ERASMUS UNIVERSITY ROTTERDAM

Erasmus School of Economics [*]

Bachelor Thesis Investing, major Quantitative Finance

Supervisor: Vladimirov, E

Second Assessor: Sluis, BP van der

GARCH implied volatility: an analysis of the risk premium

Luc van Maldegem (576765)

Date final version: Juli 1, 2024

Abstract

In this paper, the locally risk-neutral valuation relationship (LRNVR) of Duan (1995) is used to model implied volatility, as in Hao and Zhang (2013), utilising the square-root stochastic autoregressive volatility models of Meddahi and Renault (2004). This GARCH implied volatility is then extended with the modified LRNVR of Zhang and Zhang (2020) to incorporate the volatility risk premium. For the GARCH models, a GARCH(1,1) and an EGARCH(1,1) are used.

The models are applied to the volatility indices of the S&P 500, FTSE 100, DAX 30, and Nikkei 225. There appears to be a negative relationship between the persistence of shocks and the equity risk premium. The highest equity risk premiums are found for the S&P 500, followed by the FTSE 100 and DAX 30, with the lowest being the Nikkei 225. Contrary to this, the largest volatility risk premium is found for the Nikkei 225, with the S&P 500 following closely, while the DAX 30 and FTSE 100 have smaller volatility risk premiums.

Furthermore, the models are applied to different maturities of the CBOE VIX. The results suggest that as the maturity increases, the persistence increases and the equity risk premium decreases. Additionally, a lengthening in maturity results in an increase in the volatility risk premium. This indicates that for longer maturities, investors demand a smaller premium for bearing risk while being willing to pay a larger premium to hedge against an increase in risks.

[*]The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics or Erasmus University Rotterdam.

1 Introduction

In the financial world, bearing risk requires compensation. Part of this risk is explained by the volatility of returns, known as the equity risk premium. However, this volatility is not constant but varies over time. Consequently, investors also demand a volatility risk premium to compensate for the risk associated with this volatility. Carr and Wu (2008) analysed this volatility risk premium and found that the volatility risk premiums have a negative sign, indicating that investors perceive increasing market volatility as undesirable and are willing to pay a premium to hedge against it. To further analyse how this volatility risk premium is embedded into market prices, one can pursue two distinct directions. First, researchers can analyse the implied volatility, and second, they can study the difference between the variance swap rate and the realised variance[1].

This paper will highlight the first method, analysing implied volatility. By understanding and estimating implied volatility, one can analyse the equity and volatility risk premiums that are embedded in it. In this paper, the equity and volatility risk premiums will be analysed by applying a GARCH implied option pricing model to several indices and various implied volatility maturities. By understanding the properties of the equity risk premium and volatility risk premium, the ability for effective risk management and strategic decision-making in financial markets is enhanced. Furthermore, this analysis helps in understanding differences between markets by examining various indices.

The GARCH option pricing method was first introduced by Duan (1995), who applied the locally risk-neutral valuation relationship (LRNVR) to obtain a risk-neutral measure. Hao and Zhang (2013) further developed the model to include the square-root stochastic autoregressive volatility (SR-SARV) model introduced by Meddahi and Renault (2004), enabling it to effectively model implied volatility. Although the model was able to capture the equity risk premium, Hao and Zhang (2013) showed that this model was insufficient to capture the volatility risk premium. Therefore, Zhang and Zhang (2020) introduced a method that applied a modified LRNVR, which they demonstrated to be sufficient in capturing the volatility risk premium and the statistical properties of market volatility.

This paper utilises the models proposed by Hao and Zhang (2013) and Zhang and Zhang (2020) to analyse different stock indices and maturities. The indices analysed are the S&P 500, Nikkei 225, FTSE 100, and DAX 30, along with their corresponding volatility indices. Additionally, to analyse the impact of the maturity of the implied volatility, this paper examines

[1] The difference between the variance swap rate and the realised variance is not discussed in this paper. For more information on this topic, refer to Demeterfi et al. (1999) and Carr and Wu (2008).

1-day, 9-day, 1-month, 3-month, 6-month, and 1-year maturities of the CBOE VIX.

The results indicate a potential relationship between the persistence of shocks[2] and the equity risk premium: the higher the persistence, the lower the premium. Furthermore, as the maturity lengthens, the persistence increases while the equity risk premium decreases. The indices show consistent results, with higher persistence being linked to a lower equity risk premium. However, for the maturities, this reduction of the equity risk premium is paired with an increase in the volatility risk premium, indicating that investors are less sensitive to bear risk but more sensitive to a rise in risk.

The models employed in this paper utilise option pricing, the fundamentals of which were laid by Black and Scholes (1973) and Merton (1973). These models operate on no-arbitrage principles, ensuring that a fully hedged position yields a return equivalent to the risk-free rate. Extending this framework, researchers have proposed several models incorporating stochastic volatility, such as Wiggins (1987), Johnson and Shanno (1987), Hull and White (1987), and Heston (2015). The introduction of GARCH models led to their incorporation into option pricing, resulting in models such as Christoffersen et al. (2006).

Furthermore, empirical studies have shown the presence of volatility jumps, which in turn were incorporated into new models. These volatility jump models progressed from constant-volatility models, featuring constant-intensity jumps (Merton (1976) and Carr et al. (2002)), to models with stochastic-intensity jumps (Bates (2000) and Bates (2012)). Moreover, Duffie et al. (2000) presented a constant-intensity co-jump model where the spot volatility jumps synchronously and correlates with price jumps. Recent models by Andersen et al. (2015) and Carr and Wu (2020) feature co-jump models with self-exciting jump intensity.

These option pricing models all use volatility as one of the key variables. Therefore, if the price of an option is known, one can use reverse engineering to determine the volatility factor embedded in the option. Since the market prices the option, this resulting value is then referred to as market-implied volatility.

Duan (1995) states that the GARCH option pricing model has three unique features. First, it includes the risk premium embedded in the underlying asset, unlike standard preference-free option pricing models. Second, the GARCH option pricing model is non-Markovian, in contrast to traditional option pricing, which assumes the underlying asset follows a Markovian diffusion process. Third, the GARCH model can potentially account for several systematic biases commonly associated with the Black-Scholes model. These biases include the overpricing of out-of-the-money options, the underpricing of options on low-volatility securities, the underpricing

[2] A higher persistence of shocks is also in relation with a smaller impact of shocks. This coincides with a lower "News Impact Curve" (Engle, 1993)

of short-maturity options, and the U-shaped implied volatility curve relative to the exercise price (Black (1975), Gultekin et al. (1982), and Whaley (1982)).

An important assumption for option pricing is risk-neutralisation (Black and Scholes (1973), Rubinstein (1976)); however, due to the complexity of GARCH option modelling, a generalised version of risk-neutralisation is required, found Duan (1995). Therefore, he introduced the locally risk-neutral valuation relationship (LRNVR), which is different in terms of variances, as the LRNVR states that the one-step-ahead conditional variance is invariant with respect to a change to the risk-neutralised measure. This assumption is made since GARCH models are used to model the one-period-ahead conditional variance.

Hao and Zhang (2013) used the GARCH option pricing model on the S&P 500 and compared it against the CBOE VIX. First, they applied MLE with the return data to obtain parameter estimates. However, they found that the implied VIX was consistently lower than the CBOE VIX. The undervaluation had a mean error in the range of 3.47 to 3.78, which is about the same as the volatility risk premium, around 3.3 (Hao and Zhang, 2013).

Therefore, they also applied MLE with the CBOE VIX itself, in line with Fassas and Siriopoulos (2021), who found that implied volatility offers insights into future volatility that exceed the information available from past volatility. This led to a better fit; however, the parameters were distorted to match the level of the CBOE VIX, resulting in an overestimation of the equity risk premium. Additionally, they combined these two approaches, obtaining a joint maximum likelihood. These joint maximum likelihood parameters were less distorted; nonetheless, they still failed to reflect the statistical properties of the CBOE VIX. This indicates that while combining information from both returns and implied volatility can improve model fit, challenges remain in accurately capturing the equity risk premium and the underlying statistical properties.

The first to point out that GARCH option pricing under the LRNVR has poor pricing and hedging performance were Chernov and Ghysels (2000) and Christoffersen and Jacobs (2004). They highlighted the limitations of the LRNVR in accurately pricing options. Barone-Adesi et al. (2008) identified that the restrictions imposed by the LRNVR caused these issues. They took a non-parametric approach, using filtered historical innovations to address the problem. Additionally, Christoffersen et al. (2013) developed a new pricing kernel that allows for a volatility premium, providing a more flexible and accurate framework for option pricing under GARCH models.

In line with this, Zhang and Zhang (2020) proposed a modified LRNVR (mLRNVR) to address the inadequacies of the original LRNVR, which also allowed for a volatility risk premium. They proposed a risk-neutral measure that has different conditional volatilities than the real-world measure. Specifically, this modified measure is adjusted to be more persistent, enhancing its ability to capture the volatility risk premium. They found that under the mLRNVR, GARCH implied VIX fits the CBOE VIX and its statistical properties.

GARCH implied option pricing facilitates the computation of conditional volatility. However, to compare this against a volatility index, extrapolation is required, which is achieved by the introduction of the square-root stochastic autoregressive volatility (SR-SARV) model. The parametric SR-SARV model was first introduced by Andersen (1994) and subsequently enhanced to be semi-parametric by Meddahi and Renault (2004). If the underlying index follows an SR-SARV process under the LRNVR, an analytic formula for the GARCH implied VIX can be derived (Hao and Zhang, 2013).

The GARCH option pricing model can utilise several types of GARCH models. GARCH models, pioneered by Engle (1982) and Bollerslev (1986), excel in their ability to capture the properties of asset returns, such as excess kurtosis and skewness. To further enhance the GARCH model, several different versions have been proposed, such as exponential GARCH by Nelson (1991), GJR-GARCH by Glosten et al. (1993), non-linear asymmetric GARCH by Engle (1993), and component GARCH by Ding and Granger (1996).

This paper will cover the GARCH$(1,1)$ and EGARCH$(1,1)$[3] models, as well as the GARCH$(1,1)$ model under the modified LRNVR, denoted by mGARCH. The GARCH model is used as a benchmark, while the EGARCH and mGARCH models are used as extensions. The EGARCH model is chosen because Hao and Zhang (2013) found it to be the best fit among all models. In the following section, the data is discussed. The third section introduces the methodology of the models. The results are presented in the fourth section, and the final section provides a conclusion.

2 Data

In this paper, several datasets are applied[4]. There are four stock indices with their corresponding volatility indices: The NASDAQ (CBOE VIX), the Nikkei 225 (Nikkei 225 VI), the FTSE 100 (VFTSE) and the DAX (VDAX). The country's 3-month government bond rate is used as the risk-free rate for each stock index. The different indices are analysed from May 6, 2012, to May 26, 2019[5]. Furthermore, the CBOE VIX is also analysed at a forecast horizon of 1-day, 9-day, 1-month, 3-month, 6-month, and 1-year forecasts. The maturities are analysed from January 4,

[3] The implied volatility under the EGARCH model is computed using the square-root stochastic exponential autoregressive volatility (SR-SEARV) model, a slightly different version of the SR-SARV, which is consistent with the exponential form of the EGARCH model.

[4] The data was found online through public sources, such as the website of the Chicago Board Options Exchange.

[5] There are 1677 observations in the dataset of the indices

2011, until April 22, 2022, except for the 1-day maturity, which starts on May 13, 2022[6].

3 Methodology

3.1 GARCH option pricing

Consider a discrete-time economy where X_t denotes the price of the asset at time t. Duan (1995) proposed a linear GARCH process for modelling the asset and the options written on it. Furthermore, he proposed that the one-period-ahead rate of return of this asset is log-normally distributed, that is:

$$\ln(\frac{X_t}{X_{t-1}}) = r + \lambda_1\sqrt{h_t} + -\frac{1}{2}h_t + \varepsilon_t \tag{1}$$

where ε_t follows a zero mean distribution with conditional volatility h_t, given information set J_{t-1}. In this paper, the conditional volatility is modelled using either a GARCH$(1,1)$ or an EGARCH$(1,1)$ model, for which the formulas under the real-world measure P are presented in Equation (2) and Equation (3), respectively. Moreover, the risk-free rate is denoted by r, and λ_1 is considered the equity risk premium.

$$h_t = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta_1 h_{t-1} \tag{2}$$

$$\ln(h_t) = \alpha_0 + \alpha_1 \ln(h_{t-1}) + g(z_{t-1}), \quad z_{t-1} = \varepsilon_{t-1}/\sqrt{h_{t-1}} \tag{3}$$

$$g(z_{t-1}) = \alpha_1 z_{t-1} + \kappa(\mid z_{t-1} \mid - \sqrt{2/\pi}) \tag{4}$$

3.1.1 Locally Risk Neutral Valuation Relation (LRNVR)

Following Hao and Zhang (2013), a locally risk-neutral measure Q is utilised instead of the real-world probabilities P. Under Q, the expected return must equal the risk-free rate. Furthermore, the one-day-ahead variance must be the same under both measures to ensure the feasibility of the estimation, given that only under P is the conditional variance observable. Combining this with the fact that the conditional mean can be replaced with the risk-free rate results in a well-specified model that does not locally depend on preferences. However, the LRNVR cannot fully eliminate preferences; nonetheless, they are encapsulated in the risk-premium parameter λ_1. Applying the LRNVR yields Equation (5) for the one-period-ahead rate of return. The conditional variances for the GARCH$(1,1)$ and EGARCH$(1,1)$ models under the LRNVR are

[6] There are 3345 observations in the dataset for the maturities, except for the 1-day maturity, for which there are 487 observations.

described in Equation (6) and Equation (7) respectively.

$$\ln(\frac{X_t}{X_{t-1}}) = r - \frac{1}{2}h_t + \xi_t \tag{5}$$

$$h_t = \alpha_0 + \alpha_1\big(\xi_{t-1} - \lambda_1\sqrt{h_{t-1}}\big)^2 + \beta_1 h_{t-1} \tag{6}$$

$$\ln(h_t) = \alpha_0 + \beta_1 \ln(h_{t-1}) + g(u_{t-1} - \lambda_1)^7, \quad u_t = \xi_t/\sqrt{h_t} \tag{7}$$

Definition 3.1 *A pricing measure Q is said to satisfy to locally risk-neutral valuation relationship (LRNVR) if measure Q is mutually absolutely continuous with respect to measure P, $X_t/X_{t-1} \mid J_{t-1}$ distributes lognormally (under Q),*

$$E^Q[X_t/X_{t-1} \mid J_{t-1}] = e^r$$

and

$$Var^Q(\ln(X_t/X_{t-1} \mid J_{t-1}) = Var^P(\ln(X_t/X_{t-1} \mid J_{t-1})$$

almost surely.

3.1.2 modified LRNVR

Hao and Zhang (2013) found the LRNVR inadequate for capturing the statistical properties of the CBOE VIX, attributing this to the omission of the volatility risk premium. Therefore, Zhang and Zhang (2020) proposed a modified LRNVR, in which an additional parameter λ_2 is incorporated to capture the volatility risk premium and increase the persistence of the conditional volatility. The returns are computed under the mLRNVR as in Equation (5), and below, the GARCH model under the mLRNVR is described.

$$h_t = \alpha_0 + \alpha_1\big(\xi_{t-1} - \lambda_1\sqrt{h_{t-1}}\big)^2 + (\beta_1 - \sqrt{2}\alpha_1\lambda_2)h_{t-1} \tag{8}$$

3.2 Implied Volatility

The market volatility denotes the expected volatility of the underlying index over several days, commonly 21 trading days. In this paper, the implied volatility is computed using GARCH option pricing; however, the GARCH models operate on a daily frequency. To overcome this discrepancy, a daily proxy is employed. This daily proxy has the following relation: $\text{Vix}_t = \frac{1}{252}\Big(\frac{\text{VIX}_t}{100}\Big)^2$, where Vix is the daily proxy of the implied volatility, denoted by VIX.

[7] For function $g(\cdot)$, see Equation (4).

In this paper, the implied volatility is approached by taking the expected arithmetic average of the variance over the next n subperiods of the next τ_0 trading days as described in Equation (9). As the GARCH models operate on a daily frequency, $n = \tau_0$ and $t + \dfrac{\tau_0 k}{n}$ reduces to $t + k$.

$$\left(\frac{\text{VIX}_t}{100}\right)^2 = \frac{1}{n}\sum_{k=1}^{n} \text{E}_t^Q\left[\tilde{h}_{t+\dfrac{\tau_0 k}{n}}\right] \tag{9}$$

$$\text{Vix}_t = \frac{1}{n}\sum_{k=1}^{n} \text{E}_t^Q\left[\tilde{h}_{t+k}\right] \tag{10}$$

In these formulas, $\tilde{h}_s$ is the instantaneous annualised variance of the rate of return of the underlying index. If the underlying follows an SR-SARV(p) model, as developed by Meddahi and Renault (2004), under the LRNVR, then an analytic formula for the Vix can be derived. In line with the GARCH(p, p) models, an SR-SARV(p) model is proposed, where p denotes the number of lags in the VAR(p) sub-model.

Definition 3.2 *Discrete time SR-SARV(p) model*
A stationary square-integrable process $\{\varepsilon_t, t \in \mathbb{Z}\}$ 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 w.r.t. J_{t-1}, that is $E[\varepsilon_t|J_{t-1}] = 0$,

(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:

$$f_t = \text{Var}[\varepsilon_{t+1}|J_t] = e' F_t, \tag{11}$$

$$F_t = \Omega + \Gamma F_{t-1} + V_t, \quad \textit{with } E[V_t|J_{t-1}] = 0, \tag{12}$$

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

Proposition 3.1 *If the underlying follows a SR-SARV(p) process under the LRNVR Q proposed by Duan (1995), then the implied VIX at time t is affine in F_t, i.e.,*

$$Vix_t = \zeta + \Psi F_t, \quad \Psi \in \mathbb{R}^p \tag{13}$$

In particular, if $p = 1$ (then $e = 1$), the implied VIX at time t is a linear function of the conditional variance of the next period,

$$Vix_t = \zeta + \psi f_t, \quad \psi \in \mathbb{R}, \tag{14}$$

where

$$\zeta = \frac{\Omega}{1-\Gamma(1-\psi)},$$
$$\psi = \frac{1-\Gamma^n}{n(1-\Gamma)}$$

Proof: See Appendix (A).

Proposition 3.2 *Let $\xi_t, t \in Z$ be a m.d.s. with the conditional variance $h_t \equiv Var[\xi_t \mid \xi_\tau, \tau \leq t-1]$ under the LRNVR. If h_t is given by Equation* (6) *or Equation* (8), *then ξ_t is a SR-SARV*(1) *process.*
Proof: See Appendix (A).

Applying propositions 3.1 and 3.2, the Vix_t of the GARCH model can be described as:

$$\begin{aligned} Vix_t &= A + Bh_{t+1} \\ A &= \frac{\alpha_0}{1-\eta}(1-B) \\ B &= \frac{1-\eta^n}{n(1-\eta)} \\ \eta &= \alpha_1(1+\lambda^2) + \beta_1 \\ \eta^{*8} &= \alpha_1(1+\lambda^2) + \beta_1 - \sqrt{2}\alpha_1\lambda_2 \end{aligned}$$

The EGARCH model is not incorporated into the class of SR-SARV models, therefore, the square-root stochastic exponential autoregressive volatility model is introduced below.

Definition 3.3 *Discrete time SR-SEARV(1) model*
A stationary square-integrable process $\{\epsilon_t, t \in \mathbb{Z}\}$ is called a SR-SEARV(1) process with respect to a filtration $\mathcal{J}_t, t \in \mathbb{Z}$, if:

(i) ε_t is a martingale difference sequence w.r.t. $\mathcal{I}_{t-1}$, that is $\mathbb{E}[\varepsilon_t|\mathcal{J}_{t-1}] = 0$,

(ii) the logarithm of the conditional variance process f_t of ϵ_{t+1} given $\mathcal{J}_t$ is a stationary $\mathcal{J}_t$-adapted AR(1):

$$\ln f_t = \omega + \gamma \ln f_{t-1} + v_t, \quad \textit{with } v_t \textit{ i.i.d.} \tag{15}$$

where $|\gamma| < 1$.

[8] η has a slightly different form when using the mLRNVR instead of the LRNVR, denoted here by an asterisk.

Proposition 3.3 *If the underlying follows an SR-SEARV*(1) *process under the LRNVR proposed by Duan (1995), then the implied VIX is a polynomial function of the conditional variance of the next period,* h_{t+1}.

Proposition 3.4 *If* $\{\xi_t, t \in Z\}$ *is a m.d.s. under the LRNVR with the conditional variance* $h_t \equiv Var[\xi_t \mid \xi_\tau, \tau \leq t-1]$ *given by Equation* (7) *and* $u_t = \xi_t/\sqrt{h_t}$ *i.i.d., then* ξ_t *is a SR-SEARV(1) process.*

Applying propositions 3.3 and 3.4, the implied Vix for the EGARCH$(1,1)$ model is as follows:

$$Vix_t = \frac{1}{n}\left[h_{t+1} + \sum_{k=1}^{n-1}\left(\prod_{i=0}^{k-1} l_i\right) h_{t+1}^{\beta_1^k}\right]$$

$$l_i = e^{\beta_1^i(\alpha_0 - \kappa\sqrt{2/\pi})}\left\{e^{-\beta_1^i(\alpha_1-\kappa)\lambda + \frac{[\beta_1^i(\alpha_1-\kappa)]^2}{2}} N[\lambda - \beta_1^i(\alpha_1-\kappa)]\right.$$

$$\left. + e^{-\beta_1^i(\alpha_1+\kappa)\lambda + \frac{[\beta_1^i(\alpha_1+\kappa)]^2}{2}} N[\beta_1^i(\alpha_1+\kappa) - \lambda]\right\}$$

3.3 Estimation

The estimation of the parameters utilises MLE using one of three methods. The first is MLE using the returns, the second is using the volatility index and the third is using both the returns and the volatility index.

For the returns, the log-likelihood is computed as follows:

$$\ln(L_R) = -\frac{T}{2}\ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\left(\ln(h_t) + \left[\ln\left(\frac{X_t}{X_{t-1}}\right) - r - \lambda\sqrt{h_t} + \frac{1}{2}h_t\right]^2 / h_t\right) \tag{16}$$

Here, h_t is computed using the corresponding GARCH model. As for MLE with the volatility index, a difference has to be made between the implied volatility and the market volatility, as the daily return innovation determines the current price and the conditional variance of the next period. Therefore, an error term μ is added to allow for a difference. This error term follows an $i.i.d$ normal distribution with variance s^2, estimated by: $\hat{s}^2 = \text{Var}(VIX^{MKT} - VIX^{Imp})$.

$$\text{VIX}^{IND} = \text{VIX}^{IMP} + \mu, \qquad \mu \sim i.i.d.\ N(0, s^2) \tag{17}$$

$$\ln(L_V) = -\frac{T}{2}\ln(2\pi\hat{s}^2) - \frac{1}{2\hat{s}^2}\sum_{t=1}^{T}\left(VIX_t^{Mkt} - VIX_t^{Imp}\right)^2 \tag{18}$$

When optimising using both returns and the volatility index, the total log-likelihood is:

$$\ln(L_T) = \ln(L_R) + \ln(L_V)$$

When maximising the log-likelihood, the constraints under measure Q are applied, as they are stricter than the constraints under measure P. The constraints under measure Q are denoted below and are necessary to ensure stationarity.

$$\begin{aligned}
\text{GARCH under LRNVR} &= \alpha_1(1+\lambda_1^2) + \beta_1 < 1 \\
\text{GARCH under mLRNVR} &= \alpha_1(1+\lambda_1^2) + \beta_1 - \sqrt{2}\alpha_1\lambda_2 < 1 \\
\text{EGARCH under LRNVR} &= |\beta_1| < 1
\end{aligned}$$

4 Results

In this section, first the GARCH and EGARCH parameters and performance under the LRNVR are compared against the values found in Hao and Zhang (2013). Then the indices and maturities are evaluated and compared using three distinct models: GARCH$(1,1)$ and EGARCH$(1,1)$ under the LRNVR, and GARCH$(1,1)$ under the mLRNVR, denoted as the mGARCH model. These models are all maximised using both the returns and the corresponding volatility index. Additionally, in Appendix B the mGARCH and EGARCH models are plotted against their volatility indices for both the indices and the maturities.

4.1 Reproduction

To evaluate the performance of the models, they are first compared against the results from Hao and Zhang (2013)[9]. Due to potential differences in optimisation procedures and slight variations in data, the results might not be identical. Nonetheless, no significant discrepancies are found, as the largest difference is 60% of the standard deviation for α_0 in the EGARCH model optimised by both the returns and the VIX. However, the majority of differences fall within the range of 10% to 20% of the standard deviation. In Table 1, the parameter estimates under the LRNVR are presented, which are applied to the same data as in Hao and Zhang (2013). Moreover, the performance measures exhibit only slight differences. Only the kurtosis is consistently lower than reported by Hao and Zhang (2013).

[9] Hao and Zhang (2013) utilised several GARCH models; however, only the GARCH$(1,1)$ and EGARCH$(1,1)$ models are applied in this paper.

For the GARCH model, the biggest difference when maximising with the VIX instead of the returns is in the parameter λ_1, which changes from 0.0518 to 0.774, indicating a significant increase in the equity risk premium. Moreover, when estimating solely with the return, the mean error was in the range of 3.5 to 3.7, almost consistent with the volatility premium in standard deviation unit of around 3.3 (Hao and Zhang, 2013). In the EGARCH model with VIX estimation, a negative value of -0.0678 is observed for λ_1, while slightly positive values are noted for the returns and joint optimisation.

The performance of the GARCH and EGARCH models improves when estimating with the VIX, with the mean error decreasing from around 3.5 to 0.1. However, the parameters become distorted when estimating solely with the volatility index, as the equity risk premium is overvalued. This indicates that estimation using both the returns and the volatility index is superior. Additionally, the EGARCH model has the best model fit, as the errors are lower and the statistical properties are closer to those of the CBOE VIX.

		α_0	α_1	β_1	κ	λ_1	Ret	VIX	Both
GARCH									
Ret	Hao	7.245E-07	0.0632	0.9312	$\sim$	0.0523	**16097**	20552	36649
	Mald	7.123E-07	0.0638	0.9309	$\sim$	0.0529	**16059**	20566	36625
	std	(1.579E-07)	(0.0065)	(0.0068)	$\sim$	(0.0139)			
VIX	Hao	1.711E-06	0.0366	0.9387	$\sim$	0.7914	14472	**23873**	38345
	Mald	1.715E-06	0.0367	0.9390	$\sim$	0.7859	14455	**23832**	38287
	std	(4.121E-08)	(0.0009)	(0.0014)	$\sim$	(0.0300)			
Both	Hao	1.675E-06	0.0473	0.9498	$\sim$	0.2068	15872	23634	**39507**
	Mald	1.687E-06	0.0471	0.9499	$\sim$	0.2067	15832	23592	**39423**
	std	(4.652E-08)	(0.0011)	(0.0012)	$\sim$	(0.0122)			
EGARCH									
Ret	Hao	-0.1329	-0.0921	0.9854	0.1105	0.0167	**16169**	20492	36661
	Mald	-0.1354	-0.0933	0.9851	0.1119	0.0165	**16132**	20452	36584
	std	(0.0166)	(0.0078)	(0.0018)	(0.0095)	(0.0142)			
VIX	Hao	-0.0761	-0.0641	0.9895	0.0906	-0.0690	16002	**24478**	40480
	Mald	-0.0770	-0.0639	0.9894	0.0901	-0.0678	15964	**24432**	40395
	std	(0.0021)	(0.0016)	(0.0002)	(0.0019)	(0.0176)			
Both	Hao	-0.0838	-0.0614	0.9891	0.0955	0.0096	16038	24468	**40506**
	Mald	-0.0849	-0.0612	0.9890	0.0952	0.0113	15997	24424	**40420**
	std	(0.0018)	(0.0016)	(0.0002)	(0.0101)	(0.0021)			

Table 1: A comparison of the MLE parameters of Hao and Zhang (2013), referred to as "Hao", and those discussed in this paper, referred to as "Mald". The data range utilised in this study is identical to that of Hao and Zhang (2013), covering the period from January 2, 1990 to August 8, 2009[10]. The models used are the GARCH(1, 1) and the EGARCH(1, 1) under the LRNVR. The likelihood that is maximised is presented in the first column, with "Ret" denoting Returns, "VIX" referring to VIX data, and "Both" representing a joint log-likelihood of the Returns and VIX data. The maximised log-likelihood values are indicated in boldface. The standard deviation is denoted in parentheses, and since the GARCH model does not include the κ parameter, it is represented by "$\sim$".

[10]This dataset consists of 4938 observations.

Model and Data		ME	Std.Err.	MAE	MSE	RMSE	P-value	Violation of one-sigma band	Corr.Coef.	AR1	AR10	AR30	Var	Skew	Kurt
GARCH															
	Hao	3.63	3.31	4.02	24.13	4.91	0.00	7.86%	0.92	0.9943	0.9355	0.7751	70.64	3.10	17.06
Ret	Mald	3.65	3.29	4.05	24.18	4.92	0.00	7.78%	0.92	0.9949	0.9400	0.7855	71.88	3.05	13.58
	Hao	0.12	3.08	2.36	9.51	3.08	0.48	1.61%	0.93	0.9961	0.9551	0.8221	65.23	3.27	17.66
Vix	Mald	0.10	3.08	2.35	9.48	3.08	0.55	1.60%	0.93	0.9962	0.9552	0.8234	66.03	3.25	14.54
	Hao	0.26	3.23	2.39	10.47	3.24	0.13	2.30%	0.92	0.9967	0.9554	0.8162	66.80	3.27	17.95
Both	Mald	0.26	3.22	2.39	10.45	3.23	0.12	2.29%	0.92	0.9967	0.9555	0.8171	66.90	3.26	14.84
EGARCH															
	Hao	3.62	3.12	3.76	22.81	4.78	0.00	8.08%	0.94	0.9889	0.9068	0.7457	47.16	2.19	10.92
Ret	Mald	3.54	3.11	3.69	22.20	4.71	0.00	7.65%	0.94	0.9887	0.9052	0.7444	47.97	2.18	7.90
	Hao	0.00	2.73	2.10	7.45	2.73	1.00	0.75%	0.95	0.9953	0.9510	0.8295	63.13	2.17	10.53
Vix	Mald	0.02	2.73	2.09	7.44	2.73	0.89	0.75%	0.95	0.9953	0.9508	0.8301	63.06	2.16	7.49
Both	Hao	0.09	2.73	2.10	7.48	2.73	0.57	0.71%	0.95	0.9949	0.9475	0.8203	64.04	2.18	10.65
	Mald	0.09	2.73	2.09	7.46	2.73	0.58	0.77%	0.95	0.9949	0.9473	0.8211	64.26	2.17	7.61
CBOE	Hao									0.9844	0.9162	0.7846	70.65	2.06	10.26
VIX	Mald									0.9845	0.9164	0.7868	70.80	2.06	7.25

Table 2: Comparison of various statistics from the paper of Hao and Zhang (2013) and the parameters found in this paper, using the data from January 2, 1990 to August 10, 2009. The likelihood that is maximised is presented in the first column, with "Ret" denoting Returns, "VIX" referring to VIX data, and "Both" representing a joint log-likelihood of the Returns and VIX data. In the second column, "Hao" refers to values from Hao and Zhang (2013), while "Mald" indicates values from the optimisation performed in this paper. The error is determined by subtracting the implied VIX from the CBOE VIX. The mean error (ME) represents the daily average difference between the implied VIX and the CBOE VIX. The standard error (Std.Err.) measures the standard deviation of this error. The mean absolute error (MAE) calculates the daily average of the absolute differences between the implied VIX and the CBOE VIX. The mean squared error (MSE) computes the daily average of the squared differences, while the root mean squared error (RMSE) is the square root of the MSE. The P-value tests the null hypothesis that the means of the implied VIX and the CBOE VIX are equal. The violation of the one-sigma band indicates the probability that the implied VIX falls outside the one-standard-deviation range of the CBOE VIX. The correlation coefficient (Corr.Coef.) assesses the linear relationship between the implied VIX and the CBOE VIX. Additionally, autocorrelation coefficients for lags of 1, 10, and 30 days, as well as higher moments, are reported.

4.2 Indices

In Table 3, the parameters for the indices are displayed. For the S&P, the mGARCH model shows the highest log-likelihood, indicating the best fit. However, for the remaining three indices, the EGARCH model provides the highest log-likelihood. Additionally, the log-likelihood of the GARCH model is consistently close to that of the mGARCH model across all indices. However, as the mGARCH model provides a higher log-likelihood and a better fit, this model is preferred over the standard GARCH model. Therefore, in the following section, the parameters for the mGARCH models are discussed. Nonetheless, as they are similar to the GARCH model they provide insights into the GARCH model as well.

When comparing the mGARCH models, the S&P demonstrates the highest sensitivity to innovation and the least persistence, with $\alpha_1 = 0.1207$ and $\beta_1 = 0.8046$. In contrast, the Nikkei exhibits an α_1 of 0.0180 and a β_1 of 0.9623, indicating lower sensitivity and greater persistence. The DAX also shows greater persistence than the S&P, while the FTSE exhibits similar characteristics. Furthermore, the equity risk premium, λ_1, is highest for the S&P and lowest for the Nikkei, which appears to correspond with the level of persistence, exhibiting a negative relationship.

The volatility risk premium, λ_2, does not appear to be related to the other parameters. For example, the lowest values are found for the S&P and the Nikkei, with -0.2119 and -0.2355, respectively. The DAX and FTSE have higher values, with -0.1696 and -0.1370. This discrepancy suggests that the S&P and Nikkei are associated with higher levels of risk aversion among investors compared to the DAX and FTSE.

The EGARCH models all show a high level of persistence, ranging from 0.98 to 0.99. The sensitivity to innovation is greatest for the S&P 500 and FTSE, lowest for the Nikkei, and intermediate for the DAX, similar to the other models. Note that the equity risk premium is negative for the FTSE, and not significantly different from zero for the S&P 500.

In Table 4, the performance measures for the different indices and models are displayed. In line with the log-likelihoods, the EGARCH model performs worse than the GARCH models for the S&P. However, for the other indices, the mGARCH model outperforms the EGARCH model in terms of mean error but is outperformed in terms of all other error measures. This indicates that, on average, the mGARCH model provides a better fit, but the EGARCH model captures the peaks more accurately. This is due to the higher sensitivity to shocks of the EGARCH model's structure, which utilises an exponential function.

When comparing the statistical properties of the different models with their underlying volatility indices, the EGARCH model shows a higher level of autocorrelation compared to the mGARCH model, which is generally more aligned with the underlying volatility index. Furthermore, the EGARCH model has a higher variance, and a lower skewness and kurtosis, providing a better fit for most indices except for the S&P 500. However, none of the models are able to consistently capture all of the statistical properties across all indices.

Index	Model	α_0	α_1	β_1	κ	λ_1	λ_2	ML
S&P 500	Garch	3.688E-06	0.1287	0.8311	∼	0.3626	∼	14761
		(1.247E-07)	(0.0051)	(0.0064)	∼	(0.0180)	∼	
	mGARCH	3.575E-06	0.1207	0.8046	∼	0.3685	-0.2119	**14781**
		(1.211E-07)	(0.0047)	(0.0083)	∼	(0.0182)	(0.0356)	
	EGARCH	-0.1131	-0.0725	0.9857	0.1136	0.0006	∼	14126
		(0.0137)	(0.0140)	(0.0014)	(0.0083)	(0.0108)	∼	
Nikkei	Garch	2.459E-06	0.0169	0.9703	∼	0.0534	∼	13206
		(1.667E-07)	(0.0009)	(0.0017)	∼	(0.0229)	∼	
	mGARCH	2.677E-06	0.0180	0.9623	∼	0.0546	-0.2355	13218
		(1.84E-07)	(0.0010)	(0.0026)	∼	(0.0229)	(0.0521)	
	EGARCH	-0.1314	-0.0042	0.9843	0.0492	0.0543	∼	**13359**
		(0.0048)	(0.0021)	(0.0006)	(0.0018)	(0.0229)	∼	
DAX	Garch	2.188E-06	0.0483	0.9391	∼	0.2356	∼	13652
		(1.566E-07)	(0.0028)	(0.0037)	∼	(0.0202)	∼	
	mGARCH	2.255E-06	0.0480	0.9275	∼	0.2392	-0.1696	13664
		(1.597E-07)	(0.0028)	(0.0047)	∼	(0.0203)	(0.0351)	
	EGARCH	-0.1589	-0.0427	0.9818	0.0899	0.1016	∼	**13748**
		(0.0010)	(0.0009)	(0.0001)	(0.0009)	(0.0010)	∼	
FTSE	Garch	2.506E-06	0.0926	0.8802	∼	0.2781	∼	14318
		(1.351E-07)	(0.0044)	(0.0056)	∼	(0.0200)	∼	
	mGARCH	2.603E-06	0.0923	0.8616	∼	0.2789	-0.1370	14327
		(1.84E-07)	(0.0010)	(0.0026)	∼	(0.0229)	(0.0521)	
	EGARCH	-0.1144	-0.0740	0.9870	0.0918	-0.0197	∼	**14449**
		(0.0066)	(0.0036)	(0.0007)	(0.0221)	(0.0049)	∼	

Table 3: The parameters for the different indices, namely, the S&P 500, the Nikkei 225, the DAX 30, and the FTSE 100. The models are the GARCH$(1,1)$ under the LRNVR and the mLRNVR (denoted as mGARCH), and the EGARCH$(1,1)$ under the LRNVR. The data spans from 6 May 2012 to 26 May 2019, and all models are optimized using both the returns and the corresponding volatility index. The corresponding maximum log-likelihood (ML) values are provided in the final column, with the highest log-likelihood in boldface. The standard deviations of the parameters are denoted in parentheses below the parameters.

4.3 Maturity of CBOE VIX

In Table 3, the parameters for the different maturities and models are displayed. Important to note is that the 1-day maturity does not use the same data range, as it was only introduced in 2020. Therefore, the 1-day maturity will only be discussed in the last paragraph. The models are all maximised using the returns and the CBOE VIX with the corresponding maturity.

For the 9-day and 6-month maturities, the EGARCH model provides the highest log-likelihoods.

Model & Index	ME	Std.Err.	MAE	MSE	RMSE	P-value	Violation of one-sigma band	Corr.Coef.	AR1	AR10	AR30	Var	Skew	Kurt
S&P 500									0.9259	0.5588	0.2605	14.17	1.57	4.15
GARCH	0.03	1.83	1.35	3.34	1.83	0.7879	4.23%	0.8744	0.9521	0.5378	0.1949	11.36	2.05	5.21
mGARCH	0.00	1.82	1.34	3.32	1.82	0.9685	4.35%	0.8751	0.9557	0.5553	0.2026	11.22	2.01	4.94
EGARCH	-1.21	2.37	2.05	7.09	2.66	0.0000	16.34%	0.8395	0.9858	0.8021	0.4500	19.07	0.95	0.86
Nikkei									0.9654	0.8416	0.7071	16.58	0.54	-0.49
GARCH	0.09	2.85	2.32	8.14	2.85	0.4598	14.01%	0.7145	0.9944	0.9182	0.6917	9.31	0.94	0.39
mGARCH	0.04	2.85	2.33	8.14	2.85	0.7770	13.66%	0.7141	0.9938	0.9117	0.6728	9.26	0.99	0.51
EGARCH	0.06	2.63	2.14	6.90	2.63	0.6576	10.49%	0.7642	0.9950	0.9221	0.7100	9.96	0.51	-0.26
DAX									0.9576	0.7239	0.5471	20.81	1.08	1.29
GARCH	0.08	2.58	1.84	6.64	2.58	0.5770	5.72%	0.8257	0.9872	0.8509	0.5905	14.98	1.24	1.71
mGARCH	0.03	2.57	1.84	6.62	2.57	0.8247	5.61%	0.8257	0.9871	0.8503	0.5890	14.68	1.25	1.74
EGARCH	0.08	2.50	1.84	6.26	2.50	0.5701	4.59%	0.8363	0.9873	0.8535	0.5898	14.95	0.69	0.48
FTSE									0.9392	0.6521	0.3352	14.18	1.48	2.93
GARCH	-2.08	2.45	2.70	10.36	3.22	0.0000	21.41%	0.7623	0.9895	0.8320	0.5183	9.97	1.71	3.33
mGARCH	0.01	2.28	1.62	5.18	2.27	0.9612	7.39%	0.7968	0.9679	0.6358	0.2300	9.19	2.11	5.75
EGARCH	0.14	2.18	1.52	4.77	2.18	0.2510	5.90%	0.8179	0.9826	0.7902	0.4627	11.00	1.22	1.88

Table 4: The performance measures for the different indices, namely, the S&P 500, the Nikkei 225, the DAX 30, and the FTSE 100. The models are the GARCH$(1,1)$ under the LRNVR, and the mLRNVR (denoted as mGARCH), and the EGARCH$(1,1)$ under the LRNVR. The data spans from 6 May 2012 to 26 May 2019, and all models are optimised using both the returns and the corresponding volatility index. The row of the index displays the statistical properties of the volatility index. For a detailed description of the performance measures, see Table 2

Maturity	Model	α_0	α_1	β_1	κ	λ_1	λ_2	Both
1 Day	Garch	9.190E-13	0.0898	0.9029	~	0.2832	~	3751
		(2.485E-07)	(0.0079)	(0.0088)	~	(0.0363)	~	
	mGARCH	1.50E-12	0.0928	0.9108	~	0.2847	0.0931	**3752**
		(2.84E-07)	(0.0082)	(0.0095)	~	(0.0359)	(0.0408)	
	EGARCH	-0.1012	-0.0699	0.9869	0.0998	-0.0206	~	3621
		(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)	~	
9 Days	Garch	4.460E-06	0.1253	0.8536	~	0.2497	~	26580
		(1.488E-07)	(0.0034)	(0.0038)	~	(0.0140)	~	
	mGARCH	4.48E-06	0.1194	0.8183	~	0.2586	-0.2465	26582
		(1.45E-07)	(0.0032)	(0.0056)	~	(0.0140)	(0.0247)	
	EGARCH	-0.1042	0.0974	0.9845	0.1955	0.2161	~	**25606**
		(0.0009)	(0.0009)	(0.0001)	(0.0011)	(0.0009)	~	
1 Month	Garch	3.657E-06	0.0879	0.8905	~	0.2230	~	26843
		(1.196E-07)	(0.0032)	(0.0038)	~	(0.0149)	~	
	mGARCH	4.136E-06	0.0946	0.8627	~	0.2405	-0.2737	**26906**
		(1.025E-07)	(0.0007)	(0.0008)	~	(0.0009)	(0.0010)	
	EGARCH	-0.1367	-0.0635	0.9836	0.3682	-0.0218	~	26013
		(0.0025)	(0.0024)	(0.0003)	(0.0029)	(0.0030)	~	
3 Months	Garch	1.278E-06	0.0253	0.9703	~	0.1905	~	26658
		(7.450E-08)	(0.0015)	(0.0018)	~	(0.0166)	~	
	mGARCH	3.56E-06	0.0810	0.8661	~	0.2312	-0.2813	**26826**
		(1.14E-07)	(0.0029)	(0.0049)	~	(0.0150)	(0.0276)	
	EGARCH	-0.1280	0.1099	0.9828	0.1911	0.2132	~	26455
		(0.0023)	(0.0050)	(0.0003)	(0.0027)	(0.0044)	~	
6 Months	Garch	5.090E-07	0.0092	0.9906	~	0.1159	~	26693
		(1.304E-08)	(0.0001)	(0.0001)	~	(0.0001)	~	
	mGARCH	1.11E-06	0.0206	0.9616	~	0.1731	-0.4953	26999
		(5.38E-08)	(0.0009)	(0.0019)	~	(0.0169)	(0.0318)	
	EGARCH	-0.0822	-0.0314	0.9907	0.0899	0.1839	~	**27149**
		(0.0017)	(0.0028)	(0.0002)	(0.0025)	(0.0035)	~	
1 Year	Garch	2.897E-07	0.0057	0.9942	~	0.1143	~	26842
		(5.695E-09)	(0.0000)	(0.0000)	~	(0.0000)	~	
	mGARCH	4.41E-07	0.0092	0.9832	~	0.1609	-0.5591	**27088**
		(1.83E-08)	(0.0002)	(0.0006)	~	(0.0172)	(0.0310)	
	EGARCH	-0.0901	-0.0070	0.9900	0.1844	0.2117	~	26943
		(0.0017)	(0.0040)	(0.0001)	(0.0031)	(0.0040)	~	

Table 5: The parameters for the different maturities of the CBOE VIX. The models are the GARCH$(1,1)$ under the LRNVR, the GARCH$(1,1)$ under the mLRNVR (denoted as mGARCH), and the EGARCH$(1,1)$ under the LRNVR. The data spans from 4 January 2011 to 22 April 2022, except for the 1-day maturity, which starts on May 13, 2022, and all models are optimised using both the returns and the CBOE VIX with the corresponding volatility. The log-likelihood values are provided in the final column, with the highest log-likelihood in boldface. The standard deviations of the parameters are denoted in parentheses below the parameters.

In contrast, for the 1-month, 3-month, and 1-year maturities, the mGARCH model has the highest log-likelihoods. When examining the GARCH and mGARCH models, it is apparent that persistence increases and sensitivity to innovation decreases as maturity lengthens, with α_1

decreasing from 0.1194 to 0.0092 and β_1 increasing from 0.8183 to 0.9832 when lengthening the maturity from 9 days to 1 year. This decrease of α_1 is due to the fact that a single day has less impact on the expected volatility of a year than the next 9 days.

In line with the findings from the indices, the equity risk premium decreases with an increase in persistence, and therefore with an increase in maturity. Furthermore, λ_2 displays a similar effect, decreasing as maturity lengthens. This indicates that for longer maturities, investors are willing to pay a larger premium to hedge against volatility increases. Thus the equity risk premium decreases, while the volatility risk premium increases as maturity lengthens.

The EGARCH model shows a persistence parameter around 0.98 to 0.99, with the absolute innovation parameter decreasing when the maturity lengthens, except for the 3-month maturity, which has a higher absolute α_1 than the 1-month maturity.

Among the three models, the mGARCH provides the best fit, regardless of maturity, as evidenced by Table 6. The errors associated with the mGARCH model are lower, the p-value of equal means is higher, and the violation of the one-sigma band is lower for all maturities.

Contrary to the findings for the indices, the EGARCH model is not consistently the highest in terms of autocorrelation across different maturities of the CBOE VIX. Nonetheless, none of the models is able to consistently capture the autocorrelations of the underlying CBOE VIX. In terms of variance, the mGARCH model generally outperforms the EGARCH model. For skewness and kurtosis, neither model consistently outperforms the other across all maturities.

5 Conclusion

This paper analyses the GARCH implied volatility model by Hao and Zhang (2013) and the modified version by Zhang and Zhang (2020) across various indices and maturities. By applying these models, the equity and volatility risk premiums are analysed to determine their behaviour across different indices and various maturities.

First, this paper reproduces the findings of Hao and Zhang (2013) to validate the models used in the subsequent analysis. This reproduction revealed that when estimating using only the returns, the implied VIX lacked approximately the same value as the volatility risk premium, consistent with the findings of Hao and Zhang (2013). Furthermore, when estimating with the VIX data, the parameters become distorted, resulting in an overestimation of the equity risk premium. Therefore, the parameters estimated using both the returns and the VIX data are superior, as they provide a good fit to the data while remaining consistent with the theoretical expectations of the equity risk premium.

Then, the S&P 500, FTSE 100, Nikkei 225, and DAX 30 indices are analysed by

Model & Maturity	ME	Std.Err.	MAE	MSE	RMSE	P-value	Violation of one-sigma band	Corr.Coef.	AR1	AR10	AR30	Var	Skew	Kurt
1 Day									0.8622	0.7150	0.6026	61.90	1.14	0.92
GARCH	-0.21	4.04	2.69	16.35	4.04	0.6525	5.54%	0.8607	0.9824	0.8794	0.7428	42.31	0.78	-0.42
mGARCH	-0.10	4.04	2.65	16.30	4.04	0.8320	5.75%	0.8608	0.9810	0.8726	0.7343	42.28	0.80	-0.37
EGARCH	-3.06	4.06	4.46	25.79	5.08	0.0000	4.93%	0.8712	0.9893	0.8946	0.7240	32.67	0.59	-0.69
9 Days									0.9446	0.7035	0.4212	70.93	3.06	17.41
GARCH	0.27	3.74	2.58	14.09	3.75	0.1699	3.71%	0.8969	0.9811	0.7403	0.3172	63.05	5.03	39.92
mGARCH	0.12	3.74	2.59	13.98	3.74	0.5593	3.50%	0.8970	0.9823	0.7480	0.3235	62.25	5.01	39.49
EGARCH	0.17	4.74	3.29	22.45	4.74	0.3952	5.83%	0.8300	0.9874	0.8401	0.5052	57.74	3.07	16.55
1 Month									0.9673	0.7946	0.5396	50.27	2.49	11.54
GARCH	0.22	3.39	2.45	11.54	3.40	0.1899	5.11%	0.8801	0.9879	0.7887	0.3628	44.23	5.11	39.34
mGARCH	0.11	3.38	2.46	11.44	3.38	0.5031	4.90%	0.8804	0.9891	0.8006	0.3792	43.61	5.04	38.00
EGARCH	-0.36	4.43	3.13	19.71	4.44	0.0745	8.67%	0.8922	0.9686	0.6969	0.3542	87.98	3.11	19.10
3 Months									0.9822	0.8675	0.6778	39.95	1.89	6.60
GARCH	0.21	3.37	2.39	11.39	3.37	0.1541	7.20%	0.8491	0.9979	0.9304	0.6869	33.73	3.54	17.07
mGARCH	0.07	3.34	2.41	11.13	3.34	0.6244	6.31%	0.8506	0.9983	0.9422	0.7308	32.09	3.25	14.19
EGARCH	0.78	3.57	2.67	13.39	3.66	0.0000	8.61%	0.8282	0.9838	0.8275	0.5231	22.62	2.39	10.79
6 Months									0.9877	0.9043	0.7679	32.50	1.33	2.55
GARCH	-0.66	3.07	2.47	9.84	3.14	0.0000	6.46%	0.8439	0.9994	0.9737	0.8655	25.29	2.15	5.61
mGARCH	0.10	3.03	2.18	9.22	3.04	0.4613	7.00%	0.8480	0.9993	0.9736	0.8653	26.22	2.12	5.44
EGARCH	0.33	3.08	2.31	9.60	3.10	0.0098	8.43%	0.8415	0.9952	0.9060	0.6351	22.48	1.88	6.43
1 Year									0.9912	0.9337	0.8345	23.58	0.86	0.05
GARCH	0.33	2.87	2.11	8.35	2.89	0.0027	10.85%	0.8092	0.9996	0.9822	0.9072	18.08	1.65	2.91
mGARCH	0.20	2.88	2.13	8.32	2.88	0.0706	10.46%	0.8064	0.9995	0.9803	0.8975	16.88	1.80	3.61
EGARCH	0.40	3.33	2.64	11.24	3.35	0.0001	13.60%	0.7285	0.9902	0.8414	0.5089	11.50	2.27	10.25

Table 6: The performance measures for the different maturities of the CBOE VIX. The models are the GARCH$(1,1)$ under the LRNVR, and the mLRNVR (denoted as mGARCH), and the EGARCH$(1,1)$ under the LRNVR. The data spans from 4 January 2011 to 22 April 2022, except for the 1-day maturity, which starts on May 13, 2022, and all models are optimised using both the returns and the CBOE VIX with the corresponding volatility. The row of the maturity displays the statistical properties of the maturity. For a detailed description of the performance measures, see Table 2

applying three different models: the GARCH$(1,1)$ under the LRNVR and the mLRNVR, and the EGARCH$(1,1)$ under the LRNVR. This analysis shows that when the persistence of shocks is higher, the equity risk premium decreases, indicating a potential relationship between the persistence of shocks and the equity risk premium. The equity risk premium is highest for the S&P 500, followed by the FTSE 100 and then the DAX 30, while the Nikkei 225 has the lowest equity risk premium. This suggests that investors in the Nikkei 225 demand the lowest excess return for bearing risk.

Furthermore, in the GARCH model under the mLRNVR, the volatility risk premium is consistently negative, aligning with theoretical expectations. However, there does not appear to be a relation with other parameters, as the lowest values are found for the S&P 500 and the Nikkei 225, while the DAX 30 and FTSE 100 have higher values. This indicates that investors in the S&P 500 and Nikkei 225 have a higher level of risk aversion than investors in the FTSE 100 and DAX 30.

The analysis of different maturities indicates that as the maturity lengthens, the persistence of shocks increases. This is accompanied by a decrease in sensitivity to shocks, resulting in a less volatile model. Additionally, as the maturity lengthens, the equity risk premium decreases, indicating that investors are willing to accept lower compensation for bearing risk over longer periods. This again suggests a potential relationship between the persistence of shocks and the equity risk premium.

Similarly, the volatility risk premium displays a decrease as maturity lengthens. This indicates that for longer maturities, investors are willing to pay a larger premium to hedge against volatility increases. Therefore, as maturity lengthens, investors become less susceptible to bearing risk, but more susceptible to an increase in risk.

Overall, the GARCH model under the mLRNVR and the EGARCH model provide a good fit for the underlying volatility index, but they still lack the ability to fully capture its statistical properties. Therefore, more models under the mLRNVR should be applied to gain further insight into the behaviour of these models. Moreover, the relationship between the risk premiums and the persistence of shocks, as well as the possible negative relation between the equity risk premium and the volatility risk premium are topics for further research.

A Appendix A

Proof Proposition 3.1

For $k \geq 1$

$$
\begin{aligned}
E_t^Q[f_{t+k}] &= e' E_t^Q[F_{t+k}] = e' E[\Omega + \Gamma F_{t+k-1} + V_{t+k}] && (11)\&(12) \\
&= e' E_t^Q[\Omega + E_{t+k-1}^Q[\Gamma F_{t+k-1} + V_{t+k}]] \\
&= e' E_t^Q[\Omega + \Gamma F_{t+k-1}] && (12) \\
&= e' E_t^Q[\Omega + \Gamma E_t^Q[F_{t+k-1}]]
\end{aligned}
$$

Iterating:

$$E_t^Q[f_{t+k}] = e' \Big(\sum_{i=0}^{k-1} \Gamma^i \Omega + \Gamma^k F_t \Big)$$

Then with $Vix_t = \dfrac{1}{n} \sum_{k=1}^{n} E_t^Q[h_{t+k}] = \dfrac{1}{n} E_t^Q[f_{t+k}]$, we have:

$$Vix_t = \zeta + \psi F_t$$

with

$$\zeta = \frac{e'}{n} \sum_{k=1}^{n-1} \sum_{i=0}^{k-1} \Gamma^i \Omega,$$

$$\psi = \frac{e'}{n} \sum_{k=1}^{n-1} \Gamma^k,$$

which is affine in F_t. For $p = 1$, we can get VIX_t as a linear function of the conditional variance of the next period, f_t,

$$Vix_t = \zeta + \psi f_t,$$

where

$$
\begin{aligned}
\zeta &= \frac{\Omega}{1-\Gamma}(1-\psi), \\
\psi &= \frac{1-\Gamma}{n(1-\Gamma).}
\end{aligned}
$$

Proof Proposition 3.2

Let $u_t = \xi_t/\sqrt{h_t}$. From Definition 3.2, $h_t = \omega + \gamma h_{t-1} + v_{t-1}$.
If h_t is given by Equation (6), then write:

$$\omega = \alpha_0, \quad \gamma = \alpha_1(1+\lambda_1^2) + \beta_1, \quad v_{t-1} = \alpha_1 h_{t-1}(u_{t-1}^2 - 1 - 2\lambda_1 u_{t-1})$$

If h_t is given by Equation (8), then write:

$$\omega = \alpha_0, \quad \gamma = \alpha_1(1+\lambda_1^2) + (\beta_1 - \sqrt{2}\alpha_1\lambda_2), \quad v_{t-1} = \alpha_1 h_{t-1}(u_{t-1}^2 - 1 - 2\lambda_1 u_{t-1})$$

Proof Proposition 3.3

Let $e^{\gamma^i \omega} E_t^Q(e^{\gamma^i v_{t+1}}) = l_i$. Under the LRNVR Q, the expectation of the conditional variance $k \geq 1$ periods ahead can be expressed as

$$\begin{aligned} E_t^Q(f_{t+k}) &= e^{\omega} E_t^Q \left[E_{t+k-1}^Q \left(f_{t+k-1}^{\gamma} e^{v_{t+k}} \right) \right] \\ &= e^{\omega} E_t^Q(e^{v_{t+k}}) \left[E_{t+k-1}^Q \left(f_{t+k-1}^{\gamma} \right) \right] \\ &= l_0 E_t^Q \left(f_{t+k-1}^{\gamma} \right) \end{aligned}$$

For $0 \leq i \leq k-1$, we have

$$\gamma^i \ln(f_{t+k-i}) = \gamma^i \omega + \gamma^{i+1} \ln(f_{t+k-i-1}) + \gamma^i v_{t+k-i}$$

Thus,

$$\begin{aligned} E_t^Q \left(f_{t+k-i}^{\gamma^i} \right) &= e^{\gamma^i \omega} E_t^Q \left(f_{t+k-i-1}^{\gamma^{i+1}} \right) E_{t+k-i-1}^Q \left(e^{\gamma^i v_{t+k-i}} \right) \\ &= l_i E_t^Q \left(f_{t+k-i-1}^{\gamma^{i+1}} \right) \end{aligned}$$

Iterating:

$$E_t^Q(f_{t+k}) = f_t^{\gamma^k} \prod_{i=0}^{k-1} l_i$$

And the implied VIX formula is

$$Vix_t = \frac{1}{n} \left[f_t + \sum_{k=1}^{n-1} \left(\prod_{i=0}^{k-1} l_i \right) f_t^{\gamma^k} \right]$$

Proof Proposition 3.4

Let $u_t \xi / \sqrt{h_t}$. From Definition 3.3, $\ln(h_t) = \omega + \gamma \ln(h_{t-1}) + v_t$.

If h_t is given by Equation (7), then write:

$$\omega = \alpha_0, \qquad \gamma = \beta_1, \qquad v_t = \alpha_1(u_{t-1} - \lambda_1) + \kappa \left(\mid u_{t-1} - \lambda_1 \mid - \sqrt{2/\pi} \right)$$

B Appendix B

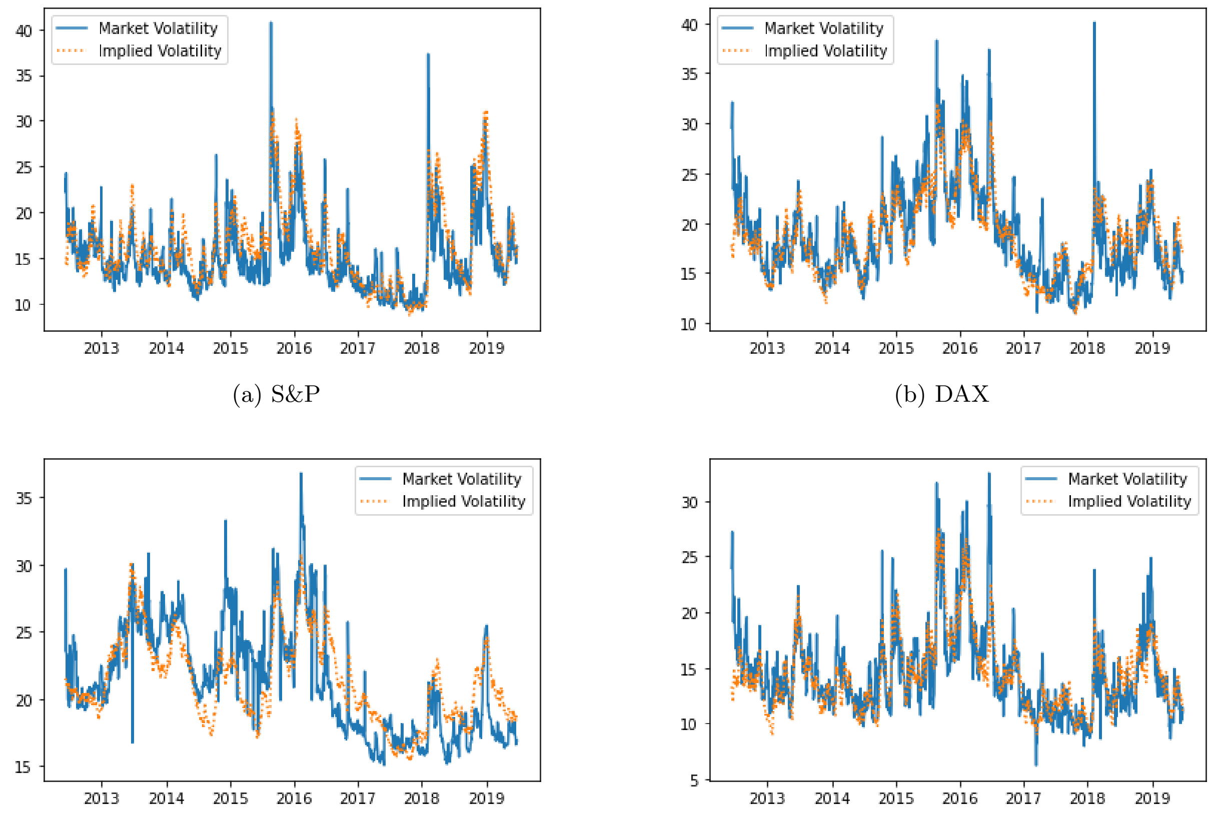

(a) S&P

(b) DAX

(c) Nikkei

(d) FTSE

Figure 1: The EGARCH model for the different indices, optimised using the returns and the VIX. The data spans from May 6, 2012 to May 26, 2019.

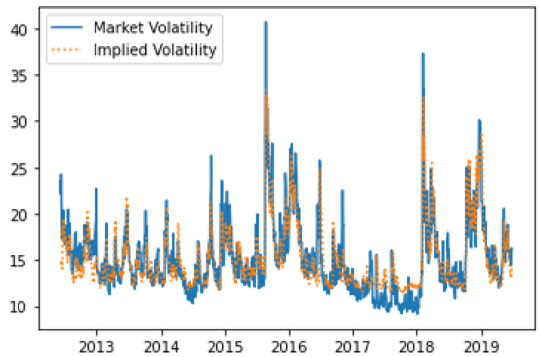

(a) S&P

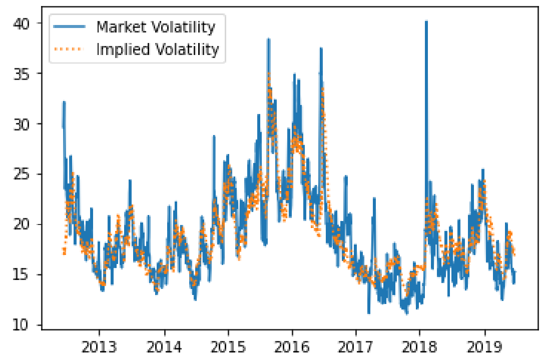

(b) Caption for DAX mod

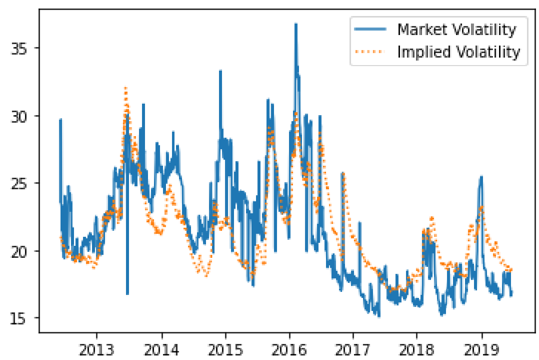

(c) Nikkei

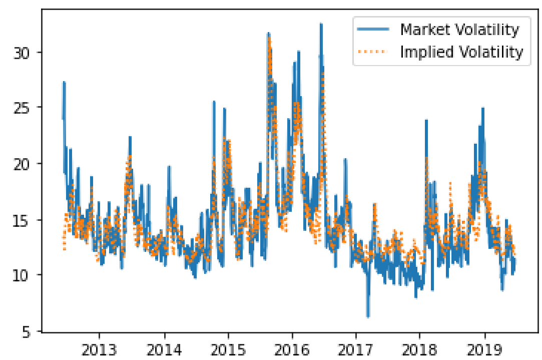

(d) FTSE

Figure 2: The mGARCH model for the different indices, optimised using the returns and the VIX. The data spans from May 6, 2021 to May 26, 2019.

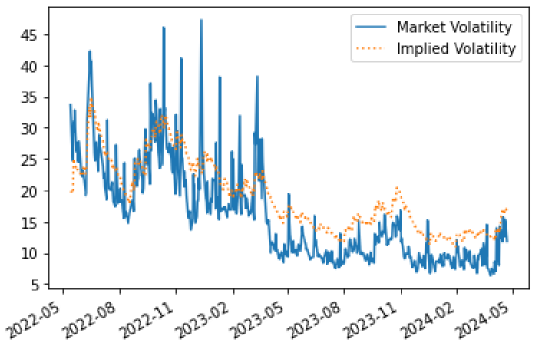

(a) Maturity of one day

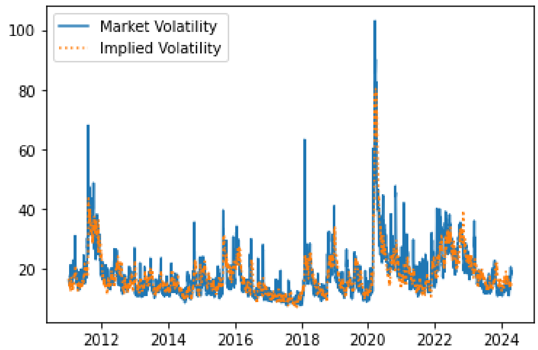

(b) Maturity of nine days

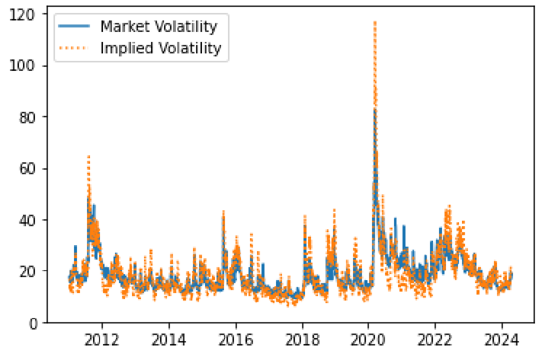

(c) Maturity of one month

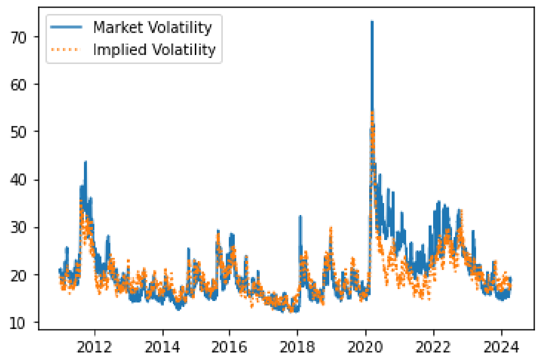

(d) Maturity of three months

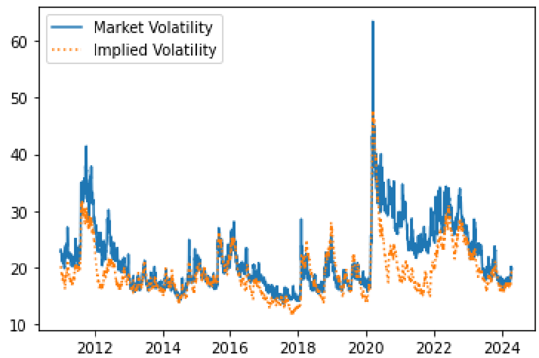

(e) Maturity of six months

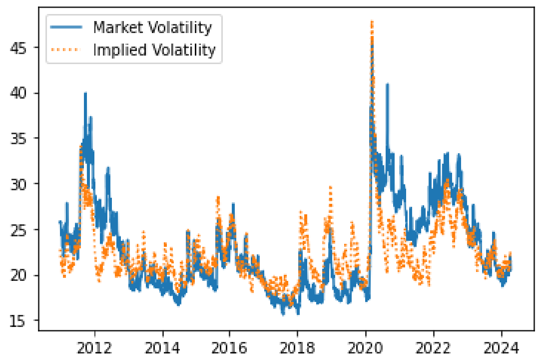

(f) Maturity of one year

Figure 3: Maturities at different time frames using the EGARCH model under the LRNVR. The data spans from 4 January 2011 to 22 April 2022, except for the 1-day maturity that starts on May 13 2022.

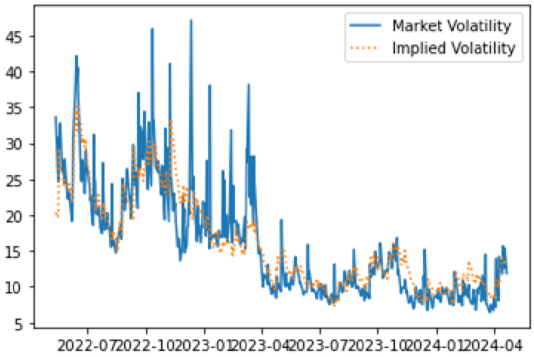

(a) Maturity of one day

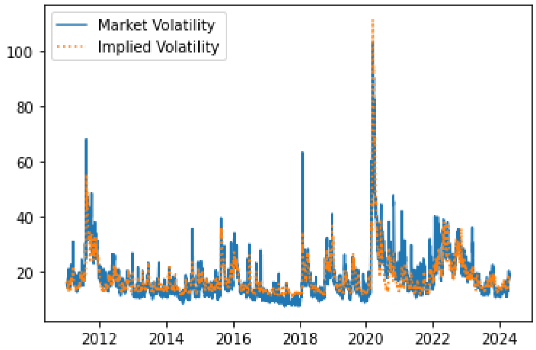

(b) Maturity of nine days

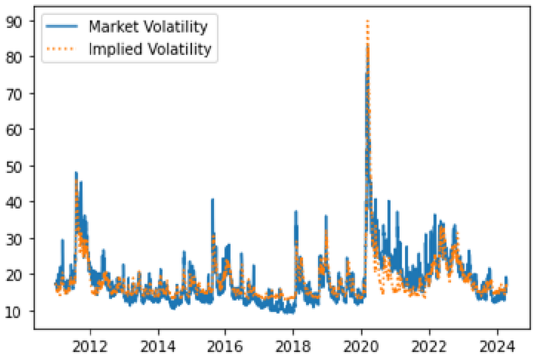

(c) Maturity of one month

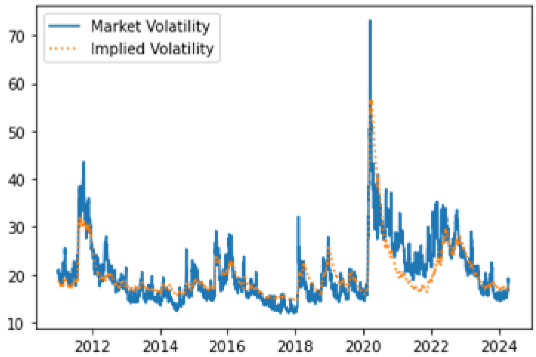

(d) Maturity of three months

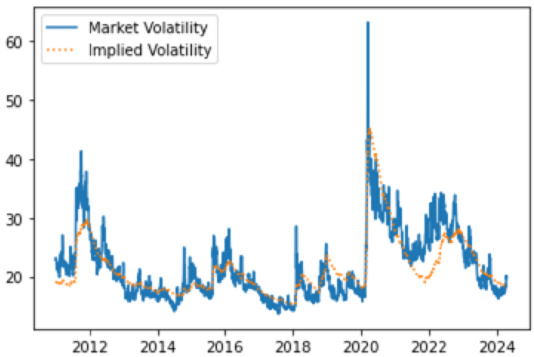

(e) Maturity of six months

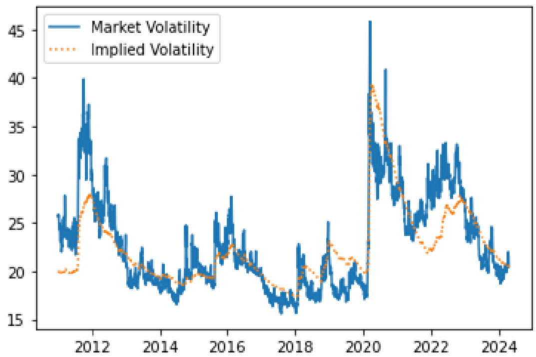

(f) Maturity of one year

Figure 4: Maturities at different time frames using the mGARCH model. The data spans from 4 January 2011 to 22 April 2022, except for the 1-day maturity that starts on May 13 2022.

Appendix C

The code used to obtain the results is structured as follows:

First there is a difference between the main and the function part. The main document is used to run and import everything, while the function document stores all of the methods used. The function document starts with performance measures and data-cleaning functions. Then follows a function that runs the GARCH models along the given time series, and functions to compute the hessian and standard deviations.

Then the GARCH option pricing models are introduced. First, the shocks are computed under measures P and Q. Then follow the functions to update the conditional variance using the different GARCH models. After this, a function is written that computes the implied volatility using the models and the conditional variance. This uses the SR-SARV models to compute the daily proxy and then a function to annualize this to obtain the implied volatility.

Finally, the likelihood functions are displayed. The return and implied volatility log-likelihood are computed as in Hao and Zhang (2013). The joint log-likelihood is differentiated in the GARCH and EGARCH models, as the return function is embedded in the method. The likelihoods are optimised using the trust-const method of SciPy, as they can optimise using the constraints necessary for the models.

References

Andersen, T. G. (1994). Stochastic autoregressive volatility: A framework for volatility modeling. *Mathematical Finance*, 4:75–102.

Andersen, T. G., Fusari, N., and Todorov, V. (2015). Parametric inference and dynamic state recovery from option panels. *Econometrica*, 83(3):1081–1145.

Barone-Adesi, G., Engle, R. F., and Mancini, L. (2008). A GARCH Option Pricing Model with Filtered Historical Simulation. *The Review of Financial Studies*, 21(3):1223–1258.

Bates, D. S. (2000). Post-'87 crash fears in the sp 500 futures option market. *Journal of Econometrics*, 94(1):181–238.

Bates, D. S. (2012). U.s. stock market crash risk, 1926–2010. *Journal of Financial Economics*, 105(2):229–259.

Black, F. (1975). Fact and fantasy in the use of options. *Financial Analysts Journal*, 31(4):36–41.

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

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31(3):307–327.

Carr, P., Geman, H., Madan, D., and Yor, M. (2002). The fine structure of asset returns: An empirical investigation. *The Journal of Business*, 75(2):305–332.

Carr, P. and Wu, L. (2008). Variance Risk Premiums. *The Review of Financial Studies*, 22(3):1311–1341.

Carr, P. and Wu, L. (2020). Option profit and loss attribution and pricing: A new framework. *The Journal of Finance*, 75(4):2271–2316.

Chernov, M. and Ghysels, E. (2000). A study towards a unified approach to the joint estimation of objective and risk neutral measures for the purpose of options valuation. *Journal of Financial Economics*, 56(3):407–458.

Christoffersen, P., Heston, S., and Jacobs, K. (2006). Option valuation with conditional skewness. *Journal of Econometrics*, 131(1):253–284.

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. and Jacobs, K. (2004). Which garch model for option valuation? *Management science*, 50(9):1204–1221.

Demeterfi, K., Derman, E., and Kamal, M. (1999). A guide to volatility and variance swaps. *Journal of derivatives*, 6(4):9–32.

Ding, Z. and Granger, C. W. (1996). Modeling volatility persistence of speculative returns: A new approach. *Journal of Econometrics*, 73(1):185–215.

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

Duffie, D., Pan, J., and Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 68(6):1343–1376.

Engle, R. F. (1982). Autoregressive conditional heteroscedasticity with estimates of the variance of united kingdom inflation. *Econometrica*, 50(4):987–1007.

Engle, Robert F. Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The Journal of Finance*, 48(5):1749–1778.

Fassas, A. P. and Siriopoulos, C. (2021). Implied volatility indices – a review. *The Quarterly Review of Economics and Finance*, 79:303–329.

Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance*, 48(5):1779–1801.

Gultekin, N. B., Rogalski, R. J., and Tinic, S. M. (1982). Option pricing model estimates: Some empirical results. *Financial Management*, 11(1):58–69.

Hao and Zhang (2013). GARCH Option Pricing Models, the CBOE VIX, and Variance Risk Premium. *Journal of Financial Econometrics*, 11(3):556–580.

Heston, S. L. (2015). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *The Review of Financial Studies*, 6(2):327–343.

Hull, J. and White, A. (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance*, 42(2):281–300.

Johnson, H. and Shanno, D. (1987). Option pricing when the variance is changing. *Journal of Financial and Quantitative Analysis*, 22(2):143–151.

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

Merton, R. C. (1973). Theory of rational option pricing. *The Bell Journal of Economics and Management Science*, 4(1):141–183.

Merton, R. C. (1976). Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics*, 3(1):125–144.

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59(2):347–370.

Rubinstein, M. (1976). The valuation of uncertain income streams and the pricing of options. *The Bell Journal of Economics*, 7(2):407–425.

Whaley, R. E. (1982). Valuation of american call options on dividend-paying stocks: Empirical tests. *Journal of Financial Economics*, 10(1):29–58.

Wiggins, J. B. (1987). Option values under stochastic volatility: Theory and empirical estimates. *Journal of Financial Economics*, 19(2):351–372.

Zhang, W. and Zhang, J. E. (2020). Garch option pricing models and the variance risk premium. 13:51.