Jan 28, 2025

Comparative Analysis of Risk Metrics for Linear Portfolios:

Application to Spot and Perpetual Contracts

Abstract

This paper presents a comprehensive framework for risk assessment in linear investment portfolios, including spot and perpetual contracts, by estimating Value at Risk (VaR) and Conditional Value at Risk (CVaR) using four methodologies. The study emphasises the importance of correctly accounting for the signs and magnitudes of net positions (long and short) within the portfolio covariance matrix to reflect economic exposure accurately. Portfolio variance and standard deviation are calculated using signed weights to capture both the direction and magnitude of positions.

Four VaR estimation methods are employed: the Normal distribution method, the Skewed Student's t-distribution method, the Cornish-Fisher expansion, and Monte Carlo simulation. CVaR, as an extension of VaR, is estimated using the Cornish-Fisher and Monte Carlo methods to quantify the average losses beyond the VaR threshold, providing a more comprehensive view of tail risks.

Backtesting results, evaluated using the Kupiec Proportion of Failures Test, indicate that the Cornish-Fisher method delivers the most accurate VaR estimates, effectively capturing skewness and kurtosis in the return distribution while accounting for asset correlations. The findings support the use of Cornish-Fisher VaR for linear portfolios and highlight the practical advantages of CVaR in stress scenarios, where understanding potential losses beyond VaR is critical.

This study underscores the significance of proper weighting in risk modelling and advocates for adaptive methods, such as exponentially weighted moving averages (EWMA), to enhance the sensitivity of risk estimates to changing market conditions. The proposed framework provides valuable insights for risk management, enabling robust assessment of potential losses in portfolios exposed to diverse market dynamics.

1. Introduction

In financial risk management, accurately estimating potential losses in investment portfolios is critical for informed decision-making, regulatory compliance, and optimising portfolio performance. Value at Risk (VaR) has emerged as a cornerstone metric for quantifying market risk, representing the maximum expected loss over a specified time horizon at a given confidence level. However, traditional VaR estimation methods often rely on simplifying assumptions, such as normally distributed returns, which may not adequately capture the complexities inherent in portfolios comprising various linear instruments, such as spot and perpetual contracts.

Linear portfolios, widely utilised across asset classes, are characterised by their direct price dependency on underlying assets, making them particularly sensitive to volatility, correlations, and market movements. These portfolios often include both long and short positions to capitalise on directional or relative price changes. This dual exposure introduces challenges in risk modelling, particularly in capturing offsetting effects and accurately reflecting the economic exposure of the portfolio.

While VaR provides an estimate of the threshold of potential losses, it does not offer insights into the magnitude of losses beyond that threshold, which is critical in stress scenarios. To address this limitation, Conditional Value at Risk (CVaR), also known as Expected Shortfall (ES), is employed as an extension. CVaR measures the average loss given that losses exceed the VaR threshold, offering a more comprehensive view of tail risks. This metric is particularly valuable for understanding the severity of potential losses during extreme market events, providing an additional layer of risk assessment for portfolio managers.

This paper proposes a robust framework for estimating both VaR and CVaR for linear portfolios by leveraging the covariance matrix with signed weights to account for the direction and magnitude of net positions. This approach ensures that both long and short positions are incorporated accurately, offering a comprehensive view of portfolio risk.

Four distinct methodologies are employed to estimate VaR:

1. **Normal Distribution Method**: Assumes normally distributed returns and calculates VaR using portfolio variance derived from the covariance matrix.
2. **Skewed Student's t-Distribution Method**: Accounts for skewness and excess kurtosis by fitting a skewed Student's t-distribution to historical returns.
3. **Cornish-Fisher Expansion**: Adjusts quantiles of the normal distribution to incorporate higher moments (skewness and kurtosis), improving accuracy in non-normal return distributions.
4. **Monte Carlo Simulation**: Generates a distribution of portfolio returns using simulated asset returns based on the covariance matrix.

For CVaR estimation, the Cornish-Fisher expansion and Monte Carlo simulation methods are extended to evaluate the expected losses beyond the VaR threshold, demonstrating their utility in quantifying tail risk.

The effectiveness of these methodologies is assessed through backtesting, employing the Kupiec Proportion of Failures Test to evaluate their ability to predict potential losses under various market scenarios. The findings highlight the strengths and limitations of each method, with the Cornish-Fisher Expansion demonstrating superior performance in capturing non-linearities and tail risks in portfolio returns.

This paper is structured as follows: **Section 2** describes the data and methodology, emphasising the importance of signed portfolio weights. **Section 3** explores the VaR and CVaR estimation methods. **Section 4** presents backtesting results and comparative analyses. Finally, **Section 5** concludes with recommendations for practitioners and implications for broader risk management applications.

2. Data and Methodology

2.1 Data

The dataset employed in this study consists of daily price data for 159 digital assets, encompassing both long and short positions across various sectors. The data spans multiple years, allowing for robust risk estimation and analysis. A 180-day rolling window is used to calculate portfolio risk metrics, with backtesting conducted over a one-year period. Daily log returns are computed to stabilise variance and account for compounding effects, enabling a more consistent evaluation of portfolio volatility.

The portfolio composition prioritises larger positions in high market capitalisation assets due to their higher liquidity and lower price volatility. This focus aligns with common practices in linear portfolio construction, where the stability of large-cap assets mitigates risks associated with low-cap assets, which are included in smaller allocations for diversification.

Despite the robustness of the dataset, several constraints merit discussion. The focus on large-cap assets may limit the generalisability of results to portfolios that heavily include small-cap assets, which exhibit higher idiosyncratic risk. Additionally, the reliance on historical data assumes that past volatility and correlations are representative of future market behaviour, an assumption that may falter during structural changes or extreme market conditions. While the dataset provides a diverse cross-section of digital assets, its ability to reflect the broader dynamics of more volatile or illiquid markets remains limited.

By concentrating on linear instruments, such as spot and perpetual contracts, this study provides a focused analysis but does not extend to portfolios with nonlinear exposures, such as those involving options or dated futures. Furthermore, the covariance structure employed, though effective over the rolling window, may not fully capture time-varying correlations during periods of market stress. These factors underline the need for cautious interpretation of results, particularly in applications beyond the specific portfolio characteristics studied.

Figure 1: Portfolio Positions

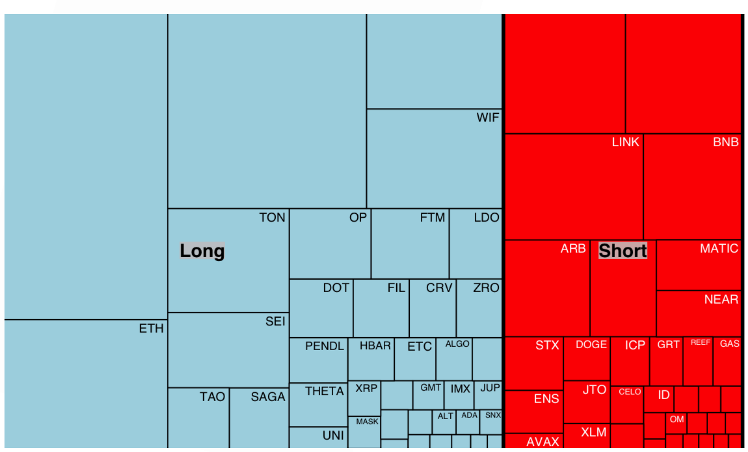

Figure 2: Portfolio Net Positions Size

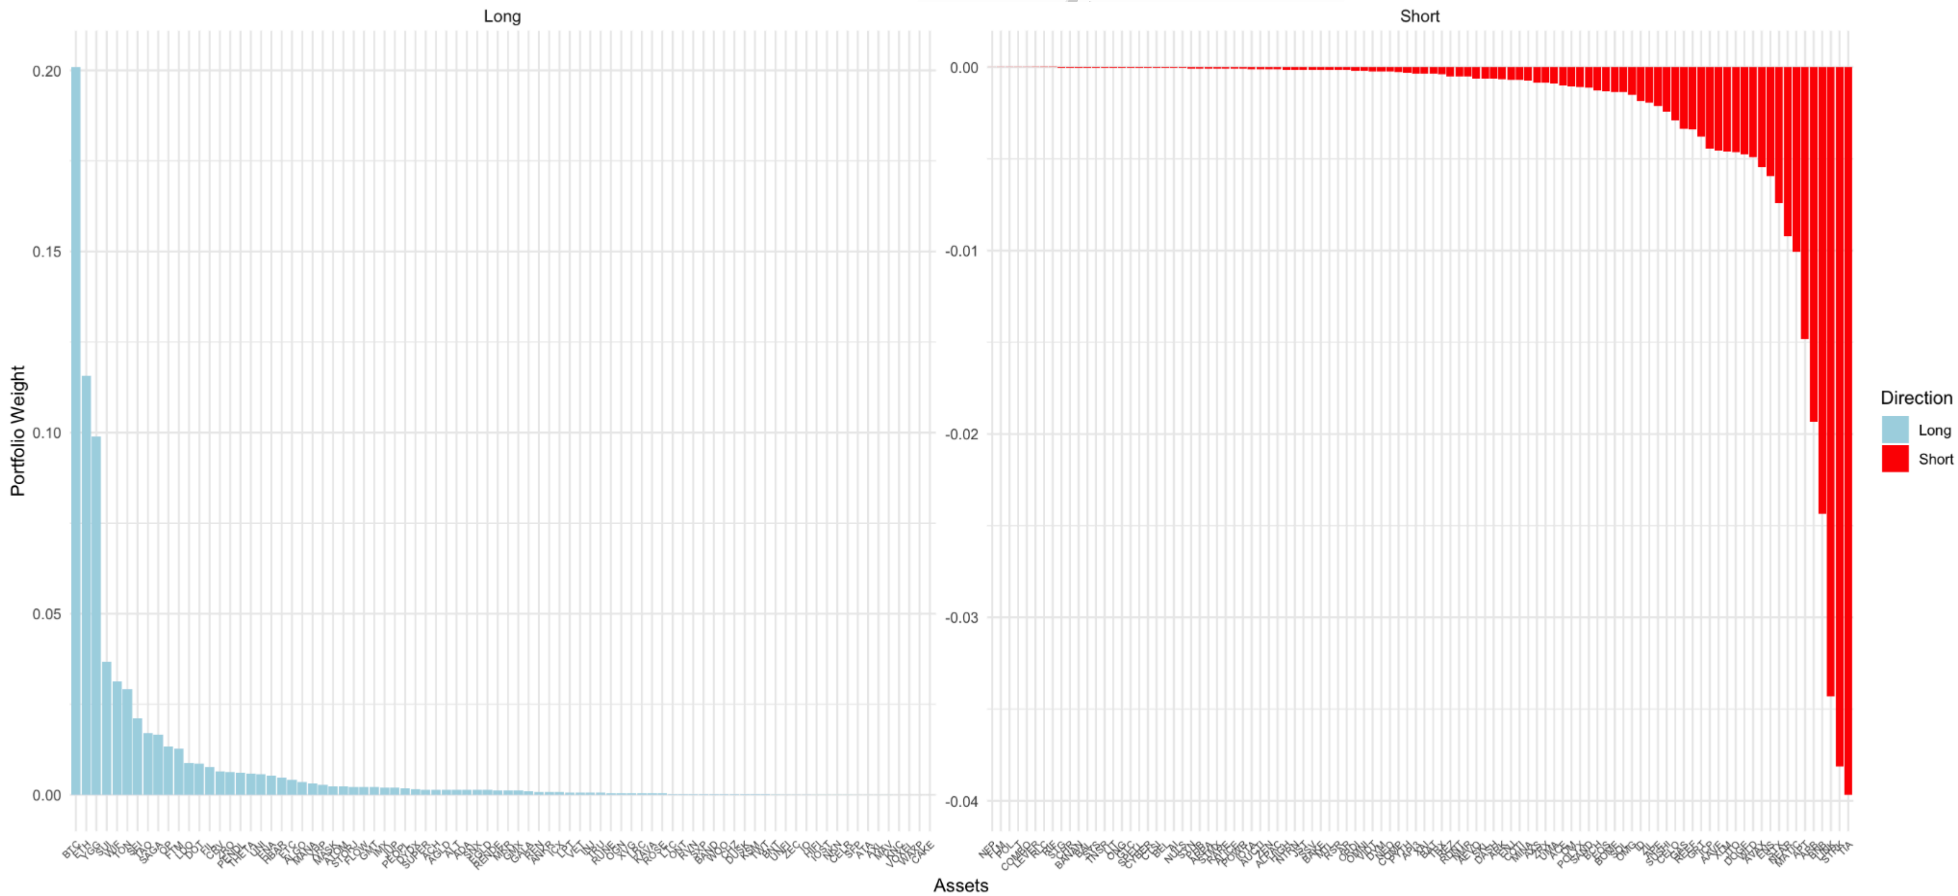

2.2 Portfolio Covariance Approach

The portfolio covariance approach forms the foundation of this study's risk modelling framework. By incorporating signed portfolio weights and the covariance matrix of asset returns, this methodology ensures that both the magnitude and direction of each position (long or short) are accurately reflected in risk estimates. This section outlines the key steps in the approach: calculating asset returns, estimating the covariance matrix, determining portfolio weights, and deriving portfolio variance and standard deviation.

2.2.1 Calculation of Asset Returns

Daily log returns are calculated for each asset to standardise the return series and stabilise variance over time. Log returns are defined as:

$$r_{i,t} = ln\left(\frac{P_{i,t}}{P_{i,t-1}}\right)$$

Where $r_{i,t}$ is the log return of asset *i* on day *t*. $P_{i,t}$ is the closing price of asset *i* on day *t* and $P_{i,t-1}$ is the closing price on the previous day. This method accounts for continuous compounding and stabilises the variance of the returns, which is beneficial for statistical analysis.

2.2.2 Estimation of the Covariance Matrix

The covariance matrix Σ of asset returns is estimated using the historical log returns:

$$\Sigma = Cov(R)$$

Where ***R*** represents the matrix of asset returns. The diagonal elements of Σ capture the variances of individual asset returns, while the off-diagonal elements represent the covariances between pairs of assets. This matrix provides a comprehensive view of the relationships between assets, enabling the modelling of diversification effects and co-movements.

2.2.3 Portfolio Weights with Signed Positions

In traditional portfolio theory, weights are determined based on the proportion of capital allocated to each asset and typically sum to one. For linear portfolios that include both long and short positions, weights must also account for the direction and magnitude of net positions to reflect true economic exposure. The portfolio weight for each asset is defined as:

The weights are calculated as:

$$\omega_i = \frac{netPosition_i}{\sum_{j=1}^{N} \left|NetPosition_j\right|}$$

Where;

- ω_i is the weight of asset *i* in the portfolio
- $NetPosition_i$ is the net position of asset *i,* positive for long positions and negative for short positions
- *N* is the total number of assets in the Portfolio

where ω_i is the weight of asset *i*, $NetPosition_i$ i represents the net position of asset *i* (positive for long positions, negative for short positions), and *N* is the total number of assets in the portfolio. This approach allows the weights to reflect offsetting effects, ensuring an accurate representation of portfolio exposure. Notably, the weights may not sum to one and can sum to zero in market-neutral configuration.

2.2.4 Calculation of Portfolio Variance and Standard Deviation

The portfolio variance is calculated using the signed weights and the covariance matrix:

$$\sigma_p^2 = w^T \Sigma w$$

Where σ_p^2 is the portfolio variance, w is the vector of portfolio weights, including signs for long and short positions and Σ is the covariance matrix of asset returns. The portfolio standard deviation. σ_p, is the square root of the variance:

$$\sigma_p = \sqrt{\sigma_p^2}$$

This calculation incorporates both the magnitude and direction of each position, capturing the hedging effects of the portfolio's structure. The standard deviation serves as a measure of the portfolio's volatility, forming a critical input for VaR and CVaR calculations.

Figure 2 presents a sample standard deviation using the signed weights and covariance matrix.

Figure 3: Histogram of Predicted Portfolio Standard Deviations

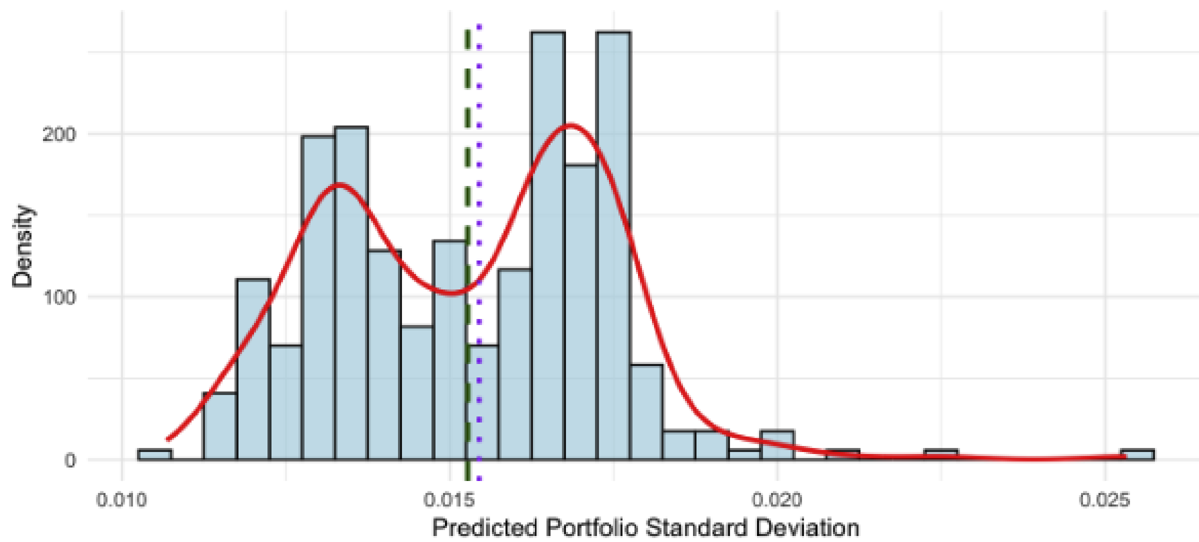

Red line: Density curve, dashed green line: Mean and dotted purple line: Median.

The portfolio covariance approach ensures a rigorous representation of economic exposure in risk estimation for linear portfolios. By leveraging signed weights and the covariance matrix, this methodology accounts for the interplay between long and short positions, providing a robust foundation for calculating portfolio volatility and assessing potential losses. However, as discussed, the use of historical covariance matrices assumes stationarity, which may not hold during periods of market stress or regime shifts, a limitation addressed in subsequent sections.

3. Value at Risk Estimation Methods

Accurately estimating Value at Risk (VaR) for linear portfolios requires methodologies that can account for the unique characteristics of portfolio return distributions, including volatility clustering, skewness, and kurtosis. In this study, four distinct approaches are employed to estimate VaR: the Normal Distribution Method, the Skewed Student's t-Distribution Method, the Cornish-Fisher Expansion, and Monte Carlo Simulation. Each method is suited to different assumptions about return behaviour and levels of computational complexity, enabling a comprehensive comparison of their performance.

While the Normal Distribution Method serves as a benchmark with its simplicity and widespread use, the other methods aim to address its limitations by incorporating asymmetry and fat tails. These enhancements allow for more precise risk assessment in portfolios where return distributions deviate from normality.

The subsequent subsections provide detailed explanations of each method, outlining their theoretical underpinnings, computational processes, and practical applications in the context of linear portfolios.

3.1 Normal Distribution Method

The Normal Distribution Method is a foundational approach to VaR estimation, based on the assumption that portfolio returns follow a normal distribution. This method is computationally straightforward, relying on the portfolio's standard deviation and the Z-score corresponding to the chosen confidence level (α) to estimate potential losses:

$$VaR_{Normal} = z_{\alpha} \times \sigma_{p}$$

Where;

- z_{α} is the Z-score for the confidence level α (e.g. $z_{0.05}$ = -1.645 for 95% confidence level).
- σ_{p} is the portfolio standard deviation as calculated from the covariance matrix.

The key advantage of this method lies in its simplicity and the ease with which it can be implemented. However, it assumes that returns are symmetrically distributed and neglects the skewness and kurtosis often observed in financial return distributions. These assumptions can lead to underestimation of tail risks, particularly in turbulent market conditions.

Despite its limitations, the Normal Distribution Method remains widely used as a benchmark for risk assessment, offering a baseline against which the performance of more sophisticated methods can be evaluated.

Figure 2 below shows the normal distribution fitted to the portfolio return data, highlighting potential deviations in the tails.

Figure 4: Empirical Portfolio Distribution vs Normal distribution

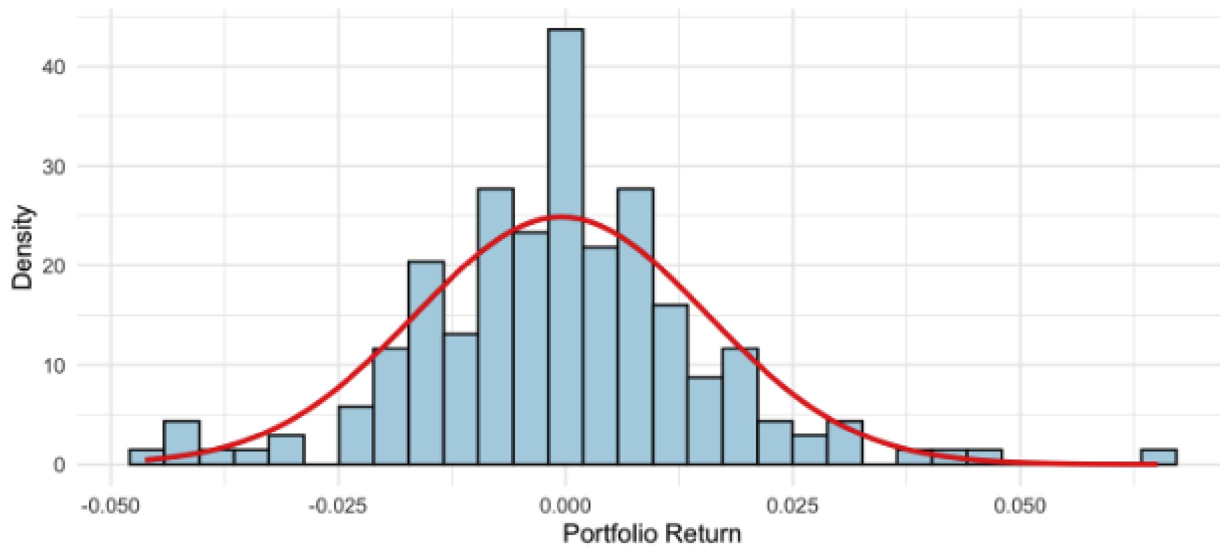

3.2 Skewed Student's t-Distribution Method

The Skewed Student's t-Distribution Method extends the traditional VaR framework by addressing the limitations of normality assumptions. Empirical evidence from financial markets demonstrates that asset returns often exhibit fat tails (kurtosis) and asymmetry (skewness), which are inadequately captured by the normal distribution. The Skewed Student's t-distribution offers a flexible model capable of fitting return distributions with these properties, providing a more realistic representation of portfolio risk.

In this method, the portfolio return distribution is approximated by a skewed Student's t-distribution, parameterised to account for skewness (γ**)** and degrees of freedom ($\mathcal{V}$) that control the thickness of the tails. VaR is calculated using the quantile function of the fitted distribution:

$$VaR_{skewed-t} = Q_{skewed-t}(\alpha)$$

where:

- $Q_{skewed-t}(\alpha)$ is the quantile of the skewed Student's t-distribution at confidence level α.

The parameters of the skewed Student's t-distribution—location, scale, skewness, and degrees of freedom—are estimated using maximum likelihood techniques, ensuring that the model fits the empirical return data as closely as possible.

This approach is particularly effective for portfolios exposed to asymmetric market conditions or extreme price movements, as it captures the tail risks that traditional methods may overlook.

However, the computational intensity of fitting a skewed t-distribution, especially for portfolios with a large number of assets, poses a challenge. Additionally, the accuracy of the results depends on the quality of parameter estimation, which can be sensitive to sample size and the chosen estimation window.

Despite these challenges, the Skewed Student's t-Distribution Method provides a significant improvement over simpler approaches in capturing the higher-order characteristics of return distributions. Its ability to model both skewness and kurtosis makes it a valuable tool for risk managers seeking a more nuanced understanding of potential losses in linear portfolios.

Figure 5: Empirical Portfolio Distribution vs Skewed T-Distribution

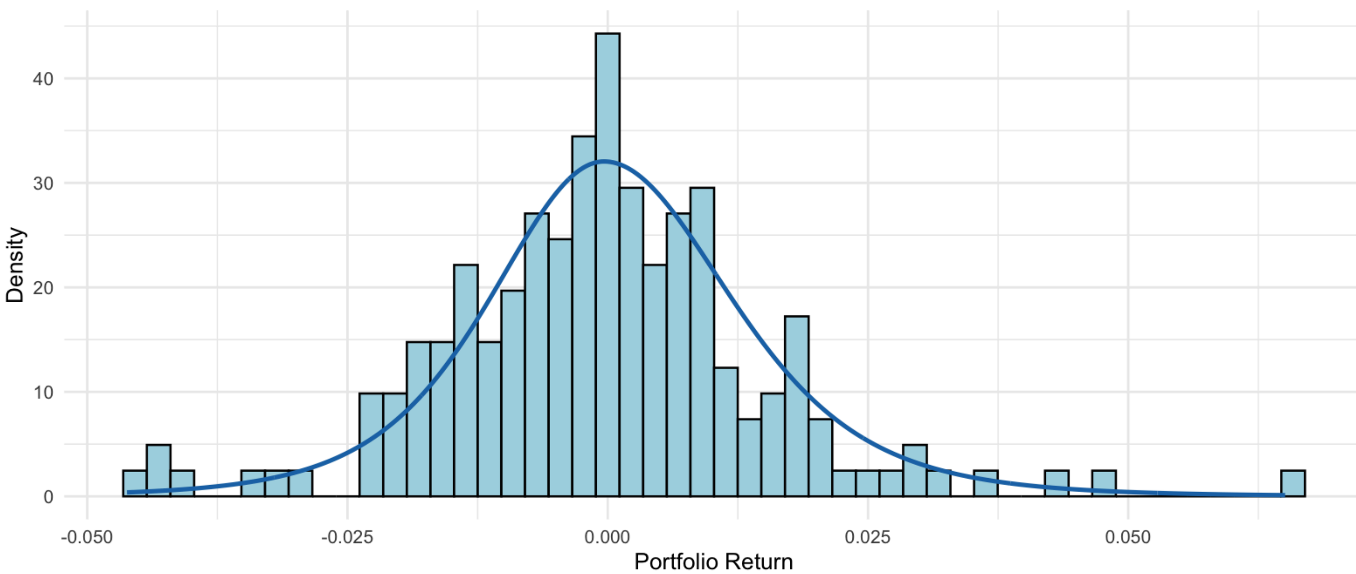

3.3 Cornish-Fisher Expansion

The Cornish-Fisher expansion offers a method for adjusting quantiles of the normal distribution to incorporate skewness (γ) and kurtosis (κ) observed in portfolio return distributions. By modifying the Z-scores typically used in VaR calculations, this approach enables the estimation of risk metrics that account for deviations from normality, providing a more accurate assessment of tail risks.

The adjusted Z-score for a given confidence level (α) is calculated as:

$$z_{\alpha}^{CF} = Z_{\alpha} + \frac{1}{6}\left(z_{\alpha}^{2} - 1\right)\gamma + \frac{1}{24}\left(z_{\alpha}^{3} - 3z_{\alpha}\right)\kappa - \frac{1}{36}\left(2z_{\alpha}^{3} - 5z_{\alpha}\right)\gamma^{2}$$

where:

- Z_{α} is the standard Z-score for confidence level α\alphaα.
- γ is the skewness of the portfolio return distribution.

- κ is the excess kurtosis (kurtosis minus 3) of the portfolio return distribution.

The VaR is then calculated as:

$$VaR_{CF} = -\left(\mu_p + z_\alpha^{CF} \sigma_p\right)$$

Where:

- μ_p is the portfolio mean return (often assumed to be negligible over short horizons).
- σ_p is the portfolio standard deviation.

This method retains the use of the covariance matrix while adjusting for the higher moments of the return distribution. Consequently, the Cornish-Fisher expansion strikes a balance between computational simplicity and the ability to capture non-normal characteristics, making it suitable for a wide range of portfolios.

3.4 Conditional Value at Risk (CVaR) Using Cornish-Fisher Expansion

The Cornish-Fisher expansion can also be extended to estimate Conditional Value at Risk (CVaR), which measures the average loss given that the VaR threshold is exceeded. This metric is particularly valuable in understanding the severity of tail risks and provides a more comprehensive perspective on potential losses during extreme market conditions.

Using the adjusted quantiles from the Cornish-Fisher expansion, CVaR is approximated as:

$$CVaR_{CF} =- (\mu_p + \sigma_p W)$$

where W represents the adjustment factor for the tail, incorporating the skewness and kurtosis parameters. By integrating the tail distribution beyond the VaR threshold, CVaR provides insights into the magnitude of extreme losses, complementing the threshold-based perspective of VaR. The following specifies the adjustment factor:

$$W = Z_{modified}\left(1 + \frac{z_\alpha \gamma}{6} + (1 - 2z_\alpha^2)\frac{\gamma^2}{36} + (-1 + z_\alpha^2\right)\frac{k}{24}$$

The Cornish-Fisher expansion's ability to incorporate skewness and kurtosis makes it a powerful tool for portfolios with return distributions that deviate significantly from normality. However, it relies on accurate estimation of higher moments, which can be sensitive to outliers and limited sample sizes. Moreover, the method assumes that these higher moments remain stable over the estimation window, an assumption that may not hold during periods of significant market volatility.

Despite these limitations, the Cornish-Fisher expansion is a practical and efficient approach for improving the accuracy of risk estimates in linear portfolios, particularly in scenarios where skewness and kurtosis are prominent.

3.4 Monte Carlo Simulation

Monte Carlo simulation involves generating a large number of hypothetical portfolio return scenarios based on the statistical properties of the assets within the portfolio. This process incorporates the covariance structure of returns and allows for the modelling of diversification effects. The steps are as follows:

1. **Simulate Asset Returns**:

   Generate random samples of asset returns using the multivariate normal distribution:

$$R_{sim} \sim N(\mu, \Sigma_{adj})$$

where:

- μ is the vector of asset mean returns (which can be assumed to be zero for simplicity).
- Σ_{adj} is the adjusted covariance matrix, ensuring positive semi-definiteness.

2. **Calculate Simulated Portfolio Returns**:

For each simulated scenario, the portfolio return is computed as:

$$r_{p,\,sim} = w^T R_{sim}$$

where w represents the vector of signed portfolio weights.

3. **Estimate VaR:**

The VaR is derived by identifying the quantile of the simulated portfolio return distribution at the chosen confidence level:

$$VaR_{MC} = -\,Quantile\left(r_{p,\,sim,},\ \alpha\right)$$

Monte Carlo simulation is highly adaptable, enabling the modelling of portfolios with non-normal return distributions and incorporating dynamic correlation structures. Its ability to generate return distributions directly from the covariance matrix ensures that the method is compatible with linear portfolios comprising spot and perpetual contracts. Furthermore, the approach naturally accommodates non-linearities and tail risks, providing precise estimates of both VaR and CVaR.

The primary drawback of Monte Carlo simulation is its computational intensity. Generating and evaluating a large number of return scenarios requires significant processing power, particularly for portfolios with numerous assets. Additionally, the method's accuracy depends on the quality of input parameters, such as the covariance matrix, and may suffer during periods of structural market shifts where historical data is less predictive.

Another limitation is that, despite its flexibility, Monte Carlo simulation typically assumes the covariance matrix remains constant throughout the estimation period. This assumption may reduce the accuracy of risk estimates during periods of rapidly changing market conditions.

Figure 6: illustrates the distribution of simulated portfolio returns and the VaR estimate.

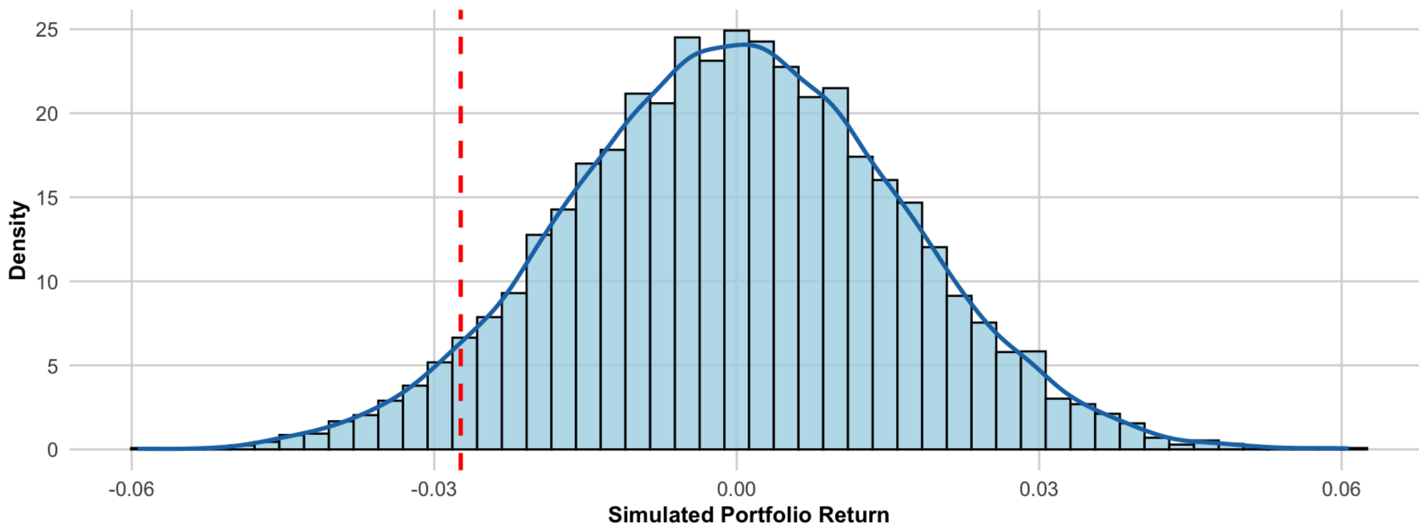

The red dashed line represents the VaR at 5 % confidence level

4. Backtesting All Models

Backtesting is an essential step in evaluating the reliability of risk estimation models. By comparing predicted losses with actual portfolio performance, backtesting assesses whether Value at Risk (VaR) and Conditional Value at Risk (CVaR) estimates accurately capture the portfolio's risk profile under historical market conditions. This validation process is crucial for ensuring the practical applicability of risk metrics in decision-making and regulatory compliance.

The backtesting framework used in this study evaluates the performance of four VaR models—Normal Distribution, Skewed Student's t-Distribution, Cornish-Fisher Expansion, and Monte Carlo Simulation—by analysing their ability to predict exceptions (instances where actual losses exceed the estimated VaR) over a one-year period. Additionally, the Conditional VaR estimates are assessed for their effectiveness in capturing average losses beyond the VaR threshold. The Kupiec Proportion of Failures (PoF) test is employed to statistically evaluate the alignment between observed and expected exceptions.

4.1 Objective of Backtesting

The primary objective of backtesting in this study is to validate the predictive performance of the VaR and CVaR models for a portfolio comprising 159 digital assets, with both long and short positions. Backtesting serves to answer two critical questions: (1) Do the models accurately estimate potential losses at the selected confidence levels (e.g., 95% and 99%)? (2) Are the observed frequencies of exceptions consistent with the expected frequencies?

To achieve these objectives, the study employs a rolling-window approach to simulate real-time risk estimation. The portfolio is assumed to be rebalanced daily, maintaining the weights derived from the net positions of assets. For each day in the backtesting period, the following steps are conducted:

1. **Model Estimation**: VaR and CVaR are calculated for each day using the four risk estimation methods and a 180-day rolling window of historical returns.
2. **Portfolio Return Calculation**: Actual portfolio returns are computed based on realised returns and portfolio weights for the same day.
3. **Exception Identification**: An exception is recorded if the actual portfolio loss exceeds the VaR estimate for the day.
4. **Performance Evaluation**: The frequency of exceptions is compared to the expected frequency at the given confidence level. For example, at a 95% confidence level, exceptions are expected on 5% of the days.

This approach ensures that the backtesting process reflects real-world risk management practices, where risk estimates are derived from historical data available at the time of prediction. The inclusion of daily rebalancing further enhances the realism of the analysis by accounting for the dynamic nature of portfolio composition.

4.2 Portfolio Composition and Data

The backtesting analysis is conducted on a portfolio comprising 159 digital assets, with both long and short positions strategically allocated to capture relative value opportunities. The portfolio reflects a diversified investment strategy, balancing exposure to large-cap assets with smaller allocations to mid- and small-cap assets. Larger positions are concentrated in high market capitalisation coins to ensure liquidity and stability, while smaller allocations in lower-cap assets introduce diversification and potential for higher returns.

Historical price data for each asset were collected to calculate daily log returns. The dataset spans a sufficient period to facilitate rolling window analysis, enabling the estimation of risk measures over time. The long-short nature of the portfolio introduces additional complexity in modelling, as traditional risk measures may not adequately capture the dynamics introduced by short positions.

4.2.1 Rolling Window Estimation

To simulate a realistic investment environment and assess the models' predictive capabilities, a rolling window approach was employed. The backtesting was conducted over a one-year period, encompassing **365 days**. For each day within this period, the VaR models were estimated using a historical window of minimum **180 days** of asset return data. A number of different windows were tried ranging from **180** days to **365** days.

This rolling window methodology ensures that the VaR estimates are based solely on information available up to the prediction date, thereby mimicking the real-time forecasting process used in practical risk management.

4.2.2 Daily Portfolio Rebalancing

To maintain consistency with real-world portfolio management practices, the portfolio is assumed to be rebalanced daily. Rebalancing ensures that portfolio weights remain aligned with the target strategy, accounting for changes in asset prices and ensuring that net positions reflect the desired economic exposure. The weights for each asset are recalculated as:

$$\omega_i = \frac{netPosition_i}{\sum_{j=1}^{N} \left|NetPosition_j\right|}$$

Where;

- ω_i is the weight of asset *i* in the portfolio
- $NetPosition_i$ is the net position of asset *i,* positive for long positions and negative for short positions
- *N* is the total number of assets in the Portfolio

This rebalancing approach allows for dynamic adjustment to market conditions and ensures that the risk estimates accurately reflect the portfolio's current composition.

Volatility estimation is a critical component of the risk modelling framework. For each day in the backtesting period, the portfolio's predicted standard deviation is computed as:

$$\sigma_{\rho} = \sqrt{w^{T} \Sigma w}$$

where w represents the vector of portfolio weights and Σ is the covariance matrix of asset returns.

4.3 Exception Identification

An exception is recorded on day *t* if the actual portfolio loss exceeds the VaR estimate for that day. Mathematically, an exception occurs if:

$$Actual\ Loss_{t} > VaR_{t}$$

where the actual loss is defined as the negative realised portfolio return:

$$Actual\ Loss_{t} = -r_{p,t}$$

For models estimating Conditional Value at Risk (CVaR), exceptions are similarly identified, with the CVaR threshold serving as the benchmark. This process enables a consistent evaluation of the models' ability to predict extreme losses accurately.

4.4 Statistical Tools for Backtesting

The Kupiec Proportion of Failures (PoF) test is employed to evaluate the accuracy of the VaR estimates. The test assesses whether the observed frequency of exceptions aligns with the expected frequency at the specified confidence level (α). The test statistic is given by:

$$LR_{PoF} = -2ln\left(\frac{(1-\hat{p})^{N-x}\hat{p}^{x}}{(1-\alpha)^{N-x}\alpha^{x}}\right)$$

where:

- *N* is the total number of observations (days).
- *x* is the number of exceptions observed.
- $\hat{p}$ = *x* / N is the observed exception rate.
- α is the expected exception rate (e.g., 5% for 95% confidence level).

The test statistic follows a chi-squared distribution with one degree of freedom. If LR_{PoF} exceeds the critical value from the chi-squared distribution, the null hypothesis that the model accurately predicts exceptions is rejected.

5. Backtesting Results

This section summarises the backtesting results of the VaR and CVaR models, evaluating their accuracy in predicting potential losses at 95% and 99% confidence levels. The performance of each model is assessed based on the frequency of exceptions and the results of the Kupiec Proportion of Failures (PoF) test. These metrics provide insight into each method's reliability and suitability for risk management in linear portfolios.

5.1 Results Overview at the 95% Confidence Level

This section summarises the backtesting of Value-at-Risk (VaR) estimates using multiple methods at both the 5% and 1% confidence levels. The performance of each model is evaluated based on the number of VaR exceptions and the Kupiec test statistics. The goal is to determine which method provides the most accurate risk estimates for the portfolio. The backtesting was conducted over **343 days**, with an expected number of **17.15** VaR exceptions at the 5% confidence level (α = 0.05) . At the 1% confidence level (α = 0.01), the expected number of exceptions is **3.43** over 343 days. The below tables show the results from both confidence levels and for all the VaR estimates considered.

At the 5% Confidence Level:

Table 1: Backtesting Results (α = 0.05)

Model	Exceptions	Expected Exceptions	P-value	Interpretation
Skewed T Distribution VaR	20	17.15	0.491	Slightly conservative. The difference is not statistically significant, suggesting acceptable model performance.
Skewed Cornish-Fisher VaR	17	17.15	0.97	Perfect alignment with expectations. Highest p-value indicates exceptional accuracy.
Skewed Cornish-Fisher EL (CVaR)	9	17.15	0.027	Significantly fewer exceptions than expected. The low p-value indicates the model is overly conservative and does not accurately predict losses.

Normal Distribution VaR	15	17.15	0.587	Slightly fewer exceptions. Acceptable model performance, though it may slightly underestimate risk.
Monte Carlo Simulation VaR	15	17.15	0.587	Same as Normal VaR: acceptable performance, with slight underestimation of risk.

The Skewed Cornish-Fisher VaR method demonstrated the best alignment with expectations, achieving the highest p-value and closely matching the expected number of exceptions. Conversely, the Skewed Cornish-Fisher CVaR method was overly conservative, significantly underpredicting exceptions, which may result in excessive capital allocation for risk management.

5.2 Results Overview at the 99% Confidence Level

At a 99% confidence level (α=0.01), the expected number of exceptions over 343 days is 3.433.433.43. Table 2 summarises the results:

Table 2: Backtesting Results at 99% Confidence Level (α=0.01)

Model	Exceptions	Expected Exceptions	P-value	Interpretation
Skewed T Distribution VaR	8	3.43	0.034	*More exceptions than expected. The p-value indicates statistical significance, and the model underestimates risk.*
Skewed Cornish-Fisher VaR	6	3.43	0.207	*More exceptions than expected, but the p-value suggests acceptable performance with slight underestimation of risk.*
Skewed Cornish-Fisher EL (CVaR)	1	3.43	0.12	*Fewer exceptions than expected. The model passes the test but is overly conservative, likely overestimating risk.*
Normal Distribution VaR	8	3.43	0.034	*More exceptions than expected. Statistically significant underestimation of risk.*
Monte Carlo Simulation VaR	8	3.43	0.034	*Same as Normal VaR: statistically significant underestimation of risk.*

At the 99% level, the Skewed Cornish-Fisher VaR method again demonstrated the best performance, balancing risk estimation accuracy and conservatism. The Skewed Student's t-Distribution, Normal, and Monte Carlo methods failed to accurately predict tail risks, resulting in statistically significant underestimation of potential losses. The Skewed Cornish-Fisher CVaR was excessively conservative, failing to meet the expected exception frequency.

5.3 Interpretation of Results

The backtesting results highlight the strengths and weaknesses of the tested methods:

1. **Best Overall Performance**:
    a. The Skewed Cornish-Fisher VaR method consistently aligned with expected exception frequencies, particularly at the 95% confidence level, and performed well at 99%.
    b. Its ability to incorporate skewness and kurtosis makes it the most reliable method for accurately capturing tail risks in linear portfolios.
2. **Conservativeness of CVaR**:
    a. The Skewed Cornish-Fisher CVaR method's overly conservative nature suggests it may not be optimal for regular risk estimation. However, its conservative approach may be advantageous for stress testing or when a buffer against extreme tail events is desirable.
3. **Underperformance of Other Methods**:
    a. The Normal Distribution VaR and Monte Carlo Simulation methods failed to adequately capture extreme tail risks, as evidenced by their poor performance at the 99% confidence level.
    b. The Skewed Student's t-Distribution VaR, while capturing fat tails, still exhibited underestimation of risk in the most extreme scenarios.

5.4 Implications for Risk Management

The findings demonstrate that VaR methods capable of capturing higher moments, such as the Skewed Cornish-Fisher approach, are more effective for portfolios exposed to asymmetric and non-normal return distributions. The conservativeness of CVaR methods highlights their utility in stress testing, while their limitations in day-to-day risk management must be carefully weighed.

6. Further Analysis: Sensitivity to Market Conditions

The backtesting results establish the Cornish-Fisher VaR method as the most reliable risk estimation approach for linear portfolios. This section explores its performance in greater detail, focusing on sensitivity to market conditions and the relative advantages of Cornish-Fisher CVaR in capturing tail risk under stress scenarios.

6.1 Sensitivity to Market Conditions

To evaluate the responsiveness of the Cornish-Fisher methods to changing market dynamics, the portfolio VaR estimates were analysed during periods of heightened volatility. A key observation was that the Cornish-Fisher VaR demonstrated a lagged response to sudden shifts in market volatility, particularly during March to May, a period of pronounced market stress. This lag arises from the 180-day rolling window used to estimate the covariance matrix and higher moments. Equal weighting of historical observations in this window diminishes the influence of more recent volatility spikes.

The graph below illustrates the relationship between actual portfolio returns and Cornish-Fisher VaR estimates. In periods where actual losses exceeded VaR estimates, subsequent increases in the VaR threshold were observed, reflecting the model's adaptive nature but also its sensitivity limitations.

Figure 7: Actual Portfolio Returns vs. Cornish-Fisher VaR Estimates Over Time (alpha = 0.01%)

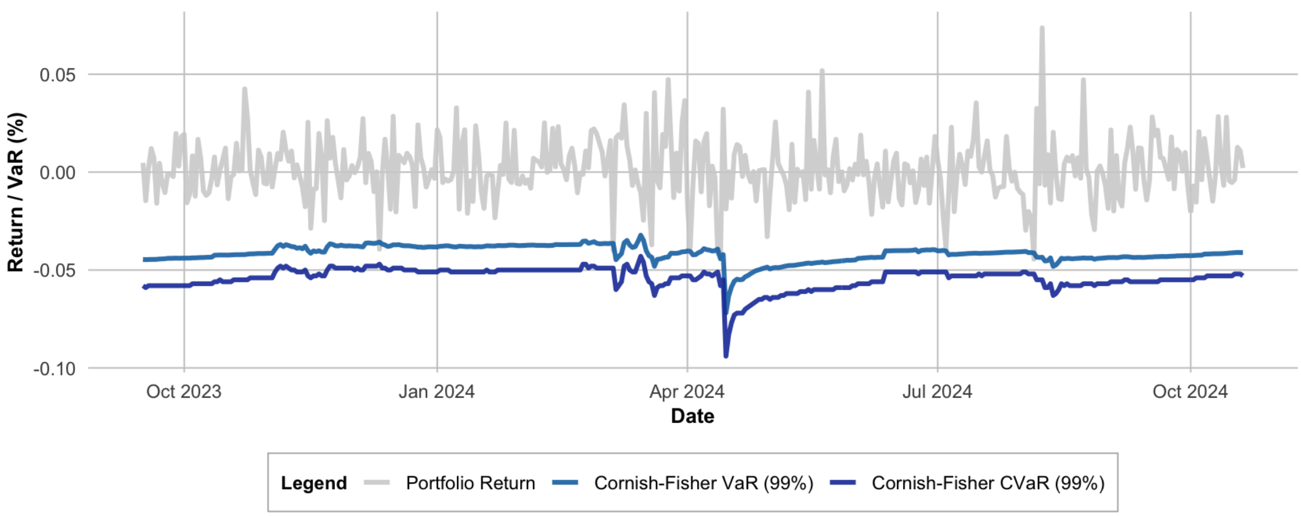

To improve responsiveness, an exponentially weighted moving average (EWMA) approach for the covariance matrix can be adopted. By assigning greater weight to recent data, EWMA enhances the model's ability to react to rapid changes in market conditions, providing a more accurate risk estimate during periods of heightened volatility.

The idea of using exponentially weighted moving averages to estimate covariance matrices is rooted in making the model adaptive to recent market changes. The classic reference for this type of approach is J.P. Morgan's **RiskMetrics** framework, which was introduced using exponentially weighted moving average models for estimating volatility and covariances. The original RiskMetrics Technical Document from 1996 provides an in-depth explanation[1].

[1] https://www.msci.com/documents/10199/5915b101-4206-4ba0-aee2-3449d5c7e95a

6.2 EWMA for Volatility and Correlation Estimation

The weighted covariance matrix using EWMA gives more recent returns a higher weight, which aligns with the idea that recent data is more indicative of current risk. This approach is widely used in financial econometrics, particularly for estimating time-varying volatilities and correlations.

The covariance matrix calculation using EWMA can be linked to **GARCH (Generalized Autoregressive Conditional Heteroskedasticity)** models, as both approaches are aimed at capturing time-varying volatility. The EWMA essentially performs a simplified GARCH(1,1) without explicitly estimating the parameters.

The EWMA weights are calculated to give more importance to recent data while decaying the importance of older observations. The formula for the EWMA weights for each observation i is:

$$\omega_i = \lambda^{(n-i)}$$

where:

- λ is the smoothing parameter (typically close to 1, like 0.94).
- n is the total number of observations (days) in the rolling window.
- i represents each observation from the most recent to the oldest.

These weights are normalised so that their sum is equal to 1:

$$\omega_i = \frac{\lambda^{(n-i)}}{\sum_{j=1}^{n} \lambda^{(n-j)}}$$

EWMA Covariance Matrix Calculation

The covariance matrix is calculated by taking into account the weighted contributions of each observation (day). The below denotes the formula for the EWMA covariance matrix:

$$\Sigma_{EWMA} = \sum_{i=1}^{n} \omega_i \cdot r_i r_i^T$$

where:

- $\omega_i = \lambda^{n-i}$ is the EWMA weight for observation *i*, with λ as the smoothing parameter (commonly 0.94),
- r_i is the vector of **centred returns** (i.e., returns with the mean subtracted) for observation i

- $r_i r_i^T$ is the **outer product** of the centred return r_i with itself.

The smoothing parameter λ controls the decay rate of the weights, with higher values giving more emphasis to recent data. This dynamic weighting improves the model's ability to track time-varying volatility and correlations, particularly during periods of market stress.

6.3 Improved Cornish-Fisher VaR with EWMA

When applied to the Cornish-Fisher framework, the EWMA covariance matrix enhances the model's responsiveness, reducing the lag in VaR estimates during volatile periods. Figure 9 compares the standard Cornish-Fisher VaR estimates with those derived using the EWMA covariance matrix, demonstrating improved alignment with realised portfolio losses in periods of heightened volatility.

Figure 8: Portfolio Returns vs. EWMA Cornish-Fisher VaR Estimates Over Time (α = 1%)

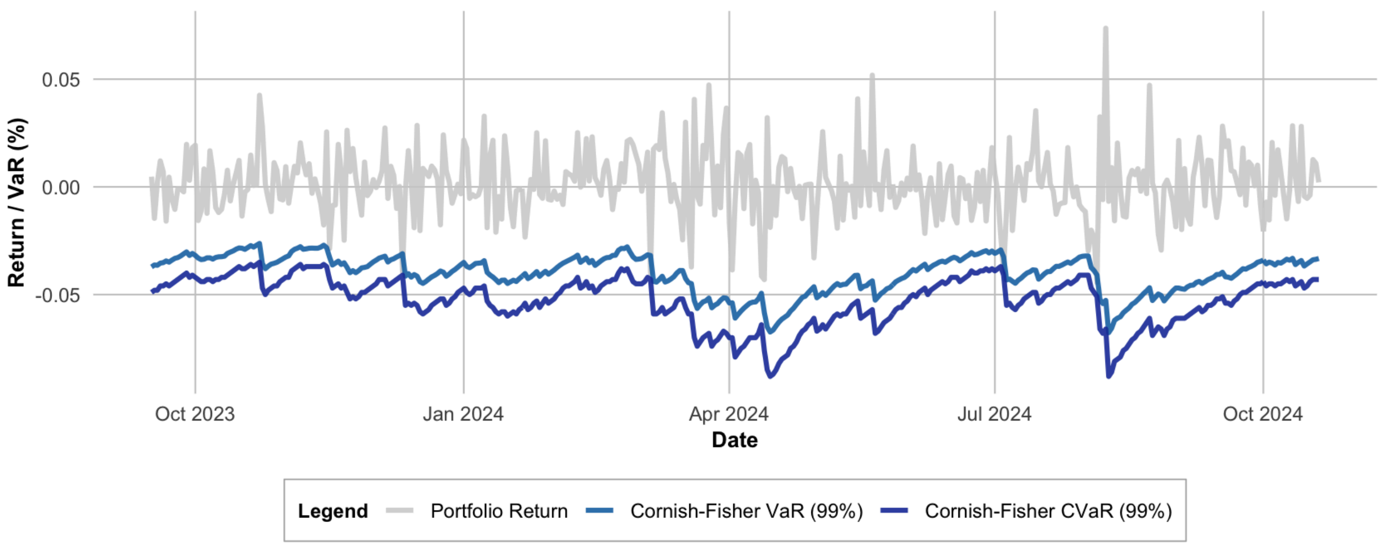

By incorporating EWMA, the model adapts dynamically to shifts in the risk environment, resulting in more accurate and timely risk estimates. This is particularly valuable for portfolios exposed to sudden market shocks, where traditional rolling window methods may understate risks.

6.4 Backtesting Results Summary EWMA VaR and EL

The below table shows the updated statistics for the Cornish-Fisher VaR and EL using the EWMA calculation method as per the above.

Table 3: Backtesting Results (α = 0.01)

Model	Exceptions	Expected Exceptions	Kupiec Test P-value	Interpretation
Skewed Cornish-Fisher VaR	4	3.43	0.76	*The model is performing quite closely to the expected number of exceptions. A high p-value indicates that the number of VaR exceptions is consistent with the expected number, suggesting that the Cornish-Fisher VaR model is well-calibrated.*
Skewed Cornish-Fisher EL (CVaR)	0	3.43	0.0086	*No exceptions, which is possibly too conservative for this confidence level.* *A low p-value indicates that the observed exceptions deviate significantly from the expected number. Since Cornish-Fisher CVaR produced 0 exceptions, it suggests that the model is overly conservative.*

7. Conclusions and Recommendations

This study evaluates and compares four methodologies for estimating Value at Risk (VaR) and Conditional Value at Risk (CVaR) in linear portfolios, including the Normal Distribution Method, Skewed Student's t-Distribution, Cornish-Fisher Expansion, and Monte Carlo Simulation. By incorporating the covariance matrix with signed portfolio weights, the analysis captures the economic exposure of portfolios with both long and short positions, enabling a comprehensive assessment of potential losses under varying market conditions.

7.1 Key Findings:

Cornish-Fisher VaR:

The Cornish-Fisher VaR consistently demonstrated superior performance across backtesting scenarios, particularly at the 95% confidence level. Its ability to incorporate skewness and kurtosis allows it to accurately capture the characteristics of non-normal return distributions, making it the most reliable method for routine risk monitoring.

Conservativeness of Cornish-Fisher CVaR:

The Cornish-Fisher CVaR method exhibited significant conservatism, underestimating the frequency of exceptions across confidence levels. While this conservatism may result in over-allocation of capital under normal conditions, it provides a valuable buffer during stress scenarios, highlighting its utility in capturing tail risks.

Limitations of Traditional Methods:

The Normal Distribution and Monte Carlo Simulation methods failed to adequately capture extreme tail risks, as evidenced by their underperformance at the 99% confidence level. The Skewed Student's t-Distribution, while addressing fat tails, also struggled to reliably estimate risk during extreme market events.

Dynamic Adaptation with EWMA:

The integration of exponentially weighted moving averages (EWMA) into the Cornish-Fisher framework enhanced the model's responsiveness to changing market conditions. By assigning greater weight to recent observations, EWMA-adjusted risk metrics improved alignment with realised portfolio losses during volatile periods.

7.2 Practical Implications

For Routine Risk Management:

The Cornish-Fisher VaR provides a reliable and efficient tool for day-to-day risk monitoring, offering accurate estimates of potential losses while maintaining regulatory compliance. Its adaptability to portfolio-level skewness and kurtosis ensures robust risk assessments across diverse market conditions.

For Stress Testing:

Cornish-Fisher CVaR is particularly well-suited for stress testing and scenario analysis, providing insights into the magnitude of losses beyond the VaR threshold. Its conservatism ensures preparedness for extreme market events, but its use should be complemented by more calibrated methods in routine operations.

Balancing Efficiency and Conservatism:

Risk managers should adopt a dual approach, using VaR for capital allocation and regulatory reporting, while leveraging CVaR for stress testing and extreme risk scenarios. The integration of adaptive methodologies, such as EWMA, further enhances the reliability of these metrics.

7.3 Future Directions

The findings underscore the importance of continued refinement in risk estimation methodologies. Future research could explore:

- The integration of non-parametric methods or machine learning models to dynamically estimate higher moments, reducing reliance on fixed-window assumptions.
- The application of time-varying covariance models, such as GARCH, to complement the EWMA framework for improved responsiveness.
- Expanding the scope to include nonlinear instruments, such as options or dated futures, to evaluate the broader applicability of these methodologies.

8. Limitations

While this study provides a comprehensive analysis of risk estimation methodologies for linear portfolios, several limitations warrant discussion. These constraints highlight areas where caution should be exercised in interpreting results and where future research could provide further insights.

8.1 Data and Portfolio Composition

1. **Concentration in Large-Cap Assets**:
   The portfolio under analysis places significant weight on high market capitalisation assets, reflecting common investment practices to prioritise liquidity and stability. While this approach enhances the robustness of the covariance matrix and reduces noise from less liquid assets, it may limit the generalisability of findings to portfolios with a heavier reliance on mid- and small-cap assets, which exhibit higher idiosyncratic risk and volatility.
2. **Historical Data Assumptions**:
   The use of historical returns assumes that past volatility, correlations, and higher moments are indicative of future market behaviour. This assumption may not hold during periods of structural market change, such as regulatory shifts, sudden macroeconomic shocks, or changes in market participant behaviour.

8.2 Methodological Constraints

1. **Fixed Estimation Window**:
   The primary analysis relies on a 180-day rolling window for estimating the covariance matrix and higher moments. While alternative windows were considered for sensitivity analysis, fixed-length windows may fail to capture abrupt changes in market conditions. The inclusion of exponentially weighted moving averages (EWMA) addresses this to some extent, but additional refinement may be required for scenarios with rapid volatility clustering.
2. **Higher-Moment Estimation Stability**:
   Methods such as the Cornish-Fisher expansion rely on accurate estimation of skewness and kurtosis. These higher moments are sensitive to sample size and outliers, particularly in shorter estimation windows. As a result, the stability of these estimates may be compromised during periods of extreme market movements, impacting the accuracy of risk metrics.
3. **Simplifications in Monte Carlo Simulation**:
   The Monte Carlo approach assumes a multivariate normal distribution for asset returns, which may oversimplify real-world return dynamics. While computationally efficient, this assumption limits the model's ability to fully account for non-linear dependencies or tail events.

8.3 Model Applicability

1. **Exclusion of Nonlinear Instruments**:
   The analysis focuses exclusively on linear portfolios comprising spot and perpetual contracts. Portfolios including nonlinear instruments, such as options or dated futures, introduce additional complexities, such as time decay and path dependency, which are not addressed in this framework.
2. **Assumptions of Portfolio Rebalancing**:
   The assumption of daily portfolio rebalancing ensures alignment with the target strategy but may not reflect real-world constraints such as transaction costs, slippage, or operational limitations. These factors could affect the practical application of the proposed methodologies.

9. References

Foundational Risk Metrics and Methodologies:

- Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica, 50*(4), 987-1007.
- Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics, 31*(3), 307-327.
- Jorion, P. (2001). *Value at Risk: The New Benchmark for Managing Financial Risk.* 2nd Edition. McGraw-Hill.

Cornish-Fisher Expansion:

- Cornish, E. A., & Fisher, R. A. (1937). Moments and Cumulants in the Specification of Distributions. *Revue de l'Institut International de Statistique, 4*(4), 307-320.

Monte Carlo Simulation:

- Glasserman, P. (2004). *Monte Carlo Methods in Financial Engineering.* Springer.

Skewed Student's t-Distribution:

- Aas, K., & Haff, I. H. (2006). The Generalized Hyperbolic Skew Student's t-Distribution. *Journal of Financial Econometrics, 4*(2), 275-309.

RiskMetrics Framework:

- J.P. Morgan & Reuters. (1996). *RiskMetrics Technical Document.* 4th Edition. Available at: https://www.msci.com.

Portfolio Theory and Covariance Matrices:

- Markowitz, H. (1952). Portfolio Selection. *The Journal of Finance, 7*(1), 77-91.
- Ledoit, O., & Wolf, M. (2004). A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices. *Journal of Multivariate Analysis, 88*(2), 365-411.

Volatility Modelling:

- Poon, S.-H., & Granger, C. W. J. (2003). Forecasting Volatility in Financial Markets: A Review. *Journal of Economic Literature, 41*(2), 478-539.
- Hull, J., & White, A. (1998). Incorporating Volatility Updating into the Historical Simulation Method for Value-at-Risk. *Journal of Risk, 1*(1), 5-19.

Applications to Digital Assets:

- Nakamoto, S. (2008). Bitcoin: A Peer-to-Peer Electronic Cash System. [Online]. Available at: https://bitcoin.org/bitcoin.pdf.
- Liu, Y., & Tsyvinski, A. (2018). Risks and Returns of Cryptocurrency. *The Review of Financial Studies, 34*(6), 2689-2727.

Stress Testing and Tail Risks:

- Kupiec, P. H. (1995). Techniques for Verifying the Accuracy of Risk Measurement Models. *The Journal of Derivatives, 3*(2), 73-84.
- Embrechts, P., Kluppelberg, C., & Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance.* Springer.

Exponential Weighting and Dynamic Models:

- RiskMetrics Group. (1996). *RiskMetrics Technical Document: Exponentially Weighted Moving Average Models for Risk Management.* Available at: https://www.msci.com.
- Tsay, R. S. (2010). *Analysis of Financial Time Series.* 3rd Edition. Wiley.