Systematic Skewness and Expected Returns under Market Proxy Inefficiency[†]

Moncef Asmar[*] Sandra Nolte[‡] Mark Shackleton[§] Viorel Roscovan[¶]

March, 2026

Abstract

This paper re-evaluates the empirical validity of the three-moment capital asset pricing model by addressing the structural and statistical challenges posed by market proxy inefficiency and higher-order moment estimation risk. Central to our inquiry is the Roll (1977) critique, which we address by evaluating systematic skewness pricing across Value-Weighted, Equally-Weighted, and Price-Weighted market proxies. We demonstrate that traditional tests of market non-linear response are frequently suppressed by systematic measurement noise, leading to an empirical attenuation of risk premia and a subsequent false rejection of the non-linear risk-return trade-off. To resolve this identification challenge, we implement a Principal Components Analysis-based latent factor correction that filters risk loadings of systematic estimation error. Our findings reveal a robust and statistically significant skewness premium across all market proxies once latent noise is addressed — a result that remains invariant to the inclusion of Fama-French size and value anomalies. Furthermore, we subject our primary specifications to Shanken (1992) correction, confirming that the retrieved skewness premium is a universal feature of the equilibrium cross-section rather than a byproduct of sampling variation. We conclude that the "skewness puzzle" prevalent in the literature is largely a structural artefact of market proxy construction, and that systematic skewness constitutes a fundamental, non-diversifiable risk factor in contemporary financial markets.

[†]The authors gratefully acknowledge the financial support provided by Invesco Quantitative Strategies. The views expressed in this paper are those of the authors and do not necessarily reflect those of Invesco Quantitative Strategies.

[*]**Corresponding author.** Lancaster University Management School, Lancaster, UK; Invesco Management S.A., Frankfurt, Germany; m.asmar@lancaster.ac.uk

[‡]Lancaster University Management School, Lancaster, UK; s.nolte@lancaster.ac.uk

[§]Lancaster University Management School, Lancaster, UK; m.shackleton@lancaster.ac.uk

[¶]Invesco Management S.A., Frankfurt, Germany; Viorel.Roscovan@invesco.com

1. Introduction

The Capital Asset Pricing Model (CAPM) has shaped financial economics for over fifty years, yet it embodies a fundamental simplification: risk is captured entirely by systematic variation with the market portfolio. This assumption, central to the theoretical work of Sharpe (1964) and Lintner (1965), has proven both elegant and durable. However, empirical evidence reveals persistent deviations from this framework. Stock returns exhibit significant negative skewness, crashes occur with greater frequency than normal distribution theory predicts, and investor behaviour suggests a preference for positive skewness that extends beyond variance considerations. These observations motivate an examination of whether financial markets price systematic skewness as a distinct risk dimension.

Theoretical arguments support this possibility. Kraus and Litzenberger (1976) demonstrate that investors with non-increasing absolute risk aversion must exhibit preferences over higher moments of the return distribution, implying that systematic skewness should command a risk premium in equilibrium. This prediction received empirical support from Harvey and Siddique (2000), who document a significant relationship between coskewness and cross-sectional expected returns. Yet the literature exhibits a troubling inconsistency: the magnitude and even the sign of the skewness premium varies substantially across studies, specifications, and market definitions. This heterogeneity suggests either that the skewness premium is genuinely unstable or that methodological factors obscure its identification. The Roll (1977) critique is particularly relevant here: since the true market portfolio is unobservable, any empirical test necessarily employs a proxy. If the estimated skewness premium is sensitive to the choice of the proxy, it may reflect index construction artefacts rather than genuine risk pricing.

This paper evaluates skewness pricing across three market proxies with distinct weighting schemes: value-weighted, equally-weighted, and price-weighted indices. If systematic skewness is a fundamental priced risk factor, the premium should be robust to proxy choice. Our approach follows the standard methodology in the higher-moment literature: we employ the squared market factor as a proxy for systematic skewness, which captures the sensitivity of individual stock returns to market non-linear response. This approach is more statistically tractable than direct skewness estimation, which suffers from substantial estimation noise.

A critical methodological innovation underpins our analysis. We recognise that market proxy inefficiency introduces systematic measurement error into our risk loading estimates. Rather than treating this noise as an impediment to inference, we explicitly model it using Principal Components Analysis. We extract the latent common factor in our estimated residuals and incorporate this factor into our asset pricing specifications. This procedure serves as an identification strategy that filters measurement error from our skewness loading estimates, thereby allowing us to isolate the true relationship between crash risk exposure and expected returns.

Our findings address three substantive dimensions. First, we establish that the skewness premium is robust across alternative market definitions once measurement error is properly controlled for. Second, we demonstrate that systematic skewness is a distinct risk dimension that is empirically separable from Fama-French factors, suggesting it captures independent variation in the cross-section of returns. Third, we show that measurement error substantially attenuates the estimated skewness premium and that the latent factor correction is essential to its identification. These results suggest that the apparent inconsistencies in the skewness pricing literature are largely attributable to methodological challenges in proxy construction rather than to genuine instability in the risk-return relationship.

The rest of the paper is organised as follows. Section 2 covers the theoretical background and explains why the three-moment CAPM matters. Section 3 explains the different indices construction. Section 4 lays out our methodology in detail, including how we use PCA to correct for measurement error. Section 5 presents our main findings on the pricing of skewness across different market definitions. Section 6 reports robustness checks. Section 7 concludes and discusses what our results mean for understanding risk and return in equity markets.

2. Literature Review

The genesis of modern asset pricing lies in the mean–variance framework of Markowitz (1952), which assumes that risk is fully characterised by the second moment of returns. Within this paradigm, the Capital Asset Pricing Model of Sharpe (1964) and Lintner (1965) posits that an asset's risk premium depends solely on its covariance with the market portfolio. However, the theoretical sufficiency of the first two moments has been long contested. Arrow (1971) and Pratt (1964) demonstrate that any investor with non-increasing absolute risk aversion must exhibit a preference for positive skewness. This implies that in a world of non-normal returns—where extreme downside events are more frequent than a Gaussian distribution suggests—the third moment must enter the equilibrium pricing relation.

Early challenges to the CAPM emerged from evidence that market beta alone could not account for the full cross-section of expected returns. Banz (1981) documented that smaller NYSE-listed firms earned systematically higher risk-adjusted returns than larger ones — a pattern he traced back several decades and attributed to a fundamental misspecification in the single-factor model. Fama and French (1992) formalised this critique by showing that, once size was controlled for, beta lost all explanatory power, and that firm size and book-to-market equity jointly characterised the cross-section of average US stock returns far more effectively than the CAPM had ever managed.

Kraus and Litzenberger (1976) provide the first formal extension of the CAPM to incorporate these preferences. Their three-moment CAPM suggests that investors require compensation not only for variance risk β but also for systematic coskewness γ. In their framework, assets that decrease in value when market skewness turns negative are considered risky and must offer higher expected returns. Whilst this theory is elegant, early empirical tests yield mixed results. Jurczenko and Maillet (2001) subsequently situate this result within a unified multi-moment asset pricing framework, providing a systematic comparison of two-, three-, and four-moment pricing models and establishing the precise utility-theoretic conditions under which the three-moment specification constitutes a valid and sufficient extension of the standard CAPM. Whilst this theoretical framework is elegant, early empirical tests yield mixed results. Friend and Westerfield (1981) find that whilst skewness is often significant, the results prove highly sensitive to the chosen sample period and asset selection, raising fundamental questions about the model's empirical stability. This sensitivity to specification choices presaged concerns that would later dominate the literature.

A fundamental challenge to these tests is the Roll (1977) critique, which argues that any test of a CAPM-like model is inherently a joint test of the model itself and the mean-variance efficiency of the market proxy. This critique is particularly acute in a three-moment setting. If the true market portfolio is unobservable, the estimated premiums for β and γ may reflect measurement error rather than genuine pricing. A value-weighted index may lead to different conclusions than an equally-weighted index about whether skewness is priced, not because the underlying risk preference differs, but because different weighting schemes capture different dimensions of systematic risk. Our study addresses this directly by systematically varying the market proxy to determine whether the γ premium is a robust economic phenomenon or an artefact of proxy selection.

The empirical methodology for testing these relations was significantly refined by Barone-Adesi (1985), who recasts the three-moment CAPM within a quadratic factor framework derived from Arbitrage Pricing Theory. By treating the market return and its squared deviation as two distinct risk factors, Barone-Adesi provides a statistically tractable approach for separating local volatility from non-linear skewness exposure. This Quadratic Market Model (QMM) remains the standard methodology for estimating systematic skewness risk in modern asset pricing research.

Critically, Barone-Adesi et al. (2004) extend their framework by exploring the relationship between firm size and systematic skewness. They document that smaller-cap portfolios exhibit significantly more negative coskewness than their large-cap counterparts. This observation raises an important identification question: is the size premium documented by Fama and French (1993) partly a reward for bearing systematic skewness concentrated in small firms? This size-skewness nexus motivates our research design. By employing 25 test portfolios double-sorted on size and book-to-market, we can evaluate whether the skewness premium persists across different size buckets and whether its magnitude varies when moving from a value-weighted market (dominated by large caps) to an equally-weighted market (which amplifies small-cap effects).

More recent empirical work has clarified the distinction between systematic and idiosyncratic skewness. Harvey and Siddique (2000) document that conditional coskewness is a priced factor, commanding a premium of approximately 3.6% per annum, and that this premium remains significant even after controlling for Fama-French factors. This evidence suggests that skewness pricing is distinct from size and value effects. However, alternative explanations have emerged from the behavioural finance literature. Mitton and Vorkink (2007) and Barberis and Huang (2008) argue that idiosyncratic skewness may be priced due to lottery-seeking behaviour among under-diversified investors, introducing a separate mechanism that could confound empirical estimates of systematic skewness pricing. Indeed, Boyer et al. (2010) examine whether preferences for idiosyncratic skewness are reflected in equilibrium prices. Using a predictive model of expected idiosyncratic skewness derived from lagged firm characteristics, they document a robust negative relation between expected skewness and subsequent returns — consistent with investors accepting lower compensation for stocks that offer lottery-like payoffs.

Tauscher and Wallmeier (2016) raise a pointed methodological concern: since the test portfolios and the SMB and HML factors in standard Fama–French tests are constructed on the same underlying sorting variables, the resulting overlap mechanically inflates the model's apparent fit. Applying resampling and split-sample methods to European equity data, they demonstrate that this overlapping bias is economically non-trivial and that conventional inference overstates the model's pricing ability. Their findings carry a broader lesson — that the persuasiveness of a factor model in time-series tests may, in part, be an artefact of how the test is designed rather than a genuine reflection of its explanatory content.

Our study bridges these two research strands — the proxy-dependency strand initiated by Roll (1977) and the size-skewness nexus explored by Barone-Adesi et al. (2004). By testing whether γ is priced across diverse market definitions whilst controlling for size and value effects, we provide a comprehensive assessment of the three-moment CAPM's empirical validity. If systematic skewness represents a fundamental dimension of equilibrium risk pricing, the γ premium should remain statistically significant and economically consistent regardless of whether the market is defined by market capitalisation, equal weighting, or price weighting. Conversely, if the premium vanishes when moving away from equal-weighted or small-cap-heavy specifications, this would suggest that apparent skewness pricing is inextricably linked to the size effect rather than reflecting genuine preferences for market non-linear response. This distinction has important implications for both theoretical understanding of investor preferences and practical applications to portfolio construction and risk management.

3. Data and variable construction

This study investigates the distributional properties of three distinct weighting schemes applied to US stock market indices from July 1963 to December 2024. By comparing these different weighting approaches, we can uncover valuable insights into risk and return characteristics, particularly when examining higher-order moments such as skewness and kurtosis. Understanding how different index construction methodologies affect these distributional properties is crucial for asset pricing research, as the choice of market portfolio can materially influence empirical results and the interpretation of factor pricing models. We focus on the S&P 500 universe, obtaining data from the CRSP database. The S&P 500 index, as a cap weighted index, tracks the 500 largest US companies by market capitalisation and serves as our primary asset pool for portfolio construction, allowing us to make clean comparisons across weighting schemes whilst maintaining consistency in the underlying security universe. Whilst value-weighted indices represent the canonical market proxy in empirical asset pricing, their exclusive reliance on capitalisation weights concentrates exposure in the largest firms and may inadequately capture the higher-moment properties of the broader market. The equal-weighted variant assigns uniform weight to all constituents, introducing a deliberate tilt towards smaller stocks. The price-weighted variant draws its justification from long-established empirical practice, replicating the construction logic of the Dow Jones Industrial Average and providing a historically grounded benchmark against which cap-weighted results can be assessed. Each of the three resulting indices then is used as a market portfolio proxy in our subsequent analysis, enabling us to assess whether systematic risk pricing — particularly coskewness risk — varies systematically with the choice of market definition.

3.1 Market proxy

The construction of market-representative assets lies at the heart of empirical asset pricing research. The choice of weighting scheme determines not only the average returns and volatility of the resulting portfolio, but also its higher-order moments, which have received increasing attention in recent pricing models. Using the historical composition of the S&P 500, we construct three different market portfolios while applying fundamentally different weighting methodologies.

The Value-Weighted Index (VWI) applies the standard market capitalisation weighting method, which is the predominant passive portfolio used in academic research and industry practice. Under this scheme, larger firms exert greater influence on index performance, reflecting their greater economic significance and the proportion of capital invested in them. The weight of a constituent $w_{i,t}$ is defined as its relative market capitalisation at the end of the previous trading period.

$$w_{i,t} = \frac{P_{i,t-1}N_{i,t-1}}{\sum_j P_{j,t-1}N_{j,t-1}}, \tag{1}$$

where $P_{i,t-1}$ and $N_{i,t-1}$ are the price and outstanding common shares of constituent i at time $t-1$, respectively. This weighting approach ensures that the index reflects the overall market portfolio as theorised in the CAPM, making it the passive benchmark for assessing systematic risk.

given the evidence in Barone-Adesi et al. (2004) that small-firm portfolios systematically exhibit negative coskewness with the market, the EWI is expected to differ materially from the VWI in its skewness characteristics, making it a theoretically motivated alternative for identifying the coskewness premium.

The Equal-Weighted Index (EWI) takes a markedly different approach, assigning identical weights to all S&P 500 constituents regardless of firm size.

$$w_{i,t} = \frac{1}{I}, \tag{2}$$

where I is the number of eligible S&P 500 stocks. This construction fundamentally amplifies the size effect, giving smaller firms substantially greater influence on index performance than in the VWI. From a portfolio management perspective, maintaining equal weights requires continuous rebalancing as relative market prices evolve. Because stock allocations naturally drift from equal weights over time, the EWI requires more frequent rebalancing than the VWI to maintain equal constituents' weights. Thus, it typically results in higher trading costs and increased portfolio turnover. This distinction has important implications for understanding the transaction costs[1] associated with index replication and the practical feasibility of equal-weighted strategies.

The Price-Weighted Index (PWI) represents a historically significant but economically less intuitive approach, assigning weights based solely on individual stock prices. The weight of a constituent $w_{i,t}$ is equal to its market price at the end of the previous trading period divided by the sum of the prices of all constituents.

$$w_{i,t} = \frac{P_{i,t-1}}{\sum_j P_{j,t-1}}. \tag{3}$$

Under this scheme, stocks with higher nominal prices receive disproportionately greater influence on index performance, regardless of their market capitalisation or economic importance. Since the expected number of shares $N_{i,t}$ in issue is uncertain and can change, it remains agnostic to this factor. This was especially true in the early days of the stock market. The PWI construction mirrors the methodology of the Dow Jones Industrial Average, the oldest major US stock index dating back to 1896, which effectively allocates one share to each stock in the index. Although price weighting is rarely justified

[1] We assume zero transaction costs in constructing and rebalancing our indices. In practice, real-world index replication incurs trading costs, market impact, and bid-ask spreads, which would be substantially higher for equal-weighted indices due to their greater rebalancing frequency.

on economic grounds — market price per share is largely a function of arbitrary historical corporate actions such as stock splits and dividend policy — it remains of historical and practical interest, not least because it requires rebalancing only upon such corporate events rather than continuously. This third weighting scheme provides a distinctive perspective that enhances our ability to analyse how different index construction methods affect higher-moment measures and whether empirical results regarding systematic risk pricing are robust to alternative market portfolio definitions.

3.2 Sample

We retrieve the historical monthly composition of the S&P 500 index from the CRSP database, including information on stocks rebalanced into the index over time.[2] This comprehensive historical record allows us to accurately reconstruct the index composition at each point in time, which is necessary for computing accurate historical returns under each weighting scheme.

For a security to be included in our indices, it needs to meet the following criteria: it must have valid prices at both the current and previous period end-dates, and it must be part of the S&P 500 universe at the end of the current period. These requirements ensure that we compute returns only for securities with complete pricing information and that we maintain consistency with the official S&P 500 constituent list at each rebalancing date.

We then extract relevant stock attributes from CRSP and apply each of the three weighting schemes to construct our indices. Dividends are assumed to be reinvested, as our objective is to construct total return indices suitable for comprehensive performance analysis. This assumption is standard in the asset pricing literature and ensures that our returns capture the full economic return to holding the index. The dataset comprises monthly observations spanning June 1963 to December 2024, providing an extensive period of over 60 years across varying market conditions — including multiple bull markets, bear markets, financial crises, and regime changes — to robustly test higher-moment models and assess the stability of pricing relationships. Monthly stock returns are computed as:

$$r_{i,t} = \frac{P_{i,t} + D_{i,t} - P_{i,t-1}}{P_{i,t-1}}, \tag{4}$$

where $P_{i,t}$ denotes the price of stock i at the end of period t, and $D_{i,t}$ represents dividends paid at time t for holding the stock i.

3.3 Summary statistics

Table 1 reports the summary statistics for the three indices based on their monthly returns, revealing notable differences in distributional properties that persist despite the common underlying security universe. These differences have important implications for understanding how the index construction affects the statistical properties available to pricing models.

Mean returns vary substantially across indices, with the EWI exhibiting the highest average return at 1.1% per month, followed by the VWI at 0.9%, whilst the PWI records the lowest at 0.6%. These differences likely reflect the differential exposure of

[2] Retrieving all historical members of the S&P 500 requires an API connection with WRDS starting from July 2020 onwards.

each index to the size premium, with equal weighting providing the greatest exposure to smaller stocks, which have historically offered higher average returns[3]. The 50 basis points difference between VWI and EWI represents a meaningful economic distinction over long investment horizons. Returns variance differs markedly across the three indices. The PWI shows the largest variability with a standard deviation of 7%, substantially exceeding both the EWI at 5% and the VWI at 4.3%. This relatively higher volatility for price weighting reflects the pronounced influence of individual stock price movements — large percentage moves in high-priced stocks can substantially move the index — and the mechanical properties of price weighting in selecting stocks regardless of their economic importance. Both the VWI and EWI display negative skewness (–0.42 and –0.34, respectively), indicating asymmetric return distributions characterised by relatively larger downside movements. This negative skewness is consistent with well-documented properties of equity returns and suggests that tail risk is relevant for investors. The PWI exhibits near-zero skewness at –0.004, a notable contrast to the other two indices, potentially reflecting the different composition effects when stocks are weighted by price rather than market value or equal allocation. All three indices exhibit elevated kurtosis, ranging from 4.70 for the VWI to 5.50 for the EWI, substantially exceeding the value of 3 expected under a normal distribution. This indicates return distributions with pronounced tail behaviour and a higher frequency of extreme observations — both large gains and large losses appear more frequently than would be expected under normality. The Jarque–Bera test of normality strongly rejects the null hypothesis for all three indices (p-value $\approx$ 0.000), providing conclusive statistical evidence of significant departures from normality (see Appendix A for kernel density estimates and Q–Q plots). These findings align with well-documented characteristics of financial returns[4] and underscore the importance of considering higher-order moments in asset pricing models, particularly when assessing whether investors demand compensation for bearing skewness and kurtosis risk.

	Mean	SD	Skew.	Kurt.	Min.	Max.	JB p	Obs.
VWI	0.009	0.043	−0.415	4.698	−0.216	0.168	0.000	738
EWI	0.011	0.050	−0.335	5.499	−0.256	0.231	0.000	738
PWI	0.006	0.070	−0.004	4.843	−0.270	0.300	0.000	738

Note: JB p denotes the p-value of the Jarque–Bera test for normality. The sample period covers monthly returns from 1963 to 2024.

Table 1. Summary statistics of monthly returns

Table 2 shows the relationships between the three indices based on monthly returns over the full sample period. A strong positive correlation of 0.95 exists between VWI and EWI, suggesting that despite their fundamentally different weighting schemes, these two indices capture largely similar underlying market movements. The price-weighted version, by contrast, exhibits noticeably lower correlations with both the VWI (0.67) and the EWI (0.71). These substantially lower correlations suggest that the price weighting mechanism

[3]The size premium was first documented by Banz (1981), who showed that smaller firms earn higher risk-adjusted returns than larger firms. Fama and French (1992) subsequently confirm this finding, demonstrating a clear inverse relationship between firm size and average monthly returns across NYSE decile portfolios.

[4]Peiró, A. (1999) documents, using distribution-free tests across international equity index returns, that financial returns exhibit significant departures from normality, predominantly characterised by excess kurtosis rather than asymmetry.

introduces a meaningful alternative market return series, one that can diverge materially from conventional market-weighted or equally-weighted approaches. This divergence is economically meaningful and suggests that the choice between price weighting and other approaches can have significant implications for empirical results.

	VWI	**EWI**	**PWI**
VWI	1.00		
EWI	0.95	1.00	
PWI	0.67	0.71	1.00

Note: Correlations are computed using monthly returns over 1963–2024.

Table 2. Correlation matrix

3.4 Test assets

To evaluate systematic skewness pricing, we employ the 25 size–value portfolios (5×5 double sort) from Kenneth French's data library as our test assets. These standard portfolios are formed each June by sorting all NYSE, AMEX, and NASDAQ stocks on the basis of size (market capitalisation) and book-to-market equity ratio, creating a comprehensive cross-section of the equity market. The use of these widely adopted benchmark portfolios provides several advantages. First, they provide extensive historical data, allowing us to conduct robust empirical tests across our full sample period. Second, their use facilitates comparison with existing asset pricing research, as these portfolios have become the standard test assets in the empirical asset pricing literature. Third, they provide a well-diversified cross-section spanning both size and value dimensions, which is essential for reliably estimating risk premiums and testing whether systematic skewness is priced across the market. By employing these established test assets, we can evaluate whether higher-moment pricing varies depending on the choice of market portfolio definition, thereby assessing the robustness of pricing relationships to alternative market specifications.

4. Methodology

We employ a rigorous multi-stage empirical framework designed to evaluate the pricing of systematic skewness within the cross-section of equity returns. Our methodology addresses the key challenges in estimating higher-order moments - traditional estimators are sensitive to market portfolio definition and the inclusion of non-traded assets. We adopt a structural factor loading approach that treats market excess returns and market non-linear response as distinct, priced risk factors. This will ensure robustness through three complementary stages. First, we run time-series regressions across the full sample to estimate factor loading estimates. Second, we implement a standard Fama-MacBeth cross-sectional procedure to identify risk premiums. This two-stage structure is deliberately chosen: Fama and French (2020) demonstrate that combining time-series loading estimates with cross-sectional pricing tests yields more reliable inference than either approach used in isolation, and that the construction of factors carries as much consequence for empirical performance as their selection. Our design is informed by this principle throughout. Finally, we add a latent factor correction using PCA to isolate any systematic measurement noise. By employing this multi-step approach across three different market proxies, we provide a comprehensive and robust assessment of whether systematic skewness earns a significant risk premium in equilibrium.

A central contribution of this study is the systematic comparison of skewness pricing across three differently weighted market portfolios: value-weighted, equal-weighted, and price-weighted. Each market proxy yields alternative estimates of market returns and market squared deviations, and thus alternative factor loadings. By repeating our estimation across all three market proxies, we assess the robustness of our findings to the choice of market portfolio. If the coskewness premium is fundamental and economically meaningful, it should materialise across different market definitions, even if the magnitude varies. Conversely, if the premium is an artefact of a particular market proxy, our results will diverge significantly across the three specifications.[5]

4.1 Quadratic Market Model

The theoretical motivation for incorporating higher-order moments into our pricing framework begins with a formal approximation of the investor's objective function. While the traditional CAPM implicitly relies on quadratic utility, Kraus and Litzenberger (1976) argue that such a form is unsuitable for risk-averse agents as it violates essential properties of rational preference. Instead, they favour utility functions — such as logarithmic, power, or negative exponential — that exhibit non-increasing absolute risk aversion. For this class of functions, the investor's expected utility of end-of-period wealth, $E[U(\tilde{W})]$, can be approximated via a Taylor series expansion around the mean wealth level $\bar{W} = E[U(\tilde{W})]$:

$$E[U(\tilde{W})] = U(\bar{W}) + \frac{U''(\bar{W})}{2!}\sigma_w^2 + \frac{U'''(\bar{W})}{3!}m_w^3 + \textit{terms of higher order} \tag{5}$$

Within this expansion, the second and third moments are formally defined as expectations of wealth deviations: $\sigma_w^2 = E[(\tilde{W} - \bar{W})^2]$ represents the variance of end-of-period wealth. $m_w = E[(\tilde{W} - \bar{W})^3]$ denotes the third central moment, or wealth skewness.

Kraus and Litzenberger (1976) note that the condition of non-increasing absolute risk aversion is a sufficient requirement to ensure $U'''(\bar{W}) > 0$. This property formally characterises a preference for positive skewness, implying that for a given level of wealth, investors favour distributions with longer right tails. However, to move from this normative utility framework to a testable empirical specification, we adopt the market-based equilibrium perspective of Barone-Adesi (1985). This shift allows us to treat systematic skewness as a non-diversifiable risk factor rather than an unobservable utility trait. To operationalise this intuition, we use a linear two-factor model, the Quadratic Market Model (QMM), as a return-generating process. Market returns and the square of the market returns are the two factors. Specifically, market non-linear response proxies for coskewness exposure. Here, $r_{p,t} = R_{p,t} - R_{F,t}$ and $r_{M,t} = R_{M,t} - R_{F,t}$ represent excess returns, over the risk-free rate of return $R_{F,t}$, for the test portfolio and the market, respectively.

$$r_{p,t} = \alpha_p + \beta_p r_{M,t} + \gamma_p (R_{M,t} - \bar{R}_M)^2 + \epsilon_{p,t}, \qquad t = 1, \dots, T \tag{6}$$

where the α_p coefficient denotes the intercept representing the portfolio-specific component of returns not captured by market factors. The term $\epsilon_{p,t}$ represents the idiosyncratic

[5] This comparative analysis directly addresses the Roll (1977) critique and strengthens our identification strategy.

risk, which is assumed to satisfy $E[\epsilon_{p,t}|r_{M,t}, (R_{M,t} - \bar{R}_M)^2] = 0$. Standard errors are computed using Newey-West (1987) procedure to account for potential heteroscedasticity and autocorrelation in $\epsilon_{p,t}$.

We intentionally center the quadratic term around the sample mean ($\bar{R}_M$) for two reasons: it isolates sensitivity to market asymmetry, secondly it eliminates the multicollinearity that would otherwise affect a raw R_M^2 term. In this setup, β_p measures traditional systematic market risk, while γ_p captures the coskewness loading. For a given asset, a negative γ_p coefficient amplifies its downside exposure during market volatility, reflecting a 'crash risk'. Conversely, a positive γ_p indicates that the asset exhibits a non-linear, convex response to market returns — its performance is disproportionately amplified during large positive market movements, providing a natural hedge against the downside asymmetry captured by negative systematic skewness.

4.2 Time-Series Estimation of Factor Loadings

For each alternative market proxy k, we estimate the factor loadings ($\hat{\alpha}_p^{(k)}$, $\hat{\beta}_p^{(k)}$ and $\hat{\gamma}_p^{(k)}$) using an ordinary least squares (OLS) for the 25 test size–value portfolios excess returns. For the risk-free rate of return, we consider one-month Treasury bill rate.

$$\begin{gathered} r_{p,t} = \alpha_p^{(k)} + \beta_p^{(k)} r_{M,t}^{(k)} + \gamma_p^{(k)} (R_{M,t}^{(k)} - \bar{R}_M^{(k)})^2 + \epsilon_{p,t}^{(k)} \\ t = 1, \ldots, T; \quad k \in \{VWI, EWI, PWI\} \end{gathered} \tag{7}$$

We correct all standard errors for heteroscedasticity and autocorrelation using the Newey-West procedure with six lags. The extended sample period (60+ years) provides substantial variation in both the level and squared deviations of market returns, ensuring reliable identification of both linear and quadratic sensitivities. Importantly, by estimating loadings over the full sample, we implicitly assume that the relationship between portfolio returns and market moments is stable across time and extended macro-economic cycles.

4.3 Cross-Sectional Pricing and Factor Identification

In the second stage, we employ the cross-sectional approach of Fama and MacBeth (1973) to estimate risk premia. For each month t, we regress portfolios' excess returns on the estimated factor loadings from the first stage. This approach suits our setting as its time-series averaging of cross-sectional estimates yields standard errors that are robust to cross-sectional dependence in returns.[6] Additionally, we assume that realised returns provide an unbiased measure of ex-ante expectations over a long sample period, as idiosyncratic shocks are expected to mean-revert to zero. This allows for the inference of equilibrium risk premia from the cross-section of average returns.

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + u_{p,t}. \tag{8}$$

The coefficients $\lambda_{1,t}$ and $\lambda_{2,t}$ capture the market risk premium and coskewness risk premium, respectively. To obtain our final estimates, we average the monthly regression coefficients across the entire sample period: $\bar{\lambda}_j = T^{-1} \sum_t \lambda_{j,t}$, for each factor $j = 0, 1, 2$.

[6] Standard errors are adjusted for the EIV bias arising from the use of estimated factor loadings in Section 6.

We compute t-statistics using Newey-West standard errors with six lags to account for serial correlation and heteroscedasticity in the computed premiums.

Furthermore, we control for the Fama-French size (SMB) and value (HML) factors to ensure that any skewness premium we identify is not contaminated by size or value effects. By explicitly controlling for these established factors, we can confidently attribute any remaining premium to genuine compensation for both systematic variance and skewness, rather than attributing it to other sources of return predictability.

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{3,t}\hat{\beta}_{p,\text{SMB}} + \lambda_{4,t}\hat{\beta}_{p,\text{HML}} + u_{p,t}. \tag{9}$$

where $\hat{\beta}_{p,\text{SMB}}$ and $\hat{\beta}_{p,\text{HML}}$ denote the estimated loadings on SMB and HML factors, whilst $u_{p,t}$ represents the cross-sectional pricing error assumed to be a zero-mean, serially uncorrelated process and orthogonal to the estimated factor loadings $\hat{\beta}_p$, $\hat{\gamma}_p$, $\hat{\beta}_{p,\text{SMB}}$, and $\hat{\beta}_{p,\text{HML}}$. Similarly, for each $j = 0, 1, 2, 3, 4$ we average out independently $\lambda_{j,t}$ to get the associated premiums.

4.4 Omitted Factors Correction

Whilst repeating tests across multiple market definition mitigates concerns about proxy invalidity, a more direct approach is to estimate the unobserved common factor that our proxies may be missing. Following Giglio and Xiu (2021), we implement a latent factor procedure to address measurement error from omitted factors. If our market proxy is an imperfect representation of the true latent wealth portfolio, the residuals from our time-series regressions in the first stage will contain a common systematic component. We construct a $T \times P$ residual matrix (where P is the number of months and N is the number of test portfolios) and apply Principal Components Analysis (PCA) to extract the first principal component PC_1. This component captures the dominant common variation in residuals, helping to mitigate - not fully eliminate - measurement error from proxy invalidity. Residual differences across market definitions reflect both the unique measurement errors inherent in each proxy and potentially genuine variation in skewness pricing across market specifications. We then respecify our cross-sectional model by including sensitivity to this latent factor, where $\hat{\beta}_{p,\text{latent}}$ denotes the loading on the latent factor, estimated by regressing each portfolio's returns on the extracted PC_1.

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{5,t}\hat{\beta}_{p,\text{latent}} + e_{p,t} \tag{10}$$

By controlling for this latent factor alongside SMB and HML, we mitigate measurement noise that arises from using imperfect market proxies. Thus, If λ_2 remains significant and economically meaningful after controlling for the latent factor and Fama-French factors, it indicates that skewness pricing reflects genuine equilibrium compensation for systematic risk. The consistency of our findings across three distinct market definitions and with latent factor adjustment demonstrates that skewness pricing is not simply an artefact of proxy choice or specification error, but rather a fundamental feature of how markets price systematic risks (i.e., variance and skewness).

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{3,t}\hat{\beta}_{p,\text{SMB}} + \lambda_{4,t}\hat{\beta}_{p,\text{HML}} + \lambda_{5,t}\hat{\beta}_{p,\text{latent}} + e_{p,t} \tag{11}$$

where $e_{p,t}$ represents the cross-sectional pricing error assumed to be a zero-mean, serially uncorrelated process and orthogonal to the estimated factor loadings $\hat{\beta}_p$, $\hat{\gamma}_p$, $\hat{\beta}_{p,\text{SMB}}$, $\hat{\beta}_{p,\text{HML}}$, and $\hat{\beta}_{p,\text{latent}}$.

5. Empirical Results

This section presents the empirical results, focusing on whether systematic skewness commands a significant risk premium across different market portfolio definitions. By testing the QMM across our three market proxies, we assess whether the skewness premium reflects fundamental risk pricing or merely an artefact of proxy construction. The first stage estimates time-series risk loadings for the 25 size and book-to-market sorted value-weighted portfolios.[7] These loadings capture each portfolio's sensitivity to the linear market factor β and to the market non-linear response γ, which proxies for systematic skewness. Table 3 reports these sensitivities alongside Newey-West t-statistics for each market proxy.

The β loadings exhibit the expected size gradient across all three market definitions - smaller firms display higher market betas, reflecting greater sensitivity to overall market movements. When using VWI (and its square) as market proxy (Panel A), small-cap portfolios show betas ranging from 0.99 to 1.29, whilst large-cap portfolios cluster near unity. The EWI proxy (Panel B) displays a similar pattern, with small-cap betas between 0.96 and 1.20 and large-cap betas ranging from 0.77 to 0.93. The PWI proxy (Panel C) produces systematically lower betas overall, reflecting its distinct weighting structure. More importantly, the γ loadings reveal pronounced patterns in systematic skewness exposure. Under the VWI proxy, smaller firms exhibit highly significant negative γ loadings, with t-statistics frequently exceeding $|3.0|$. For instance, the small-cap/low-value portfolio displays $\gamma = -1.83$ ($t = -3.15$), indicating disproportionate sensitivity to market downturns. As firm size increases, negative γ loadings diminish substantially. The largest firms display near-zero or even positive loadings - the big-cap/low-value portfolio has $\gamma = 0.37$ ($t = 2.05$), suggesting this portfolio hedges against crash risk. This size gradient in coskewness loadings directly corroborates the structural finding of Barone-Adesi et al. (2004), who demonstrate that small-firm portfolios systematically exhibit negative coskewness with the market whilst large-firm portfolios display positive coskewness — a pattern they interpret as evidence that smaller firms disproportionately amplify portfolio skewness risk.

The EWI produces broadly similar patterns, though with attenuated magnitudes. Small-cap portfolios still exhibit negative γ loadings ranging from -0.52 to -0.78, roughly half those under the VWI, with most remaining statistically significant at conventional levels. This attenuation reflects equal weighting's reduced exposure to extreme market movements driven by large-cap stocks. The size gradient persists — as firms get larger, γ loadings move closer to zero and lose significance — but the effect is less pronounced than under the VWI. The attenuation of the size–coskewness gradient under the EWI is consistent with the equal-weighted proxy's reduced exposure to the large-cap dynamics that drive the positive coskewness end of the distribution.

The PWI presents a notably different pattern - small size portfolios still show negative γ, but surprisingly, large size ones also display significant negative γ under PWI. This departure is consistent with the price-weighting scheme's high-priced stock concentration distorting the market's higher-moment structure relative to capitalisation- and equal-weighted alternatives, and suggests that the size–coskewness relationship documented by Barone-Adesi et al. (2004) is not invariant to market proxy construction — a result that directly motivates the multi-proxy design adopted in the present study.

[7] We also present the estimates for the 25 size and book-to-market sorted equally-weighted portfolios as robustness check.

Panel A: Value-Weighted Index (VWI)

Size \ Value	Alpha ($\alpha \times 10^2$)					Beta (β)					Gamma (γ)				
	Low	2	3	4	High	Low	2	3	4	High	Low	2	3	4	High
Small	−0.120	0.430*	0.440**	0.690***	0.860***	1.294***	1.115***	1.035***	0.961***	0.992***	−1.825***	−1.692***	−1.485***	−1.615***	−1.815***
	(−0.45)	(1.86)	(2.39)	(3.71)	(4.19)	(23.95)	(21.86)	(22.27)	(19.60)	(18.43)	(−3.15)	(−3.11)	(−3.43)	(−3.23)	(−3.41)
2	−0.010	0.390**	0.500***	0.530***	0.640***	1.302***	1.120***	1.030***	0.981***	1.093***	−1.089**	−1.369***	−1.272***	−1.100***	−1.514***
	(−0.07)	(2.46)	(3.63)	(3.72)	(3.62)	(27.98)	(28.65)	(25.41)	(24.10)	(18.44)	(−2.38)	(−3.55)	(−3.57)	(−2.75)	(−3.30)
3	−0.040	0.320**	0.310***	0.390***	0.520***	1.253***	1.087***	1.000***	0.991***	1.072***	−0.815**	−0.882***	−0.722**	−0.539	−0.840**
	(−0.24)	(2.50)	(2.59)	(3.19)	(3.10)	(31.10)	(33.45)	(28.44)	(25.59)	(22.26)	(−2.01)	(−2.82)	(−2.37)	(−1.61)	(−2.04)
4	0.000	0.130	0.240**	0.330***	0.350**	1.179***	1.079***	1.022***	0.997***	1.091***	−0.167	−0.530**	−0.586*	−0.297	−0.655
	(0.01)	(1.22)	(2.22)	(3.06)	(2.23)	(40.11)	(40.09)	(32.61)	(26.59)	(22.78)	(−0.60)	(−2.16)	(−1.69)	(−0.67)	(−1.37)
Big	−0.060	0.020	0.080	0.040	0.100	1.041***	0.966***	0.908***	0.941***	1.028***	0.369**	−0.054	−0.156	−0.240	−0.018
	(−0.88)	(0.28)	(0.94)	(0.37)	(0.61)	(59.58)	(50.31)	(39.67)	(23.25)	(19.28)	(2.05)	(−0.32)	(−0.63)	(−0.40)	(−0.03)

Note: This table reports the time-series regression results for the Quadratic Market Model (QMM):

$$r_{p,t} = \alpha_p^{(k)} + \beta_p^{(k)} r_{M,t}^{(k)} + \gamma_p^{(k)} (R_{M,t}^{(k)} - \bar{R}_M^{(k)})^2 + \epsilon_{p,t}^{(k)} \quad (7)$$
$$t = 1, \ldots, T; \quad k \in \{VWI, EWI, PWI\}$$

where $r_{p,t}$ is the excess return of the p-th value-weighted (VW) size and book-to-market sorted portfolio, $r_{M,t}^{(k)}$ is the excess return for the and k-th market proxy. Panel A reports results for the $k = VWI$ market proxy. Panels B and C (following page) provide resutls for $k = EWI$ and $k = PWI$. Alpha (α) values are multiplied by 100. Newey–West t-statistics (6 lags) are in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 3. Time-Series Risk Loadings: Sensitivity to Systematic Skewness

Panel B: Equally-Weighted Index (EWI)

	Alpha ($\alpha \times 10^2$)					**Beta** (β)					**Gamma** (γ)				
Size \ Value	Low	2	3	4	High	Low	2	3	4	High	Low	2	3	4	High
Small	−0.400*	0.090	0.140	0.370**	0.520***	1.203***	1.073***	1.017***	0.961***	1.012***	−0.784**	−0.516	−0.516*	−0.544	−0.690*
	(−1.71)	(0.47)	(0.98)	(2.46)	(3.00)	(23.80)	(21.58)	(26.86)	(23.75)	(24.04)	(−2.30)	(−1.47)	(−1.94)	(−1.59)	(−1.73)
2	−0.200	0.140	0.240***	0.280***	0.350***	1.190***	1.070***	1.006***	0.962***	1.091***	−0.562*	−0.605***	−0.523**	−0.395	−0.661**
	(−1.13)	(1.19)	(2.68)	(2.66)	(2.67)	(26.33)	(34.37)	(38.65)	(34.08)	(28.22)	(−1.78)	(−3.01)	(−2.44)	(−1.50)	(−2.27)
3	−0.180	0.130	0.130*	0.220***	0.290**	1.117***	1.024***	0.956***	0.959***	1.063***	−0.460*	−0.381**	−0.333*	−0.251	−0.362
	(−1.11)	(1.40)	(1.94)	(2.76)	(2.38)	(28.80)	(44.57)	(44.04)	(43.01)	(35.20)	(−1.70)	(−2.15)	(−1.90)	(−1.20)	(−1.26)
4	−0.040	0.030	0.090	0.210**	0.140	1.015***	0.987***	0.960***	0.944***	1.057***	−0.238	−0.444***	−0.328**	−0.196	−0.286
	(−0.26)	(0.53)	(1.47)	(2.49)	(1.21)	(30.63)	(54.25)	(63.27)	(41.92)	(37.34)	(−0.83)	(−3.75)	(−2.00)	(−0.59)	(−0.87)
Big	0.060	0.030	0.090	−0.010	0.020	0.804***	0.813***	0.775***	0.840***	0.926***	−0.170	−0.280	−0.376***	−0.267	−0.041
	(0.43)	(0.40)	(1.29)	(−0.13)	(0.13)	(29.70)	(43.52)	(40.80)	(28.93)	(24.20)	(−0.58)	(−1.30)	(−2.71)	(−0.69)	(−0.11)

Panel C: Price-Weighted Index (PWI)

	Alpha ($\alpha \times 10^2$)					**Beta** (β)					**Gamma** (γ)				
Size \ Value	Low	2	3	4	High	Low	2	3	4	High	Low	2	3	4	High
Small	0.580**	0.940***	0.910***	1.090***	1.260***	0.609***	0.547***	0.511***	0.488***	0.510***	−0.828**	−0.588**	−0.493*	−0.485*	−0.542*
	(2.01)	(4.00)	(4.60)	(5.71)	(5.78)	(10.57)	(10.95)	(11.15)	(10.73)	(10.54)	(−2.48)	(−2.03)	(−1.87)	(−1.82)	(−1.93)
2	0.680***	0.950***	0.980***	0.960***	1.120***	0.610***	0.547***	0.519***	0.490***	0.544***	−0.533*	−0.577**	−0.461*	−0.345	−0.492*
	(2.99)	(5.43)	(6.36)	(6.32)	(5.94)	(11.37)	(11.85)	(12.06)	(11.61)	(10.51)	(−1.75)	(−2.13)	(−1.88)	(−1.47)	(−1.71)
3	0.690***	0.920***	0.830***	0.990***	1.080***	0.579***	0.531***	0.494***	0.485***	0.537***	−0.551*	−0.484**	−0.360*	−0.444*	−0.424
	(3.42)	(6.10)	(6.28)	(6.54)	(6.12)	(11.56)	(12.28)	(12.37)	(11.48)	(11.30)	(−1.97)	(−2.01)	(−1.69)	(−1.97)	(−1.63)
4	0.790***	0.840***	0.880***	1.000***	0.970***	0.526***	0.499***	0.493***	0.474***	0.522***	−0.489**	−0.579**	−0.520**	−0.487**	−0.446*
	(4.39)	(5.94)	(6.73)	(7.51)	(5.50)	(11.71)	(11.93)	(12.00)	(10.87)	(10.63)	(−2.06)	(−2.40)	(−2.07)	(−2.14)	(−1.77)
Big	0.780***	0.700***	0.700***	0.680***	0.740***	0.420***	0.399***	0.383***	0.411***	0.450***	−0.527***	−0.445**	−0.414**	−0.459*	−0.274
	(5.34)	(5.90)	(5.63)	(4.79)	(4.09)	(11.55)	(11.50)	(10.90)	(9.82)	(9.81)	(−2.64)	(−2.16)	(−1.99)	(−1.80)	(−1.15)

Table 3. Time-Series Risk Loadings: Sensitivity to Systematic Skewness (Continued)

We now move to the second step of the empirics where the focus is shifted to evaluate whether the market price of systematic skewness is robust to the construction of the market proxy. Table 4 presents the Fama-MacBeth risk premia estimates (λ) across all three market definitions, both for the baseline QMM (Panel A) and for the PCA corrected specification (Panel B).

Panel A: Baseline Quadratic Market Model

Market Proxy	Intercept (λ_0)	Market Risk (λ_β)	Skewness Risk (λ_γ)
VWI	1.708***	−1.024**	−0.146*
	(3.51)	(−2.10)	(−1.74)
EWI	0.773**	−0.071	−0.095
	(2.33)	(−0.18)	(−0.48)
PWI	0.743*	0.994	1.025***
	(1.90)	(1.24)	(2.88)

Panel B: PCA-Augmented Quadratic Market Model

Market Proxy	Intercept (λ_0)	Market Risk (λ_β)	Skewness Risk (λ_γ)	Latent Risk (λ_L)
VWI	2.242***	−1.648***	0.265***	2.760***
	(4.40)	(−3.21)	(2.94)	(3.47)
EWI	0.593***	0.102	−0.206	−0.269
	(2.61)	(0.36)	(−0.97)	(−0.52)
PWI	1.150***	−0.571	1.164***	2.306
	(3.51)	(−0.71)	(3.24)	(1.34)

Note: This table reports second-stage Fama–MacBeth (1973) cross-sectional risk premia. Panel A presents the baseline QMM estimates according to:

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + u_{p,t} \tag{8}$$

Panel B presents the augmented model including latent factors identified via Principal Component Analysis (PCA):

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{5,t}\hat{\beta}_{p,\text{latent}} + e_{p,t} \tag{10}$$

The cross-section consists of 25 value-weighted (VW) size and book-to-market sorted portfolios. All coefficients are reported in percent per month. t-statistics are reported in parentheses. Asterisks ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 4. Fama–MacBeth Regressions: Systematic Skewness and Latent Risk Premia

The baseline results in Panel A reveal substantial variation in the skewness premium across market definitions. Under the VWI, the skewness premium is negative but statistically insignificant ($\lambda_\gamma = -0.15\%$, $t = -1.74$), suggesting that investors do not demand meaningful compensation for bearing systematic skewness. The EWI produces a similarly insignificant estimate ($\lambda_\gamma = -0.09\%$, $t = -0.48$). By contrast, the PWI yields a large and highly significant positive premium ($\lambda_\gamma = 1.02\%$, $t = 2.88$), indicating that portfolios with greater exposure to crash risk earn substantially higher returns. This divergence across proxies initially appears to support the Roll's critique - if the skewness premium depends so heavily on proxy choice, it may simply reflect measurement error rather than genuine equilibrium pricing. However, Panel B tells a different story. Once we control for the latent factor extracted from the first stage residuals, the skewness premium becomes positive and significant under the VWI ($\lambda_\gamma = 0.26\%$, $t = 2.94$). This suggests that measurement noise, rather than a lack of systematic pricing, drives the baseline insignificance. The latent factor itself commands a substantial premium ($\lambda_{latent} = 2.76\%$, $t = 3.47$), reflecting compensation for exposure to the unobserved component of the market portfolio. Under the PWI, the skewness premium remains large and significant ($\lambda_\gamma = 1.16\%$, $t = 3.24$), with an economically meaningful magnitude. However, the EWI continues to show an insignificant skewness premium even after the PCA correction ($\lambda_\gamma = -0.21\%$, $t = -0.97$), suggesting that equal weighting may obscure systematic skewness pricing through its dilution of large-cap effects.

Overall, the pattern across panels suggests two key findings. First, systematic skewness is priced in equilibrium, but proxy invalidity can obscure this relationship in conventional tests. Second, the magnitude of the premium varies substantially across market definitions - PWI yields premiums roughly four times larger than VWI. This variation likely reflects differences in how each weighting scheme captures systematic skewness rather than pure measurement error, suggesting that the economic content of skewness risk depends meaningfully on market portfolio construction.

Adding Factor Controls

A rigorous test of the QMM requires demonstrating that systematic skewness is not merely a proxy for the well-known size and value anomalies. Table 5 introduces Fama-French controls (SMB and HML) to assess whether the skewness premium persists after accounting for these established factors.

The findings in Table 5 confirm that systematic skewness remains a distinct and priced risk factor even after controlling for size and value effects. In Panel A, the skewness premium under VWI becomes positive and highly significant ($\lambda_\gamma = 0.30\%$, $t = 3.33$), in contrast to its insignificance in the baseline specification without controls. This pattern suggests that omitting the Fama-French factors confounds the skewness premium. The EWI produces a similar positive premium ($\lambda_\gamma = 0.29\%$, $t = 1.97$), though with weaker statistical significance, whilst the PWI maintains a large and significant premium ($\lambda_\gamma = 0.57\%$, $t = 2.75$). Importantly, the SMB premium is statistically insignificant across all market proxy specifications, suggesting that the size premium documented by Banz (1981) and incorporated into the Fama–French (1992) framework may be partially absorbed once systematic skewness is explicitly priced. However, the HML premium survives the inclusion of coskewness, consistent with the value effect representing a distinct and independent source of priced risk. Indeed, the HML premium remains significant across all three specifications (t-statistics around 2.2–2.4), indicating that value effects persist alongside coskewness pricing. Panel B shows that the PCA correction has relatively modest effects once Fama-French controls are included. For VWI, the skewness premium remains positive and significant ($\lambda_\gamma = 0.27\%$, $t = 2.93$), with the latent factor contributing only marginally ($\lambda_{latent} = 0.23\%$, $t = 1.32$). This suggests that much of the measurement error addressed by the latent factor is already captured by SMB and HML

Panel A: Baseline QMM with Fama–French Controls

Market Proxy	Intercept (λ_0)	λ_β	λ_γ	λ_{SMB}	λ_{HML}
VWI	1.507***	−0.956***	0.303***	0.119	0.308**
	(5.80)	(−3.23)	(3.33)	(1.00)	(2.23)
EWI	1.129***	−0.399	0.292**	0.115	0.321**
	(5.28)	(−1.46)	(1.97)	(0.97)	(2.32)
PWI	0.976***	−0.214	0.569***	0.126	0.333**
	(4.63)	(−0.40)	(2.75)	(1.07)	(2.41)

Panel B: PCA-Augmented QMM with Fama–French Controls

Market Proxy	Intercept (λ_0)	λ_β	λ_γ	λ_{SMB}	λ_{HML}	λ_L
VWI	1.571***	−1.063***	0.275***	0.148	0.305**	0.229
	(5.62)	(−3.21)	(2.93)	(1.26)	(2.21)	(1.32)
EWI	1.216***	−0.382	0.376***	0.182	0.278**	−0.666***
	(5.55)	(−1.40)	(2.63)	(1.56)	(2.04)	(−2.80)
PWI	1.156***	2.668***	0.398*	0.167	0.300**	−7.623***
	(5.19)	(2.74)	(1.95)	(1.43)	(2.21)	(−3.23)

Note: This table reports second-stage Fama–MacBeth (1973) risk premia while controlling for the Fama–French (1993) size (SMB) and value (HML) factors. Panel A presents the baseline specification:

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{3,t}\hat{\beta}_{p,\text{SMB}} + \lambda_{4,t}\hat{\beta}_{p,\text{HML}} + u_{p,t} \tag{9}$$

Panel B includes latent factors identified via PCA:

$$r_{p,t} = \lambda_{0,t} + \lambda_{1,t}\hat{\beta}_p + \lambda_{2,t}\hat{\gamma}_p + \lambda_{3,t}\hat{\beta}_{p,\text{SMB}} + \lambda_{4,t}\hat{\beta}_{p,\text{HML}} + \lambda_{5,t}\hat{\beta}_{p,\text{latent}} + e_{p,t} \tag{11}$$

The cross-section comprises 25 value-weighted (VW) size and book-to-market sorted portfolios. All coefficients are in percent per month. t-statistics are in parentheses. Asterisks ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 5. Fama–MacBeth Regressions: Impact of Fama–French Controls

controls. The skewness premium strengthens slightly for EWI ($\lambda_\gamma = 0.38\%$, $t = 2.63$), whilst the PWI produces a somewhat attenuated premium ($\lambda_\gamma = 0.40\%$, $t = 1.95$). Interestingly, the latent factor premiums for EWI and PWI are large, negative, and significant, potentially reflecting the interaction between equal/price weighting and the unobserved market component. Overall, λ_γ is positive and statistically significant across all market proxy specifications, and remains so after augmenting the model with Fama–French controls and PCA-extracted latent factors. This is consistent with the broader skewness pricing narrative of Boyer et al. (2010), who document a significant relationship between expected skewness and returns in the cross-section, and extends their evidence to the systematic component of skewness.

6. Robustness Checks

To ensure that the documented pricing of systematic skewness is a fundamental market feature rather than an artefact of specific empirical configurations, we conduct an extensive series of robustness tests. These checks address the sensitivity of our findings to test asset construction, the precision of risk premia estimates under stringent econometric adjustments, and the stability of results across alternative specifications.

6.1 Alternative Test Asset Weighting

Our primary results use value-weighted test portfolios, which naturally place greater emphasis on large firms. To assess whether skewness pricing depends on this choice, we repeat our analysis using equally-weighted test assets. Tables 6 and 7 report the baseline and Fama-French augmented results, respectively, for equally-weighted portfolios.

The results confirm that coskewness pricing is not driven by test asset weighting. In Table 6, Panel A shows patterns similar to the value-weighted test assets: the baseline VWI and EWI produce insignificant skewness premiums, whilst the PWI yields a positive but marginally significant estimate ($\lambda_\gamma = 0.74\%$, $t = 1.49$). Panel B demonstrates that the PCA correction substantially strengthens the skewness premium under the PWI ($\lambda_\gamma = 1.50\%$, $t = 3.33$), whilst the VWI and EWI show modest improvements in significance. Table 7 reveals that adding Fama-French controls produces statistically significant skewness premiums across all specifications. Moreover, with PCA correction and factor controls (Panel B), the skewness premium reaches $\lambda_\gamma = 0.46\%$ ($t = 3.43$) for EWI, demonstrating that systematic skewness pricing emerges consistently once we properly account for measurement error and factor controls. The robustness across test asset weighting schemes strengthens our conclusion that skewness represents a fundamental dimension of risk pricing.

6.2 Shanken Standard Error Correction

A significant concern in multi-stage asset pricing models is the errors-in-variables (EiV) problem - the second-stage cross-sectional regressions treat first-stage factor loadings as known quantities rather than estimated parameters. This sampling variation can lead to understated standard errors and inflated $t - stats$. To address this, we implement the Shanken (1992) correction, which adjusts standard errors to account for estimation uncertainty in the factor loadings. Table 8 reports the results for our full specification across all three market proxies.

Panel A: Baseline Quadratic Market Model

Market Proxy	Intercept (λ_0)	Market Risk (λ_β)	Skewness Risk (λ_γ)
VWI	1.854***	−1.045***	−0.115
	(4.46)	(−2.70)	(−1.21)
EWI	1.123***	−0.332	−0.051
	(2.99)	(−0.79)	(−0.21)
PWI	1.262***	−0.241	0.737
	(2.88)	(−0.37)	(1.49)

Panel B: PCA-Augmented Quadratic Market Model

Market Proxy	Intercept (λ_0)	Market Risk (λ_β)	Skewness Risk (λ_γ)	Latent Risk (λ_L)
VWI	2.136***	−1.462***	0.182	2.130**
	(4.54)	(−3.15)	(1.43)	(1.97)
EWI	1.508***	−0.732**	0.196	0.495
	(4.81)	(−2.52)	(1.16)	(0.97)
PWI	2.131***	−2.950***	1.503***	4.708**
	(4.70)	(−2.87)	(3.33)	(2.34)

Note: This table reports second-stage Fama–MacBeth (1973) risk premia for robustness. Unlike the previous analysis, the test assets $r_{p,t}$ are the excess returns of the 25 equally-weighted (EW) size and book-to-market sorted portfolios. Panel A presents the baseline estimates (Eq. 8), and Panel B presents the PCA-augmented specification (Eq. 10). All coefficients are reported in percent per month. t-statistics are in parentheses. Asterisks ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 6. Fama–MacBeth Regressions: Robustness Check with Equally-Weighted Test Assets

Panel A: Baseline QMM with Fama–French Controls

Market Proxy	Intercept (λ_0)	λ_β	λ_γ	λ_{SMB}	λ_{HML}
VWI	1.562***	−0.998***	0.197**	0.131	0.345**
	(5.65)	(−3.50)	(2.13)	(0.97)	(2.49)
EWI	1.407***	−0.711**	0.251	0.087	0.358***
	(4.68)	(−2.49)	(1.56)	(0.67)	(2.59)
PWI	1.415***	−1.130**	0.568*	0.114	0.385***
	(5.87)	(−2.14)	(1.92)	(0.88)	(2.84)

Panel B: PCA-Augmented QMM with Fama–French Controls

Market Proxy	Intercept (λ_0)	λ_β	λ_γ	λ_{SMB}	λ_{HML}	λ_L
VWI	1.715***	−1.142***	0.210**	0.236**	0.311**	−0.526*
	(6.54)	(−4.09)	(2.30)	(2.02)	(2.27)	(−1.78)
EWI	1.863***	−1.165***	0.457***	0.208*	0.296**	−0.643**
	(7.04)	(−4.09)	(3.43)	(1.76)	(2.19)	(−2.37)
PWI	1.225***	1.219	0.116	0.239*	0.354***	−5.757**
	(5.11)	(1.05)	(0.44)	(1.93)	(2.64)	(−2.21)

Note: This table reports second-stage Fama–MacBeth (1973) risk premia for the 25 equally-weighted (EW) size and book-to-market sorted portfolios. Panel A presents the baseline specification (Eq. 9) including Fama–French (1993) controls. Panel B augments this model with latent factors identified via PCA (Eq. 11). All coefficients are reported in percent per month. t-statistics are in parentheses. Asterisks ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Table 7. Fama–MacBeth Regressions: Robustness with Equally-Weighted Assets and Controls

	VWI Proxy			EWI Proxy			PWI Proxy		
Parameter	Prem.	t_{NW}	t_{Sh}	Prem.	t_{NW}	t_{Sh}	Prem.	t_{NW}	t_{Sh}
Intercept	1.571***	5.62	4.37	1.216***	5.55	4.43	1.156***	5.19	4.16
λ_β	−1.063**	−3.21	−2.49	−0.382	−1.40	−1.12	2.668**	2.74	2.19
λ_γ	0.275**	2.93	2.27	0.376**	2.63	2.10	0.398	1.95	1.57
λ_{SMB}	0.148	1.26	0.98	0.182	1.56	1.24	0.167	1.43	1.14
λ_{HML}	0.305*	2.21	1.72	0.278	2.04	1.63	0.300*	2.21	1.77
λ_L	0.229	1.32	1.02	−0.666**	−2.80	−2.23	−7.623***	−3.23	−2.59
Adj. Factor (c)	1.660			1.571			1.559		

Note: This table reports the second-stage Fama–MacBeth (1973) risk premia for the full PCA-augmented model, incorporating the Shanken (1992) correction for the errors-in-variables (EIV) bias. t_{NW} denotes the Newey–West (1987) adjusted t-statistics (6 lags), while t_{Sh} provides the EIV-corrected t-statistics. Significance asterisks ***, **, and * are determined strictly by the Shanken t-statistics at the 1%, 5%, and 10% levels, respectively.

Table 8. Fama–MacBeth Regressions with Shanken (1992) EIV Correction

The results demonstrate that the skewness premium survives this stringent adjustment. For the VWI and EWI, the skewness premium remains statistically significant after the Shanken correction, with adjusted t-statistics of 2.27 and 2.10, respectively. Whilst the adjustment factors (c) are substantial—ranging from 1.56 to 1.66—due to the high-dimensional covariance matrix of factors, the persistence of significant premiums confirms that skewness pricing is robust to estimation uncertainty in factor loadings. For PWI, the skewness premium remains positive but loses marginal significance ($t_{Shanken} = 1.57$), though the economic magnitude remains meaningful at 0.40% per month. Moreover, the latent factor premium remains highly significant for the EWI and PWI specifications even after the Shanken adjustment ($t_{Shanken} = -2.23$ and -2.59, respectively), confirming that our PCA procedure successfully isolates a priced systematic component that traditional market proxies fail to capture. The HML premium also persists across all three specifications with Shanken adjusted t-statistics around 1.6–1.8, demonstrating that value effects remain distinct from skewness pricing.

7. Conclusion

This study investigates whether investors actually price systematic skewness—the risk that their portfolio crashes when the market crashes. We test this question across three different market proxies to see if the skewness premium is a real feature of how markets work or just an artefact of how we measure the market. Our key finding is simple: coskewness is priced, and it matters. Importantly, we find that the coskewness premium is positive and statistically significant across all market proxy specifications, and remains so after augmenting the model with Fama–French controls and PCA-extracted latent factors.

Our results paint a clear picture. Low value stocks, i.e., growth, are hit harder during market crashes than large, low growth stocks. This is what we would expect theoretically — investors should demand higher returns for bearing crash risk. But this skewness premium exists independently of the size and value effects. When we control for those factors, the skewness premium persists. This tells us that investors care about crash risk for its own sake, not just because small and growth stocks happen to crash more often.

A practical insight emerges from comparing different market indices. Even though it is subject to corporate action rebalancing, it is the PWI that captures the most stable skewness premium than the traditional value-weighted index. EW and PW indices both place lower weight on the effect of large-cap stocks, allowing us to see how the broader market moves.

On a methodological level, our work directly addresses the Roll (1977) critique, which posits that unobservable true reference portfolios render any empirical asset pricing test inherently a joint test of the model and the efficiency of the chosen proxy. We resolve this identification challenge by implementing a three-stage regression model that extracts systematic variation omitted by traditional market proxies. Firstly, our empirical design follows the methodological recommendations of Fama and French (2020), who demonstrate that combining time-series loading estimates with cross-sectional pricing regressions yields more reliable inference than either approach alone. Secondly, by controlling for this latent factor, we effectively filter measurement error from our risk loading estimates, allowing the skewness premium to emerge with statistical clarity. The significance of the latent factor premium across all specifications suggests that measurement error, rather than absence of genuine pricing, obscured the skewness premium in prior research. One limitation worth acknowledging concerns the use of the 25 Fama–French size and book-to-market sorted portfolios as test assets. As Tauscher and Wallmeier (2016) demonstrate, when test portfolios and factors share the same sorting variables, a mechanical overlap bias can inflate the model's apparent explanatory fit. Whilst the $\lambda_\gamma$is not a construction input for the test portfolios — and thus no direct mechanical overlap arises — we acknowledge that coskewness is empirically correlated with firm size, as documented by Barone-Adesi et al. (2004). The significance of λ_γ should therefore be interpreted with this in mind, and we address this concern further in the robustness analysis.

Our findings indicate that asset returns reflect compensation for exposure to non-linear market risk that standard two-moment models fail to capture. Systematic skewness represents a priced risk dimension that is empirically distinct from size and value effects, even after controlling for Fama-French factors. These results contribute to our understanding of the cross-sectional determinants of expected returns and underscore the importance of accounting for distributional preferences beyond mean and variance in empirical asset pricing research.

References

[1] Adcock, C. J. (2003). Asset pricing and portfolio selection based on the multivariate skew-student distribution. *Annals of Operations Research*, 116(1), 201–223.

[2] Adcock, C. J., & Shutes, K. (2000). Portfolio selection based on the multivariate skew-normal distribution. In *Advances in Quantitative Methods for Financial Markets* (pp. 45–62). Springer, Berlin.

[3] Arrow, K. J. (1971). *Essays in the Theory of Risk-Bearing*. Markham Publishing Company, Chicago.

[4] Athayde, G. M. de, & Flôres, R. G. Jr. (1997). Finding a maximum skewness portfolio—a general solution to three-moments portfolio choice. *Journal of Economic Dynamics and Control*, 21(9-10), 1529–1543.

[5] Athayde, G. M. de, & Flôres, R. G. Jr. (2000). Incorporating skewness and kurtosis in portfolio optimization: A multivariate approach. *Brazilian Review of Econometrics*, 20(1), 3–28.

[6] Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics*, 9(1), 3–18.

[7] Barberis, N., & Huang, M. (2008). Stocks as lotteries: The implications of probability weighting for security prices. *American Economic Review*, 98(5), 2066–2100.

[8] Barone-Adesi, G. (1985). Arbitrage equilibrium with skewed asset returns. *Journal of Financial and Quantitative Analysis*, 20(3), 299–313.

[9] Barone-Adesi, G., Gagliardini, P., & Urga, G. (2004). Testing asset pricing models with coskewness. *Journal of Business and Economic Statistics*, 22(4), 474–485.

[10] Boyer, B., Mitton, T., and Vorkink, K. (2010). Expected idiosyncratic skewness. *Review of Financial Studies*, 23(1), 169–202.

[11] Dittmar, R. F. (2002). Nonlinear pricing kernels, kurtosis preference, and evidence from the cross-section of equity returns. *The Journal of Finance*, 57(1), 369–403.

[12] Fama, E. F., & French, K. R. (1992). The cross-section of expected stock returns. *The Journal of Finance*, 47(2), 427–465.

[13] Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, 33(1), 3–56.

[14] Fama, E. F., and French, K. R. (2020). Comparing cross-section and time-series factor models. *Review of Financial Studies*, 33(5), 1891–1926.

[15] Fama, E. F., & MacBeth, J. D. (1973). Risk, return, and equilibrium: Empirical tests. *Journal of Political Economy*, 81(3), 607–636.

[16] Giglio, S., & Xiu, D. (2021). Asset pricing with omitted factors. *Journal of Political Economy*, 129(7), 1947–1990.

[17] Harvey, C. R., & Siddique, A. (2000). Conditional skewness in asset pricing tests. *The Journal of Finance*, 55(3), 1263–1295.

[18] Jurczenko, E. and Maillet, B. (2001). The 3-CAPM: theoretical foundations and an asset-pricing model comparison in a unified framework. In Dunis, C. L., Timmermann, A. and Moody, J. (eds.), *Developments in Forecast Combination and Portfolio Choice*. John Wiley & Sons, Chichester.

[19] Jurczenko, E., & Maillet, B. (2015). *Theories and Models for Financial Markets.* John Wiley & Sons, Hoboken, NJ.

[20] Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica*, 47(2), 263–291.

[21] Karehnke, P. (2024). Systematic skewness and stock returns. *The Review of Asset Pricing Studies*, 14(4), 578–612.

[22] Kon, S. J. (1984). Models of stock returns—a comparison. *The Journal of Finance*, 39(1), 147–165.

[23] Kraus, A., & Litzenberger, R. H. (1976). Skewness preference and the valuation of risk assets. *The Journal of Finance*, 31(4), 1085–1100.

[24] Langlois, H. (2020). Measuring skewness premia. *Journal of Financial Economics*, 135(2), 399–424.

[25] Leland, H. E. (1998). Agency costs, risk management, and capital structure. *The Journal of Finance*, 53(4), 1213–1243.

[26] Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *The Review of Economics and Statistics*, 47(1), 13–37.

[27] Markowitz, H. (1952). Portfolio selection. *The Journal of Finance*, 7(1), 77–91.

[28] Mills, T. C. (1995). Modelling skewness and kurtosis in the London Stock Exchange FT-SE index return distributions. *Journal of the Royal Statistical Society: Series D (The Statistician)*, 44(3), 323–332.

[29] Mitton, T., & Vorkink, K. (2007). Equilibrium underdiversification and the preference for skewness. *The Review of Financial Studies*, 20(4), 1255–1288.

[30] Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica*, 34(4), 768–783.

[31] Newey, W. K., & West, K. D. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3), 703–708.

[32] Peiró, A. (1999). Skewness in financial returns. *Journal of Banking & Finance*, 23(6), 847–862.

[33] Peiró, A. (2002). Skewness in individual stocks at different investment horizons. *Quantitative Finance*, 1(3), 365–373.

[34] Pratt, J. W. (1964). Risk aversion in the small and in the large. *Econometrica*, 32(1-2), 122–136.

[35] Premaratne, G., & Tay, A. S. (2002). How should we interpret evidence of time varying conditional skewness? (Working Paper No. 2002-01). National University of Singapore, Department of Economics.

[36] Roll, R. (1977). A critique of the asset pricing theory's tests Part I: On past and potential testability of the theory. *Journal of Financial Economics*, 4(2), 129–176.

[37] Rubinstein, M. (1973). The fundamental theorem of parameter-preference security valuation. *Journal of Financial and Quantitative Analysis*, 8(1), 61–69.

[38] Scott, R. C., & Horvath, P. A. (1980). On the direction of preference for moments of higher order than the variance. *The Journal of Finance*, 35(4), 915–919.

[39] Shanken, J. (1992). On the estimation of beta-pricing models. *The Review of Financial Studies*, 5(1), 1–33.

[40] Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance*, 19(3), 425–442.

[41] Tauscher, K., and Wallmeier, M. (2016). Portfolio overlapping bias in tests of the Fama–French three-factor model. *European Financial Management*, 22(3), 367–393.

A. Appendix

This appendix provides visual assessments to complement the Jarque-Bera tests presented in the main analysis. While the formal tests already reject normality, these plots offer a visual representation of how the distributions deviate from the normal.

A.1 Distribution Plots

Figure 1 shows the distribution of returns for each of the three market indices alongside their kernel density estimates. The visual inspection confirms the departures from normality identified by the statistical tests.

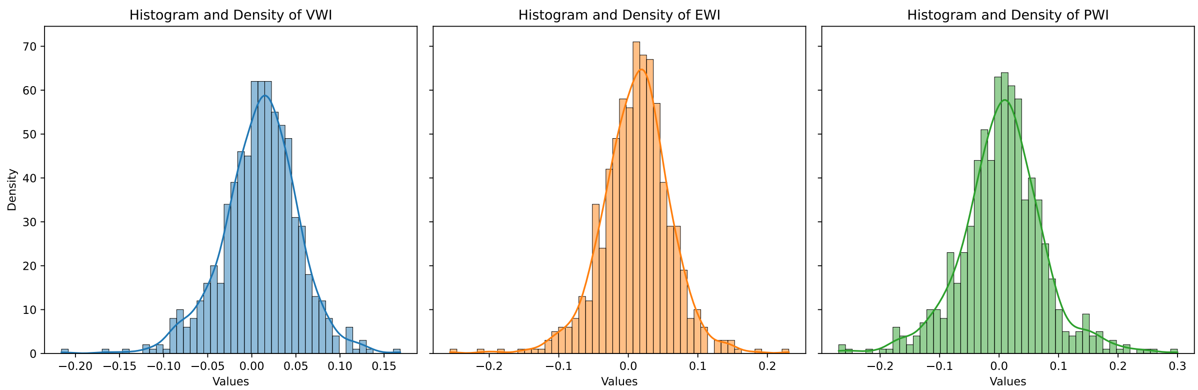

Figure 1. Distribution of returns across VWI, EWI, and PWI

A.2 Q-Q Plots

Figure 2 plots the sample quantiles against theoretical normal quantiles. Under normality, the points should fall along the diagonal line. The clear deviations, particularly in the tails, visually confirm the rejection of normality.

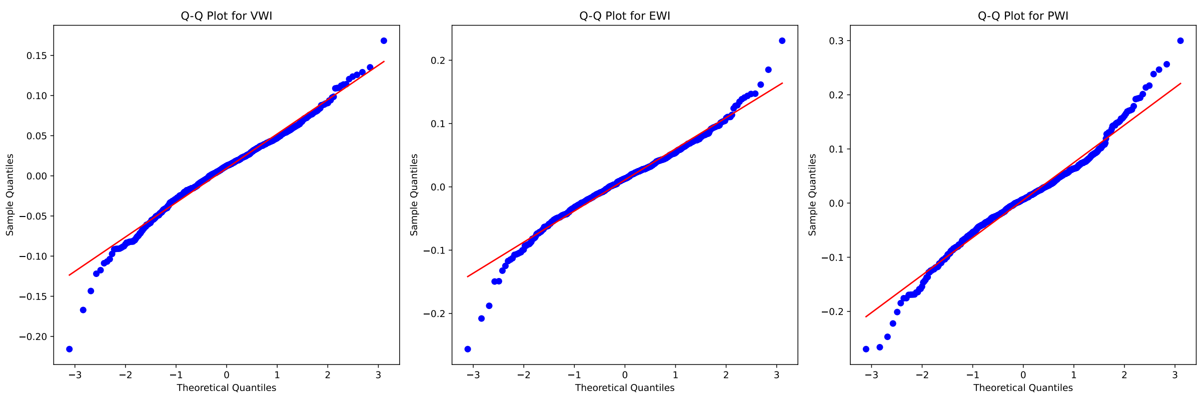

Figure 2. Q-Q plots for normality assessment