R}^n$, scatter matrix $\Sigma \succeq 0$, and characteristic generator ψ , written

$$X \sim \mathcal{E}_n(\mu, \Sigma, \psi),$$

if its characteristic function has the form

$$\varphi_X(t) = e^{it^\top \mu} \psi(t^\top \Sigma t).$$

When a density exists and $\Sigma \succ 0$, it can be written

$$f_X(x) = |\Sigma|^{-1/2} g\left((x - \mu)^\top \Sigma^{-1} (x - \mu)\right)$$

for a suitable density generator g .

A useful stochastic representation is

$$X \stackrel{d}{=} \mu + RAU, \tag{8}$$

where $AA^\top = \Sigma$, U is uniform on the unit sphere $\mathbb{S}^{n-1}$, $R \geq 0$ is a radial random variable, and R is independent of U . This representation separates radial tail behavior from angular portfolio geometry; see Fang et al. (1990) and McNeil et al. (2015).

Definition 3.2 (Finite elliptical mixture). A finite mixture of K elliptical distributions has CDF and density

$$F_X(x) = \sum_{k=1}^K \pi_k F_k(x), \quad f_X(x) = \sum_{k=1}^K \pi_k f_k(x),$$

where $\pi_k > 0$ and $\sum_k \pi_k = 1$, and component k is elliptical with parameters $(\mu_k, \Sigma_k, \psi_k)$.

Mixtures provide a simple way to represent heterogeneous volatility and tail regimes. The components need not share the same location, scatter, or tail generator.

4 Canonical quadratic-form representation

We first state a general reduction and then specialize it to the intrinsically quadratic securities portfolio.

4.1 General quadratic loss

Consider

$$L(X) = c + b^\top X + X^\top H X, \quad H = H^\top. \quad (9)$$

With $X = \mu + RAU$ from (8),

$$L = c_0 + Rd^\top U + R^2 U^\top B U, \quad (10)$$

where

$$c_0 = c + b^\top \mu + \mu^\top H \mu, \quad d = A^\top (b + 2H\mu), \quad B = A^\top H A.$$

Thus, conditional on the direction $U = u$, the loss is a scalar quadratic polynomial in the radial variable R .

Theorem 4.1 (Radial-angular representation). *Let L be given by (9) and X admit representation (8). Then for every $\ell \in \mathbb{R}$,*

$$F_L(\ell) = \mathbb{E}_U \left[\int_0^\infty \mathbf{1}\{a(U)r^2 + d(U)r + c_0 \leq \ell\} dF_R(r) \right], \quad (11)$$

where $a(u) = u^\top B u$ and $d(u) = d^\top u$.

Proof. Condition on U in (10), use independence of R and U , and integrate the indicator with respect to F_R . $\square$

The theorem covers indefinite quadratic forms and therefore delta-gamma portfolios. For the long-only securities problem, however, a substantially simpler representation is available.

4.2 Intrinsic quadratic securities portfolio

Assume henceforth $a_i > 0$ and define $W = \text{diag}(a_1, \dots, a_n)$. Let

$$Y = W^{1/2}(X + \mathbf{1}), \quad Q := Y^\top Y = (X + \mathbf{1})^\top W (X + \mathbf{1}). \quad (12)$$

Then Proposition 2.3 gives

$$L^Q = \frac{P_0 - Q}{2}. \quad (13)$$

The entire VaR/ES problem therefore reduces to the lower tail of a positive quadratic form Q .

5 Value-at-Risk and Expected Shortfall

5.1 Definitions

For a loss L and confidence level $p \in (0, 1)$, define

$$\text{VaR}_p(L) = \inf\{\ell \in \mathbb{R} : \mathbb{P}(L \leq \ell) \geq p\}. \quad (14)$$

The general Expected Shortfall definition is

$$\text{ES}_p(L) = \frac{1}{1-p} \int_p^1 \text{VaR}_u(L) du. \quad (15)$$

If L has a continuous distribution, this reduces to

$$\text{ES}_p(L) = \mathbb{E}[L \mid L > \text{VaR}_p(L)].$$

5.2 A joint VaR–ES theorem for the intrinsic quadratic portfolio

Let

$$F_Q(q) = \mathbb{P}(Q \leq q), \quad M_Q(q) = \mathbb{E}[Q \mathbf{1}\{Q \leq q\}].$$

Theorem 5.1 (Quadratic VaR and ES). *Assume $a_i > 0$, Q in (12) has a continuous strictly increasing CDF on the relevant range, and $p \in (0, 1)$. Let*

$$\tau_p := F_Q^{-1}(1-p). \quad (16)$$

Then

$$\boxed{\text{VaR}_p(L^Q) = \frac{P_0 - \tau_p}{2}} \quad (17)$$

and

$$\boxed{\text{ES}_p(L^Q) = \frac{P_0}{2} - \frac{M_Q(\tau_p)}{2(1-p)}}. \quad (18)$$

Proof. Since $L^Q = (P_0 - Q)/2$ is strictly decreasing in Q ,

$$\mathbb{P}(L^Q > \ell) = \mathbb{P}(Q < P_0 - 2\ell).$$

At $\ell = \text{VaR}_p(L^Q)$ the upper loss-tail probability equals $1-p$, hence $F_Q(P_0 - 2\ell) = 1-p$ and (17) follows. For ES, continuity gives

$$\begin{aligned} \text{ES}_p(L^Q) &= \frac{1}{1-p} \mathbb{E} \left[\frac{P_0 - Q}{2} \mathbf{1}\{Q \leq \tau_p\} \right] \\ &= \frac{P_0}{2} - \frac{1}{2(1-p)} \mathbb{E}[Q \mathbf{1}\{Q \leq \tau_p\}], \end{aligned}$$

which is (18). □

Remark 5.2 (Why the factor 1/2 matters). Both (17) and (18) inherit the coefficient 1/2 from

the Taylor term $X_i^2/2$. Writing the problem through Q makes this normalization explicit and prevents a multiplicative factor from being lost during changes of variables.

Corollary 5.3 (Bounds). *Under the assumptions of Theorem 5.1,*

$$\text{VaR}_p(L^Q) \leq \frac{P_0}{2}, \quad \text{ES}_p(L^Q) \leq \frac{P_0}{2}.$$

Remark 5.4 (Existence of ES). For the long-only intrinsic quadratic loss, $L^Q \leq P_0/2$, so upper-tail ES is finite whenever the relevant conditional expectation is well-defined; in fact the tail contribution is bounded. For a *general* quadratic loss (9), however, $|L| = O(\|X\|^2)$, and a sufficient condition is $\mathbb{E}\|X\|^2 < \infty$. In the multivariate Student- t case this requires $\nu > 2$. The distinction is important: the structured long-only model has a bounded loss tail, whereas an indefinite delta-gamma form need not.

6 Finite mixtures

Suppose X follows a K -component mixture with weights π_k . Define component quadratic-form distributions

$$F_{Q,k}(q) = \mathbb{P}(Q \leq q \mid Z = k), \quad M_{Q,k}(q) = \mathbb{E}[Q \mathbf{1}\{Q \leq q\} \mid Z = k].$$

Theorem 6.1 (Mixture decomposition). *For every $q \geq 0$,*

$$F_Q(q) = \sum_{k=1}^K \pi_k F_{Q,k}(q), \tag{19}$$

and

$$M_Q(q) = \sum_{k=1}^K \pi_k M_{Q,k}(q). \tag{20}$$

Consequently, τ_p is the solution of

$$\sum_{k=1}^K \pi_k F_{Q,k}(\tau_p) = 1 - p, \tag{21}$$

and

$$\text{ES}_p(L^Q) = \frac{P_0}{2} - \frac{1}{2(1-p)} \sum_{k=1}^K \pi_k M_{Q,k}(\tau_p). \tag{22}$$

Proof. Both identities follow from the law of total probability/expectation with respect to the latent mixture indicator Z . The VaR and ES formulas then follow from Theorem 5.1. $\square$

The result is useful because mixture complexity enters only through the component CDFs and lower partial moments. No ad hoc “mixture VaR” averaging is valid: the root must be solved after mixing the CDFs.

7 Gaussian and Student- t specializations

7.1 Gaussian factors

Suppose

$$X \sim N_n(\mu, \Sigma).$$

Then

$$Y = W^{1/2}(X + \mathbf{1}) \sim N_n(m, \Omega), \quad m = W^{1/2}(\mu + \mathbf{1}), \quad \Omega = W^{1/2}\Sigma W^{1/2}.$$

Let $\Omega = U\Lambda U^\top$ with $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n)$ and set $\delta = U^\top m$. If $Z_j \stackrel{iid}{\sim} N(0, 1)$, then

$$Q \stackrel{d}{=} \sum_{j=1}^n \lambda_j \left(Z_j + \frac{\delta_j}{\sqrt{\lambda_j}} \right)^2, \quad (23)$$

for positive eigenvalues. Thus Q follows a generalized noncentral chi-square distribution. Its CDF can be computed by characteristic-function inversion such as the methods of [Imhof \(1961\)](#) or [Davies \(1980\)](#). This route is typically more accurate and simpler than high-dimensional hyperspherical cubature in the Gaussian case.

The lower partial moment can be evaluated numerically as

$$M_Q(q) = \int_0^q t dF_Q(t) = qF_Q(q) - \int_0^q F_Q(t) dt, \quad (24)$$

where the second equality follows by integration by parts. Hence once a stable routine for F_Q is available, ES requires only a one-dimensional numerical integration.

Corollary 7.1 (Gaussian VaR/ES computation). *Under Gaussian risk factors, compute τ_p from $F_Q(\tau_p) = 1 - p$ using a generalized noncentral chi-square CDF. Then*

$$\text{VaR}_p = \frac{P_0 - \tau_p}{2},$$

and with (24),

$$\text{ES}_p = \frac{P_0}{2} - \frac{1}{2(1-p)} \left[\tau_p(1-p) - \int_0^{\tau_p} F_Q(t) dt \right].$$

7.2 Student- t factors

Let

$$X \sim t_{\nu,n}(\mu, \Sigma),$$

where Σ is the scale matrix. Use the normal scale-mixture representation

$$X \stackrel{d}{=} \mu + \sqrt{\frac{\nu}{S}} Z, \quad Z \sim N_n(0, \Sigma), \quad S \sim \chi_\nu^2, \quad (25)$$

with $Z \perp S$. Conditional on $S = s$,

$$Y \mid S = s \sim N_n\left(m, \frac{\nu}{s}\Omega\right),$$

where $m = W^{1/2}(\mu + \mathbf{1})$ and $\Omega = W^{1/2}\Sigma W^{1/2}$.

Theorem 7.2 (Student- t quadratic-form representation). *Let $F_{Q|s}$ and $M_{Q|s}$ denote the Gaussian quadratic-form CDF and lower partial moment conditional on $S = s$. Then*

$$F_Q(q) = \int_0^\infty F_{Q|s}(q) f_{\chi_\nu^2}(s) ds, \quad (26)$$

$$M_Q(q) = \int_0^\infty M_{Q|s}(q) f_{\chi_\nu^2}(s) ds. \quad (27)$$

Therefore VaR and ES follow from Theorem 5.1 after one-dimensional mixing over s .

Proof. Apply the law of total probability and total expectation to the scale mixture (25). $\square$

This formulation has two advantages over a direct special-function expansion. First, it separates the Gaussian quadratic-form engine from the tail-thickness parameter ν . Second, it extends immediately to finite mixtures of Student- t components by applying Theorem 6.1.

8 Numerical implementation

8.1 Root-and-moment algorithm

The theoretical results suggest the following generic procedure.

Algorithm 1 Quadratic VaR and ES for a long-only securities portfolio

Require: Exposures $a_i > 0$, distributional parameters, confidence level p .

- 1: Set $P_0 = \sum_i a_i$, $W = \text{diag}(a_i)$.
 - 2: Construct a numerical routine for $F_Q(q)$ where $Q = (X + \mathbf{1})^\top W (X + \mathbf{1})$.
 - 3: Solve $F_Q(\tau_p) = 1 - p$ for $\tau_p \geq 0$ by bisection or safeguarded Newton.
 - 4: Set $\text{VaR}_p = (P_0 - \tau_p)/2$.
 - 5: Compute $M_Q(\tau_p) = \mathbb{E}[Q \mathbf{1}\{Q \leq \tau_p\}]$ directly or via (24).
 - 6: Set $\text{ES}_p = P_0/2 - M_Q(\tau_p)/[2(1 - p)]$.
 - 7: Evaluate a local-validity diagnostic using Proposition 2.5.
-

8.2 Gaussian engine

For Gaussian components, generalized chi-square algorithms based on inversion of the characteristic function are natural. The classical methods of Imhof (1961) and Davies (1980) remain standard references. In modern implementations one should monitor absolute and relative tail-error tolerances because the relevant probability is $1 - p$, which can be small.

8.3 Student- t engine

For Student- t components, use Gauss–Laguerre or adaptive quadrature on the χ_ν^2 mixing variable in (26)–(27). Each quadrature node calls the Gaussian engine with covariance scaled by ν/s . This is computationally stable and avoids deriving separate hypergeometric expressions for each degree of freedom.

8.4 Elliptical radial-angular engine

For a general elliptical generator, the stochastic representation (8) and Theorem 4.1 suggest integration over the sphere and radial law. Fully symmetric spherical rules such as those studied by Genz (2003) can be used, and the transformations in Kamdem and Genz (2008) remain relevant for indefinite quadratic regions. For the positive long-only form, Monte Carlo or randomized quasi-Monte Carlo over the sphere combined with one-dimensional radial integration is also straightforward.

8.5 Monte Carlo validation

Regardless of the analytic engine, we recommend validating the reported VaR/ES against direct simulation of three losses:

$$L^L = -a^\top X, \quad L^Q = -a^\top X - \frac{1}{2}X^\top W X, \quad L^E = \sum_i a_i(1 - e^{X_i}).$$

The difference $L^Q - L^E$ is a direct diagnostic of Taylor approximation error, while $L^L - L^E$ measures the incremental value of the quadratic correction.

9 Reproducible simulation study

9.1 Design

We consider a normalized long-only portfolio with $P_0 = 1$ and monetary weights

$$a = (0.50, 0.30, 0.20)^\top.$$

The location vector is

$$\mu = (0.0005, 0.0004, 0.0003)^\top,$$

and marginal volatilities are (0.04, 0.05, 0.06). The correlation matrix is

$$C = \begin{pmatrix} 1 & 0.4 & 0.3 \\ 0.4 & 1 & 0.5 \\ 0.3 & 0.5 & 1 \end{pmatrix},$$

so $\Sigma = \text{diag}(\sigma)C \text{diag}(\sigma)$.

We use 800,000 simulated observations and a fixed random seed (20260930). Three distributions are considered:

- (a) Gaussian $N(\mu, \Sigma)$;
- (b) standardized multivariate Student- t_5 with covariance Σ ;
- (c) a 90/10 mixture of the Gaussian component and a Student- t_5 component with covariance 2.25Σ .

For each model we compute empirical VaR and ES at 95% and 99% for L^L , L^Q , and L^E .

Table 1: Monte Carlo comparison of linear, quadratic, and exact losses

Distribution	Method	VaR ₉₅	ES ₉₅	VaR ₉₉	ES ₉₉
Gaussian	Linear	0.05970	0.07503	0.08461	0.09687
	Quadratic	0.05745	0.07163	0.08053	0.09160
	Exact	0.05751	0.07174	0.08066	0.09180
Student- t_5	Linear	0.05650	0.08118	0.09470	0.12504
	Quadratic	0.05429	0.07643	0.08904	0.11457
	Exact	0.05436	0.07667	0.08929	0.11528
Mixture	Linear	0.06174	0.08073	0.09053	0.11325
	Quadratic	0.05927	0.07647	0.08570	0.10505
	Exact	0.05935	0.07665	0.08588	0.10553

Losses are expressed as fractions of initial portfolio value. Results are Monte Carlo estimates based on 800,000 observations with seed 20260930. The Student- t scale is adjusted so that its covariance equals Σ .

9.2 Accuracy of the quadratic approximation

Relative to exact revaluation, the quadratic approximation is extremely accurate in this experiment. The relative errors (quadratic minus exact, divided by exact) are approximately: The quadratic correction therefore captures almost all of the nonlinearity generated by the

Table 2: Relative error of quadratic approximation versus exact revaluation (percent)

Distribution	VaR ₉₅	ES ₉₅	VaR ₉₉	ES ₉₉
Gaussian	-0.11	-0.16	-0.17	-0.22
Student- t_5	-0.14	-0.31	-0.28	-0.62
Mixture	-0.13	-0.24	-0.21	-0.46

exponential price map at the chosen horizon. The linear approximation is systematically more conservative in this design because it omits the positive second-order contribution to P&L, equivalently the negative quadratic term in the loss.

The Student- t and mixture models increase the gap between 95% and 99% ES, as expected under heavier tails. Yet even in these cases the local quadratic approximation remains close to exact revaluation for the selected volatility scale. This should not be extrapolated mechanically to much longer horizons: Proposition 2.3 implies that the quadratic model cannot represent losses beyond $P_0/2$.

9.3 Reproducibility pseudocode

The simulation can be reproduced with the following operations:

1. Generate X under the selected multivariate law.
2. Compute $L^L = -a^\top X$.
3. Compute $L^Q = -\sum_i a_i (X_i + X_i^2/2)$.
4. Compute $L^E = -\sum_i a_i (e^{X_i} - 1)$.

5. Estimate VaR_p by the empirical p -quantile and ES_p by the mean of losses above that quantile.

In an empirical implementation, the same workflow can be used with estimated locations, covariance matrices, degrees of freedom, and mixture probabilities.

10 Empirical application: an eleven-stock French equity portfolio

10.1 Data and portfolio construction

We apply the framework to weekly closing prices for eleven French equities: Air Liquide, Schneider Electric, Airbus, Sanofi, LVMH, TotalEnergies, BNP Paribas, AXA, Drone Volt, Vinci, and Danone. The sample runs from 2 September 2019 to 14 September 2026. Continuously compounded weekly returns are defined by

$$X_{i,t} = \log P_{i,t} - \log P_{i,t-1},$$

yielding $T = 367$ return observations for each asset.

To keep the empirical comparison transparent and independent of an optimization criterion, we use a weekly rebalanced equal-weight portfolio,

$$w_i = \frac{1}{11}, \quad i = 1, \dots, 11.$$

Portfolio value is normalized to $P_0 = 1$, so all reported losses, VaR, and ES can be read directly as percentages of initial portfolio value.

For each observed return vector X_t , we compute the three losses

$$L_t^L = - \sum_{i=1}^{11} w_i X_{i,t}, \tag{28}$$

$$L_t^Q = - \sum_{i=1}^{11} w_i \left(X_{i,t} + \frac{1}{2} X_{i,t}^2 \right), \tag{29}$$

$$L_t^E = - \sum_{i=1}^{11} w_i (e^{X_{i,t}} - 1). \tag{30}$$

The last expression is the exact one-period mark-to-market loss for a portfolio rebalanced to equal weights at the beginning of each week.

The descriptive statistics provide strong evidence against a Gaussian cross-sectional return model. Several assets exhibit substantial negative skewness, and excess kurtosis is particularly large for TotalEnergies, Drone Volt, Airbus, Vinci, and AXA. Because fourth moments are highly sensitive to a small number of extreme observations, Table 4 reports a direct robustness diagnostic for the most leptokurtic large-cap names and for the CAC 40 benchmark. This distinction is important: a large full-sample kurtosis may reflect episodic crisis observations rather than a stable unconditional distributional feature.

Table 3: Descriptive statistics of weekly log returns for the eleven-stock portfolio and CAC 40 benchmark

Asset	Mean (%)	S.D. (%)	Minimum (%)	Maximum (%)	Skewness	Excess kurtosis
Air Liquide	0.179	2.815	-15.733	10.214	-0.550	3.991
Schneider Electric	0.360	4.033	-22.681	13.102	-0.705	2.948
Airbus	0.120	5.451	-31.993	35.643	-0.405	11.229
Sanofi	-0.021	3.270	-19.353	10.402	-0.957	4.890
LVMH	0.022	4.112	-14.296	17.259	0.010	1.592
TotalEnergies	0.151	4.262	-35.591	25.464	-1.070	16.700
BNP Paribas	0.244	4.985	-19.514	19.283	-0.705	3.113
AXA	0.201	4.078	-26.592	20.167	-0.669	7.487
Drone Volt	-1.067	12.562	-39.548	101.394	2.661	16.672
Vinci	0.033	3.815	-26.152	18.091	-0.851	8.409
Danone	-0.078	2.862	-13.063	10.302	-0.421	3.538
CAC 40	0.112	2.760	-22.073	10.159	-1.743	13.079

Note: $N = 367$ weekly observations for every series. Mean, standard deviation, minimum, and maximum are expressed in percentage points. Skewness and excess kurtosis use bias-corrected sample estimators.




Figure 1: Normalized weekly closing prices of the eleven French equities, 2 September 2019–14 September 2026. Each price series is normalized to 100 at the first observation, which removes differences in nominal price levels and makes relative price dynamics directly comparable.

Table 4: Robustness of weekly excess kurtosis to crisis and extreme observations

Asset	Full sample	Excl. Feb.–May 2020	Drop 1	Drop 2	Drop 3	Drop 5
TotalEnergies	16.70	1.75	6.07	1.97	1.19	0.39
Airbus	11.23	12.12	8.68	6.39	4.05	3.09
AXA	7.49	4.56	3.81	2.56	2.21	1.76
Vinci	8.41	2.68	3.91	2.86	2.30	1.56
CAC 40	13.08	2.72	3.66	2.38	1.77	1.05

Notes: Entries are unbiased Fisher excess-kurtosis estimates of weekly log returns. The full sample contains 367 observations from September 2019 to September 2026. “Excl. Feb.–May 2020” removes the 17 weekly observations dated from 3 February through 25 May 2020. “Drop k ” removes, separately for each series, the k observations with the largest absolute log returns $|X_t|$ and recomputes excess kurtosis on the remaining sample.




Figure 2: Weekly log returns of the equally weighted eleven-stock portfolio and the CAC 40 benchmark. The figure highlights volatility clustering and the exceptional observations around the 2020 market disruption.

The robustness exercise substantially changes the interpretation of the raw fourth moments. TotalEnergies is the clearest case: its excess kurtosis falls from 16.70 to 1.75 when the February–May 2020 window is removed, to 6.07 after removing only its single largest absolute weekly return, and to 1.97 after removing the two largest observations. AXA, Vinci, and the CAC 40 display the same broad pattern, although with different magnitudes. Their full-sample tail thickness is therefore strongly concentrated in a small set of crisis weeks. Airbus is different. Excluding February–May 2020 raises its excess kurtosis from 11.23 to 12.12 because a very large positive weekly return on 1 June 2020 remains outside the exclusion window; nevertheless, trimming its largest absolute observations reduces excess kurtosis monotonically to 3.09 after five observations are removed. These results support heavy-tailed and mixture specifications, but they also caution against interpreting a single full-sample kurtosis estimate as a time-invariant structural parameter. Figures 1–2 complement these moments by showing that the extreme-return episodes are localized in time rather than uniformly distributed across the sample.

10.2 Historical VaR and Expected Shortfall

We first use the empirical distribution of the 367 weekly loss observations. For confidence level p , historical VaR is the empirical p -quantile and historical ES is the average loss weakly above that empirical quantile.

The quadratic approximation is exceptionally close to exact revaluation. At 99%, for example,

$$\text{VaR}_{0.99}^Q = 7.418\%, \quad \text{VaR}_{0.99}^E = 7.433\%,$$

while

$$\text{ES}_{0.99}^Q = 12.544\%, \quad \text{ES}_{0.99}^E = 12.620\%.$$

The corresponding relative errors are about -0.20% for VaR and -0.60% for ES. In contrast, the linear approximation yields relative errors of about $+7.25\%$ for 99% VaR and $+8.80\%$ for 99% ES.

Table 5: Historical weekly VaR and ES: linear, quadratic, and exact losses

Method	Measure	95%	97.5%	99%
Linear	VaR	4.487	6.101	7.972
	ES	7.658	9.751	13.731
Quadratic	VaR	4.305	5.779	7.418
	ES	7.148	8.994	12.544
Exact	VaR	4.311	5.792	7.433
	ES	7.176	9.040	12.620

Note: entries are percentages of initial portfolio value. At 95%, 97.5%, and 99%, the empirical tail contains approximately 19, 10, and 4 observations, respectively.

Hence the second-order term is empirically material even though the quadratic approximation itself remains very accurate.

10.3 Parametric elliptical models

We next estimate two parametric benchmarks using the full eleven-dimensional return vector.

For the Gaussian benchmark, the location and covariance matrices are estimated by their maximum-likelihood sample analogues.

For the multivariate Student- t model we estimate

$$X_t \sim t_\nu(\mu, \Sigma)$$

by maximum likelihood, using the standard scale-matrix parameterization. The estimated degrees of freedom are

$$\hat{\nu} = 5.11,$$

indicating substantial tail thickness. The location and scale matrix are estimated jointly by the usual fixed-point/EM representation of the multivariate Student likelihood.

Finally, to illustrate the finite-mixture extension, we estimate a two-component full-covariance Gaussian mixture. The estimated component probabilities are approximately

$$\hat{\pi}_1 = 0.879, \quad \hat{\pi}_2 = 0.121.$$

Each Gaussian component is elliptical, so this specification belongs directly to the finite mixture-of-ellipticals framework of [Section 6](#).

For each fitted distribution we generate 10^6 return vectors with random seed 20260930 and evaluate L^L , L^Q , and L^E on exactly the same simulated vectors.

Two findings are worth emphasizing. First, the quadratic approximation tracks exact revaluation extremely closely under all three fitted specifications. Second, a heavier-tailed Student law need not produce a larger VaR at every confidence level when its scale is estimated jointly from the data. Here, Student VaR is below the Gaussian VaR even at 99%, whereas Student ES is larger:

$$\text{VaR}_{0.99}^{E,t} = 6.280\% < \text{VaR}_{0.99}^{E,N} = 6.663\%,$$

Table 6: Parametric weekly VaR and ES for the eleven-stock portfolio

Model	Method	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
Gaussian	Linear	4.964	6.236	5.922	7.074	7.040	8.072
	Quadratic	4.716	5.894	5.608	6.665	6.641	7.571
	Exact	4.728	5.911	5.624	6.686	6.663	7.598
Student- <i>t</i>	Linear	3.979	5.748	5.105	7.021	6.728	8.878
	Quadratic	3.764	5.345	4.797	6.467	6.252	8.055
	Exact	3.775	5.370	4.814	6.506	6.280	8.121
Gaussian mixture	Linear	4.766	5.938	5.645	6.713	6.686	7.635
	Quadratic	4.575	5.667	5.400	6.386	6.370	7.229
	Exact	4.582	5.678	5.410	6.399	6.384	7.247

Note: entries are percentages of initial portfolio value. Monte Carlo size is 10^6 for each fitted model. Student-*t* parameters are estimated from the full eleven-dimensional weekly return sample. The mixture is a two-component Gaussian mixture with full covariance matrices.

but

$$\text{ES}_{0.99}^{E,t} = 8.121\% > \text{ES}_{0.99}^{E,N} = 7.598\%.$$

This illustrates why tail-risk assessment should not rely on VaR alone. The Student fit assigns more probability to both the center and the remote tails, and ES reveals the greater severity of those remote-tail losses.

10.4 Risk-minimizing allocations

As a portfolio-construction extension, we also compare three long-only allocations on the same eleven-asset universe: the global minimum-variance portfolio, the minimum 99% VaR portfolio, and the minimum 99% ES portfolio. The constraints are

$$\mathbf{1}^\top w = 1, \quad w_i \geq 0, \quad i = 1, \dots, 11.$$

The optimization uses the one-step-ahead conditional covariance matrix from the companion EWMA elliptical fit. The VaR and ES objectives use the corresponding corrected 99% location-scale tail coefficients. This allocation exercise is deliberately kept separate from the main comparison between linear, intrinsic-quadratic, and exact exponential revaluation: its purpose is to show how the same eleven-stock data translate into economically different risk-minimizing holdings.

The allocations reveal three robust features. First, TotalEnergies and Danone jointly account for roughly 57–59% of each optimized portfolio, reflecting their diversification contribution under the conditional covariance forecast. Second, BNP Paribas and Vinci receive zero weight in all three long-only solutions, whereas Drone Volt receives only a very small allocation because of its exceptionally high volatility and tail risk. Third, the minimum-VaR and minimum-ES portfolios are close but not identical: ES places slightly less weight on Air Liquide and slightly more on Danone and Drone Volt, which is consistent with ES penalizing the severity of losses beyond the quantile rather than only the quantile boundary.

For reference, the annualized volatility of the three optimized portfolios is approximately 11.28%, 11.31%, and 11.30%, respectively. Their model-implied 99% VaR values are approxi-

Table 7: Long-only optimal allocations for the eleven-stock portfolio (%)

Asset	Minimum variance	Minimum VaR ₉₉	Minimum ES ₉₉
Air Liquide	0.00	2.46	2.00
Schneider Electric	10.38	11.55	11.35
Airbus	11.24	11.15	11.16
Sanofi	6.09	5.61	5.70
LVMH	9.53	8.76	8.90
TotalEnergies	30.77	30.70	30.71
BNP Paribas	0.00	0.00	0.00
AXA	2.77	3.29	3.20
Drone Volt	1.42	0.53	0.69
Vinci	0.00	0.00	0.00
Danone	27.81	25.95	26.29
Total	100.00	100.00	100.00

Note: all three portfolios impose full investment and no short selling. VaR and ES are minimized at the 99% confidence level under the same conditional elliptical-EWMA risk forecast; figures are rounded to two decimal places.

mately 3.91%, 3.90%, and 3.90%, and the corresponding ES values are approximately 4.76%, 4.76%, and 4.76%. These differences are economically small because, under a common location-scale elliptical specification and the same conditional covariance matrix, variance, VaR, and ES are closely related risk objectives. The main value of Table 7 is therefore not a large change in aggregate risk, but the comparison of the asset-level allocations induced by the alternative criteria.

10.5 Empirical approximation errors

Table 8 reports the relative error of the quadratic and linear approximations against exact historical revaluation.

Table 8: Relative approximation error against exact historical revaluation (%)

Approximation	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
Linear	4.09	6.71	5.34	7.87	7.25	8.80
Quadratic	-0.13	-0.40	-0.22	-0.51	-0.20	-0.60

The error pattern is monotone in economic significance: the omission of the second-order price transformation becomes increasingly costly as the confidence level moves deeper into the tail. Conversely, the residual third-order error of the quadratic approximation remains below one percent for both VaR and ES in the present weekly application. This provides direct empirical support for the local approximation analysis in Section 2.

11 Daily-to-weekly horizon validation across major equity indices

The local Taylor analysis in Section 2 predicts that approximation error should increase as the return horizon becomes longer and the relevant tail becomes more extreme. We test this

implication directly using ten years of observed index data rather than relying on a square-root-of-time scaling rule. The data comprise daily and weekly close-to-close log returns for the S&P 500, NASDAQ Composite, and CAC 40 from 25 September 2016 to 25 September 2026. The weekly observations are constructed directly from weekly closing levels (week ending Friday); they are not obtained by multiplying daily risk estimates by $\sqrt{5}$. The daily samples contain 2,513 observations for the two U.S. indices and 2,559 observations for the CAC 40; synchronization across the three markets leaves 2,489 daily observations. Each weekly series contains 521 returns.

For each index return X_t , we compute

$$L_t^L = -X_t, \quad L_t^Q = -\left(X_t + \frac{1}{2}X_t^2\right), \quad L_t^E = 1 - e^{X_t}.$$

For the equally weighted three-index portfolio, the transformations are applied componentwise before aggregation:

$$L_{p,t}^L = -\frac{1}{3} \sum_{i=1}^3 X_{i,t}, \quad L_{p,t}^Q = -\frac{1}{3} \sum_{i=1}^3 \left(X_{i,t} + \frac{1}{2}X_{i,t}^2\right),$$

$$L_{p,t}^E = \frac{1}{3} \sum_{i=1}^3 (1 - e^{X_{i,t}}).$$

This componentwise exact revaluation is essential: in general, $\frac{1}{3} \sum_i e^{X_i} \neq \exp(\frac{1}{3} \sum_i X_i)$. Historical VaR is the empirical quantile and ES is the average of the corresponding upper-tail losses, using the same convention as in [Section 10](#).

11.1 Descriptive statistics, prices, and returns

Before comparing tail-risk approximations, [Table 9](#) documents the unconditional return environment at both horizons. All three markets are negatively skewed and strongly leptokurtic, with especially large daily excess kurtosis for the S&P 500 and CAC 40 and persistent weekly tail thickness for the CAC 40. These diagnostics make the Gaussian square-root-of-time benchmark deliberately demanding and motivate the direct horizon-by-horizon comparison that follows.

Table 9: Descriptive statistics of daily and weekly log returns, 2016–2026

Frequency	Market	N	Mean (%)	S.D. (%)	Ann. S.D. (%)	Min (%)	Max (%)	Skew.	Ex. kurt.
Daily	S&P 500	2513	0.051	1.141	18.118	-12.765	9.089	-0.677	16.755
	NASDAQ	2513	0.065	1.398	22.190	-13.149	11.478	-0.433	8.739
	CAC 40	2559	0.024	1.097	17.417	-13.098	8.056	-0.911	14.390
Weekly	S&P 500	521	0.244	2.389	17.224	-16.228	11.424	-0.904	7.352
	NASDAQ	521	0.313	2.820	20.337	-13.513	10.062	-0.409	2.119
	CAC 40	521	0.115	2.539	18.310	-22.142	10.164	-1.655	13.248

Note: returns are close-to-close log returns in percentage points. Annualized standard deviations use $\sqrt{252}$ for daily returns and $\sqrt{52}$ for weekly returns. Weekly returns are computed directly from week-ending closing levels rather than by scaling daily returns.




Figure 3: Normalized closing-price indices for the S&P 500, NASDAQ Composite, and CAC 40, September 2016–September 2026. Each series is set equal to 100 at the beginning of the sample.




Figure 4: Daily close-to-close log returns for the S&P 500, NASDAQ Composite, and CAC 40, 2016–2026. The pronounced volatility bursts, especially around 2020, are inconsistent with a homoskedastic Gaussian benchmark.




Figure 5: Weekly close-to-close log returns for the S&P 500, NASDAQ Composite, and CAC 40, 2016–2026. Weekly aggregation reduces some high-frequency variation but preserves large tail events.

11.2 Daily horizon

Table 10 reports the daily VaR and ES estimates. At all three confidence levels, the quadratic values are visually almost indistinguishable from exact revaluation. The difference between the linear and exact losses is already visible in the 99% tail, especially for NASDAQ.

Table 10: Daily historical VaR and ES: linear, quadratic, and exact revaluation (%)

Market	Method	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
S&P 500	Linear	1.670	2.802	2.395	3.656	3.342	4.894
	Quadratic	1.656	2.752	2.366	3.573	3.286	4.750
	Exact	1.656	2.753	2.366	3.575	3.286	4.754
NASDAQ	Linear	2.237	3.388	3.040	4.225	4.016	5.335
	Quadratic	2.212	3.321	2.994	4.123	3.936	5.173
	Exact	2.212	3.322	2.995	4.125	3.937	5.176
CAC 40	Linear	1.604	2.655	2.112	3.494	3.401	4.761
	Quadratic	1.591	2.610	2.090	3.419	3.343	4.628
	Exact	1.592	2.610	2.090	3.421	3.343	4.631
Equal-weight portfolio	Linear	1.562	2.584	2.166	3.341	3.059	4.575
	Quadratic	1.542	2.535	2.142	3.265	2.995	4.440
	Exact	1.542	2.536	2.142	3.267	2.996	4.443

Note: losses and risk measures are percentages of initial value. The portfolio is equally weighted across the S&P 500, NASDAQ Composite, and CAC 40 and uses synchronized trading dates.

To quantify the approximation quality, define for $M \in \{L, Q\}$ and $\rho \in \{\text{VaR}, \text{ES}\}$

$$\text{RE}_\rho^M(p) = 100 \times \frac{\rho_p^M - \rho_p^E}{\rho_p^E}.$$

Positive values denote overestimation relative to exact revaluation. As shown in Table 11, the quadratic error is below 0.08% in absolute value for every daily VaR/ES entry, whereas the linear error reaches approximately 3.1% for NASDAQ 99% ES.

Table 11: Daily relative errors against exact revaluation (%)

Market	Approx.	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
S&P 500	Linear	0.837	1.810	1.202	2.254	1.680	2.946
	Quadratic	-0.005	-0.032	-0.010	-0.046	-0.019	-0.072
NASDAQ	Linear	1.122	1.996	1.528	2.427	2.022	3.061
	Quadratic	-0.008	-0.034	-0.016	-0.049	-0.027	-0.075
CAC 40	Linear	0.804	1.728	1.060	2.155	1.710	2.813
	Quadratic	-0.004	-0.030	-0.007	-0.042	-0.019	-0.067
Equal-weight portfolio	Linear	1.294	1.885	1.103	2.289	2.129	2.958
	Quadratic	-0.011	-0.034	-0.008	-0.048	-0.020	-0.074

11.3 Weekly horizon

The weekly experiment provides a sharper test because the return magnitudes are larger and the Taylor remainder scales cubically to first order. Table 12 shows that the quadratic approximation remains remarkably accurate even at 99%, while the linear approximation becomes materially more conservative.

Table 12: Weekly historical VaR and ES: linear, quadratic, and exact revaluation (%)

Market	Method	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
S&P 500	Linear	3.329	5.703	4.761	7.376	6.135	10.101
	Quadratic	3.273	5.498	4.648	7.052	5.946	9.535
	Exact	3.274	5.504	4.650	7.063	5.950	9.558
NASDAQ	Linear	4.147	6.204	5.203	7.628	7.853	10.052
	Quadratic	4.061	5.984	5.068	7.306	7.544	9.528
	Exact	4.062	5.990	5.070	7.316	7.552	9.546
CAC 40	Linear	3.720	6.230	5.039	8.116	7.089	11.388
	Quadratic	3.651	5.965	4.912	7.688	6.838	10.604
	Exact	3.652	5.976	4.914	7.707	6.844	10.645
Equal-weight portfolio	Linear	3.263	5.516	4.789	7.056	6.039	9.633
	Quadratic	3.202	5.305	4.654	6.725	5.855	9.062
	Exact	3.203	5.312	4.657	6.738	5.859	9.087

Note: weekly returns are directly observed close-to-close log returns from weekly closing levels; no $\sqrt{5}$ scaling is used.

11.4 Direct weekly risk versus $\sqrt{5}$ scaling of daily risk

A separate question from the Taylor approximation is whether one-day tail-risk measures can be extrapolated to one week by the conventional square-root-of-time rule. Under independent, identically distributed Gaussian returns with negligible mean, both volatility-based VaR and ES satisfy the approximation

$$\rho_{p,1w} \approx \sqrt{5} \rho_{p,1d}, \quad \rho \in \{\text{VaR}, \text{ES}\}.$$

Table 13: Weekly relative errors against exact revaluation (%)

Market	Approx.	VaR ₉₅	ES ₉₅	VaR _{97.5}	ES _{97.5}	VaR ₉₉	ES ₉₉
S&P 500	Linear	1.674	3.617	2.400	4.434	3.099	5.683
	Quadratic	-0.019	-0.113	-0.038	-0.157	-0.064	-0.235
NASDAQ	Linear	2.088	3.565	2.624	4.260	3.978	5.295
	Quadratic	-0.029	-0.098	-0.046	-0.133	-0.105	-0.192
CAC 40	Linear	1.871	4.258	2.541	5.312	3.587	6.981
	Quadratic	-0.023	-0.178	-0.043	-0.251	-0.085	-0.386
Equal-weight portfolio	Linear	1.867	3.839	2.821	4.725	3.065	6.002
	Quadratic	-0.025	-0.136	-0.058	-0.190	-0.063	-0.281

To assess this approximation without confounding it with Taylor error, we use the *exact exponential revaluation* at both horizons. Define

$$SR_p^\rho = \frac{\rho_{p,1w}^E}{\sqrt{5} \rho_{p,1d}^E}, \quad \rho \in \{\text{VaR}, \text{ES}\}.$$

A ratio equal to one indicates exact square-root-of-time scaling; $SR < 1$ means that $\sqrt{5}$ scaling overstates directly observed weekly risk, whereas $SR > 1$ means that it understates weekly risk.

Table 14: Directly observed weekly VaR/ES versus $\sqrt{5}$ scaling of daily exact risk

Market	Level	Weekly VaR	$\sqrt{5}$ Daily VaR	SR^{VaR}	Weekly ES	$\sqrt{5}$ Daily ES	SR^{ES}
S&P 500	95%	3.274	3.702	0.884	5.504	6.155	0.894
	97.5%	4.650	5.291	0.879	7.063	7.994	0.883
	99%	5.950	7.349	0.810	9.558	10.629	0.899
NASDAQ	95%	4.062	4.946	0.821	5.990	7.428	0.806
	97.5%	5.070	6.696	0.757	7.316	9.225	0.793
	99%	7.552	8.803	0.858	9.546	11.575	0.825
CAC 40	95%	3.652	3.559	1.026	5.976	5.837	1.024
	97.5%	4.914	4.673	1.052	7.707	7.649	1.008
	99%	6.844	7.476	0.915	10.645	10.354	1.028
Equal-weight portfolio	95%	3.203	3.448	0.929	5.312	5.670	0.937
	97.5%	4.657	4.790	0.972	6.738	7.304	0.922
	99%	5.859	6.698	0.875	9.087	9.935	0.915

Note: all entries are based on exact exponential revaluation. Weekly risk is estimated directly from weekly close-to-close log returns. The scaled columns equal $\sqrt{5}$ times the corresponding daily historical VaR or ES. Ratios are observed weekly risk divided by scaled daily risk.

The square-root-of-time rule is therefore materially less accurate than the quadratic revaluation itself. For the S&P 500, NASDAQ, and the equally weighted portfolio, the scaling ratios are below one at all reported confidence levels, indicating systematic overstatement of observed weekly tail risk. At the 99% level, for example, the portfolio ratios are $SR^{VaR} = 0.875$ and $SR^{ES} = 0.915$, corresponding to scaled values about 14.3% and 9.3% above the directly observed weekly VaR and ES, respectively. The NASDAQ displays the largest departures from one at 95% and 97.5%.

The CAC 40 behaves differently. Its scaling ratios exceed one for 95% and 97.5% VaR and for ES at all three confidence levels, although the 99% VaR ratio falls below one. Hence square-root-of-time scaling can either overstate or understate weekly tail risk depending on the market and tail functional. This non-uniformity is consistent with the empirical presence of volatility clustering, temporal dependence, non-Gaussian tails, and aggregation effects, and it

shows why daily and weekly risk should be treated as distinct empirical horizons rather than as mechanically rescaled versions of the same object.

11.5 Weekend information and weekday heterogeneity

The conventional $\sqrt{5}$ rule also suppresses an important calendar-time distinction. A standard trading week contains approximately five trading sessions but spans seven calendar days. In particular, the Monday close-to-close return embeds information accumulated during the non-trading interval following the previous market close. This is closely related to the classical weekend-effect literature, which documents systematic differences between Monday and other weekday returns (French, 1980; Keim and Stambaugh, 1984). For the present purpose, the issue is not whether a persistent mean-return anomaly exists, but whether the distribution of the Monday return is sufficiently different from that of the other trading days to invalidate the homogeneous-i.i.d. approximation underlying simple square-root-of-time scaling.

Let C_t and O_t denote the close and open on trading day t . For a Monday observation $t \in \mathcal{M}$, define the close-to-close Monday return

$$r_t^{Mon} = \log\left(\frac{C_t}{C_{t^-}}\right),$$

where t^- denotes the immediately preceding trading day. We decompose it as

$$r_t^{Mon} = g_t + r_t^{intra}, \quad g_t = \log\left(\frac{O_t}{C_{t^-}}\right), \quad r_t^{intra} = \log\left(\frac{C_t}{O_t}\right).$$

Thus g_t measures the non-trading information gap carried into the Monday open, while r_t^{intra} measures the subsequent Monday intraday return. When a Friday is a holiday, C_{t^-} is the last available pre-Monday trading close, so g_t should be interpreted as a non-trading-gap proxy rather than a literal Friday-to-Monday return.

Table 15: Monday versus Tuesday–Friday daily log-return distributions, 2016–2026

Market	Subsample	N	Mean (%)	S.D. (%)	Minimum (%)	Skewness	Excess kurtosis
S&P 500	Monday	469	0.076	1.240	-12.765	-2.825	29.913
	Tue.–Fri.	2,044	0.045	1.118	-9.994	-0.010	12.227
NASDAQ	Monday	469	0.122	1.473	-13.149	-1.943	16.457
	Tue.–Fri.	2,044	0.052	1.380	-9.910	-0.018	6.576
CAC 40	Monday	503	-0.022	1.185	-8.764	-0.790	11.422
	Tue.–Fri.	2,056	0.035	1.075	-13.098	-0.943	15.400

Note: Returns are close-to-close daily log returns. Monday is identified by the calendar weekday of the closing observation. The Tuesday–Friday benchmark pools the remaining standard trading days. Moments are bias-corrected sample statistics; percentages are percentage points.

The evidence rejects the idea that all five daily increments entering a trading week are distributionally homogeneous. Monday volatility exceeds the Tuesday–Friday standard deviation for all three indices, and the difference is accompanied by much stronger downside asymmetry and tail thickness for the two U.S. indices. The decomposition in Table 16 is even more revealing: the non-trading gap has excess kurtosis of approximately 40.0 for the S&P 500 and 14.5 for NASDAQ, far above the corresponding Monday intraday values of 6.9 and 5.8. For the CAC 40,

Table 16: Decomposition of Monday returns into non-trading gap and intraday components

Market	Component	N	Mean (%)	S.D. (%)	Skewness	Excess kurtosis
S&P 500	Non-trading gap g_t	469	0.014	0.687	-3.612	39.992
	Monday intraday r_t^{intra}	469	0.062	0.814	-0.941	6.884
NASDAQ	Non-trading gap g_t	469	0.019	1.013	-2.212	14.541
	Monday intraday r_t^{intra}	469	0.104	1.029	-0.829	5.791
CAC 40	Non-trading gap g_t	503	-0.029	0.917	-2.085	13.867
	Monday intraday r_t^{intra}	503	0.007	0.788	0.309	7.496

Note: $g_t = \log(O_t/C_{t-})$ uses the Monday open and the previous available trading close; $r_t^{intra} = \log(C_t/O_t)$. Their sum equals the Monday close-to-close log return. The gap includes all information arriving during the non-trading interval and is therefore a proxy for weekend-information accumulation, not a pure causal weekend effect.

the gap is also strongly negatively skewed and leptokurtic, although Tuesday–Friday close-to-close returns retain substantial tail thickness.

This weekday heterogeneity provides a concrete mechanism through which weekly tail risk can depart from $\sqrt{5}$ scaling. Writing the weekly return schematically as

$$R_w = g_{Mon} + r_{Mon}^{intra} + r_{Tue} + r_{Wed} + r_{Thu} + r_{Fri},$$

its variance is

$$\text{Var}(R_w) = \sum_{j=1}^5 \text{Var}(r_j) + 2 \sum_{i < j} \text{Cov}(r_i, r_j),$$

with the Monday component itself split into a non-trading gap and an intraday increment. The textbook relation $\text{Var}(R_w) \approx 5\sigma_d^2$ therefore requires not only weak serial dependence but also approximate homogeneity of the daily marginal distributions. The observed Monday and non-trading-gap diagnostics show that this condition is empirically demanding. They help explain why the scaling ratios in Table 14 can deviate materially from one, although they should not be interpreted as a causal decomposition of the scaling error: volatility clustering, cross-day dependence, and non-Gaussian aggregation remain additional channels.

11.6 Horizon effect and interpretation

Three empirical conclusions emerge. First, the quadratic approximation is effectively exact at the daily horizon for VaR and ES up to 99%: its absolute relative error never reaches one tenth of one percent. Second, the error rises at the weekly horizon, exactly as predicted by the third-order remainder, but remains economically small. The largest quadratic discrepancy in Table 13 is the CAC 40 weekly 99% ES error, approximately -0.386% . Third, the omission of the quadratic term is increasingly costly in the tail. For weekly 99% ES, the linear approximation overstates exact risk by approximately 5.68% for the S&P 500, 5.29% for NASDAQ, 6.98% for the CAC 40, and 6.00% for the equally weighted portfolio.

This pattern has a simple analytical explanation. For a negative log return x ,

$$1 - e^x = -x - \frac{x^2}{2} - \frac{x^3}{6} - \dots.$$

The linear loss $-x$ omits the positive correction $x^2/2$ in the price expansion and therefore

exceeds the exact percentage loss in the relevant left tail. The quadratic approximation restores this dominant curvature term; its remaining discrepancy is third order. As return magnitudes rise from one day to one week, the third-order residual grows, but much more slowly than the first-order misspecification incurred by ignoring curvature altogether over the empirical range considered here.

The direct daily-to-weekly comparison therefore provides an empirical validation domain for the intrinsic-quadratic approach: over these three major equity indices and the synchronized equally weighted portfolio, second-order revaluation remains highly accurate through the weekly 99% tail. At the same time, Table 14 shows that the conventional $\sqrt{5}$ rule can generate much larger errors than the quadratic Taylor approximation itself. This conclusion should not be extrapolated mechanically to monthly horizons, 99.9% tails, or crisis scenarios involving exceptionally large single-period moves; those remain settings in which full exponential revaluation is preferable.

12 Model risk, tail validity, and practical diagnostics

12.1 The quadratic map is not globally monotone in negative returns

For one asset,

$$L^Q(x) = -a \left(x + \frac{x^2}{2} \right) = \frac{a}{2} - \frac{a}{2}(x+1)^2.$$

The quadratic loss increases as x falls from 0 toward -1 , but decreases again when $x < -1$. By contrast, the exact loss $a(1 - e^x)$ increases monotonically as $x \rightarrow -\infty$. This is not a defect of the algebra; it is the unavoidable consequence of using a local Taylor polynomial outside its validity region.

A practical implementation should therefore report, together with quadratic VaR/ES, at least one of the following diagnostics:

- (i) the simulated probability $\mathbb{P}(\|X\|_\infty > \delta)$ for a prespecified local region;
- (ii) the deterministic error bound ε_δ from Proposition 2.1;
- (iii) the empirical discrepancy between quadratic and exact revaluation in a Monte Carlo benchmark.

12.2 Student- t log returns

A multivariate Student- t model is attractive for heavy-tailed risk-factor changes, but direct Student- t modeling of log returns has a well-known implication: positive exponential moments do not exist. This is not an obstacle for computing the upper tail of the *loss* $L^E = \sum_i a_i(1 - e^{X_i})$, which is bounded above by P_0 in a long-only portfolio, but it does make some unconditional price moments economically problematic. For pricing applications or long horizons, tempered or exponentially tilted heavy-tailed models may be preferable.

12.3 Short positions

If some $a_i < 0$, then W is indefinite and Proposition 2.3 no longer applies. The appropriate framework is the general quadratic representation (9) and Theorem 4.1. Tail existence conditions must then be checked explicitly. This is also the natural route for portfolios combining long and short cash equity positions.

12.4 Comparison with delta–gamma VaR

The two quadratic problems should not be conflated:

	Intrinsic securities quadratic	Delta–gamma derivatives quadratic
Origin of quadraticity	e^{X_i} price mapping	second derivative of derivative price
Quadratic matrix	holdings/exposures	gamma matrix
Cross terms	absent in basic cash-equity expansion	generally present
Definiteness	long-only gives negative-semidefinite loss curvature	may be indefinite
Natural geometry	ellipsoid/ball after completion of squares	ellipsoid or hyperboloid

This distinction explains why the positive-form reduction of Theorem 5.1 is much simpler than a fully general delta–gamma treatment.

13 Extensions

Several extensions follow naturally.

13.1 Conditional elliptical risk

We now formulate the one-step-ahead version of the quadratic VaR–ES result. Let $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ be a filtered probability space. At date t , let $a_t = (a_{1,t}, \dots, a_{n,t})^\top$ denote the vector of monetary exposures, define

$$W_t := \text{diag}(a_{1,t}, \dots, a_{n,t}), \quad P_t := \sum_{i=1}^n a_{i,t},$$

and let X_{t+1} denote the vector of log returns over $(t, t+1]$. The one-step intrinsic quadratic loss is

$$L_{t+1}^Q = -a_t^\top X_{t+1} - \frac{1}{2} X_{t+1}^\top W_t X_{t+1} = \frac{P_t}{2} - \frac{1}{2} Q_{t+1}, \quad (31)$$

where

$$Q_{t+1} := (X_{t+1} + \mathbf{1})^\top W_t (X_{t+1} + \mathbf{1}). \quad (32)$$

For a generic loss L_{t+1} , define the conditional VaR by the $\mathcal{F}_t$ -measurable conditional quantile

$$\text{VaR}_{p,t}(L_{t+1}) := \text{ess inf} \left\{ \ell \in L^0(\mathcal{F}_t) : \mathbb{P}(L_{t+1} \leq \ell \mid \mathcal{F}_t) \geq p \right\}, \quad (33)$$

and define conditional Expected Shortfall by

$$\text{ES}_{p,t}(L_{t+1}) := \frac{1}{1-p} \int_p^1 \text{VaR}_{u,t}(L_{t+1}) du, \quad (34)$$

whenever the conditional quantile integral is finite.

Assumption 13.1 (Conditional elliptical predictive law). For a fixed confidence level $p \in (0, 1)$, suppose that, almost surely:

(C1) $a_{i,t} > 0$ for every i , and a_t is finite and $\mathcal{F}_t$ -measurable;

(C2) conditional on $\mathcal{F}_t$,

$$X_{t+1} \mid \mathcal{F}_t \sim \mathcal{E}_n(\mu_t, \Sigma_t, \psi_t), \quad (35)$$

where μ_t and Σ_t are $\mathcal{F}_t$ -measurable, $\Sigma_t \in \mathbb{S}_{++}^n$, and the conditional generator ψ_t is specified so that (35) defines a proper predictive distribution;

(C3) the conditional CDF

$$F_{Q,t}(q) := \mathbb{P}(Q_{t+1} \leq q \mid \mathcal{F}_t) \quad (36)$$

is continuous and strictly increasing in a neighbourhood of its $(1-p)$ -quantile;

(C4) the conditional quantile

$$\tau_{p,t} := F_{Q,t}^{-1}(1-p) \quad (37)$$

is finite.

Define the conditional lower partial first moment

$$M_{Q,t}(q) := \mathbb{E}[Q_{t+1} \mathbf{1}\{Q_{t+1} \leq q\} \mid \mathcal{F}_t]. \quad (38)$$

Theorem 13.2 (Conditional quadratic VaR and Expected Shortfall). *Under Assumption 13.1, the one-step-ahead intrinsic quadratic loss in (31) satisfies, almost surely,*

$$\boxed{\text{VaR}_{p,t}(L_{t+1}^Q) = \frac{P_t - \tau_{p,t}}{2}} \quad (39)$$

and

$$\boxed{\text{ES}_{p,t}(L_{t+1}^Q) = \frac{P_t}{2} - \frac{M_{Q,t}(\tau_{p,t})}{2(1-p)}}. \quad (40)$$

Moreover,

$$\text{VaR}_{p,t}(L_{t+1}^Q) \leq \frac{P_t}{2}, \quad \text{ES}_{p,t}(L_{t+1}^Q) \leq \frac{P_t}{2} \quad a.s. \quad (41)$$

No finite unconditional second moment of X_{t+1} is required for (40): because $Q_{t+1} \geq 0$,

$$0 \leq Q_{t+1} \mathbf{1}\{Q_{t+1} \leq \tau_{p,t}\} \leq \tau_{p,t},$$

so $M_{Q,t}(\tau_{p,t}) < \infty$ whenever $\tau_{p,t} < \infty$.

Proof. Conditional on $\mathcal{F}_t$, P_t and W_t are fixed and $L_{t+1}^Q = (P_t - Q_{t+1})/2$ is a strictly decreasing affine transformation of Q_{t+1} . Hence, for every $\mathcal{F}_t$ -measurable threshold ℓ ,

$$\mathbb{P}(L_{t+1}^Q > \ell \mid \mathcal{F}_t) = \mathbb{P}(Q_{t+1} < P_t - 2\ell \mid \mathcal{F}_t).$$

By continuity of $F_{Q,t}$ at $\tau_{p,t}$, the conditional upper-tail probability is $1 - p$ when $P_t - 2\ell = \tau_{p,t}$, which yields (39). The same monotonicity implies

$$\{L_{t+1}^Q \geq \text{VaR}_{p,t}(L_{t+1}^Q)\} = \{Q_{t+1} \leq \tau_{p,t}\} \quad \text{a.s.}$$

Therefore, using conditional continuity so that the tail probability equals $1 - p$,

$$\begin{aligned} \text{ES}_{p,t}(L_{t+1}^Q) &= \frac{1}{1-p} \mathbb{E} \left[\frac{P_t - Q_{t+1}}{2} \mathbf{1}_{\{Q_{t+1} \leq \tau_{p,t}\}} \mid \mathcal{F}_t \right] \\ &= \frac{P_t}{2} - \frac{1}{2(1-p)} \mathbb{E}[Q_{t+1} \mathbf{1}_{\{Q_{t+1} \leq \tau_{p,t}\}} \mid \mathcal{F}_t], \end{aligned}$$

which proves (40). Since $Q_{t+1} \geq 0$, (41) follows immediately. $\square$

Remark 13.3 (Role of the conditional elliptical specification). The identities (39)–(40) are geometric consequences of the long-only quadratic portfolio map and therefore hold for any conditional predictive law satisfying the stated continuity conditions. Conditional ellipticity supplies a tractable model for computing $F_{Q,t}$ and $M_{Q,t}$. In particular, Σ_t may be generated by EWMA, multivariate GARCH or BEKK, DCC, or realized-covariance methods. A stochastic-volatility specification is covered only when the resulting one-step predictive distribution, conditional on the chosen information set, remains elliptical; integrating out latent volatilities need not preserve ellipticity in general.

Remark 13.4 (Time-varying density representation). If the conditional elliptical law admits a density

$$f_t(x \mid \mathcal{F}_t) = c_t |\Sigma_t|^{-1/2} g_t \left((x - \mu_t)^\top \Sigma_t^{-1} (x - \mu_t) \right),$$

then the static sphere–ball representation carries over pointwise in t after replacing (μ, Σ, g, W) by $(\mu_t, \Sigma_t, g_t, W_t)$. Thus all model dynamics enter the risk forecast through the time-varying objects $F_{Q,t}$, $M_{Q,t}$, and $\tau_{p,t}$, while the VaR–ES transformation itself is unchanged.

Dynamic mixtures. Replace constant π_k by filtered regime probabilities $\pi_{k,t}$. This produces a Markov-switching or latent-regime quadratic VaR/ES model without changing the mixture decomposition in Theorem 6.1.

General positive quadratic maps. The methodology extends immediately to $L = c - Q/2$ whenever Q is a positive quadratic form. The cash-equity portfolio is then a financially interpretable special case.

Capital allocation. Once ES_p is available, one may study marginal or Euler allocations with respect to monetary exposures a_i . Because the threshold τ_p also depends on a , differentiation requires care; this is a promising extension for risk budgeting and bancassurance capital allocation.

14 Conclusion

This paper reconstructs the quadratic portfolio VaR/ES problem around a simple observation: a portfolio that is linear in security prices is intrinsically nonlinear in log returns. The second-order exponential expansion creates a quadratic loss without invoking derivative Greeks. Under long-only positions, completing the square turns that loss into an affine decreasing transform of a positive quadratic form. This immediately yields compact joint formulas for VaR and ES in terms of a lower quadratic-form quantile and a lower partial first moment.

The reformulation clarifies several points that were less transparent in the original treatment. The coefficient $1/2$ is structurally unavoidable; the restriction $VaR < P_0/2$ reflects the maximum of the quadratic approximation rather than an arbitrary numerical assumption; and the quadratic approximation is local, not a globally valid model of catastrophic equity losses. These observations are important when the risk-factor distribution is heavy-tailed.

For Gaussian factors, generalized noncentral chi-square algorithms provide a modern computational route. For Student- t factors, the Gaussian scale-mixture representation reduces the problem to one-dimensional mixing. Finite mixtures follow by linearity of component probabilities and truncated moments. A simulation study shows that, at moderate horizons, the intrinsic quadratic approximation can track exact exponential revaluation very closely and improve materially on a linear approximation.

The main practical recommendation is therefore not to replace exact revaluation blindly by a quadratic formula, but to use the quadratic framework as an analytically tractable local risk model accompanied by explicit approximation diagnostics. This interpretation preserves the original contribution while making the method compatible with modern standards of tail-risk modeling and model validation.

A Additional proofs and formulas

A.1 Quantile reversal under a decreasing affine map

Let $L = a - bQ$ with $b > 0$ and continuous F_Q . Then

$$\mathbb{P}(L \leq \ell) = \mathbb{P}\left(Q \geq \frac{a - \ell}{b}\right) = 1 - F_Q\left(\frac{a - \ell}{b}\right).$$

Setting this equal to p gives

$$F_Q\left(\frac{a - VaR_p}{b}\right) = 1 - p,$$

whence $VaR_p = a - bF_Q^{-1}(1 - p)$. Theorem 5.1 is the case $a = P_0/2$ and $b = 1/2$.

A.2 Lower partial moment identity

For nonnegative Q and $q \geq 0$,

$$M_Q(q) = \int_{[0,q]} t dF_Q(t).$$

Integration by parts yields

$$M_Q(q) = qF_Q(q) - \int_0^q F_Q(t) dt,$$

provided $F_Q(0^-) = 0$. This identity is numerically useful because it transforms a truncated first moment into a one-dimensional integral of the CDF.

A.3 A quantile-error implication of the local CDF bound

Suppose F_Q is continuously differentiable and has density bounded below by $m > 0$ on an interval containing the relevant quantile. Proposition 2.5 implies heuristically that a small CDF perturbation η translates into a quantile perturbation of order η/m , in addition to the deterministic horizontal error ε_δ . A rigorous finite-sample statement can be obtained by inverting the two CDF inequalities with left- and right-continuous generalized inverses. This provides a practical route to choosing δ : balance the Taylor error ε_δ against the tail probability $\mathbb{P}(A_\delta^c)$.

B Notation summary

Symbol	Meaning
P_0	initial portfolio value
$a_i = \theta_i S_i(0)$	initial monetary exposure in security i
X_i	log return of security i
L^E	exact loss
L^L	linear loss approximation
L^Q	intrinsic quadratic loss approximation
$W = \text{diag}(a)$	exposure matrix
$Q = (X + \mathbf{1})^\top W (X + \mathbf{1})$	positive quadratic form
p	confidence level
$\tau_p = F_Q^{-1}(1 - p)$	lower quadratic-form quantile associated with upper loss tail
$M_Q(q)$	$\mathbb{E}[Q \mathbf{1}\{Q \leq q\}]$

References

- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228. [doi:10.1111/1467-9965.00068](https://doi.org/10.1111/1467-9965.00068).
- Brummelhuis, R., Córdoba, A., Quintanilla, M., and Seco, L. (2002). Principal component Value-at-Risk. *Mathematical Finance*, 12(1), 23–43.

- Cardenas, J., Fruchard, E., Koehler, E., Michel, C., and Thomaseau, I. (1997). VaR: One step beyond. *Risk*, 10(10), 72–75.
- Davies, R. B. (1980). Algorithm AS 155: The distribution of a linear combination of χ^2 random variables. *Applied Statistics*, 29(3), 323–333. doi:10.2307/2346911.
- Duffie, D., and Pan, J. (1997). An overview of Value at Risk. *The Journal of Derivatives*, 4(3), 7–49.
- Fang, K.-T., Kotz, S., and Ng, K.-W. (1990). *Symmetric Multivariate and Related Distributions*. Chapman & Hall, London.
- French, K. R. (1980). Stock returns and the weekend effect. *Journal of Financial Economics*, 8(1), 55–69. doi:10.1016/0304-405X(80)90021-5.
- Keim, D. B., and Stambaugh, R. F. (1984). A further investigation of the weekend effect in stock returns. *The Journal of Finance*, 39(3), 819–835. doi:10.1111/j.1540-6261.1984.tb03675.x.
- Genz, A. (2003). Fully symmetric interpolatory rules for multiple integrals over hyperspherical surfaces. *Journal of Computational and Applied Mathematics*, 157, 187–195.
- Imhof, J. P. (1961). Computing the distribution of quadratic forms in normal variables. *Biometrika*, 48(3/4), 419–426.
- Sadefo Kamdem, J. (2003a). Value-at-Risk and Expected Shortfall for quadratic portfolio of securities with mixture of elliptic distributed risk factors. arXiv:cs/0310043; later presented in *Computing in Economics and Finance 2004*, Society for Computational Economics.
- Sadefo Kamdem, J. (2003b). Value-at-Risk and Expected Shortfall for linear portfolios with elliptically distributed risk factors. arXiv:math/0309211.
- Sadefo Kamdem, J., and Genz, A. (2008). Approximation of multiple integrals over hyperboloids with application to a quadratic portfolio with options. *Computational Statistics & Data Analysis*, 52(7), 3389–3407. doi:10.1016/j.csda.2007.12.006.
- Landsman, Z. M., and Valdez, E. A. (2003). Tail conditional expectations for elliptical distributions. *North American Actuarial Journal*, 7(4), 55–71. doi:10.1080/10920277.2003.10596118.
- McNeil, A. J., Frey, R., and Embrechts, P. (2015). *Quantitative Risk Management: Concepts, Techniques and Tools*, 2nd ed. Princeton University Press.
- J.P. Morgan/Reuters (1996). *RiskMetrics Technical Document*, 4th ed. J.P. Morgan, New York.
- Rockafellar, R. T., and Uryasev, S. (2000). Optimization of Conditional Value-at-Risk. *The Journal of Risk*, 2(3), 21–41. doi:10.21314/JOR.2000.038.
- Rockafellar, R. T., and Uryasev, S. (2002). Conditional Value-at-Risk for general loss distributions. *Journal of Banking & Finance*, 26(7), 1443–1471.