Erasmus University Rotterdam
Erasmus School of Economics
Bachelor Thesis Econometrics and Operational Research

GARCH option pricing models: a multivariate VIX estimation approach

Bart Overduin (569799)

Supervisor:	Vladimirov, E
Second assessor:	van der Sluis, BP
Date final version:	1st July 2024

Abstract

This study investigates the impact of using multiple VIX indices with varying maturities on the performance of GARCH models for volatility modeling. We incorporate VIX indices with maturities of 9 days, 1 month, 3 months, 6 months, and 1 year, and estimate the CBOE VIX with a new multivariate estimation approach. The models evaluated are GARCH(1,1), EGARCH(1,1), CGARCH(1,1), and ACGARCH(1,1), and are estimated under both the locally risk-neutral valuation relationship (LRNVR) and modified LRNVR (mLRNVR) framework. The newly introduced multivariate VIX estimation approach generally does not yield better results than the traditional VIX estimation approach, except for minor improvements in the CGARCH and ACGARCH models when using both returns and VIX data. The EGARCH model under the mLRNVR, using the traditional estimation approach with returns and VIX data, still provides the best fit for the CBOE VIX. Additionally, a comprehensive residual analysis of the EGARCH model indicates that normality is firmly rejected for the error terms of both estimation methods.

1 Introduction

Modeling volatility in financial markets is a well-researched topic since it allows investors and institutions to manage their risk more effectively. For example, forecasting this index accurately allows for more precise option prices. A well-known measurement for volatility is the Chicago Board Options Exchange (CBOE) VIX. This index measures the expected volatility over the next 30 days and plays a crucial role in the pricing of VIX derivatives and options. In addition to the VIX, the CBOE provides volatility indices with varying maturities.

Previous studies on GARCH option pricing models have primarily focused on the CBOE VIX, which gives insight into the expected volatility for the upcoming month. However, considering that incorporating a volatility measure enhances model estimation, it is possible that incorporating information from multiple VIX indices could improve the estimation of the CBOE VIX.

In this paper, we aim to address the question: How does the use of multiple VIX indices affect the performance of GARCH models compared to using a single VIX index? We answer this question by considering a multivariate approach to estimate the CBOE VIX with different VIX maturities for the GARCH models. More specifically, we consider the following maturities for the VIX: 9 days, 1 month, 3 months, 6 months and 1 year. We evaluate the performance of the indices with four GARCH models: the GARCH(1,1)-model, the EGARCH(1,1)-model, the CGARCH(1,1)-model, and the ACGARCH(1,1)-model. We estimate the models under the LRNVR and the mLRNVR.

Generally, the multivariate VIX estimation method does not yield improvements in fitting the CBOE VIX. The error terms are only slightly smaller for the CGARCH and ACGARCH models under the LRNVR when estimated with both returns and VIX data. Still, the implied VIX of these models struggles to capture the statistical properties of the CBOE VIX. Additionally, a comprehensive residual analysis reveals that both the single VIX estimation and the multivariate VIX estimation methods produce error terms for which normality is firmly rejected. This highlights the inherent challenges of accurately modeling the CBOE VIX.

The remainder of this proposal is as follows: Section 2 gives an overview of previous research regarding GARCH option pricing models. Section 3 shows the data and descriptive statistics. We describe our methods in Section 4. The results are discussed in Section 5, and we conclude our findings in Section 6.

2 Literature Review

The compensation for the volatility risk premium has been extensively studied since the end of the last century. Volatility in financial markets is time-varying, as shown by Schwert (1988). GARCH-type models, initially proposed by Engle and Bollerslev (1986), are well-known for accommodating time-varying volatility.

Option pricing models aim to model the volatility risk premium. A noteworthy framework is developed by Black and Scholes (1973), introducing closed-form solutions for option valuation. However, one of the assumptions of this model is constant volatility and neglects that volatility is time-varying. Duan (1995) developed a completely new framework and introduced GARCH option pricing models, which allow for time-varying volatility. He argued that, under certain assumptions, options can be priced under a risk-neutral relationship (LRNVR) when the price of an asset follows a GARCH process.

Since then, further research has extended the literature of GARCH option pricing models under the LRNVR. For example, Heston and Nandi (2000) derived a closed-form option valuation formula for a spot asset which follows a GARCH process. Additionally, various extensions of the standard GARCH model have also been brought under the LRNVR, see e.g. Hao and Zhang (2013). Many studies have used the LRNVR to price VIX derivatives and options, see e.g. Wang et al. (2017) and Tong and Huang (2021).

Hao and Zhang (2013) incorporated both returns and VIX data for model estimation. To accommodate both time series, they allowed for a discrepancy between the CBOE VIX and the implied VIX from the GARCH models. Kanniainen et al. (2014) extended this idea by including an autoregressive term in their VIX estimation approach. However, Zhang and Zhang (2020) found that the extra autoregressive term is insignificant for most GARCH models. For this reason, we adopt the simpler approach of Hao and Zhang (2013) for our analysis.

Furthermore, Hao and Zhang (2013) argued that the GARCH option pricing models under the LRNVR fail to incorporate the volatility risk premium, since these models are only able to capture the equity risk premium. They found that the implied VIX of the GARCH option pricing models is underestimated when only returns are used, and that the estimated parameters become distorted if the models are jointly estimated with the returns and the VIX.

To address these limitations, Zhang and Zhang (2020) modified Duan's LRNVR, introducing the modified LRNVR (mLRNVR). They found that the mLRNVR can capture the volatility risk premium and recommend using the mLRNVR in combination with the EGARCH(1,1) model with return and VIX data. Another way to capture the volatility risk premium is by incorporating jump risk in the GARCH models.Christoffersen et al. (2012) introduced a new class of models, which allow for dynamic volatility and jump intensity.

Parallel to these studies, Tong (2024) introduces a generalized LRNVR, where the specification of a pricing kernel is not necessary. Variance dependent pricing kernels are another commonly used method to bring GARCH models under the risk-neutral measure. The concept of variance dependent pricing kernels for GARCH models has been introduced by Christoffersen et al. (2013). Byun et al. (2015) have expanded the literature on this topic by introducing a model with a variance dependent pricing kernel and the inclusion of jump risk.

The primary motivation for our research is based on Kanniainen et al. (2014), who found an improvement when the CBOE VIX is considered for GARCH option pricing models. However, this study only considered the VIX index with a 30-day maturity and does not evaluate VIX indices with different maturities. Furthermore, the LRNVR is used for the GARCH models, which do not capture the volatility risk premium, only the equity risk premium. We extend the literature by incorporating information from multiple VIX indices with different maturities under the LRNVR and mLRNVR.

Additionally, we bring a new GARCH model under the LRNVR: the ACGARCH(1,1)-model. The ACGARCH (Asymmetric Component GARCH) model extends the CGARCH (Component GARCH) model by adding an asymmetry component. The CGARCH model decomposes volatility into long-run and short-run components, capturing both the long-term and transient short-term effects of shocks on volatility. The ACGARCH model has an additional asymmetry term, allowing for different impacts of positive and negative shocks on volatility.

3 Data

The dataset for this study consists of the S&P 500 closing price, the risk-free rate, and various CBOE VIX measures with different maturities. Our sample starts from January 4, 2011, to 29 December, 2023, resulting in 3269 observations. The starting point of the dataset is determined by the introduction of the VIX with 9-day maturity, which the Chicago Board Options Exchange has provided since January 2011.

The data are obtained from three different sources. The closing price of the S&P 500 is obtained from the CRSP. The risk-free rate is obtained from the Federal Reserve Economic Data. Similar to Hao and Zhang (2013), we use the daily 3-month Treasury Bills secondary market rate as the risk-free rate. The time series of the CBOE VIX indices are obtained directly from the CBOE.

The VIX indices are popular measures of expected future volatility, calculated as a weighted average of options of the S&P 500. More information about the calculation of the VIX indices can be found on the website of the CBOE [1]. The CBOE provides VIX indices with varying maturities. The maturities of the VIX indices that are considered in this paper consist of 9 days, 1 month, 3 months, 6 months, and 1 year. Henceforth, we denote these variables as VIX9D, VIX1M, VIX3M, VIX6M, and VIX1Y. Summary statistics for every variable can be found in Table 1. These statistics provide an overview of the data distribution, including mean, variance, skewness, kurtosis, and autocorrelations. Additionally, Figure 1 displays the S&P 500 closing price over time in panel 1a, and the VIX indices over the sample period in panel 1b.

To address missing observations of the risk-free rate, we opt for a simple approach where we fill the missing observation with the value of the previous day. This approach is based on the idea that the risk-free rate is very stable over time and the daily fluctuations are minimal.

[1] https://www.cboe.com/tradable_products/vix/

Table 1: Summary statistics of the S&P 500, risk-free rate and VIX indices

Data	Mean	Variance	Skewness	Kurtosis	Max	Min	AR1	AR10	AR30
Return S&P500 (%)	0.05	1.21	−0.51	16.09	9.38	−11.98	−0.1276	−0.0286	0.0008
Risk-free rate	1.00	2.20	1.72	5.00	5.36	−0.05	0.9986	0.9847	0.9492
VIX9D	17.69	72.78	3.22	22.57	106.66	7.10	0.9398	0.6859	0.4003
VIX1M	18.30	51.06	2.45	14.24	82.69	9.14	0.9671	0.7926	0.5355
VIX3M	20.18	40.83	1.90	9.89	72.98	11.85	0.9815	0.8580	0.6661
VIX6M	21.62	32.98	1.30	5.42	61.11	13.75	0.9873	0.9009	0.7605
VIX1Y	22.58	22.31	0.79	2.84	41.93	15.56	0.9917	0.9383	0.8375

This Table shows summary statistics for S&P 500, risk-free rate, and VIX indices from January 4, 2011, to December 29, 2023. The first column denotes the name of the variables, columns 2 to 7 show basic statistical properties, and the last three columns showcase the autocorrelation coefficients for lag 1, 10, and 30, respectively.

Figure 1: The S&P 500 and VIX indices from January 2011 to December 2023

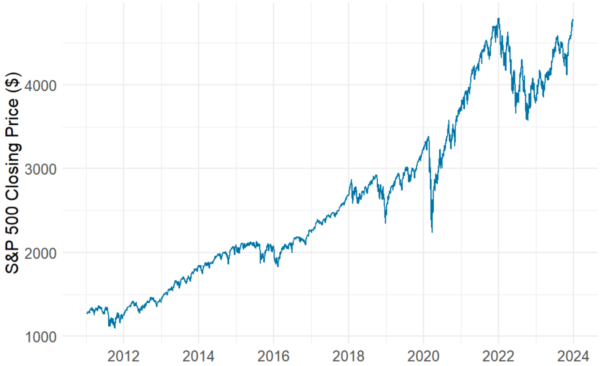

(a) The closing price of the S&P 500

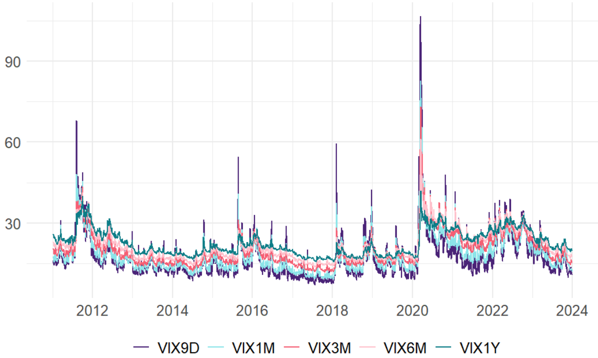

(b) VIX indices with different maturities

4 Methodology

In this section, we define the GARCH models and the estimation methods. We commence by explaining the return dynamics for the GARCH(1,1)-model, which will be the benchmark model. After that, we specify the formulas for the other GARCH models. We then derive formulas for the implied VIX of the GARCH models, and finish with the estimation methods.

4.1 Return dynamics of the GARCH(1,1)-model

4.1.1 Physical measure

We follow Duan (1995) and assume that the returns of the asset X follow a log normal distribution under the physical pricing measure $\mathbb{P}$ as

$$\ln \frac{X_t}{X_{t-1}} = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \epsilon_t, \tag{1}$$

where X_t is the price of asset X at time t, r is the risk-free rate, λ is the asset risk premium, and ϵ_t follows a GARCH(p, q) process with mean zero and h_t as the conditional variance

$$\epsilon_t \mid \mathcal{I}_{t-1} \sim N\left(0, h_t\right) \quad \text{under measure } \mathbb{P}$$

$$h_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j h_{t-j}, \tag{2}$$

where $\mathcal{I}_t$ is the information set containing all the information up to time t; and with the restrictions $\alpha_0 \geq 0$, $\alpha_i \geq 0$ for $i = 1, 2, \ldots, q$ and $\beta_j \geq 0$ for $j = 1, 2, \ldots, p$. In this paper, we only consider the GARCH(1,1)-model, and equation 2 becomes

$$h_t = \alpha_o + \alpha_1 \epsilon_{t-1}^2 + \beta_1 h_{t-1}. \tag{3}$$

4.1.2 Risk-neutrality under the LRNVR

In order to account for the heteroskedasticity of the returns, Duan (1995) proposed the locally risk-neutral valuation relationship (LRNVR). The risk-neutral pricing measure $\mathbb{Q}$ of the asset returns is denoted as

$$\ln \frac{X_t}{X_{t-1}} = r - \frac{1}{2} h_t + \xi_t, \tag{4}$$

$$\xi_t \mid \mathcal{I}_{t-1} \sim N(0, h_t) \quad \text{under measure } \mathbb{Q}$$

$$h_t = \alpha_o + \alpha_1 (\xi_{t-1} - \lambda h_{t-1}) + \beta_1 h_{t-1}. \tag{5}$$

4.1.3 Risk-neutrality under the mLRNVR

Zhang and Zhang (2020) modifies the LRNVR, intending to capture the volatility risk premium. The asset returns satisfy the mLRNVR if it has the following form:

$$\ln \frac{X_t}{X_{t-1}} = r - \frac{1}{2} h_t + \xi_t, \tag{6}$$

$$h_{t-1} = \alpha_0 + \alpha_1 \left(\xi_{t-1} - \lambda_1 \sqrt{h_{t-1}} \right)^2 + \left(\beta_1 - \sqrt{2} \alpha_1 \lambda_2 \right) h_{t-1}, \tag{7}$$

where λ_1 is the lambda of the LRNVR and λ_2 is the volatility risk premium of the asset. We assume a negative volatility risk premium in 7, consistent with empirical findings of Bollerslev et al. (2009) and Carr and Wu (2009). The main difference between the mLRNVR and the LRNVR is the persistence parameter, which is now $\beta_1 - \sqrt{2}\alpha_1\lambda_2$ in equation 7, whereas the persistence parameter in equation 5 is β_1. The theoretical justification for the adjustment can be found in Zhang and Zhang (2020). The introduction of the second lambda does not require additional data. The same dataset used for estimating the GARCH models under the LRNVR is also applied to estimate the GARCH models under the mLNRVR.

4.1.4 Other models

In addition to the GARCH(1,1)-model, we also consider other models. More specifically, we consider the exponential GARCH(1,1) (EGARCH) of Nelson (1991), the component GARCH(1,1) (CGARCH) and asymmetric component GARCH(1,1) (ACGARCH) of Lee and Engle (1993). The formulas for these models are shown below:

EGARCH$(1,1)$:

Physical measure :

$$\ln h_t = \alpha_0 + \beta_1 \ln h_{t-1} + \alpha_1 z_{t-1} + \kappa \left(|z_{t-1}| - \sqrt{2/\pi} \right), \quad z_t = \epsilon_t / \sqrt{h_t}, \tag{8}$$

LRNVR :

$$\ln h_t = \alpha_0 + \beta_1 \ln h_{t-1} + \alpha_1 \left(u_{t-1} - \lambda \right) + \kappa \left(|u_{t-1} - \lambda| - \sqrt{2/\pi} \right), \tag{9}$$

$$u_t = \xi_t / \sqrt{h_t},$$

where ξ_t is a m.d.s. with conditional variance h_t under the risk-neutral measure.

CGARCH$(1,1)$:

Physical measure :

$$h_t - q_t = \alpha_1 \left(\epsilon_{t-1}^2 - q_{t-1} \right) + \beta_1 \left(h_{t-1} - q_{t-1} \right), \tag{10}$$

$$q_t = \alpha_0 + \rho q_{t-1} + \phi \left(\epsilon_{t-1}^2 - h_{t-1} \right),$$

LRNVR :

$$h_t - q_t = \alpha_1 \left[\left(\xi_{t-1} - \lambda \sqrt{h_{t-1}} \right)^2 - q_{t-1} \right] + \beta_1 \left(h_{t-1} - q_{t-1} \right), \tag{11}$$

$$q_t = \alpha_0 + \rho q_{t-1} + \phi \left[\left(\xi_{t-1} - \lambda \sqrt{h_{t-1}} \right)^2 - h_{t-1} \right],$$

ACGARCH$(1,1)$:

Physical measure :

$$h_t - q_t = \left(\alpha_1 + \theta \mathbf{1}_{\{\epsilon_{t-1} \leq 0\}} \right) \left(\epsilon_{t-1}^2 - q_{t-1} \right) + \beta_1 \left(h_{t-1} - q_{t-1} \right), \tag{12}$$

$$q_t = \alpha_0 + \rho q_{t-1} + \phi \left(\epsilon_{t-1}^2 - h_{t-1} \right),$$

LRNVR :

$$h_t - q_t = \left(\alpha_1 + \theta \mathbf{1}_{\{\epsilon_{t-1} \leq 0\}} \right) \left[\left(\xi_{t-1} - \lambda \sqrt{h_{t-1}} \right)^2 - q_{t-1} \right] + \beta_1 \left(h_{t-1} - q_{t-1} \right), \tag{13}$$

$$q_t = \alpha_0 + \rho q_{t-1} + \phi \left[\left(\xi_{t-1} - \lambda \sqrt{h_{t-1}} \right)^2 - h_{t-1} \right].$$

We evaluate the GARCH and the EGARCH model also under the mLRNVR, as Zhang and Zhang (2020) derived formulas for these models under the mLRNVR. The formulas for the GARCH and EGARCH models can be found by substituting $\beta_1 = \beta_1 - \sqrt{2}\alpha_1\lambda_2$ in the corresponding formulas.

4.2 GARCH implied VIX

The CBOE VIX measures the market volatility and is calculated with SPX call and put options prices. The GARCH implied VIX can be calculated as follows:

$$\left(\frac{\text{VIX}_t^{\text{Imp}}\left(\tau_0\right)}{100}\right)^2 = \frac{1}{n}\sum_{k=1}^{n} E_t^Q\left(h_{t+\frac{\tau_0 k}{n}}\right) \tag{14}$$

where $\text{VIX}_t^{\text{Imp}}\left(\tau_0\right)$ is the model implied VIX, and τ_0 is equal to the amount of trading days. Note that τ_0 changes for the different VIX maturities. Since the CBOE VIX is denoted in calendar days, we assume $\tau_0 = 7, 21, 63, 121, 252$ days for VIX9D, VIX1M, VIX3M, VIX6 and VIX1Y, respectively. These transformations are commonly used in literature. Since we only use closing time data for the underlying asset, we assume that τ_0 is equal to n. We calculate the daily variance of the S&P 500 with

$$V_t = \frac{1}{n}\sum_{k=1}^{n} E_t^Q\left(h_{t+k}\right), \tag{15}$$

where $V_t = \frac{1}{252}\left(\frac{\text{VIX}_t^{\text{Imp}}(\tau_0)}{100}\right)^2$ is a proxy for the daily variance.

Following Hao and Zhang (2013), we consider SR-SARV(p) models (Meddahi and Renault, 2004) to calculate the implied VIX as a conditional mean of future variance.

Definition 1: *(Discrete time SR-SARV(p) model (Meddahi and Renault, 2004)). A stationary square-integrable process $\{\epsilon_t,\ t \in \mathbb{Z}\}$ is called a SR-SARV(p) process with respect to a filtration $J_t,\ t \in \mathbb{Z}$, if:*

(i) ϵ_t is a martingale difference sequence with respect to J_{t-1}, that is $\mathbb{E}[\epsilon_t \mid J_{t-1}] = 0$,
(ii) the conditional variance process f_t of ϵ_{t+1} given J_t is a marginalization of a stationary J_t-adapted VAR(1) of dimension p:

$$f_t \equiv \text{Var}[\epsilon_{t+1} \mid J_t] = e' F_t, \tag{16}$$

$$F_t = \Omega + \Gamma F_{t-1} + V_t, \quad \text{with} \quad \mathbb{E}[V_t \mid J_{t-1}] = 0, \tag{17}$$

where $e \in \mathbb{R}^p$, $\Omega \in \mathbb{R}^p$ and the eigenvalues of Γ have modulus less than one.

Hao and Zhang (2013) show that the GARCH(1,1), CGARCH(1,1) and EGARCH(1,1) models are SR-SARV(p) processes. Now we show that the ACGARCH(1,1) model is also an SR-SARV(p) process. It is important to note that we assume that $\epsilon_t/\sqrt{h_t}$ and $\xi_t/\sqrt{h_t}$ are i.i.d. under both the physical and risk-neutral measure for the proof to hold.

Proposition 1: *Let $\{\xi_t, t \in Z\}$ be a m.d.s. with the conditional variance $h_t \equiv \text{Var}\left[\xi_t \mid \xi_\tau, \tau \leq t-1\right]$ under the LRNVR. If $u_t = \xi_t/\sqrt{h_t}$ is i.i.d., the ACGARCH(1,1) model (13) is a SR-SARV(2) process.*

Proof. See Appendix A.1.

We calculate the implied VIX based on the methods derived in Hao and Zhang (2013). The corresponding equations can be found in Appendix A.2.

4.3 Estimation

We estimate the GARCH models using maximum likelihood estimation with three distinct approaches: considering only returns data, only VIX data, and a joint estimation using both returns and VIX data. The first approach, using returns data, is estimated under the physical measure $\mathbb{P}$. The log-likelihood function ($\ln L_R$) for this approach is

$$\ln L_R = -\frac{T}{2}\ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\left\{\ln(h_t) + \left[\ln(X_t/X_{t-1}) - r - \lambda\sqrt{h_t} + \frac{1}{2}h_t\right]^2 / h_t\right\}, \tag{18}$$

where h_t is updated by the respective GARCH process and T is the amount of observations.

For the second approach, we consider an estimation using the VIX. The approach of Hao and Zhang (2013) is used as the benchmark approach and is denoted in this paper as the single VIX estimation method. This approach allows a difference between the actual CBOE VIX index, and the implied VIX by specifying

$$\mathrm{VIX}^{Mkt} = \mathrm{VIX}^{Imp} + \mu, \quad \mu \; i.i.d. \; N(0, s^2), \tag{19}$$

where s^2 is calculated as the sample variance of $\hat{s}^2 = \mathrm{VIX}^{Mkt} - \mathrm{VIX}^{Imp}$. The reason for including μ is that we also consider a joint estimation with returns and VIX, but the innovation z_t influences both the returns and the VIX. To accommodate this, we allow a difference between these variables. The log-likelihood function based on the VIX ($\ln L_V$) is formulated as

$$\ln L_V = -\frac{T}{2}\ln\left(2\pi\hat{s}^2\right) - \frac{1}{2\hat{s}^2}\sum_{t=1}^{T}\left(\mathrm{VIX}_t^{Mkt} - \mathrm{VIX}_t^{Imp}\right)^2. \tag{20}$$

The log-likelihood function of the joint estimation is a sum of $\ln L_R$ and $\ln L_V$, and is specified as

$$\ln L_T = \ln L_R + \ln L_V. \tag{21}$$

4.3.1 The multivariate VIX estimation approach

The new approach to VIX estimation, which extends the methodology of Hao and Zhang (2013), considers multiple VIX indices simultaneously. Similar to the approach of Hao and Zhang (2013), we allow for a difference between the CBOE VIX and implied VIX by defining

$$\mathrm{VIX}_i^{Mkt} = \mathrm{VIX}_i^{Imp} + \mu_i, \quad i \in \text{9D, 1M, 3M, 6M, 1Y}, \tag{22}$$

which can be written in vector form as

$$\mathrm{VIX}^{Mkt} = \mathrm{VIX}^{Imp} + \mu, \quad \mu \text{ i.i.d.} N(0, \Sigma), \tag{23}$$

where $\mu = (\mu_{9\mathrm{D}}, \mu_{1\mathrm{M}}, \mu_{3\mathrm{M}}, \mu_{6\mathrm{M}}, \mu_{1\mathrm{Y}})^T$ and Σ is calculated as the sample variance

$$\hat{\Sigma} = (\mathrm{VIX}^{Mkt} - \mathrm{VIX}^{Imp})(\mathrm{VIX}^{Mkt} - \mathrm{VIX}^{Imp})', \tag{24}$$

where VIX^{Mkt} and VIX^{Imp} are now the vectors from equation 23. This results in a multivariate log-likelihood function, which can be specified as

$$\ln L_V = -\frac{pT}{2}\log(2\pi) - \frac{T}{2}\log(|\hat{\Sigma}|) - \frac{1}{2}\sum_{i=1}^{T}\left(\text{VIX}_t^{Mkt} - \text{VIX}_t^{Imp}\right)' \hat{\Sigma}^{-1}\left(\text{VIX}_t^{Mkt} - \text{VIX}_t^{Imp}\right) \quad (25)$$

where p is equal to the length of x_i which is 5. However, we only consider VIX9D, VIX1M, and VIX3M for the CGARCH and ACGARCH models to reduce computation time, which results in $p = 3$.

We maximize these likelihoods subject to the stationary constraints under the risk-neutral measure $\mathbb{Q}$ for each model, since these constraints are stricter than the constraints under the physical measure $\mathbb{P}$. The stationary constraints can be found in Appendix A.3.

5 Results

In this section, we analyze the parameter estimates and properties of the implied VIX for the GARCH models. First, we present the results of the replication, followed by the outcomes under the LNRVR framework. Subsequently, we analyze the estimation results under the mLRNVR framework. Finally, we conduct a residual analysis of the EGARCH model, which demonstrates the best performance under both the LRNVR and mLRNVR frameworks.

5.1 Replication

The replication results, based on data from January 2, 1990, to August 10, 2009, are largely consistent with those reported by Hao and Zhang (2013). Only the CGARCH model shows slight differences in parameters. Table 2 presents the parameter estimates, while Table 3 shows the implied VIX characteristics. The parameter estimates are nearly identical for GARCH and EGARCH, indicating a successful replication of those models. Additionally, the log-likelihood values for each model are similar to those in the original paper, supporting the robustness of the replication. The slight differences in parameter estimates and log-likelihoods can be explained by different stopping criteria in the maximum likelihood estimation and differences in the dataset. For example, there are some missing observations in the CBOE VIX and the risk-free rate time series, but the treatment of missing observations is not discussed in Hao and Zhang (2013).

The CGARCH model shows the most notable differences in parameter estimates compared to the original study. When using returns data or both datasets, the parameter estimates remain very similar. However, when considering the VIX data, the parameters β_1, and ϕ are significantly larger. Meanwhile, the persistence parameter for the long-run volatility, ρ, is significantly lower. Despite these differences, the log-likelihood values are consistent with the log-likelihood values reported by Hao and Zhang (2013), indicating that the overall model performance is not substantially affected.

Table 2: The maximum likelihood estimates of the replication

Model & Data	α_0	α_1	β_1	κ	ρ	ϕ	λ	Log-likelihoods Total	Returns	VIX
GARCH										
Returns	7.1075e-07	0.0637	0.9310	-	-	-	0.0531	36653	**16073**	20580
	(1.7900e-07)	(0.0070)	(0.0075)	-	-	-	(0.0144)			
VIX	1.7146e-06	0.0366	0.9390	-	-	-	0.7890	38303	14456	**23846**
	(3.9149e-08)	(0.0009)	(0.0012)	-	-	-	(0.0207)			
Both	1.6919e-06	0.0471	0.9499	-	-	-	0.2080	**39453**	15844	23609
	(4.3923e-08)	(0.0011)	(0.0012)	-	-	-	(0.0118)			
EGARCH										
Returns	-0.1356	-0.0930	0.9851	0.0167	-	-	0.1119	36700	**16145**	20525
	(0.0181)	(0.0079)	(0.0019)	(0.0101)	-	-	(0.0143)			
VIX	-0.0788	-0.0614	0.9894	0.0904	-	-	-0.0299	40437	15993	**24444**
	(0.0017)	(0.0015)	(0.0002)	(0.0021)	-	-	(0.0006)			
Both	-0.0838	-0.0611	0.9891	0.0111	-	-	0.0954	**40451**	16014	24436
	(0.0019)	(0.0015)	(0.0002)	(0.0022)	-	-	(0.0037)			
CGARCH										
Returns	2.0862e-07	0.0482	0.9129	-	0.9984	0.0241	0.0537	36456	**16080**	20376
	(1.2261e-07)	(0.0095)	(0.0156)	-	(0.0012)	(0.0081)	(0.0135)			
VIX	2.9439e-06	0.0040	0.9878	-	0.9380	0.0507	0.8752	38433	14179	**24254**
	(8.0759e-08)	(0.0003)	(0.0005)	-	(0.0029)	(0.0012)	(0.0246)			
Both	1.1028e-06	0.0636	0.8651	-	0.9971	0.0388	0.2653	**39699**	15872	23826
	(3.0740e-07)	(0.0148)	(0.0354)	-	(0.0028)	(0.0015)	(0.0151)			

This table displays the maximum likelihood estimates using returns, VIX or both, replicating the results of the GARCH, EGARCH, and CGARCH models of Hao and Zhang (2013). The standard errors are given in parentheses and are calculated with the inverse of the Hessian matrix. The bold values of the log-likelihood denote that this log-likelihood has been maximized.

The implied VIX characteristics for the GARCH and EGARCH are presented in Table 3. Overall, the results correspond closely to the results in Hao and Zhang (2013), with nearly identical error terms and autocorrelation values. Additionally, the variance, skewness, and kurtosis are very similar, indicating that the implied VIX properties are consistent with Hao and Zhang (2013). The p-values of the replication are lower than those reported in the original paper, which can be explained by the assumption of equal variance between the CBOE VIX and implied VIX in Hao and Zhang (2013). Important to note is that the statistical properties of the actual CBOE VIX also differ slightly from the original paper, which indicates differences in the original dataset.

Table 3: Implied VIX characteristics of the replication

Model & Data	ME	Std.Err.	MAE	MSE	RMSE	P-value	Corr.Coef.	AR1	AR10	AR30	Variance	Skewness	Kurtosis
GARCH													
Returns	3.60	3.32	4.01	23.98	4.90	0.0000	0.93	0.9943	0.9350	0.7740	71.87	3.08	16.92
VIX	0.12	3.08	2.36	9.50	3.08	0.3597	0.93	0.9962	0.9551	0.8223	65.68	3.25	17.54
Both	0.26	3.22	2.39	10.46	3.23	0.0425	0.92	0.9967	0.9553	0.8158	66.89	3.26	17.85
EGARCH													
Returns	3.59	3.14	3.73	22.73	4.77	0.0000	0.94	0.9887	0.9047	0.7424	47.02	2.17	10.83
VIX	0.00	2.73	2.10	7.46	2.73	0.9890	0.95	0.9953	0.9506	0.8282	63.28	2.15	10.46
Both	0.09	2.73	2.10	7.48	2.74	0.4601	0.95	0.9949	0.9473	0.8199	64.17	2.17	10.58
CGARCH													
Returns	3.67	3.12	3.93	23.19	4.82	0.0000	0.93	0.9941	0.9438	0.8289	64.71	2.63	13.04
VIX	0.12	2.84	2.19	8.06	2.84	0.3361	0.94	0.9927	0.9361	0.8049	63.50	2.98	15.65
Both	0.22	3.09	2.32	9.58	3.09	0.0779	0.93	0.9917	0.9372	0.8232	65.94	3.01	15.53
CBOE VIX								0.9844	0.9161	0.7849	70.76	2.06	10.24

This table shows statistical properties and differences between the CBOE VIX and implied VIX of the GARCH models. The error is calculated as the difference between the CBOE VIX and the implied VIX. The ME displays the daily average error. The standard error (Std.Err.) displays the standard deviation of the error term. The MAE displays the daily mean absolute error. The MSE displays the daily mean squared error, and the RMSE displays the daily root mean squared error. The P-value corresponds to the P-value for the null hypothesis that the average of the CBOE VIX and implied VIX are equal, and is calculated with a t-test. The correlation coefficient (Corr.Coef.) displays the correlation between the CBOE VIX and the implied VIX. The AR1, AR10, and AR30 display the autocorrelation coefficients for lag 1, 10, and 30. The last columns display additional statistical properties of the CBOE VIX index and implied VIX.

5.2 Estimation with different VIX measures under the LRNVR

This subsection analyses the differences in performance between the single VIX estimation method and the multivariate VIX estimation method, during the period of using data from January 4, 2011, to December 29, 2023. Notably, the lambda parameter, which captures the equity risk premium, changes substantially depending on the estimation approach. The parameter estimates are presented in Table 4.

Table 4: The maximum likelihood estimates under the LRNVR

									Log-likelihoods		
Model & Data	α_0	α_1	β_1	κ	ρ	ϕ	θ	λ	Total	Returns	VIX
Panel A: Single VIX estimation											
Returns	3.9412e-06	0.1872	0.7815	-	-	-	-	0.1109	25594	**10879**	14715
	(5.2539e-07)	(0.0157)	(0.0167)	-	-	-	-	(0.0150)			
VIX	2.5771e-06	0.0439	0.9026	-	-	-	-	0.9999	24856	9270	**15586**
	(8.9203e-08)	(0.0014)	(0.0067)	-	-	-	-	(0.0588)			
Both	3.6407e-06	0.0867	0.8935	-	-	-	-	0.1990	**26166**	10766	15400
	(1.2260e-07)	(0.0031)	(0.0038)	-	-	-	-	(0.0147)			
EGARCH											
Returns	-0.3000	-0.1398	0.9677	0.1949	-	-	-	0.0531	26045	**10932**	15113
	(0.3280)	(0.0494)	(0.0314)	(0.0196)	-	-	-	(0.2572)			
VIX	-0.1432	-0.0953	0.9814	0.1226	-	-	-	-0.1293	26733	10787	**15946**
	(0.0053)	(0.0029)	(0.0006)	(0.0053)	-	-	-	(0.0026)			
Both	-0.1722	-0.0890	0.9794	0.1393	-	-	-	0.0067	**26797**	10859	15938
	(0.0047)	(0.0031)	(0.0005)	(0.0047)	-	-	-	(0.0013)			
CGARCH											
Returns	5.6780e-07	0.1586	0.7716	-	0.9941	0.0256	-	0.1102	25690	**10887**	14803
	(3.7653e-07)	(0.0164)	(0.0220)	-	(0.0031)	(0.0168)	-	(0.0173)			
VIX	2.5576e-06	0.0021	0.9920	-	0.8899	0.0724	-	0.9999	26039	9375	**16664**
	(2.6362e-07)	(0.0002)	(0.0006)	-	(0.0039)	(0.0034)	-	(0.0372)			
Both	2.7196e-07	0.1298	0.8024	-	0.9978	0.0150	-	0.3368	**26983**	10753	16230
	(3.8079e-08)	(0.0039)	(0.0054)	-	(0.0002)	(0.0012)	-	(0.0171)			
ACGARCH											
Returns	6.5208e-06	0.0239	0.9714	-	0.9371	0.1675	5.8148e-09	0.1135	26011	**10891**	15121
	(9.7322e-06)	(0.0821)	(0.0646)	-	(0.0157)	(0.0432)	(0.0389)	(0.0353)			
Both	2.6160e-07	0.1255	0.8093	-	0.9976	0.0141	3.9900e-05	0.3582	**26991**	10737	16254
	(3.8334e-08)	(0.0094)	(0.0081)	-	(0.0002)	(0.0017)	(0.0047)	(0.0260)			
Panel B: Multiple VIX estimation											
GARCH											
VIX	5.6000e-06	0.0581	0.8594	-	-	-	-	0.9999	98587	9215	**89372**
	(3.2673e-07)	(0.0014)	(0.0088)	-	-	-	-	(0.0685)			
Both	5.7418e-06	0.0905	0.8746	-	-	-	-	0.3283	**99611**	10629	88981
	(1.2454e-07)	(0.0026)	(0.0032)	-	-	-	-	(0.0124)			
EGARCH											
VIX	-0.2852	-0.0951	0.9647	0.1092	-	-	-	-0.1609	100346	10664	**89682**
	(0.0037)	(0.0026)	(0.0004)	(0.0031)	-	-	-	(0.0014)			
Both	-0.3000	-0.0903	0.9644	0.1123	-	-	-	-0.0362	**100410**	10750	89660
	(0.0213)	(0.0044)	(0.0025)	(0.0167)	-	-	-	(0.0010)			
CGARCH											
VIX	1.5107e-07	0.0990	0.7596	-	0.9936	0.0022	-	0.9751	61860	9498	**52363**
	(3.8118e-08)	(0.0130)	(0.0206)	-	(0.0007)	(0.0009)	-	(0.0773)			
Both	3.9260e-07	0.1432	0.7739	-	0.9961	0.0156	-	0.3991	**62699**	10699	52000
	(1.5216e-08)	(0.0041)	(0.0063)	-	(0.0002)	(0.0006)	-	(0.0071)			
ACGARCH											
Both	3.8698e-06	0.0012	0.9878	-	0.9542	0.1148	1.4071e-02	0.2870	**63189**	10770	52419
	(8.0873e-08)	(0.0007)	(0.0003)	-	(0.0014)	(0.0035)	(0.0007)	(0.0103)			

This table displays the maximum likelihood estimates using returns, VIX or both datasets, showcasing the results of the GARCH, EGARCH, CGARCH, and ACGARCH models. The VIX is estimated with equation 20 in Panel A and equation 25 in Panel B. The standard errors are given in parentheses and are calculated with the inverse of the Hessian matrix. The bold values of the log-likelihood denote that this log-likelihood has been maximized.

For the GARCH model, lambda increases from 0.1109 (returns used) to 0.9999 (single VIX estimation or multivariate VIX estimation) and reaches the upper bound. When returns and VIX data are used, lambda increases to 0.1990 (single VIX estimation) or 0.3283 (multivariate VIX estimation), based on the estimation method for $\ln L_v$. The persistence parameter β_1 becomes lower with the multivariate VIX estimation method. For VIX data, it goes from 0.9026 (single VIX estimation) to 0.8595 (multivariate VIX estimation), and for both datasets, it goes from 0.8935 (single VIX estimation) to 0.8594 (multivariate VIX estimation). Meanwhile, α_0 and α_1 increase when the multivariate VIX estimation method is used.

The EGARCH model only reports a significant equity premium when returns and VIX data are considered using the single VIX estimation method. Moreover, the EGARCH model reports a negative risk premium (-0.1293) when VIX data and the single VIX estimation method are used. Furthermore, when using the multivariate VIX estimation method, the model shows a risk premium of -0.1609 for VIX data and -0.0362 for returns and VIX data. All the other parameters decrease slightly when the multivariate VIX estimation is used, for VIX data and returns and VIX data.

For the CGARCH and ACGARCH models, we see similar results as for the GARCH model: the lambda parameter increases substantially when only VIX data is considered. Furthermore, for the CGARCH model, ρ drops substantially from 0.0724 (single VIX estimation) to 0.0022 (multivariate VIX estimation), when VIX data is considered.

Estimating the newly introduced ACGARCH model poses challenges because the asymmetry parameter θ becomes unidentifiable when only VIX data are used. When the single VIX estimation method is used, the asymmetric component is not significantly different from zero. Meanwhile, when the multivariate VIX estimation method is used, the coefficient is significantly different from zero. Furthermore, the log-likelihood values for the ACGARCH model are higher than those of the CGARCH model, which is anticipated, since the ACGARCH model is a nested model of the CGARCH model. However, the improvements in log-likelihood are very slim.

5.2.1 Implied VIX characteristics under the LRNVR

Table 5 displays the properties of the implied VIX for various models and estimation methods. A notable observation is that the multivariate VIX estimation method mostly results in higher mean errors compared to the benchmark approach of Hao and Zhang (2013). Additionally, the single VIX estimations generally replicate the autocorrelations and statistical properties of the CBOE VIX more accurately.

The models estimated with returns struggle to estimate the CBOE VIX well, since the mean error is substantially higher than when the other datasets are used. This means that the models tend to underestimate the CBOE VIX when returns data is used, which is in line with the results of Hao and Zhang (2013). The mean errors are very similar for every model, with the minimum error being 1.78 obtained by the EGARCH model, and the maximum 2.51, obtained by the GARCH model. Furthermore, the statistical properties of the implied VIX differ substantially from the actual properties of the CBOE VIX.

Models estimated exclusively with VIX data tend to produce the lowest errors, which is expected since the VIX is directly targeted. However, the parameters become distorted in the process, which is shown in table 4. Estimating with both returns and VIX data mostly yields slightly higher errors than using only VIX data.

For the GARCH model, targeting the VIX directly with VIX data estimated with the single VIX estimation method yields the lowest errors. Additionally, the correlation coefficient (0.89) is the best compared to the other estimates of the GARCH model. The statistical properties of the implied VIX are still a bit different than those of the CBOE VIX, since the implied VIX has a relatively lower variance and higher skewness and kurtosis. The multivariate VIX estimation produces negative mean errors, suggesting that the implied VIX is overestimated. Consequently, this leads to a higher RMSE, and the statistical properties are slightly less accurate than those obtained with the single VIX estimation.

Table 5: Implied VIX characteristics under the LNRVR

Model & Data	ME	Std.Err.	MAE	MSE	RMSE	P-value	Corr.Coef.	AR1	AR10	AR30	Variance	Skewness	Kurtosis
Panel A: Single VIX estimation													
GARCH													
Returns	1.90	3.74	2.92	17.60	4.20	0.0000	0.86	0.9658	0.6658	0.2652	49.33	5.44	50.51
VIX	0.08	3.26	2.38	10.63	3.26	0.5608	0.89	0.9898	0.8369	0.4525	43.49	4.20	29.74
Both	0.23	3.44	2.51	11.90	3.45	0.0942	0.88	0.9885	0.7938	0.3686	44.83	5.02	40.83
EGARCH													
Returns	1.78	3.24	2.64	13.67	3.70	0.0000	0.89	0.9690	0.7016	0.3358	35.48	2.48	15.78
VIX	0.01	2.92	2.18	8.52	2.92	0.9653	0.91	0.9883	0.8290	0.4773	42.66	2.56	15.45
Both	0.08	2.93	2.18	8.57	2.93	0.5520	0.91	0.9865	0.8129	0.4536	43.45	2.55	15.61
CGARCH													
Returns	2.45	2.90	2.78	14.42	3.80	0.0000	0.91	0.9139	0.7488	0.4562	41.35	4.36	36.11
VIX	0.10	2.34	1.63	5.49	2.34	0.4385	0.95	0.9450	0.7957	0.4832	43.61	3.66	27.75
Both	0.28	2.66	1.86	7.17	2.68	0.0327	0.93	0.9796	0.7777	0.4907	47.05	3.92	31.46
ACGARCH													
Returns	2.13	2.83	2.55	12.57	3.55	0.0000	0.92	0.9706	0.7360	0.4378	43.17	4.43	37.69
Both	0.28	2.64	1.85	7.06	2.66	0.0373	0.93	0.9295	0.7808	0.4898	46.80	3.93	31.35
Panel B: Multiple VIX estimation													
GARCH													
VIX	-1.15	3.54	3.10	13.84	3.72	0.0000	0.87	0.9814	0.7562	0.3056	30.46	5.27	45.06
Both	-0.93	3.51	2.96	13.17	3.63	0.0000	0.87	0.9849	0.7582	0.3115	35.66	5.58	49.19
EGARCH													
VIX	-0.99	3.72	3.13	14.84	3.85	0.0000	0.91	0.9839	0.7939	0.4303	18.09	2.54	15.74
Both	-0.80	3.67	3.00	14.08	3.75	0.0000	0.91	0.9830	0.7897	0.4250	18.95	2.46	15.09
CGARCH													
VIX	0.13	2.42	1.69	5.89	2.43	0.3413	0.94	0.9410	0.7554	0.4588	43.62	3.90	31.37
Both	0.03	2.66	1.91	7.08	2.66	0.8071	0.93	0.9762	0.7718	0.4989	47.05	3.98	31.93
ACGARCH													
Both	-0.25	2.49	1.71	6.28	2.51	0.0594	0.94	0.9394	0.7574	0.4796	52.23	3.91	31.02
CBOE VIX								0.9671	0.7925	0.5356	51.07	2.45	15.31

This table shows statistical properties and differences of the CBOE VIX maturities and implied VIX of the GARCH models. The error is calculated as the difference between the CBOE VIX indices and the implied VIX. The ME displays the daily average error. The standard error (Std.Err.) displays the standard deviation of the of the error term. The MAE displays the daily mean absolute error. The MSE displays the daily mean squared error, and the RSME displays the daily root mean squared error. The P-value corresponds to the P-value for the null hypothesis that the average of the CBOE VIX index and implied VIX are equal, and is calculated with a T-test. The correlation coefficient (Corr.Coef.) displays the correlation between the CBOE VIX index and the implied VIX. The AR1, AR10 and AR30 display the autocorrelation coefficients for lag 1, 10 and 30. The last columns display additional statistical properties of the CBOE VIX index and implied VIX.

For the EGARCH model, the single VIX estimation using the VIX dataset also provides the best performance. It has a high correlation coefficient (0.91) and low error metrics. The implied VIX closely mimics the CBOE VIX with autocorrelations (AR1: 0.9883, AR10: 0.8290, AR30: 0.4773). Furthermore, the implied VIX of the EGARCH model is able to replicate the moments of the CBOE VIX well, with skewness values of 2.56 (implied) versus 2.45 (actual) and kurtosis values of 15.45 (implied) versus 15.31 (actual). However, the variance of the implied VIX remains slightly lower. Similar to the GARCH model, the mean errors are negative for the multivariate VIX estimation method, indicating an overestimation of the CBOE VIX.

Most importantly, the EGARCH model with both returns and VIX data is the best-performing model, as it replicates the statistical properties of the CBOE VIX more accurately than any other model, despite not having the lowest RMSE. In contrast, other models struggle to accurately capture the skewness and kurtosis of the CBOE VIX.

For the CGARCH model, the single VIX estimation using the VIX dataset achieves the highest correlation coefficient (0.95) and the lowest error metrics among all models. While the implied VIX's autocorrelations are close to the actual values, the statistical properties of the implied VIX differ. The skewness and the kurtosis are higher than the actual values. The multivariate VIX estimation slightly improves the RMSE when both datasets are considered, reducing it from 2.68 to 2.66. The mean errors are close to zero, in contrast to the GARCH and EGARCH model with the multivariate approach. This might be due to excluding VIX6M and VIX1Y in the estimation to reduce computation time.

For the ACGARCH model, the error terms become lower for the multivariate VIX estimation. There is still a slight overestimation of the implied VIX, as the mean error is negative. The variance of the implied VIX (52.23) closely matches the variance of the CBOE VIX (51.07), performing the best of all models. However, the autocorrelations are slightly lower for every lag. A visualization of the implied VIX of the ACGARCH model using both returns and VIX data is shown in Figure 2. The left panels display the CBOE VIX and the implied VIX, and the right panels display the ratio between the implied VIX and the CBOE VIX.

Figure 2: The comparison between the CBOE VIX and the implied VIX of the ACGARCH model under the LRNVR estimated with both datasets

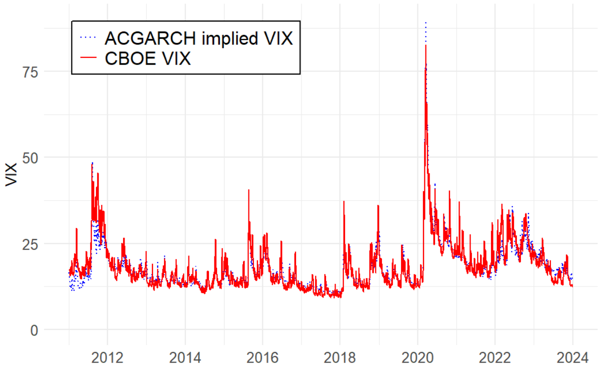

(a) Single VIX estimation: implied VIX

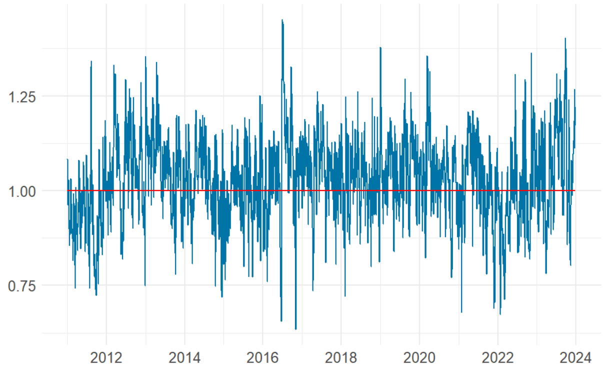

(b) Single VIX estimation: implied VIX ratio

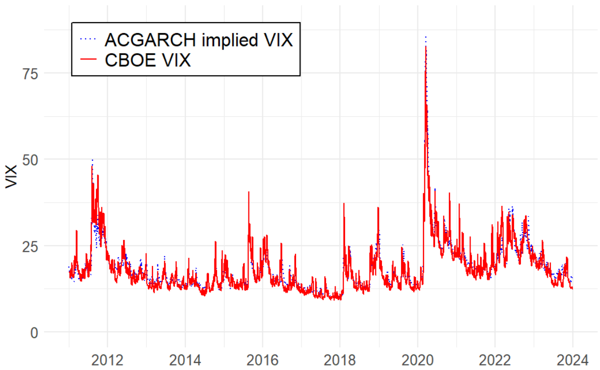

(c) Multivariate VIX estimation: implied VIX

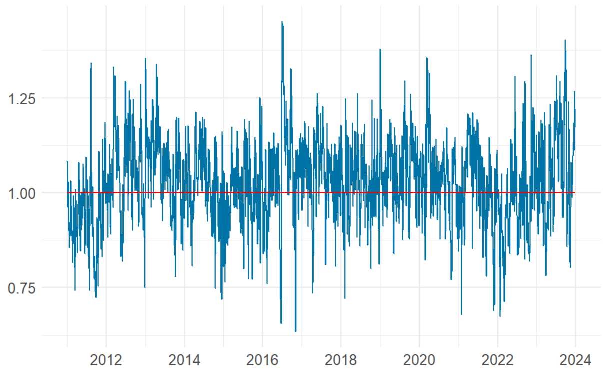

(d) Multivariate VIX estimation: implied VIX ratio

5.3 Estimation with different VIX measures under the mLRNVR

Next, we examine the results of the GARCH and EGARCH models under the mLRNVR. The parameter estimates and implied VIX characteristics are shown in Table 6 and 7, respectively. The original lambda of the LRNVR is denoted as λ_1, while the additional lambda parameter introduced for the mLRNVR is denoted as λ_2. The results for the returns data are not included here, as they are already provided in Table 4, as λ_2 is zero under the physical measure. The most important observation is that λ_1 still reaches high values for GARCH when VIX data is used, even under the mLRNVR. Another key observation is that the multivariate VIX estimation method tends to overestimate the CBOE VIX.

Table 6: The maximum likelihood estimates under the mLRNVR

							Log-likelihoods		
Model & Data	α_0	α_1	β_1	κ	λ_1	λ_2	Total	Returns	VIX
				Panel A: Single VIX estimation					
GARCH									
VIX	2.577e-06	0.0439	0.8753	-	0.9999	-0.4404	24856	9270	**15586**
	(8.9886e-08)	(0.0015)	(0.0047)	-	(0.0487)	(0.0152)			
Both	3.6058e-06	0.0854	0.8604	-	0.1998	-0.2803	**26165**	10765	15400
	(1.2080e-07)	(0.0032)	(0.0079)	-	(0.0150)	(0.0259)			
EGARCH									
VIX	-0.1556	-0.0867	0.9990	0.1263	-0.0353	-0.1484	26783	10832	**15950**
	(0.0047)	(0.0031)	(0.0008)	(0.0031)	(0.0031)	(0.0007)			
Both	-0.1722	-0.0890	0.9844	0.1393	0.0067	-0.0395	**26797**	10859	15938
	(0.0052)	(0.0035)	(0.0005)	(0.0043)	(0.0148)	(0.0009)			
				Panel B: Multiple VIX estimation					
GARCH									
VIX	5.5996e-06	0.0581	0.8510	-	0.9999	-0.1024	98587	9215	**89372**
	(1.3097e-07)	(0.0023)	(0.0034)	-	(0.0167)	(0.0122)			
Both	5.7182e-06	0.0903	0.8304	-	0.3276	-0.3495	**99611**	10630	88981
	(1.2136e-07)	(0.0025)	(0.0044)	-	(0.0151)	(0.0051)			
EGARCH									
VIX	-0.2853	-0.0951	0.9895	0.1092	-0.1606	-0.1839	100346	10664	**89682**
	(0.0032)	(0.0024)	(0.0011)	(0.0032)	(0.0064)	(0.0084)			
Both	-0.3160	-0.0904	0.9737	0.1212	-0.0210	-0.0865	**100418**	10764	89654
	(0.0040)	(0.0031)	(0.0006)	(0.0031)	(0.0167)	(0.0013)			

This table displays the maximum likelihood estimates using VIX or both returns and VIX data, showcasing the results under the mLRNVR. The VIX is estimated with equation 20 in Panel A and 25 in Panel B. The standard errors are given in parentheses and are calculated with the inverse of the Hessian matrix. The bold values of the log-likelihood denote that this log-likelihood has been maximized.

Panel A of Table 6 presents the parameter estimates obtained using the single VIX estimation approach. In the GARCH model, the newly introduced parameter λ_2, designed to capture the volatility risk premium, is significant for both datasets. Specifically, λ_2 is -0.4404 when only VIX data is used and -0.2803 when both returns and VIX data are included. When VIX data is considered under the mLRNVR, the equity risk premium λ_1 remains high, which contrasts with the findings of Zhang and Zhang (2020), where λ_1 is significantly different from zero and falls within a normal range. This difference is due to different time periods. [2]

For the EGARCH model, λ_2 remains significant for both datasets, with values of -0.1484 (using VIX data) and -0.0395 (using both datasets). Meanwhile, the equity risk premium parameter λ_1 is negative when VIX data is used, and not significantly different from zero when both datasets are considered.

The parameter estimates for the multivariate VIX estimation method are denoted in Panel B of Table 6. The variance risk premium λ_2 is significantly different than zero for all models and datasets. Consistent with the single VIX estimation method, the GARCH model shows high λ_1 values when only VIX data is used. When both VIX and returns data are considered, λ_1 is significantly different from zero (0.3276). Similar to the results under the LNRVR, the

[2] The replication of the results of Zhang and Zhang (2020) was successful, and yields the same results as reported in the paper.

estimation with the EGARCH model results in a negative λ_1 for every dataset.

5.4 Implied VIX characteristics under the mLRNVR

The implied VIX characteristics are shown in Table 7. The lowest MSE (8.50) is obtained by the EGARCH model, estimated with the single VIX estimation method using VIX data. Panel A of Table 7 illustrates that the EGARCH model consistently fits the CBOE VIX well under the mLRNVR framework, regardless of the dataset used. Nevertheless, it is important to keep in mind that the parameter estimates are distorted when only VIX data is used, therefore estimation with returns and VIX data might be preferred. Interestingly, when the multivariate VIX estimation method is used, the GARCH and EGARCH models tend to overestimate the CBOE VIX, resulting in a negative mean error. Still, the EGARCH is better at fitting the CBOE VIX considering the moments and autocorrelations of the implied VIX.

Table 7: Implied VIX characteristics under the mLRNVR

Model & Data	ME	Std.Err.	MAE	MSE	RMSE	P-value	Corr.Coef.	AR1	AR10	AR30	Variance	Skewness	Kurtosis
Panel A: Single VIX estimation													
GARCH													
VIX	0.08	3.26	2.38	10.63	3.26	0.5608	0.89	0.9898	0.8369	0.4525	43.49	4.20	29.74
Both	0.23	3.45	2.51	11.92	3.45	0.0985	0.88	0.9883	0.7915	0.3653	44.79	5.04	41.14
EGARCH													
VIX	0.01	2.92	2.17	8.50	2.92	0.9556	0.91	0.9882	0.8276	0.4750	42.75	2.54	15.31
Both	0.08	2.93	2.18	8.57	2.93	0.5521	0.91	0.9865	0.8129	0.4536	43.45	2.55	15.61
Panel B: Multiple VIX estimation													
GARCH													
VIX	-1.15	3.54	3.10	13.84	3.72	0.0000	0.87	0.9814	0.7561	0.3054	30.46	5.27	45.05
Both	-0.92	3.51	2.95	13.15	3.63	0.0000	0.87	0.9849	0.7589	0.3124	35.70	5.57	49.08
EGARCH													
VIX	-1.00	3.72	3.13	14.84	3.85	0.0000	0.91	0.9839	0.7939	0.4302	18.09	2.54	15.74
Both	-0.81	3.67	3.01	14.12	3.76	0.0000	0.91	0.9821	0.7822	0.4173	19.00	2.46	15.16
CBOE VIX								0.9671	0.7925	0.5356	51.07	2.45	15.31

This table shows statistical properties and differences of the CBOE VIX maturities and implied VIX of the GARCH models. The error is calculated as the difference between the CBOE VIX indices and the implied VIX. The ME displays the daily average error. The standard error (Std.Err.) displays the standard deviation of the error term. The MAE displays the daily mean absolute error. The MSE displays the daily mean squared error, and the RMSE displays the daily root mean squared error. The P-value corresponds to the P-value for the null hypothesis that the average of the CBOE VIX index and implied VIX are equal, and is calculated with a T-test. The correlation coefficient (Corr.Coef.) displays the correlation between the CBOE VIX index and the implied VIX. The AR1, AR10, and AR30 display the autocorrelation coefficients for lag 1, 10, and 30. The last columns display additional statistical properties of the CBOE VIX index and implied VIX.

Figure 3 shows the comparison between the CBOE VIX and the implied VIX, estimated with the EGARCH model under the mLRNVR using both datasets. The left panels display the CBOE VIX and the implied VIX, and the right panels display the ratio between the implied VIX and the CBOE VIX. The ratio for the single VIX estimation, shown in panel 3b, is nicely centered around 1. Meanwhile, the overestimation of the CBOE VIX is visible in panel 3d, where the ratio is above 1 for a long period.

Figure 3: The comparison between the CBOE VIX and the implied VIX of the EGARCH model under the mLRNVR estimated with both datasets

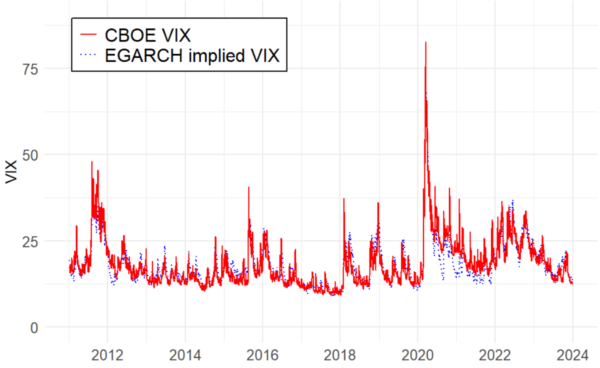

(a) Single VIX estimation: implied VIX

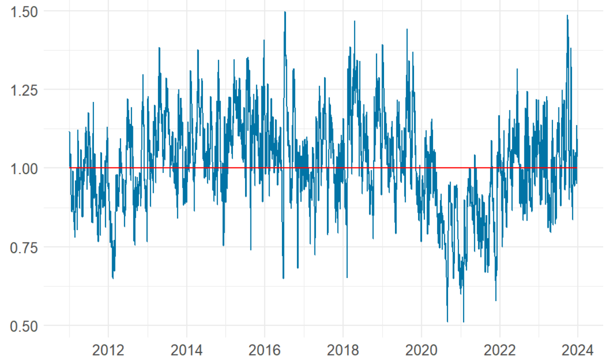

(b) Single VIX estimation: implied VIX ratio

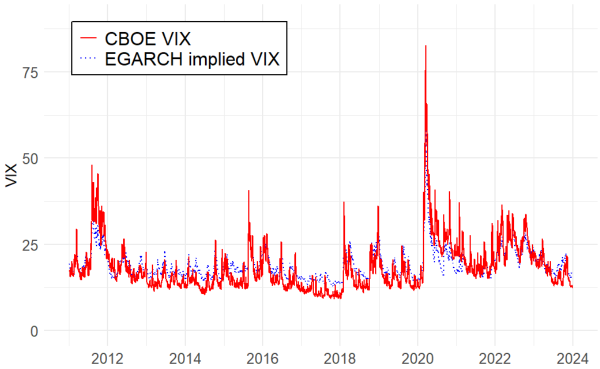

(c) Multivariate VIX estimation: implied VIX

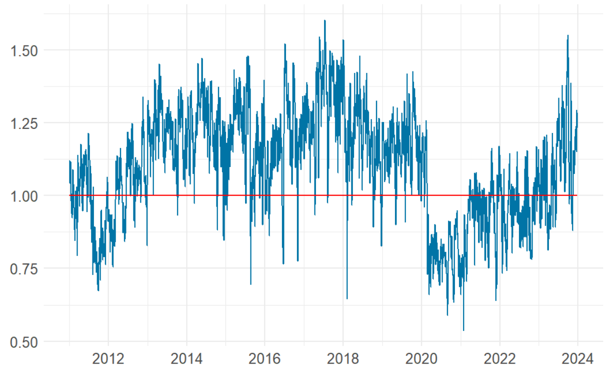

(d) Multivariate VIX estimation: implied VIX ratio

5.5 Residual analysis of the EGARCH model

Since the EGARCH model is the best-performing model under both the LRNVR and mLRNVR, we analyze the characteristics of the error term u_t. The best-performing model means that the implied VIX replicates the statistical properties of the CBOE VIX most accurately. The error term of the EGARCH model is already standardized, as $u_t = \xi_t/\sqrt{h_t}$. An important observation is that the single VIX estimation method produces higher Jarque-Bera test statistics than the multivariate VIX estimation method, yet normality is firmly rejected for both models.

The moments of the error terms are shown in table table 8. The first three columns report moments of the standardized residuals, estimated under the LRNVR. The fourth and the fifth columns report the Jarque-Bera test statistic and its corresponding p-value, respectively. The last five columns report the same statistics, instead estimated under the mLRNVR. Notably, the implied characteristics under the mLRNVR closely resemble those observed under the LRNVR.

Table 8: Characteristics of the standardized residuals of the EGARCH model

	LRNVR						mLRNVR					
Estimation Method & Data	Mean	Skewness	Variance	Kurtosis	JB statistic	P-value	Mean	Skewness	Variance	Kurtosis	JB statistic	P-value
Single VIX estimation												
VIX	0.17	-0.76	0.79	5.44	1130.78	0.0000	0.07	-0.76	0.80	5.41	1108.71	0.0000
both	0.03	-0.76	0.81	5.42	1112.24	0.0000	0.03	-0.76	0.81	5.42	1112.22	0.0000
Multivariate VIX estimation												
VIX	0.20	-0.70	0.68	5.28	977.49	0.0000	0.20	-0.70	0.68	5.28	977.32	0.0000
both	0.07	-0.71	0.70	5.24	955.54	0.0000	0.06	-0.70	0.71	5.21	935.09	0.0000

This table shows the statistical properties of the residuals of the EGARCH model, estimated with different datasets. The mean, skewness, variance, kurtosis, Jarque-Bera statistic, and its corresponding p-value are provided for each model configuration.

Every mean is close to zero, which is expected for standardized residuals. Using both datasets to estimate the EGARCH model yields a smaller mean than only using VIX data. This shows that the error terms when both returns and VIX data are used, are slightly better centered than only using VIX data. The variances should be close to 1, but the values are somewhat lower for every estimation method. Interestingly, the multivariate VIX estimation approach results in lower variance values than the single VIX estimation approach. All of the skewness values are negative, indicating that the tail is more on the left side of the distribution. Furthermore, the skewness values do not differ a lot across the models, as the lowest value is -0.76 and the highest value is -0.70. Every kurtosis value is higher than 3, meaning that the distribution has fatter tails and a sharper peak than a standard normal distribution. This means that the standardized residuals for every model contain more extreme values than a standard normal distribution.

The Jarque-Bera test statistic is high for all models, suggesting that the error terms are not standard normally distributed. The single VIX estimation method does result in slightly higher Jarque-Bera test statistics than the multivariate VIX estimation method, suggesting that the standardized residuals deviate slightly more from normality. Still, these high Jarque-Bera test statistics result in low p-values for both estimation methods, firmly rejecting the null hypothesis and implying non-normality in the standardized residuals for all models.

Figure 4 illustrates the distribution of the standardized residuals under the LRNVR, estimated with both datasets. The red line in the density plots represents the standard normal distribution. The higher kurtosis is visible for both estimation methods since both density plots have a higher peak and fatter tails. The multivariate VIX estimation exhibits slightly lighter tails and is marginally less left-skewed.

Figure 4: Density plots of the standardized residuals under the LRNVR, estimated with the EGARCH model using both datasets

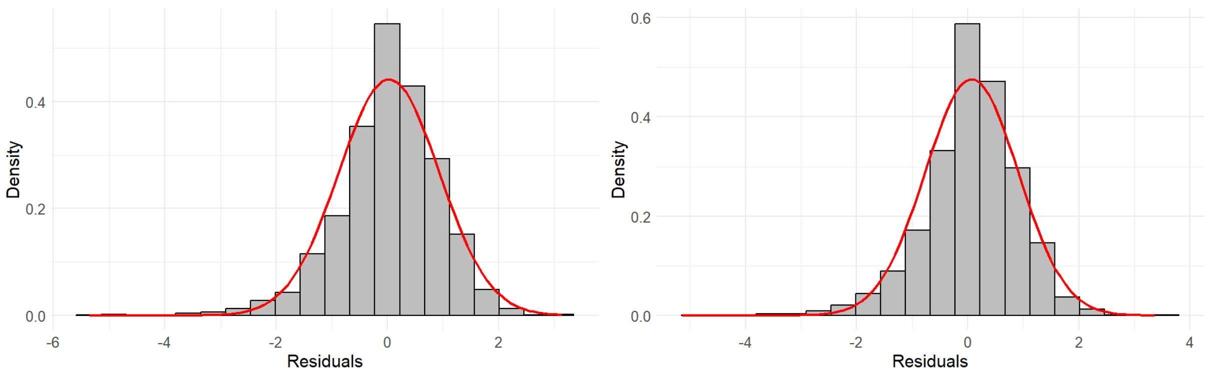

(a) Single VIX estimation (b) Multivariate VIX estimation

The first 20 autocorrelations of these models are shown in Figure 5. Both estimation methods result in largely uncorrelated residuals. Only lag 15 exceeds the confidence bounds in panel 5b. Overall, the majority of the autocorrelations are within the confidence bounds, suggesting that the residuals are not correlated.

Figure 5: Autocorrelations of the standardized residuals under the LRNVR, estimated with the EGARCH model using both datasets

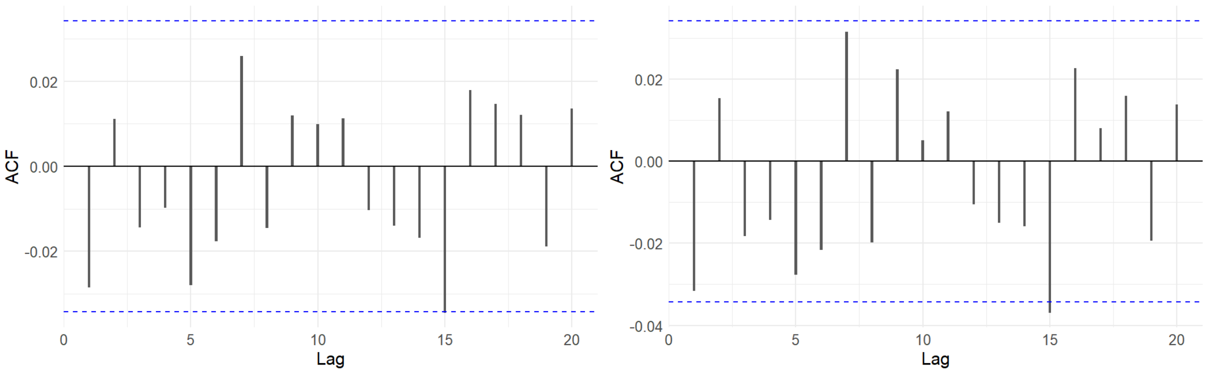

(a) Single VIX estimation (b) Multivariate VIX estimation

Figure 6 shows the behavior of the residual over time. The residuals for both estimation methods are centered around zero, indicating that there is no systematic bias in the residuals. Furthermore, the spread of the residuals appears to be consistent, suggesting homoscedasticity in the standardized residuals. The overall behavior of the residuals over time is quite similar between the two estimation methods, with no major differences in the observed characteristics.

Figure 6: Standardized residuals over time under the LRNVR, estimated with the EGARCH model using both datasets

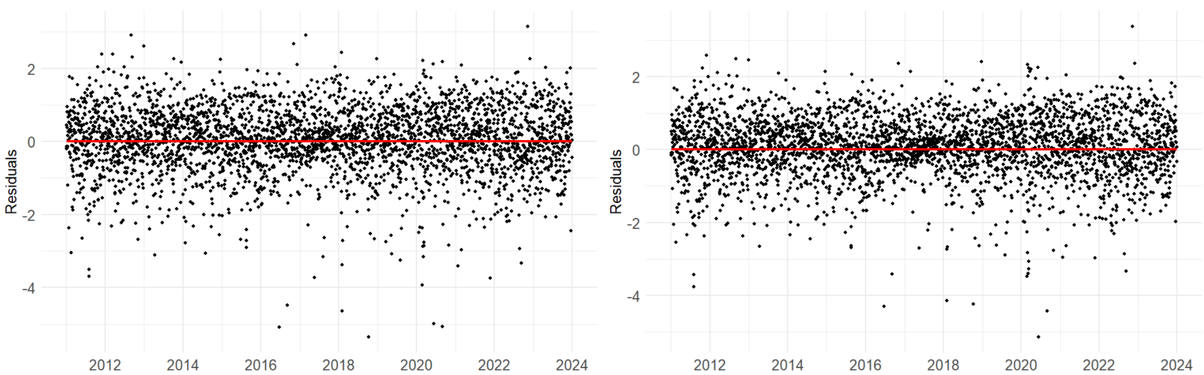

(a) Single VIX estimation (b) Multivariate VIX estimation

The residual analysis for the models under the mLRNVR yields results that are very similar to those observed for the LRNVR models. The distributions of the standardized residuals, autocorrelations, and the behavior of the residuals over time exhibit similar characteristics. The standardized residuals have a higher kurtosis, minimal autocorrelation, and are consistent over time. Given the similarity in results, detailed graphs for the mLRNVR models are provided in the appendix A.4.

6 Conclusion

In this paper, we examined whether a multivariate estimation approach with multiple VIX measures could enhance the estimation of the CBOE VIX. We analyzed the new approach with GARCH option pricing models under the framework of Duan (1995) and calculated the implied VIX based on the methodology of Hao and Zhang (2013). Additionally, we evaluated the approach under the mLRNVR framework of Zhang and Zhang (2020). We analyzed four different types of GARCH models under the LRNVR and two GARCH models under the mLRNVR.

Under the LRNVR, the single VIX estimation method is generally preferred for most GARCH models. However, the exceptions are the CGARCH and ACGARCH models when both VIX and returns data are used for estimation. Similar to the results of Hao and Zhang (2013), parameters become distorted to fit the CBOE VIX when only VIX data is considered for the estimation. Nevertheless, the GARCH implied VIX for the CGARCH and ACGARCH models still fails to match the statistical properties of the CBOE VIX. Furthermore, we see that the multivariate VIX estimation method tends to overestimate the CBOE VIX for the GARCH and EGARCH models, resulting in larger errors.

Under the mLRNVR, the single VIX estimation method yields better results than the multivariate VIX estimation method. In terms of fit and parameter estimates, the EGARCH model estimated with the single VIX estimation method using both returns and VIX data yields the best results in terms of fit. Similar to the estimation under the LRNVR, the GARCH and EGARCH models tend to overestimate when the multivariate VIX estimation method is used.

The analysis of the error terms of the EGARCH model under the LRNVR and mLRNVR reveals that the single VIX estimation method produces higher Jarque-Bera test statistics than the multivariate VIX estimation method, yet normality is firmly rejected for both models.

To address the research question, our findings suggest that the multivariate VIX estimation method generally does not provide a superior fit for the CBOE VIX. There is a slight improvement only for the CGARCH and ACGARCH models when both returns and VIX data are used. However, these models still struggle to replicate the autocorrelations and moments of the CBOE VIX. Consequently, our findings align with those of Zhang and Zhang (2020). We recommend estimating the CBOE VIX using the EGARCH model with the single VIX estimation method, incorporating both returns and VIX data, as it yields the best fit.

Further research could explore the out-of-sample performance of these models and evaluate their predictive power. Additionally, ARIMA-type models could be considered to model the VIX directly, allowing for a comparison between GARCH and ARIMA models in estimating and forecasting the CBOE VIX.

References

Black, F. and Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of political economy*, 81(3):637–654.

Bollerslev, T., Tauchen, G., and Zhou, H. (2009). Expected stock returns and variance risk premia. *The Review of Financial Studies*, 22(11):4463–4492.

Byun, S. J., Jeon, B. H., Min, B., and Yoon, S.-J. (2015). The role of the variance premium in jump-garch option pricing models. *Journal of Banking & Finance*, 59:38–56.

Carr, P. and Wu, L. (2009). Variance risk premiums. *The Review of Financial Studies*, 22(3):1311–1341.

Christoffersen, P., Heston, S., and Jacobs, K. (2013). Capturing option anomalies with a variance-dependent pricing kernel. *The Review of Financial Studies*, 26(8):1963–2006.

Christoffersen, P., Jacobs, K., and Ornthanalai, C. (2012). Dynamic jump intensities and risk premiums: Evidence from s&p500 returns and options. *Journal of Financial Economics*, 106(3):447–472.

Duan, J.-C. (1995). The garch option pricing model. *Mathematical finance*, 5(1):13–32.

Engle, R. F. and Bollerslev, T. (1986). Modelling the persistence of conditional variances. *Econometric reviews*, 5(1):1–50.

Hao, J. and Zhang, J. E. (2013). Garch option pricing models, the cboe vix, and variance risk premium. *Journal of Financial Econometrics*, 11(3):556–580.

Heston, S. L. and Nandi, S. (2000). A closed-form garch option valuation model. *The review of financial studies*, 13(3):585–625.

Kanniainen, J., Lin, B., and Yang, H. (2014). Estimating and using garch models with vix data for option valuation. *Journal of Banking & Finance*, 43:200–211.

Lee, G. G. and Engle, R. F. (1993). A permanent and transitory component model of stock return volatility. *Available at SSRN 5848*.

Meddahi, N. and Renault, E. (2004). Temporal aggregation of volatility models. *Journal of Econometrics*, 119(2):355–379.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica: Journal of the econometric society*, pages 347–370.

Schwert, G. W. (1988). Why does stock market volatility change over time? NBER Working Paper w2798, National Bureau of Economic Research. Available at SSRN: `https://ssrn.com/abstract=447222`.

Tong, C. (2024). Pricing cboe vix in non-affine garch models with variance risk premium. *Finance Research Letters*, page 105115.

Tong, C. and Huang, Z. (2021). Pricing vix options with realized volatility. *Journal of Futures Markets*, 41(8):1180–1200.

Wang, T., Shen, Y., Jiang, Y., and Huang, Z. (2017). Pricing the cboe vix futures with the heston–nandi garch model. *Journal of Futures Markets*, 37(7):641–659.

Zhang, W. and Zhang, J. E. (2020). Garch option pricing models and the variance risk premium. *Journal of Risk and Financial Management*, 13(3):51.

A Appendix

A.1 Proof of Proposition 1

Let $\{\xi_t, t \in Z\}$ be a m.d.s. with the conditional variance $h_t \equiv \mathrm{Var}[\xi_t \mid \xi_\tau, \tau \leq t-1]$ under the LRNVR. If $u_t = \xi_t/\sqrt{h_t}$ is i.i.d., the ACGARCH(1,1) model (13) is a SR-SARV(2) process.
We can rewrite the ACGARCH(1,1) model as $h_t = e'F_t$ and $F_t = \Omega + \Gamma F_{t-1} + V_{t-1}$ with

$$e = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad F_t = \begin{pmatrix} h_t \\ q_t \end{pmatrix}, \quad \Omega = \alpha_0 \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad \Gamma = \begin{pmatrix} \alpha_1 + \beta_1 + (\phi + \alpha_1)\lambda^2 + \theta S & \rho - \alpha_1 - \beta_1 - \theta S \\ \phi\lambda^2 & \rho \end{pmatrix},$$

$$V_t = h_{t-1}(u_{t-1}^2 - 2\lambda u_{t-1} - 1) \begin{pmatrix} \phi + \alpha_1 \\ \phi \end{pmatrix} + \theta h_{t-1} \left[(u_{t-1} - \lambda)^2 1\,(u_{t-1} < \lambda) - S \right] \begin{pmatrix} 1 \\ 0 \end{pmatrix}.$$

If $u_t = \xi_t/\sqrt{h_t}$ is i.i.d. N(0,1), we get $S = \left[\frac{\lambda}{\sqrt{2\pi}} e^{-\frac{\lambda^2}{2}} + \left(1 + \lambda^2\right) N(\lambda) \right]$. Since $E_{t-2}^Q(u_{t-1}) = 0$ and $E_{t-2}^Q(u_{t-1}^2) = 1$, we have $E_{t-2}^Q(V_{t-1}) = 0$.

A.2 Implied VIX formulas

We follow the approach of Hao and Zhang (2013), and calculate the implied VIX under the properties of SR-SARV models Meddahi and Renault (2004). Hao and Zhang (2013) show that the proxy V_t can be calculated as a linear function of the conditional variance of the next period, f_t and specify

$$V_t = \xi + \psi f_t. \tag{26}$$

Hao and Zhang (2013) derive the following implied VIX formulas for the GARCH models under the LRNVR:
GARCH(1,1):

$$V_t = A + Bh_{t+1}, \tag{27}$$

where

$$A = \frac{\alpha_0}{1-\eta}(1-B),$$
$$B = \frac{1-\eta^n}{n(1-\eta)},$$
$$\eta = \alpha_1(1+\lambda^2) + \beta_1.$$

EGARCH(1,1):

$$V_t = \frac{1}{n}\left(h_{t+1} + \sum_{k=1}^{n-1}\left(\prod_{i=0}^{k-1}\iota_i\right)h_{t+1}^{\beta_1^k}\right), \tag{28}$$

where ι_i is given by

$$\begin{aligned}\iota_i =& e^{\beta_1^i\left(\alpha_0-\kappa\sqrt{2/\pi}\right)}\left\{e^{-\beta_1^i(\alpha_1-\kappa)\lambda+\frac{\left[\beta_1^i(\alpha_1-\kappa)\right]^2}{2}}N\left[\lambda-\beta_1^i\left(\alpha_1-\kappa\right)\right]\right.\\ &\left.+e^{-\beta_1^i(\alpha_1+\kappa)\lambda+\frac{\left[\beta_1^i(\alpha_1+\kappa)\right]^2}{2}}N\left[\beta_1^i\left(\alpha_1+\kappa\right)-\lambda\right]\right\}\end{aligned} \tag{29}$$

if u_t is i.i.d. standard normal, which is assumed.

The implied VIX formulas for the GARCH and EGARCH models under the mLRNVR can be found by substituting $\beta_1 = \beta_1 - \sqrt{2}\alpha_1\lambda_2$.

CGARCH(1,1):

First, Hao and Zhang (2013) write the CGARCH(1,1) model as a SR-SARV(2) model and denote:

$h_t = e'F_t$ and $F_t = \Omega + \Gamma F_{t-1} + V_{t-1}$ with

$$e = \begin{pmatrix}1\\0\end{pmatrix}, \quad F_t = \begin{pmatrix}h_t\\q_t\end{pmatrix}, \quad \Omega = \alpha_0\begin{pmatrix}1\\1\end{pmatrix}, \quad \Gamma = \begin{pmatrix}\alpha_1+\beta_1+(\phi+\alpha_1)\lambda^2 & \rho-\alpha_1-\beta_1\\ \phi\lambda^2 & \rho\end{pmatrix},$$

$$V_t = h_{t-1}(u_{t-1}^2 - 2\lambda u_{t-1} - 1)\begin{pmatrix}\phi+\alpha_1\\ \phi\end{pmatrix}.$$

Subsequently, Hao and Zhang (2013) calculate the implied VIX with a numerical approach with

$$E_t^Q\left(f_{t+k}\right) = e'\left(\sum_{i=0}^{k-1}\Gamma^i\Omega + \Gamma^k F_t\right) \tag{30}$$

and equation 15. Note that f_{t+k} is h_{t+k}, but we use the notation of Meddahi and Renault (2004) in equation 30. We use the same approach to calculate the implied VIX for the ACGARCH(1,1) model, with the corresponding matrices and parameters given in proof A.1. A detailed approach of the derivation of the implied VIX formula in equation 26 can be found in Hao and Zhang (2013).

A.3 Stationairy constraints for the GARCH models under the LRNVR

GARCH(1,1):

$$\alpha_1(1+\lambda^2) + \beta_1 \leq 1$$

EGARCH(1,1):

$$|\beta_1| \leq 1$$

CGARCH(1,1):
The eigenvalues of the coefficient matrix

$$\Gamma = \begin{pmatrix} \alpha_1 + \beta_1 + (\phi + \alpha_1)\lambda^2 & \rho - \alpha_1 - \beta_1 \\ \phi\lambda^2 & \rho \end{pmatrix}$$

have a modules less than 1.

ACGARCH(1,1):
The eigenvalues of the coefficient matrix

$$\Gamma = \begin{pmatrix} \alpha_1 + \beta_1 + (\phi + \alpha_1)\lambda^2 + \theta S & \rho - \alpha_1 - \beta_1 - \theta S \\ \phi\lambda^2 & \rho \end{pmatrix}$$

have a modules less than 1.

A.4 Residual graphs under the mLRNVR

Figure 7: Density plots of the standardized residuals under the mLRNVR, estimated with the EGARCH model using both datasets

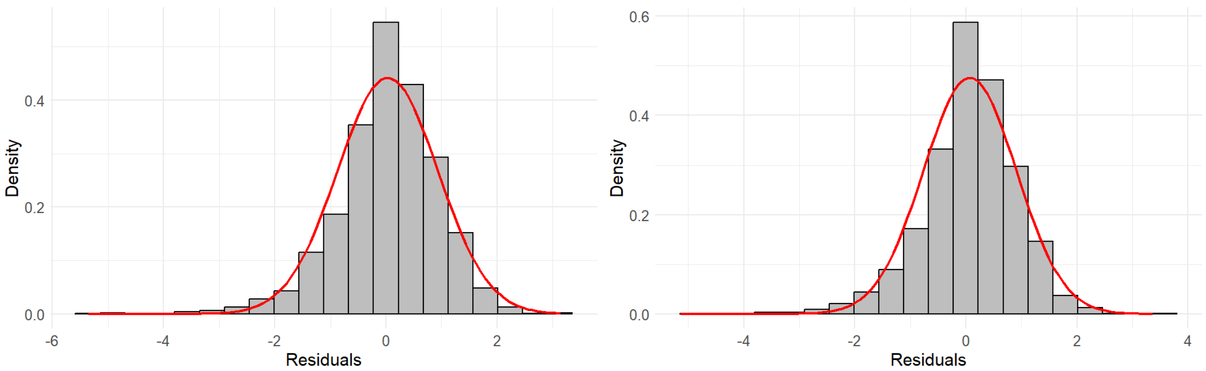

(a) Single VIX estimation (b) Multivariate VIX estimation

Figure 8: Autocorrelations of the standardized residuals under the mLRNVR, estimated with the EGARCH model using both datasets

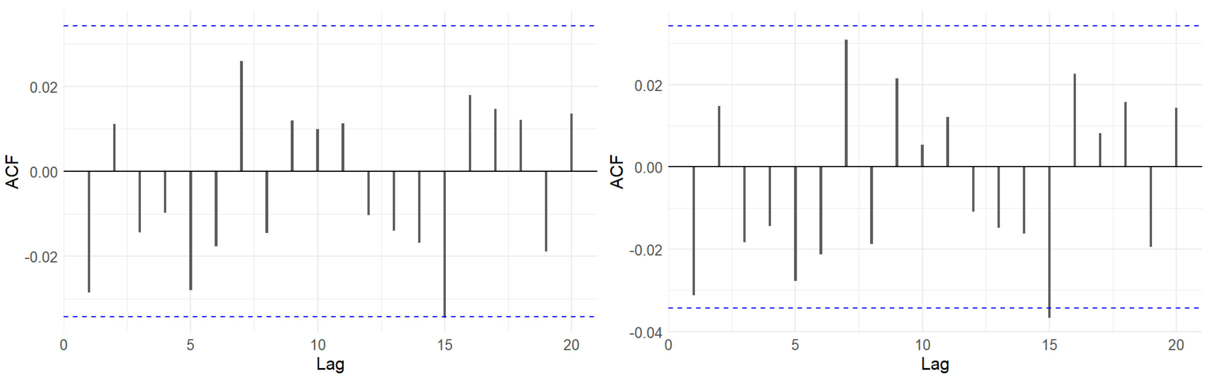

(a) Single VIX estimation (b) Multivariate VIX estimation

Figure 9: Autocorrelations of the standardized residuals under the mLRNVR, estimated with the EGARCH model using both datasets

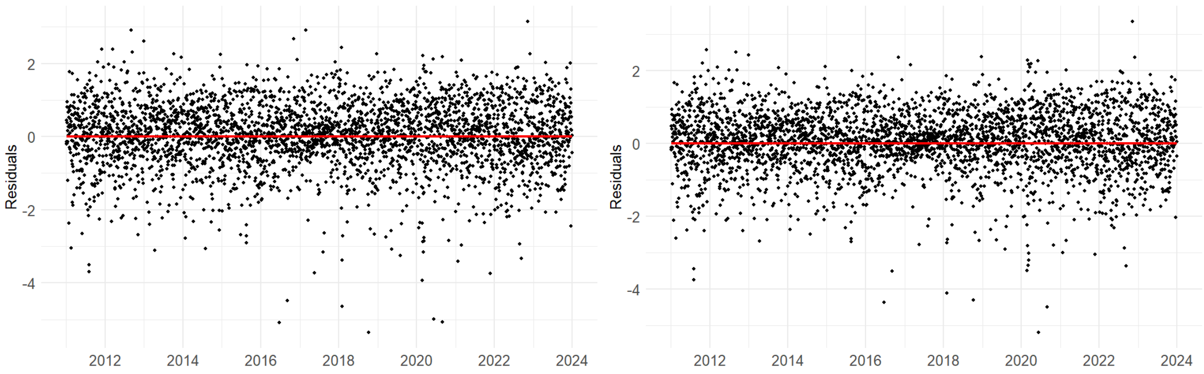

(a) Single VIX estimation (b) Multivariate VIX estimation

B Code Appendix

This appendix lists all the code and data files used to compute the results presented in this paper, along with a short explanation of each file.

- Data for the replication are located in **Replication_Dataset**.
- Data for the estimation for the time period of January 4, 2011 to 29 December, 2023 are located in **Dataset**.
- All results are computed using MATLAB code and saved to the excel file **Results_Thesis**. Each GARCH model has three associated files, where '[GARCH_MODEL]' should be replaced with the specific model name. The files used for estimation are:
  - **[GARCH_MODEL]_main.m** is the main execution file for estimating the model. The data is loaded here, sets up the initial parameters, and calls the other files for the minimization of the negative log-likelihood. This file should be executed to obtain the results.
  - **[GARCH_MODEL]_negLL.m** calculates the negative log-likelihood and returns it as output. It calls the **[GARCH_MODEL]_filter.m** file to obtain the volatility estimates, and calculates the log-likelihood based on returns data, VIX data, or both datasets.
  - **[GARCH_MODEL]_filter.m** calculates the volatility for a given set of parameters.
  - **nonlcon.m** contains non-linear constraints for optimizing the log-likelihood function and is an additional file specifically for the CGARCH and ACGARCH models.
- Important to note that the feature to choose between LNRVR and mLRNVR in the '[GARCH_MODEL]_main.m' has not been implemented yet. Therefore, the models under the mLRNVR have an additional folder that estimates the models under the mLRNVR.

- The following R-scripts are used to construct the graphs and to compute summary stats. Note that datasets should be imported manually.
  - **Summary_stats.rmd** computes the summary stats, shown in the Data section
  - **Graphs CBOE VIX.rmd** plots the graphs for every Figure in the paper and computes the results for Table 8.