Economics Letters 250 (2025) 112298

Contents lists available at ScienceDirect

Economics Letters

journal homepage: www.elsevier.com/locate/ecolet

Estimating volatility-of-volatility: A comparative analysis

Jianglei Yuan [a], Dehong Liu [a], Carl R. Chen [b], Mingye Ma [c],*

[a] School of Economics and Management, Beijing Jiaotong University, China
[b] Department of Economics and Finance, University of Dayton, United States
[c] Department of Economics, University of Southampton, United Kingdom

ARTICLE INFO

Keywords:
Volatility-of-volatility
VVIX index
Prediction

ABSTRACT

This paper compares four volatility-of-volatility estimates with the CBOE's VVIX index for prediction accuracy. Hybrid estimates, combining historical and model-based components, show closer alignment with VVIX. Practical limitations are briefly noted.

1. Introduction

Volatility-of-volatility (VV) is a key parameter in stochastic volatility models, capturing higher-moment risks in asset returns (Heston, 1993). Recent research suggests that VV risk is priced in various markets, including index (Hollstein and Prokopczuk, 2017; Huang et al., 2019; Kaeck, 2018; Branger et al., 2017) and crude oil markets (Roh et al., 2021). The VVIX index, constructed by the Chicago Board Options Exchange (CBOE), is a widely used estimate of VV (Zang et al., 2017; Cheng, 2019).[1] VVIX, however, adopts the VIX's complex formula, which limits its use beyond the S&P500 due to extensive data requirements.[2] This limitation highlights the need for simpler VV estimates/measures applicable to diverse financial markets. We compare four VV estimates adaptable to individual stocks and assess their alignment with VVIX. Identifying the VV estimate closest to VVIX can enhance VV research and applications beyond the S&P500.

This study examines four estimates of volatility-of-volatility (VV), categorized into three groups. First, the model-based estimate, Risk-Neutral Volatility-of-Volatility (RNVV), proposed by Carr and Wu (2020), relies on the theoretical curvature of the implied volatility surface to estimate VV. Second, the hybrid estimates, including Volatility-of-Volatility in the EWMA Model (EMVV) by Roh et al. (2021) and Moving Average Implied Volatility-of-Volatility (MIVV) by Baltussen et al. (2018), combine statistical frameworks with model-derived inputs, leveraging both historical data and implied volatility. Third, the historical measure, Statistical Volatility-of-Volatility (STVV), utilizes realized volatility data (instead of implied) to provide a purely statistical approach.

2. Estimates of volatility-of-volatility

The model-based estimate assumes stochastic volatility processes, unlike historical measures derived statistically from data with fewer assumptions. Hybrid estimates blend theoretical and historical inputs to estimate VV.

2.1. Model-based estimate

Risk-Neutral Volatility-of-Volatility (RNVV)

Carr and Wu (2020) propose a top-down valuation framework in which the curvature coefficient of the option-implied volatility surface is used to estimate risk-neutral VV. This framework assumes that changes in at-the-money (ATM) implied volatility reflect similar changes in other implied volatilities across the volatility surface. Empirically, the RNVV is estimated by regressing the difference between the implied volatility at a given strike price (IV_t) and the ATM implied volatility ($IV_{t,ATM}$) on moneyness variables (k_1, k_2), as follows:

$$IV_t - IV_{t,ATM} = \alpha + \beta k_1 + \gamma k_2 + \epsilon, \tag{1}$$

where: γ represents the curvature coefficient, which serves as the estimate of risk-neutral VV. k_1 and k_2 are upward-adjusted moneyness and downward-adjusted moneyness, respectively. The formulas for k_1 and k_2 are:

$$k_1 = \ln\left(\frac{S_T}{S_t}\right) + \frac{(T-t)\cdot IV_t^2}{2}, \quad k_2 = \ln\left(\frac{S_T}{S_t}\right) - \frac{(T-t)\cdot IV_t^2}{2},$$

* Corresponding author.
E-mail addresses: 22110173@bjtu.edu.cn (J. Yuan), dhliu@bjtu.edu.cn (D. Liu), rchen1@udayton.edu (C.R. Chen), Mingye.Ma@soton.ac.uk (M. Ma).

[1] VVIX estimates the expected volatility of the 30-day forward price of VIX, a benchmark index to estimate the market's expectation of future volatility. More details are available at https://www.cboe.com/tradable_products/vix/ (VIX) and https://www.cboe.com/us/indices/dashboard/vvix/ (VVIX).

[2] For example, the VIX formula requires dense strikes and a forward index level. We detail the issue of sparse strikes in the discussion.

https://doi.org/10.1016/j.econlet.2025.112298
Received 4 February 2025; Received in revised form 18 March 2025; Accepted 19 March 2025
Available online 28 March 2025

where S_t is the stock price on trading day t, and S_T is the stock price on expiration date T.

Referring to Carr and Wu (2020), we use the restricted least squares method to constrain the curvature coefficient γ to be non-negative. RNVV assumes a well-behaved implied volatility surface with uniform ATM changes across strikes, an assumption that may fail in stressed markets. Fig. 1 shows deviations during such periods.

2.2. Hybrid estimates

Volatility-of-volatility in the EWMA model (EMVV)

Roh et al. (2021) provide an alternative approach to capture the dynamic nature of volatility fluctuations. The EMVV is calculated as:

$$EMVV_t = \lambda EMVV_{t-1} + (1-\lambda) r_{t-1}^2, \tag{2}$$

where: $\lambda = 0.94$ is the smoothing parameter,[3] which assigns greater weight to more recent data; $EMVV_0 = 0$ is the initial value; r_t represents the log return of the VIX index at time t.

This method leverages the forward-looking nature of the VIX index to estimate future VV. The EWMA framework inherently prioritizes recent observations, allowing the method to adapt quickly to changing market conditions. However, due to its reliance on historical volatility components, the EMVV may be disproportionately influenced by large shocks or abrupt movements in the recent past. Such disturbances can lead to the overestimation of future VV, particularly in highly volatile markets.

Moving Average Implied Volatility-of-Volatility (MIVV)

Baltussen et al. (2018) introduce the Moving Average Implied Volatility-of-Volatility (MIVV) as a hybrid estimate that combines the forward-looking properties of implied volatility with historical volatility patterns. It is calculated as the moving average of squared deviations of implied volatilities:

$$MIVV_t = \frac{1}{n}\sum_{i=1}^{n}\left(IV_{t-i} - \bar{IV}_t\right)^2, \tag{3}$$

where the mean implied volatility $\bar{IV}_t$ is defined as:

$$\bar{IV}_t = \frac{1}{n}\sum_{i=1}^{n} IV_{t-i}.$$

Here, IV_{t-i} represents the implied volatilities of options observed over the past n periods at time t. This method emphasizes historical deviations in implied volatility, capturing both short-term fluctuations and broader trends in VV.

2.3. Historical measure

Statistical Volatility of Volatility (STVV)

Unlike hybrid or model-based approaches, STVV relies purely on historical data and avoids assumptions about forward-looking dynamics. Historical variance HV_t is calculated as:

$$HV_t = \frac{1}{n-1}\sum_{i=1}^{n}\left(r_{i,t} - \bar{r}_t\right)^2,$$

where $r_{i,t}$ represents the log-returns of the underlying asset, and $\bar{r}_t$ is the average return over the period.

STVV is then computed as the moving variance of HV_t:

$$STVV_t = \frac{1}{n-1}\sum_{i=1}^{n}\left(HV_{t-i} - \bar{HV}_t\right)^2, \tag{4}$$

where $\bar{HV}_t$ is the average historical variance over the moving window.

This measure is entirely backward-looking, capturing VV through observed market fluctuations without relying on implied volatility or other theoretical models. The delay in response due to the lack of incorporating market expectations may cause this measure to underestimate VV during periods of rapidly increasing volatility.

Table 1
Descriptive statistics for volatility-of-volatility estimates.

Method	Obs.	Mean	SD	Min	Median	Max
VVIX	4176	0.893	0.349	0.025	0.809	4.309
RNVV	4176	1.366	1.263	0.000	1.139	7.071
EMVV	4176	0.571	1.266	0.000	0.185	17.554
MIVV	4176	0.143	0.195	0.006	0.089	2.475
STVV	4176	0.150	0.159	0.012	0.109	1.611

Notes: VVIX is normalized based on the value from CBOE, i.e., $VVIX = (VVIX_{CBOE}/100)^2$ to be comparable with other estimates.

3. Empirical results

3.1. Data and descriptive statistics

This study employs S&P500 option data from OptionMetrics (March 6, 2006–December 30, 2022). Standard filters, excluding deep out-of-the-money and illiquid options (Roh et al., 2021; Baltussen et al., 2018), yield a dataset of 2,573,032 calls and 3,850,449 puts.

Fig. 1 illustrates the time series of VV estimates derived from the four methods alongside the VVIX index, highlighting their evolution over the sample period. Table 1 provides descriptive statistics for these methods, summarizing their distributional characteristics.

3.2. Measuring deviations

We assess the predictive accuracy of these four VV estimates by analyzing their deviations from VVIX across volatility levels, using both *point estimates* for numerical accuracy and *trend estimates* for directional alignment.

3.2.1. Point estimates

Table 2 provides detailed metrics for point estimates across different volatility regimes (low, medium, and high), and for the full sample. The metrics reported include the following: The Mean Absolute Error (MAE) quantifies the average magnitude of errors, emphasizing overall deviations without disproportionately penalizing outliers. The Mean Squared Error (MSE) emphasizes larger deviations by squaring individual errors, offering insights into extreme variations. The Mean Absolute Percentage Error (MAPE) expresses errors as a percentage of observed values, enhancing interpretability and comparability across methods. Lastly, the Weighted Mean Absolute Percentage Error (WMAPE) adjusts for the relative importance of individual observations, providing a balanced assessment of error magnitude.

Table 2 presents several critical insights. First, estimates incorporating historical data generally outperform model-based estimates, which rely more heavily on forward-looking expectations. Specifically, MIVV and STVV exhibit lower deviations across the full sample, with particularly strong performance in medium- and high-volatility periods. In contrast, under low-volatility conditions, EMVV demonstrates the smallest deviation, highlighting its relative accuracy in stable market environments.

Second, for most estimates (excluding RNVV), point-to-point deviations from VVIX increase with rising volatility. This trend is evident across Panels A, B, and C, where the errors in low-volatility regimes remain consistently lower than those observed in medium- and high-volatility periods. This pattern underscores the inherent difficulty in accurately estimating VV during heightened uncertainty.

[3] The value $\lambda = 0.94$ is adopted from Roh et al. (2021), and 0.94 is also considered as a benchmark value for market practitioners (Longerstaey and Zangari, 1996).

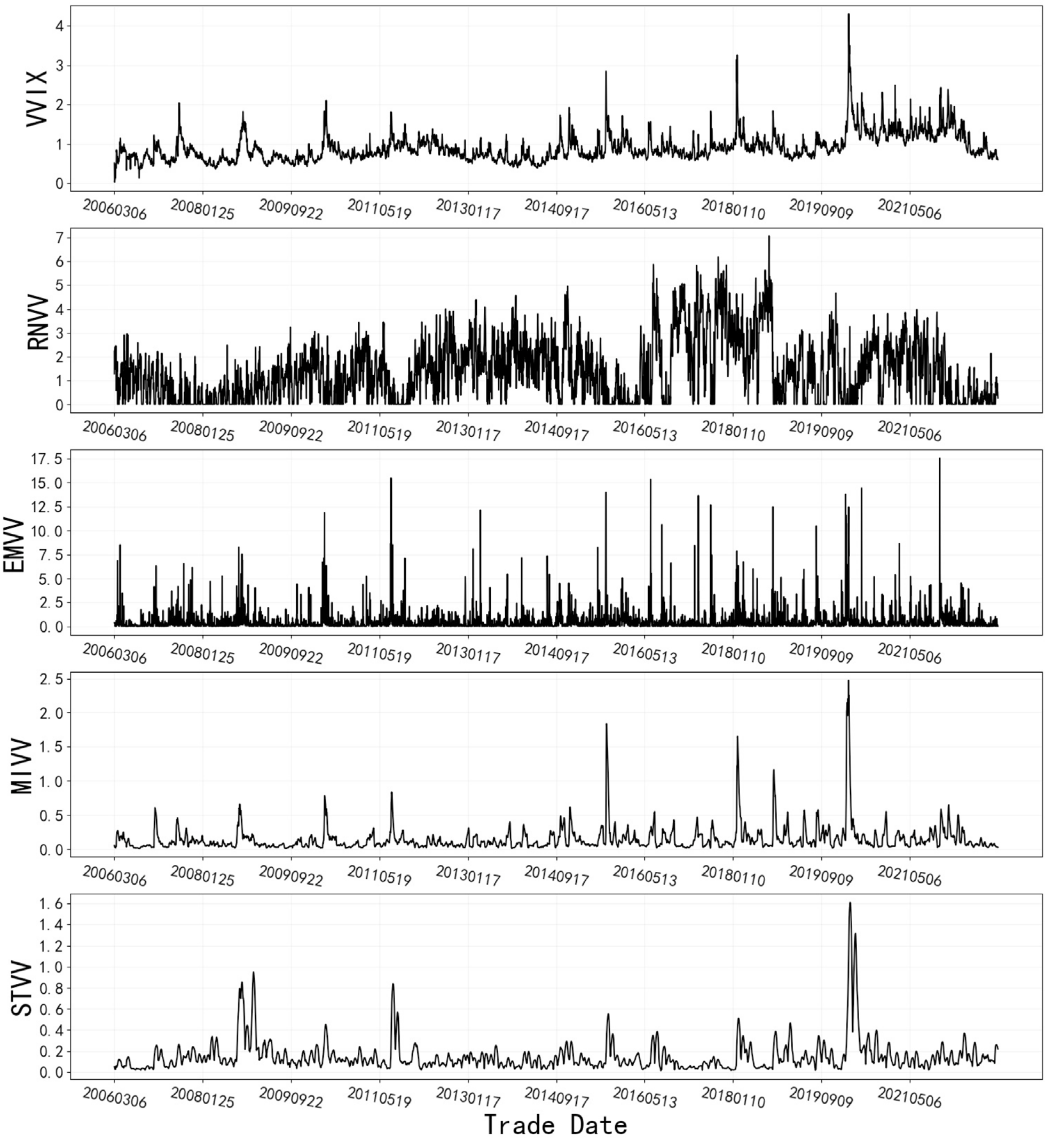

Fig. 1. Time series trends of volatility-of-volatility estimates. *Notes:* VVIX is normalized based on the value from CBOE, i.e., $VVIX = (VVIX_{CBOE}/100)^2$ to be comparable with other estimates.

Interpretation of RNVV results should be made with caution. Due to the volatility coefficients being constrained to be non-negative, a substantial proportion of values that would otherwise be negative are truncated to zero. Specifically, in the full sample, 20.43% of the RNVV estimates are zero, with 5.16%, 19.89%, and 36.28% zeros in the low-, medium-, and high-volatility subsamples, respectively. This uneven distribution across subsamples complicates comparisons of RNVV point deviations across different volatility regimes and raises concerns about its applicability in stressed market conditions.

3.2.2. Correlation analysis

The correlation analysis aims to evaluate the trend alignment. Two key methodologies are employed: Spearman correlation and copula-based analysis. Spearman correlation ranks data points, making it robust to outliers and applicable when variables exhibit non-linear dependencies. Copula, on the other hand, provides a deeper understanding of dependency structures, capturing non-linear and tail dependencies. We adopted three popular Archimedean copulas — Clayton, Gumbel, and Frank — since we expect asymmetric and nonlinear relationships, as well as potential differences in tail dependency (Genest and Rivest, 1993).[4] By visualizing the joint distributions of VV methods and VVIX, copulas show how closely these estimates align with VVIX under varying conditions.

Spearman

Table 3 reports the Spearman correlation coefficients for all estimates.[5] First, hybrid estimates (EMVV and MIVV) exhibit higher positive correlations with VVIX across most regimes, indicating stronger alignment with the trends in VVIX. Specifically, MIVV achieves the highest correlation with VVIX in all volatility regimes, underscoring its robustness in capturing VV. STVV also maintains moderate positive correlations across all subsamples.

[4] The Clayton copula emphasizes lower tail dependence, making it suitable for analyzing co-movements during periods of low volatility. The Gumbel copula captures upper tail dependence, focusing on relationships during high volatility. The Frank copula exhibits symmetric dependence with no specific emphasis on either tail.

[5] We cautiously interpret the Spearman correlation results for subsamples. Since the subsamples were divided based on VIX values, the timeline is disrupted, resulting in a loss of temporal continuity.

Table 2
Pointwise deviations across volatility regimes.

	MAE	MSE	MAPE	WMAPE
Panel A: Low volatility				
RNVV	1.658	4.083	2.473	2.307
EMVV	0.587	0.562	0.826	0.817
MIVV	0.613	0.405	0.852	0.853
STVV	0.627	0.423	0.867	0.872
Panel B: Medium volatility				
RNVV	0.853	1.091	0.984	0.939
EMVV	0.786	1.395	0.853	0.865
MIVV	0.768	0.648	0.850	0.846
STVV	0.779	0.676	0.847	0.858
Panel C: High volatility				
RNVV	0.809	0.943	0.791	0.769
EMVV	0.961	2.660	0.882	0.914
MIVV	0.873	0.902	0.848	0.830
STVV	0.823	0.858	0.773	0.783
Panel D: Full sample				
RNVV	1.106	2.04	1.417	1.24
EMVV	0.778	1.539	0.853	0.871
MIVV	0.751	0.652	0.850	0.842
STVV	0.743	0.652	0.829	0.832

Notes: Panels A, B, and C show deviations for three equally-split subsamples based on VIX representing low, middle, and high market volatility, respectively. Panel D presents the results of the full sample. The measures include the mean absolute error (MAE), $MAE = \frac{1}{n}\sum_{t=1}^{n}|VV_t - VVIX_t|$; mean squared error (MSE), $MSE = \frac{1}{n}\sum_{t=1}^{n}(VV_t - VVIX_t)^2$; mean absolute percentage error (MAPE), $MAPE = \frac{1}{n}\sum_{t=1}^{n}\frac{|VV_t - VVIX_t|}{VVIX_t}$; and weighted mean absolute percentage error (WMAPE), $WMAPE = \frac{\sum_{t=1}^{n}|VV_t - VVIX_t|}{\sum_{t=1}^{n}VVIX_t}$. Where VV_t is the respective VV estimate, $VVIX_t$ is the normalized VVIX, n is the number of observations.

Table 3
Spearman correlation coefficients (%).

	VVIX	RNVV	EMVV	MIVV	STVV
Panel A: Low volatility					
VVIX	100	20.21	14.49	25.44	0.92
RNVV	20.21	100	−10.16	0.23	−7.57
EMVV	14.49	−10.16	100	2.44	0.88
MIVV	25.44	0.23	2.44	100	27.48
STVV	0.92	−7.57	0.88	27.48	100
Panel B: Middle volatility					
VVIX	100	6.81	27.93	41.32	1.57
RNVV	6.81	100	−20.35	−16.71	−4.64
EMVV	27.93	−20.35	100	18.36	0.25
MIVV	41.32	−16.71	18.36	100	22.72
STVV	1.57	−4.64	0.25	22.72	100
Panel C: High volatility					
VVIX	100	3.23	35.94	56.91	21.51
RNVV	3.23	100	−10.43	−15.84	−1.46
EMVV	35.94	−10.43	100	34.17	8.79
MIVV	56.91	−15.84	34.17	100	31.19
STVV	21.51	−1.46	8.79	31.19	100
Panel D: Full sample					
VVIX	100	−11.64	31.61	44.39	21.23
RNVV	−11.64	100	−21.44	−17.49	−24.91
EMVV	31.61	−21.44	100	20.85	9.64
MIVV	44.39	−17.49	20.85	100	28.58
STVV	21.23	−24.91	9.64	28.58	100

Notes: The table reports Spearman correlation coefficients in percentage.

In contrast, RNVV shows negative correlations with VVIX in the full sample (−11.64%). This suggests that RNVV may be less reliable in replicating the directional behavior of VVIX. Additionally, the truncated zero values in RNVV further complicate its interpretability.

Copula

Table 4 presents the Archimedean copulas results. We also include a goodness-of-fit metric, AIC, to provide insights into which copula best captures the dependency structure. Several findings emerge from the table. First, RNVV exhibits weak or non-meaningful dependency with VVIX across both tails and symmetric relationships. This is evidenced by the invalid θ values for the Clayton ($\theta = -0.15$) and Gumbel ($\theta = 0.92$) copulas, as well as the negative θ value for the Frank copula ($\theta = -0.75$). These findings are consistent with the negative Kendall's $\tau = -0.08$, indicating a lack of meaningful alignment between RNVV and VVIX.

Table 4
Copulas results and goodness-of-fit.

Method	Clayton		Frank		Gumbel		Kendall τ
	θ	AIC	θ	AIC	θ	AIC	
RNVV	−0.15	−17 411	−0.75	−17 431	0.92	−17 298	−0.08
EMVV	0.54	−6260	1.99	−6061	1.27	−6356	0.21
MIVV	0.89	−6858	3.00	−6652	1.44	−6801	0.31
STVV	0.33	−7153	1.31	−7056	1.17	−7140	0.14

Notes: The table reports copula parameters θ, AIC (Akaike Information Criterion), and Kendall's τ values for three copula families (Clayton, Frank, and Gumbel) fitted to the dependency structure between VVIX and four VV estimates. θ represents the copula-specific parameter: for Clayton ($\theta > 0$), higher values indicate stronger lower-tail dependence; for Gumbel ($\theta \geq 1$), higher values indicate stronger upper-tail dependence; and for Frank ($\theta \neq 0$), it models symmetric dependencies. Lower AIC values indicate a better fit. Kendall's τ quantifies the overall monotonic dependency, directly comparable between different copulas families.

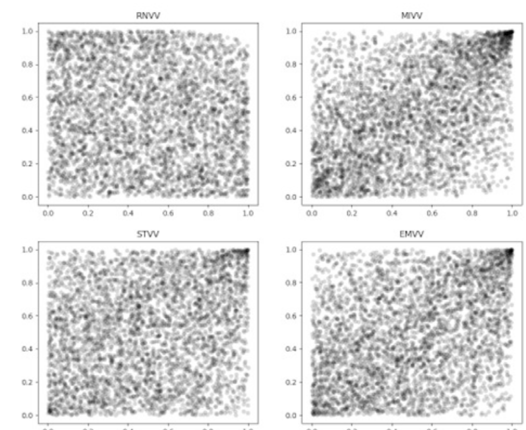

Fig. 2. Gumbel copula scatter plots.

Second, hybrid estimates demonstrate the strongest alignment with VVIX, with the highest Kendall's $\tau = 0.31$ for MIVV and $\tau = 0.21$ for EMVV. An interesting difference between MIVV and EMVV emerges from the AIC values: the lowest AIC for MIVV is achieved with the Clayton copula, indicating stronger lower-tail dependencies, while the lowest AIC for EMVV is achieved with the Gumbel copula, highlighting higher-tail co-movement. Despite the overall alignment being stronger for MIVV, EMVV may hold practical importance for studying VV during periods of high market volatility, while MIVV may be more suitable for analyzing VV in low-volatility markets. Third, STVV shows moderate alignment ($\tau = 0.14$) with similar AIC values for the Gumbel and Clayton copulas, indicating partial sensitivity to both upper-tail and lower-tail dependencies. Fig. 2 visualizes the Gumbel copula scatter plots, illustrating the dependency structures.[6] The visual patterns align with the findings in Table 4, confirming the stronger upper-tail dependencies captured by the Gumbel copula, particularly for MIVV and EMVV.

[6] Scatter plots for other copula families (e.g., Clayton, Frank) are available upon request.

4. Discussion and conclusion

VVIX represents a market-implied proxy for VIX variability rather than a direct statistical measure of realized VV. The VVIX calculation follows an approximation formula identical to that of VIX. Therefore, it inherits limitations from the VIX framework, particularly its applicability to other markets with sparse strikes. Theoretically, risk-neutral variance calculation requires a continuous range of strike prices from zero to infinity. In practice, the VIX employs a discrete subset of out-of-the-money S&P 500 options, assuming this finite set approximates the continuous integral. While this may hold for the S&P 500's dense strikes, it is less reliable for assets with sparser strikes, such as individual stocks. Heston et al. (2022) demonstrate that strike sparsity can distort the VIX's volatility estimates, while Heston and Todorov (2023) reveal that the standard VIX formula may yield negative implied volatility values when applied to other markets. Furthermore, Li (2018) notes that the VIX structure tends to overestimate implied variances, a minor issue for VIX but more pronounced for VVIX due to sparser strikes.

We conclude by attending to the practical challenges posed by sparse strikes for the four VV estimates examined. For RNVV, the variance of implied volatility is estimated through the curvature of the implied volatility surface, a process that relies on multiple, closely spaced strikes. Thus, sparse strikes can compromise the precision of this estimate. To address this issue, Carr and Wu (2020) employ interpolation to a fixed moneyness grid to limit the impact. Similarly, applying hybrid methods like EMVV and MIVV to individual stocks requires implied volatility as inputs, which could also be vulnerable to sparse strikes.[7] In contrast, STVV circumvents these challenges entirely by relying on realized volatility rather than option data. When implied volatility is employed, Heston et al. (2022) propose a promising practical solution that modifies the VIX methodology by incorporating a refined version of Simpson's rule. This adaptation improves convergence with sparse data, thereby enhancing the applicability of VV estimates in contexts where strike density is limited.

To summarize, this study compares four volatility-of-volatility estimates with VVIX using S&P500 data. Overall, hybrid estimates, especially MIVV, align best with VVIX.

Data availability

The authors do not have permission to share data.

References

Baltussen, G., Van Bekkum, S., Van Der Grient, B., 2018. Unknown unknowns: Uncertainty about risk and stock returns. J. Financ. Quant. Anal. 53 (4), 1615–1651.

Branger, N., Hülsbusch, H., Kraftschik, A., 2017. The volatility-of-volatility term structure. In: Paris December 2017 Finance Meeting EUROFIDAI-AFFI.

Carr, P., Wu, L., 2020. Option profit and loss attribution and pricing: A new framework. J. Financ. 75 (4), 2271–2316.

Cheng, I.-H., 2019. The VIX premium. Rev. Financ. Stud. 32 (1), 180–227.

Genest, C., Rivest, L.-P., 1993. Statistical inference procedures for bivariate Archimedean copulas. J. Amer. Statist. Assoc. 88 (423), 1034–1043.

Heston, S.L., 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. Rev. Financ. Stud. 6 (2), 327–343.

Heston, S.L., Jones, C.S., Khorram, M., Li, S., Mo, H., 2022. Seasonal momentum in option returns. Available At SSRN 4167231.

Heston, S.L., Todorov, K., 2023. Exploring the variance risk premium across assets. Available At SSRN 4373509.

Hollstein, F., Prokopczuk, M., 2017. How aggregate volatility-of-volatility affects stock returns. Rev. Asset Pricing Stud. 8 (2), 253–292.

Huang, D., Schlag, C., Shaliastovich, I., Thimme, J., 2019. Volatility-of-volatility risk. J. Financ. Quant. Anal. 54 (6), 2423–2452.

Kaeck, A., 2018. Variance-of-variance risk premium. Rev. Financ. 22 (4), 1549–1579.

Li, A., 2018. Improvements to the CBOE VIX calculation formula. Available At SSRN 2991829.

Longerstaey, J., Zangari, P., 1996. Risk metrics. (Accessed 14 March 2025).

Roh, T., Tourani-Rad, A., Xu, Y., Zhao, Y., 2021. Volatility-of-volatility risk in the crude oil market. J. Futur. Mark. 41 (2), 245–265.

Zang, X., Ni, J., Huang, J.-Z., Wu, L., 2017. Double-jump diffusion model for VIX: evidence from VVIX. Quant. Finance 17 (2), 227–240.

[7] OptionMetrics provides a widely used estimate of implied volatility for individual stocks, derived through Black–Scholes inversion. This calculation relies on interpolated values, such as at-the-money implied volatility, and is thus susceptible to sparse strikes, which may diminish interpolation accuracy.