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

Capturing the Variance Risk Premium at Different Time Horizons using (Realized) GARCH Option Pricing Models

Denny van Daalen (605573)

Supervisor:	E. Vladimirov
Second assessor:	B.P. van der Sluis
Date final version:	1st July 2024

Abstract

In this paper, investigated is whether (Realized) GARCH option pricing models are able to price the CBOE Volatility Index (VIX) at different time horizons. Firstly, the results of Hao and Zhang (2013) are replicated, where the implied VIX is calculated under the locally risk-neutral valuation relationship. The same findings are reported, wherein GARCH models are unable to capture the variance risk premium. Secondly, this research is extended by using an exponentially affine stochastic discount factor for the Realized GARCH model to price the VIX at a nine (VIX9D) and one (VIX1D) day time horizon. The reported results show that the Realized GARCH option pricing model outperforms the GARCH option pricing model for the VIX9D when the CBOE VIX9D is included in estimation. The Realized GARCH option pricing model does not show major fitting improvements regarding the VIX1D, but its coefficients are not distorted to fit the CBOE VIX1D as is for the GARCH model.

1 Introduction

In the field of financial markets, the variance risk premium is known to serve as a measure for the compensation investors require for bearing the risk of fluctuating volatility. This risk, known as volatility risk, was introduced, since a lot of research had shown that volatility is not constant over time. For the pricing of financial options, it meant that the Black-Scholes model by Black and Scholes (1973) was inappropriate for capturing the variance risk premium. Instead, stochastic volatility models and multiple GARCH models would become more adequate.

As most risks in the investing world require a premium, it is captivating to research the price of this volatility risk. Logically, given that volatility risk is adequately compensated in financial markets, the premium that is demanded from investors should be reflected in the prices of options, which are heavily dependent on expected volatility. In the literature, this is often measured by the difference between the realized variance and (risk-neutral) expected or implied variance, which can be synthetically constructed using a portfolio of options (Carr & Wu, 2009).

The existence of the variance risk premium has been studied from multiple perspectives. One of the first questions that arises is the sign of the risk premium. The negative correlation between the returns of indices and its corresponding volatility implies a negative variance risk premium. This is confirmed by Bakshi and Kapadia (2003), who found evidence of a negative risk premium by examining the statistical characteristics of portfolios that implement delta-hedging. Furthermore, Carr and Wu (2009) constructed synthetic option portfolios to capture the volatility risk premium of multiple indices and showed that the risk premium is negative on average.

The forces and possible implementations of the volatility risk premium are researched more recently. Todorov (2010) investigated the forces behind the temporal variations in the volatility risk premium over time. He found that the time-varying changes are mostly driven by large market jumps. Carr and Wu (2016) and Bollerslev, Gibson and Zhou (2011) show that the extracted variance risk premium has significant predictive power of future index returns. Moreover, Prokopczuk and Simen (2014) find that the (adjusted) volatility risk premium outperforms all alternative models in forecasting volatility.

In this paper, it is investigated whether multiple GARCH option pricing models can adequately capture the volatility risk premium. For a long time, GARCH models have been widely used in financial applications, as it is able to adequately capture the stylized facts of financial time series (i.e. volatility clustering and a high kurtosis). Besides the simple univariate GARCH model (Engle & Bollerslev, 1986), other more sophisticated models are also implemented, such as the TGARCH (Glosten, Jagannathan & Runkle, 1993), AGARCH (Engle & Ng, 1993) and Realized GARCH (Hansen, Huang & Shek, 2012).

Using GARCH models in the pricing of options under a locally risk-neutral valuation relationship (LRNVR) was introduced by Duan (1995). Under certain assumptions on the utility function, he constructed an equilibrium equation with which options can be priced. One of the principal differences with the regular GARCH models is the inclusion of the equity risk premium in the function of the option price. In this paper, the concept of the LRNVR is utilized in capturing the volatility risk premium.

Since the publication of Duan (1995), the LRNVR has been further investigated in the option pricing literature. Hao and Zhang (2013) studied the difference between the GARCH implied VIX and the CBOE VIX. They found that the GARCH implied VIX is consistently lower compared to the CBOE VIX when using just the asset returns. When the CBOE VIX is included in estimation[1], the GARCH implied VIX still cannot capture the CBOE VIX from a statistical perspective. Following this paper, Zhang and Zhang (2020) showed that using a modified version of the LRNVR[2] improves the fitting of the GARCH implied VIX on the CBOE VIX substantially.

Other alternatives of the GARCH option pricing model under the LRNVR is a GARCH model with a variance-dependent pricing kernel that allows for a variance premium, which was introduced by Christoffersen, Heston and Jacobs (2013). Furthermore, option pricing models which include the realized variance have also shown its potential in recent literature (Huang, Wang & Hansen, 2017). In fact, Hansen, Huang, Tong and Wang (2024) showed that the Realized GARCH model outperforms all other GARCH models in capturing the variance risk premium.

In this paper, the results of Hao and Zhang (2013) are replicated first. The main research question therein entails whether the GARCH implied VIX under the LRNVR is able to fit the CBOE VIX. This is investigated using index returns, but also the CBOE VIX itself. The relevance of this research is mostly scientific, as investigating which methods and models prices the volatility risk premium best is a captivating pursuit. There are no predictions made in this article and often ex-post information is used in estimation, making it not useful for practical applications. As the research of Hao and Zhang (2013) stated that the GARCH implied VIX under LRNVR is not able to capture the variance risk premium adequately, there is still some room for additional research in this sphere.

Besides replicating the results of Hao and Zhang (2013), this research extends the current literature by trying to capture the variance risk premium at different time horizons. Specifically, the more recent variants of the VIX, the VIX9D and VIX1D, are investigated. These indices contain the expected volatility for a shorter time period, and have not been researched often in the literature yet. Therefore, in this paper, it is also studied whether certain GARCH option pricing models are able to price the CBOE VIX9D and VIX1D. As the time horizons are shorter, it could be interesting to see whether realized measures are able to improve the pricing performance through the Realized GARCH option pricing model. Hansen et al. (2024) already found that the Realized GARCH model outperforms other GARCH models for the regular VIX, but in this paper, it is researched whether this is also the case for the VIX at shorter time horizons.

Concerning the VIX9D and VIX1D, literature has shown that the VIX term structure could be used

[1]The CBOE VIX is included in estimation through the form of a constructed measurement error between the implied VIX and CBOE VIX. More on this in Section 3

[2]The main modification in this model entails that the conditional variances under different measures (P and Q, more on this in Section 3.) are constructed to be different.

for predicting future stock returns (Aharon & Dimpfl, 2022) and variance risk premia (Johnson, 2017). Moreover, the VIX indices at a lower time horizon have also been used in forecasting the regular VIX. For example, the VIX9D could be used as a lower bound for the normal VIX (Jiang & Lazar, 2022). Lastly, there has been some skepticism surrounding the VIX1D, because it has a highly predictable intraday pattern and day-of-the-week effect, which coincides with an overnight bias (Albers & Kestner, 2024). This is something to keep in mind when determining which daily return format is used (close-to-close, open-to-close etc.). In this paper, this is tackled by only using close-to-close data.

The reported results show that, under the LRNVR, multiple GARCH models are unable to capture the variance risk premium when only returns are used. When the CBOE VIX is used in estimation, the coefficient representing the equity risk premium is distorted to fit the CBOE VIX. In regard to the VIX at lower time horizons, we find that the Realized GARCH options pricing model outperforms the GARCH option pricing model for modelling the VIX9D. Regarding the VIX1D, the Realized GARCH model does not show substantial fitting improvements, but its coefficients are not distorted as is for the GARCH model.

This paper is structured as follows. In Section 2, the data, its sources and its time periods are mentioned. This is followed by the financial framework in Section 3. The numerical results are given in Section 4 and the paper is concluded in Section 5. In the Appendix, additions to the financial framework and supplementary results are shown, together with an explanation of the used programming code.

2 CBOE VIX Indices

In this section, we shortly discuss the CBOE VIX indices that measure the expected (annualized) volatility implied by option prices. The VIX30D (often referred to as the regular VIX) was the first index solely based on volatility and has become one of the principal measures of market sentiment. As the name says, this index reflects the market's expectation of volatility for the next 30 calendar days (which equals approximately 21 trading days). VIX^2 reflects the variance swap rate with a duration of 30 calendar days, which can be interpreted as the risk-neutral expected variance. After the introduction of the VIX30D, various alternative VIX indices were introduced that cover a different time horizon. These include the VIX1D, VIX9D and VIX3M, which entail the expected implied volatility for the next day, nine days and three months, respectively.

In this paper, we will make use of CBOE VIX data during different time periods. This is done, because of the fact that these indices were not created at the same time. Therefore, the data can be divided into three sections, each corresponding to the VIX30D, VIX9D and VIX1D, respectively. Moreover, besides the different time periods, other time series from different sources are used per time period.

2.1 VIX30D

The VIX30D (often referred to as just 'VIX') is the expected volatility for the next 30 calendar days. The data for this index consists of three different time series. Firstly, the daily closing prices of the S&P 500 index, which are retrieved from CRSP. Secondly, the daily CBOE VIX time series (again closing prices) will be implemented. This data is obtained from the CBOE website. Thirdly, the daily 3-month U.S. treasury bill yields are used to serve as the risk-free rate. The yields are taken from the Federal Reserve website (FRED) and are transformed to be de-annualized. For all these time series, the period ranges from January 2nd 1990 to August 10th 2009.

2.2 VIX9D

The VIX9D is the expected volatility for the next 9 calendar days. This index is newer compared to the normal VIX, and therefore the period for this index runs from January 3th 2011 to November 24th 2021. The retrieved data for this period includes the CBOE VIX9D, the daily closing prices of the S&P 500 index and the risk-free rates. Additionally, for the Realized GARCH model, the daily realized variances of the S&P 500 at a 5-minute frequency are used. These are obtained from the Realized Library at the Oxford-Man Institute.

2.3 VIX1D

The VIX1D is the expected volatility for the next trading day. This index was introduced in May 2022. Because of that, the period for this index runs from May 16th 2022 to May 31st 2024. Again, the implemented time series are the CBOE VIX1D, the S&P 500 prices, the risk-free rates and the 5-minute realized variance. For this case, the realized variance is retrieved from the site of the University of Chicago Booth School of Business (Dacheng Xiu)[3].

3 Financial Framework

In this section, the models and the corresponding methods are discussed. As mentioned in Section 1, these will be largely identical to that of Hao and Zhang (2013). Specifically, Sections 3.1 to 3.3 serve to explain the methods that try to replicate the results of Hao and Zhang (2013), while Sections 3.4 and 3.5 extends this research.

3.1 GARCH option pricing measures

For the GARCH option pricing models, two measures are considered, namely the physical measure P and the risk-neutral measure Q (LRNVR of Duan (1995)). Under the physical measure P, the returns are modeled to have a conditional lognormal distribution:

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

where X_t is the closing price of the asset at day t, r_t the risk-free rate and λ is the unit risk premium. Moreover, ϵ is assumed to follow a GARCH(p,q) (Engle & Bollerslev, 1986) process:

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

$$\epsilon_t|\phi_{t-1}, \sim N(0, h_t) \tag{3}$$

where ϕ_t is the information set up to day t. This model is largely similar to a regular GARCH(p,q) model, except that the risk premium and volatility is included in the return equation. Under measure P,

[3]This data can be found at https://dachxiu.chicagobooth.edu/#research

the GARCH option pricing using only returns is implemented.

Under the measure Q (LRNVR), certain assumptions are made on the utility function, resulting in the following risk-neutral relationship:

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

$$h_t = \alpha_0 + \sum_{i=1}^{q}\alpha_i(\xi_{t-i} - \lambda\sqrt{h_{t-i}})^2 + \sum_{j=1}^{p}\beta_j h_{t-j}, \tag{5}$$

$$\xi_t|\phi_{t-1}, \sim N(0, h_t). \tag{6}$$

This risk-neutral relationship is utilized when the VIX is used in estimation, because the VIX holds forward-looking expectations.

3.2 GARCH implied VIX

The VIX contains the volatility expectations of the S&P 500 for the following 21 trading days. Mathematically, this can be written as

$$\left(\frac{VIX_t}{100}\right)^2 = E_t^Q\left[\frac{1}{\tau}\int_t^{t+\tau}\tilde{h}_s\,ds\right], \tag{7}$$

where $\tilde{h_s}$ is the annualized variance of the S&P 500 and τ is the amount of trading days, which in this case is 21. In this paper, the GARCH implied VIX is computed as the expected arithmetic average of the variance:

$$\left(\frac{VIX_t}{100}\right)^2 = \frac{1}{n}\sum_{k=1}^{n}E_t^Q\left[\tilde{h}_{t+\frac{\tau k}{n}}\right]. \tag{8}$$

Because daily data is used, the implied VIX equations looks as follows under measure Q.

$$Vix_t = \frac{1}{n}\sum_{k=1}^{n}E[h_{t+k}], \tag{9}$$

where n is the amount of calendar or trading days and $Vix_t = \frac{1}{252}(\frac{VIX_t}{100})^2$ is a proxy for the daily VIX.

For the calculation of the expected conditional mean of future variance, the square-root stochastic autoregressive volatility (SR-SARV) model is implemented (Meddahi & Renault, 2004). This is defined as follows.

Theorem 1 *A stationary process $\{\epsilon_t, t \in Z\}$ is defined as a SR-SARV(p) process if:*

1. *ϵ_t is a martingale difference process with respect to the information set Γ_{t-1}, i.e. $E[\epsilon_t|\Gamma_{t-1}] = 0$, and*

2. *the conditional variance process h_t of ϵ_t given Γ_{t-1} is a marginalization of a stationary VAR(1) model with dimension p:*

$$h_t = var[\epsilon_t|\Gamma_{t-1}] = e'H_t, \tag{10}$$

$$H_t = \Omega + \Lambda H_{t-1} + V_t, \text{ with } E[V_t|\Gamma_{t-1}] = 0, \tag{11}$$

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

Then, given that the S&P 500 is a SR-SARV(1) process under the LRNVR, the implied VIX can be defined as follows (proof can be found in Hao and Zhang (2013)).

$$Vix_t = \zeta + \psi h_t, \text{ where} \tag{12}$$

$$\zeta = \frac{\Omega}{1-\Lambda}(1-\psi), \tag{13}$$

$$\psi = \frac{1-\Lambda^n}{n(1-\Lambda)}. \tag{14}$$

This expression is obtained by writing out the expectation of volatility ($E[h_{t+k}]$) under measure Q and using the SR-SARV(1) characteristics as described in Theorem 1. As shown in the papers of Hao and Zhang (2013) and Meddahi and Renault (2004), all different GARCH models are variants of the SR-SARV model. For the TGARCH and AGARCH models, they take the following form. (GARCH model specification is already described in Equation 2 and 5).

TGARCH(1,1) :

$$\text{Physical measure : } h_t = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \theta \epsilon_{t-1}^2 \mathbb{1}(\epsilon_{t-1} < 0) + \beta h_{t-1}, \tag{15}$$

$$\text{LRNVR : } h_t = \alpha_0 + (\xi_{t-1} - \lambda\sqrt{h_{t-1}})^2 \Big[\alpha_1 + \theta \mathbb{1}(\xi_{t-1} - \lambda\sqrt{h_{t-1}})\Big] + \beta h_{t-1}. \tag{16}$$

AGARCH(1,1) :

$$\text{Physical measure : } h_t = \alpha_0 + \alpha_1(\epsilon_{t-1} - \kappa\sqrt{h_{t-1}})^2 + \beta h_{t-1}, \tag{17}$$

$$\text{LRNVR : } h_t = \alpha_0 + \alpha_1(\xi_{t-1} - \lambda\sqrt{h_{t-1}} - \kappa\sqrt{h_{t-1}})^2 + \beta h_{t-1}. \tag{18}$$

Using this specification, the daily GARCH implied VIX can be calculated linearly with the variance of the previous period and the estimated GARCH coefficients. The exact formulas of the GARCH implied VIX can be found in Appendix A. Moreover, the proof of how the GARCH models can be modified to fulfill the SR-SARV characteristics can be found in the paper of Hao and Zhang (2013) and Meddahi and Renault (2004).

3.3 Estimation under LRNVR

Estimation will be done in three different ways. In all methods, we make use of maximum likelihood. The first estimation entails only the index returns (so not the VIX) under the physical measure P as in Equation 2. The log-likelihood function for all different GARCH models is as follows.

$$\ln\ L_R = -\frac{T}{2}\ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\left\{\ln(h_t) + \left[\ln(X_t/X_{t-1}) - r_t - \lambda\sqrt{h_t} + \frac{1}{2}h_t\right]^2 / h_t\right\}, \tag{19}$$

where h_t is the iteratively updated variance process of the different GARCH models.

The second approach is based on the utilization of the returns combined with the market VIX, as this could contain additional information about how the underlying return series evolves. As the return innovation ϵ_t controls both the asset returns through Equation 1 and the GARCH implied VIX through h_t, we allow for a disparity between the market (CBOE) VIX and the GARCH implied VIX through an error term:

$$VIX_{CBOE} = VIX_{IMP} + \nu,\ \nu \sim i.i.d.N(0, s^2), \tag{20}$$

where, variance s^2 is estimated by $\hat{s}^2 = var(VIX_{CBOE} - VIX_{IMP})$. The log-likelihood for the market VIX is then as follows.

$$\ln\ L_V = -\frac{T}{2} ln(2\pi\hat{s}^2) - \frac{1}{2\hat{s}^2}\sum_{t=1}^{T}(VIX_{CBOE,t} - VIX_{IMP,t})^2. \tag{21}$$

To estimate the coefficients using both the returns and the market VIX, the joint likelihood function is maximized, that is

$$\ln\ L_T = \ln\ L_R + \ln\ L_V. \tag{22}$$

The third and last estimation involves just the CBOE VIX without the index returns (i.e. maximizing $ln\ L_V$). Following all the estimations, the corresponding implied VIX is analyzed and tested in terms of fitting of the CBOE VIX. Moreover, to start the variance processes, the first iteration is equal to the variance of the index returns over the whole sample period. Lastly, all GARCH coefficients are estimated under the corresponding stationary constraints:

$$\text{GARCH(1,1)} : \alpha_1(1 + \lambda^2) + \beta < 1, \tag{23}$$

$$\text{TGARCH(1,1)} : \alpha_1(1 + \lambda^2) + \beta + \theta\left[\frac{\lambda}{\sqrt{2\pi}} \exp^{-\frac{\lambda^2}{2}} + (1 + \lambda^2)N(\lambda)\right] < 1, \tag{24}$$

$$\text{AGARCH(1,1)} : \alpha_1\left[1 + (\lambda + \kappa)^2\right] + \beta < 1. \tag{25}$$

3.4 Realized GARCH

The GARCH models in the previous sections use only returns as input for its volatility modeling. However, by now it is well known that realized measures (realized variance in particular) are more accurate in modeling volatility compared to asset returns. To make use of this fact in pricing the VIX, the Realized GARCH framework is implemented. This model combines returns and realized measures, and therefore offers a more flexible way of forecasting volatility. The Realized GARCH framework was introduced by Hansen et al. (2012) and improved by Hansen and Huang (2016).

To keep it in the same context as the models in Section 3.2, this model is again approached under two measures, namely the physical measure P and the risk-neutral measure Q. Moreover, the equity risk premium is again represented by λ. The Realized GARCH model under the physical measure P looks as follows (Hansen et al., 2024).

$$\ln \frac{X_t}{X_{t-1}} = r_t + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}z_t, \tag{26}$$

$$\ln(h_{t+1}) = \alpha_0 + \beta\, ln(h_t) + \tau(z_t) + \gamma\sigma u_t, \tag{27}$$

$$\ln(p_t) = \omega + \phi\, ln(h_t) + \delta(z_t) + \sigma u_t, \tag{28}$$

where z_t is the standardized return (i.e. $\epsilon_t/\sqrt{h_t}$) and assumed to be distributed standard normally (just as u_t). The functions $\tau(z_t) = \tau_1 z + \tau_2(z^2 - 1)$ and $\delta(z_t) = \delta_1 z + \delta_2(z^2 - 1)$ are so-called leverage functions, that try to capture the correlation between returns and its volatility. Lastly, p_t is the realized measure and σ represents the volatility of volatility.

Before moving on to the model under the risk-neutral measure Q, it is important to choose the right way of relating this risk-neutral measure to the physical measure P. In the literature, there have been two principal routes to take. Firstly, the LRNVR by Duan (1995), which is already used in Section 3.1 to 3.3. Secondly, the variance-dependent stochastic discount factor (SDF) by Christoffersen et al. (2013) is another option. However, the model described above consists of a dual shock structure, z_t and u_t, which makes the LRNVR and SDF unusable in this case. Instead, Hansen et al. (2024) propose the exponentially affine SDF.

The stochastic process of this SDF, is defined by

$$M_{t+1} = \frac{\exp(-\lambda z_{t+1} - \zeta u_{t+1})}{E_t^P\Big[\exp(-\lambda z_{t+1} - \zeta u_{t+1})\Big]} = \exp\Bigg\{ -\lambda z_{t+1} - \zeta u_{t+1} - \frac{1}{2}(\lambda^2 + \zeta^2)\Bigg\}, \tag{29}$$

which satisfies the necessary condition of $E_t^P[M_{t+1}X_{t+1}] = E_t^Q[X_{t+1}]$. Essentially, this term M_{t+1} links the physical measure to the risk-neutral measure. In addition, the moment generating function under the risk-neutral measure Q can be described as

$$MGF(s_1, s_2) = E_t^Q[\exp(s_1 z_{t+1} + s_2 u_{t+1}] = E_t^P[M_{t+1}\exp(s_1 z_{t+1} + s_2 u_{t+1}] \tag{30}$$

$$= \exp[-s_1\lambda - s_2\zeta + \frac{1}{2}(s_1^2 + s_2^2)]. \tag{31}$$

This function is equal to $E_t^P[\exp(s_1 z_{t+1}^* + s_2 u_{t+1}^*]$, with $z_{t+1}^* = z_{t+1} + \lambda$ and $u_{t+1}^* = u_{t+1} + \zeta$. Then, the Realized GARCH model under measure Q looks as follows.

$$\ln\frac{X_t}{X_{t-1}} = r_t + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}z_t^*, \tag{32}$$

$$\ln(h_{t+1}) = \tilde{\alpha_0} + \beta\, ln(h_t) + \tilde{\tau}(z_t^*) + \gamma\sigma u_t^*, \tag{33}$$

$$\ln(p_t) = \tilde{\omega} + \phi\, ln(h_t) + \tilde{\delta}(z_t^*) + \sigma u_t^*, \tag{34}$$

where $\tilde{\alpha_0} = \alpha_0 - \tau_1\lambda + \tau_2\lambda^2 - \gamma\sigma\zeta$, $\tilde{\tau}(z_t^*) = z_t^*(\tau_1 - 2\tau_2\lambda) + \tau_2(z_t^{2^*} - 1)$, $\tilde{\omega} = \omega - \delta_1\lambda + \delta_2\lambda^2 - \sigma\zeta$ and $\tilde{\delta}(z_t^*) = z_t^*(\delta_1 - 2\delta_2\lambda) + \delta_2(z_t^{2^*} - 1)$.

The estimation is done through Maximum Likelihood, and in the same manner as in the GARCH models under LRNVR (Equations 19 to 22). For the Realized GARCH model, there are two shocks (z_t and u_t), which causes there to be two likelihood functions, which are maximized under measure P:

$$\ln\ L_R = -\frac{T}{2}ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\Bigg\{\ln(h_t) + \Bigg[\ln(X_t/X_{t-1}) - r_t - \lambda\sqrt{\frac{1}{2}h_t}\Bigg]^2/h_t\Bigg\}, \tag{35}$$

$$\ln\ L_X = -\frac{T}{2}ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\Bigg\{\ln(\sigma^2) + \Bigg[\ln(p_t) - \alpha_0 - \beta\ln(h_t) - \delta(z_t)\Bigg]^2/\sigma^2\Bigg\}. \tag{36}$$

Besides these likelihoods, we also include a likelihood for the VIX pricing errors, which originates from the same specification as in Equation 20. That likelihood looks as follows.

$$\ln\ L_V = -\frac{T}{2}\ln(2\pi) - \frac{1}{2}\sum_{t=1}^{T}\Bigg\{\ln(s^2) + (VIX_{CBOE,t} - VIX_{IMP,t})^2/s^2\Bigg\}. \tag{37}$$

Lastly, all these functions are combined in a total likelihood equation which will be maximized.

$$\ln\ L_T = \ln\ L_R + \ln\ L_X + \ln\ L_V. \tag{38}$$

The coefficients are estimated under the stationary constraint $|\beta| < 1$. Moreover, $\lambda, \gamma, \sigma, \tau_2, \delta_2 > 0$ and $\zeta, \tau_1, \delta_1 < 0$ to ensure that the equity risk premium is positive and the volatility premium is negative.

3.5 Implied VIX formulas (VIX9D and VIX1D)

The Realized GARCH implied VIX formula that is needed in Equation 37 is largely provided by Hansen et al. (2024). However, they studied whether the Realized GARCH option pricing model is able to adequately price the VIX30D (often referred to as the regular VIX). In our paper, we extend this research by incorporating a more recent version of the VIX, namely the VIX9D. This index has a lower time horizon, making it more volatile and more dependent on realized measures. The Realized GARCH implied VIX9D is given by (Hansen et al., 2024)

$$Vix_t^{RG} = \frac{1}{7}\left[h_{t+1} + \sum_{k=2}^{7}\left(\prod_{i=0}^{k-2} F_i\right) h_{t+1}^{\beta^{k-1}}\right], \tag{39}$$

where Vix_t^{RG} is the daily proxy of the VIX9D and $F_i = (1-2\beta^i\tau_2)^{-\frac{1}{2}} exp\Big\{\beta^i(\tilde{\alpha_0}-\tau_2)+\frac{1}{2}\beta^{2i}\Big[\frac{\tilde{\tau_1}}{1-2\beta^i\tau_2}+\gamma^2\sigma^2\Big]\Big\}$. This equation is obtained by writing out the expectation $E_t[h_{t+k}]$ under measure Q, while using the variance process in Equation 33.

As a benchmark for this implied VIX9D, the regular univariate GARCH model under the LRNVR is used. For this model, the implied VIX formula now is equal to

$$Vix_t = \frac{1}{7}\sum_{k=1}^{7} E[h_{t+k}], \tag{40}$$

which can be computed using the same specified VIX formulas as in the appendix, but with a different time horizon.

Another VIX index is the VIX1D, which contains the expected volatility for just one day. The index was introduced in 2022 and has been quite volatile on a daily basis. The calculation of the implied VIX1D simplifies to $Vix_t = E[h_{t+1}]$. For the univariate GARCH model, this is given by

$$Vix_t = E_t[h_{t+1}] = E_t[\alpha_0 + \alpha_1\epsilon_t^2 + \beta h_t] = \hat{\alpha_0} + \hat{\alpha}_1\epsilon_t^2 + \hat{\beta}h_t = \hat{h}_{t+1}, \tag{41}$$

where the estimated coefficients are calculated using the known likelihood functions. Note that this one-day-ahead expectation consists of only observable variables, which makes the inclusion of the Realized GARCH model even more interesting. The Realized GARCH implied VIX1D can be retrieved in the same manner as above, but using Equations 27 and 28.

4 Numerical Results

In this section, the results of the coefficient estimations and the corresponding implied VIX are reported. In addition, graphs of the CBOE VIX and implied VIX are shown to provide a visualization of the results. All coding was done through the programming language MATLAB. Moreover, the used programming code can be found as an attachment to this paper. Section 4.1 consists of the replication results of Hao and Zhang (2013), while Section 4.2 and 4.3 extends the literature through the Realized GARCH option pricing models for the VIX at lower time horizons.

4.1 GARCH models under LRNVR

The estimation results of the various GARCH models are given in Table 1. The coefficients are estimated during the period January 2nd 1990 to August 10th 2009. Firstly, looking at the univariate GARCH(1,1) model, it can be noted that λ, which represents the equity risk premium, increases substantially when the VIX is used. It increases from 0.0534 (only returns) to 0.2065 (returns and VIX) to 0.7877 (only VIX). Moreover, the coefficient β, representing the persistence of the conditional variance, also grows from 0.9310 (only returns) to 0.9389 (only VIX) to 0.9500 (returns and VIX). The higher values for λ and β when using the CBOE VIX both create a rise in the long-run variance of the SR-SARV(1) process under measure Q, which will also be shown in Table 2. Lastly, in the GARCH model, all coefficients are significant.

Regarding the TGARCH model, the patterns in the results are somewhat similar to those in the GARCH model. Again, the equity risk premium λ increases from 0.0228 (only returns) to 0.0798 (returns and VIX) to 0.4269 (only VIX). Furthermore, β increases as well. The coefficient θ, representing the effect of negative return error terms on the variance, decreases from 0.1104 (only returns) to 0.0612 (returns and VIX) to 0.0432 (only VIX). This means that when the implied VIX is forced to fit the CBOE VIX, the variance is less influenced by negative error terms, which is also exemplified by the increase in α_1. For the TGARCH model, all coefficients are significant.

Lastly, the estimations for AGARCH model show that its λ does not differ much between the case when only returns are used (0.0154) and the case where both returns and VIX is used (0.0159). While this value increases to 0.1998 when only the VIX is used, it is important to note that all equity premiums in the AGARCH are not significantly different from zero. Moreover, there are substantial increases in β when the VIX is used in estimation, but this is accompanied by a decrease in κ.

With these estimates of the coefficients, we can construct an implied VIX for all the different GARCH models using either just returns, just the VIX or both. The results of those are shown in Table 2, where they are contrasted to the CBOE VIX.

The first seven columns of Table 2 report certain statistical characteristics of the difference between the CBOE VIX and an implied VIX by a certain GARCH model. It can be seen that, for all GARCH models, the average difference (Mean Error) is substantially higher when only returns are used. This implies that, on average, the implied VIX (only returns) is lower than the CBOE VIX, which can also be seen in Figure 2. In the case of only returns, the TGARCH implied VIX performs worst with a mean error of 3.76[4]. Looking at the scenarios where the CBOE VIX is included in estimation, the mean error

[4]The values of the average differences for the case where only returns are used in estimation (3.46-3.76) are consistent with the actual variance premium from empirical studies. This implies that GARCH option pricing models under measure P are not able to capture this variance premium.

	α_0	α_1	β	θ	κ	λ
GARCH						
Returns	7.1036e−7 (1.6069e−7)	0.0637 (0.0065)	0.9310 (0.0069)	-	-	0.0534 (0.0139)
VIX	1.7182e−6 (0.0397e−6)	0.0366 (0.0009)	0.9390 (0.0014)	-	-	0.7868 (0.0296)
Both	1.6853e−6 (0.0471e−6)	0.0470 (0.0011)	0.9500 (0.0012)	-	-	0.2065 (0.0113)
TGARCH						
Returns	1.0808e−6 (0.1728e−6)	0.0018 (0.0049)	0.9317 (0.0069)	0.1104 (0.0120)	-	0.0228 (0.0125)
VIX	1.6192e−6 (0.0410e−6)	0.0043 (0.0024)	0.9524 (0.0017)	0.0432 (0.0032)	-	0.4269 (0.0403)
Both	1.5305e−6 (0.0410e−6)	0.0039 (0.0013)	0.9597 (0.0010)	0.0612 (0.0015)	-	0.0798 (0.0131)
AGARCH						
Returns	1.1516e−6 (0.1962e−6)	0.0563 (0.0051)	0.8795 (0.0116)	-	0.9999 (0.0667)	0.0154 (0.0147)
VIX	1.7159e−6 (0.0380e−6)	0.0366 (0.0009)	0.9389 (0.0016)	-	0.5880 (0.0836)	0.1998 (0.8352)
Both	1.7111e−6 (0.0424e−6)	0.0392 (0.0009)	0.9352 (0.0017)	-	0.7689 (0.0331)	0.0159 (0.0135)

Table 1: Estimation results for multiple GARCH models using returns, VIX or both. The meaning of the parameters can be found in the methodology and the standard errors are in parentheses.

decreases to almost zero, which is not surprising, as within its estimation, the implied VIX is forced to fit the CBOE VIX. These results are all similar among the different GARCH models.

The results in the remaining columns mostly result in the same conclusion. In the sixth column, a t-test is performed to test whether the mean of the implied VIX and the mean of the CBOE VIX are equal. The corresponding p-values are reported and for the cases where only returns are used in estimation, the means are significantly different, while it is the opposite for the estimation with the VIX. However, the correlation between the implied VIX and the CBOE VIX are consistent among all GARCH models and all estimation inputs.

The reported values in the eight to last column entail the autocorrelations and higher moments of the implied VIX. As a comparison, these characteristics are also reported for the CBOE VIX at the bottom. For the univariate GARCH model, the autocorrelations are mostly overestimated, especially when VIX is used in estimation. For the TGARCH and AGARCH model, the autocorrelations when only returns are used in estimation are largely underestimated, while the autocorrelations with VIX being included in estimation are again overestimated. Regarding the second moment (variance), it is noted that for all models, including the VIX in estimation results in a underestimation of the variance. In the case of only returns, the GARCH and AGARCH model overestimate the variance, while the TGARCH model performs best. Lastly, the third and fourth moments (skewness and kurtosis) are substantially overestimated in all models and estimation inputs, implying that these models cannot capture certain stylized facts of the CBOE VIX.

In Figure 2 in Appendix B, the graphs of the implied VIX versus the CBOE VIX are shown for the case with only returns, and the case with both returns and the VIX as estimation input. This visualization confirms the previous results, as it is clear that, when only returns are used, the implied VIX suffers from underestimation. Conversely, the time series for the implied VIX using both returns and the VIX in estimation, seem to have a decent fit with the CBOE VIX.

Combining all results, we see that, if the VIX data is not considered as estimation input, the implied VIX is substantially lower than the CBOE VIX for all GARCH models. When the VIX is considered in estimation, the implied VIX fits the CBOE VIX relatively well, but its estimated parameters are distorted to fit the VIX. In particular, all GARCH models report an equity risk premium (i.e. λ) that is too high when VIX data is considered. Moreover, none of the GARCH models seem to be able to fit the higher moments and autocorrelations of the CBOE VIX adequately. This conclusion, as the estimated coefficients and implied VIX results, is consistent with that of Hao and Zhang (2013).

	ME	Std. Err.	MAE	MSE	RMSE	P-value	Corr. Coef.	AR1	AR10	AR30	Variance	Skewness	Kurtosis
GARCH													
Returns	3.60	3.28	3.99	23.68	4.87	0.0000	0.92	0.9943	0.9349	0.7739	71.88	3.09	16.92
VIX	0.12	3.13	2.39	9.79	3.13	0.4877	0.93	0.9962	0.9550	0.8221	65.70	3.25	17.55
Both	0.26	3.19	2.38	10.26	3.20	0.1175	0.93	0.9967	0.9555	0.8162	66.93	3.26	17.83
TGARCH													
Returns	3.76	3.41	4.12	25.74	5.07	0.0000	0.92	0.9898	0.9070	0.7312	69.61	3.21	17.83
VIX	0.11	3.11	2.37	9.66	3.11	0.5201	0.93	0.9961	0.9564	0.8306	65.50	3.20	17.03
Both	0.27	3.13	2.34	9.85	3.14	0.1086	0.93	0.9960	0.9546	0.8259	67.60	3.21	17.15
AGARCH													
Returns	3.46	3.41	3.86	23.58	4.86	0.0000	0.92	0.9896	0.9089	0.7265	73.44	3.28	18.71
VIX	0.12	3.13	2.39	9.79	3.13	0.4874	0.93	0.9962	0.9551	0.8221	65.69	3.25	17.55
Both	0.27	3.14	2.38	9.90	3.15	0.1007	0.93	0.9958	0.9517	0.8134	67.44	3.27	17.79
CBOE VIX								0.9844	0.9160	0.7848	70.79	2.06	10.24

Table 2: The fit of the different GARCH implied VIXes on the CBOE VIX and various statistical properties for the GARCH models during the period from January 2nd 1990 to August 8th 2009. The difference between the CBOE VIX and an implied VIX is defined as the error. ME is the average error on a daily basis. The standard error (Std. Err.) is defined as the standard deviation of the error on a daily basis. The mean absolute error (MAE) is defined as the mean daily absolute error. The mean squared error (MSE) calculates the daily mean squared error. The root mean squared error is the square root of the MSE. The p-value is given for the null hypothesis that the means of the CBOE VIX and the implied VIX are equivalent. The correlation coefficient is defined as the linear correlation between the CBOE VIX and the implied VIX. Moreover, autocorrelation coefficients of the implied VIX with a lag of 1,10 and 30 are computed. Lastly, certain higher moments of the implied VIX are reported.

4.2 Realized GARCH implied VIX9D

After the successful replication of Hao and Zhang (2013), we extend the research with the Realized GARCH implied VIX at lower time horizons. In this subsection, the results of the VIX9D are discussed. The period regarding the reported values runs from January 3th 2011 to November 24th 2021. In Table 3, the estimated coefficients for the Realized GARCH and univariate GARCH models are shown. For the Realized GARCH model, two cases are reported, namely the one that excludes the CBOE VIX9D in estimation (i.e. maximizing $\ln\ L_R$ and $\ln\ L_X$ in Equation 38), and the one that includes it. Between the two variants, most coefficients are similar, except for γ, ω and ζ, which are significantly different. Moreover, it is noted that, on average, the standard errors are lower for the situation where VIX9D is included in estimation.

Regarding the univariate GARCH model (serving as a benchmark), the equity risk premium λ is substantially higher for the case where VIX9D is included. This pattern is consistent with the findings for the regular VIX9D in Section 4.1. Furthermore, the persistence parameter β is relatively low compared to other time periods, implying that there was less volatility clustering between January 3th 2011 and November 24th 2021.

	α_0	α_1	β	λ	τ_1	τ_2	δ_1	δ_2	γ	ω	ϕ	σ	ζ
RealGARCH													
No VIX	-0.69 (0.086)	-	0.99 (0.024)	0.06 (0.022)	-0.13 (0.015)	0.06 (0.005)	-0.17 (0.022)	0.11 (0.007)	0.12 (0.019)	-3.75 (1.039)	0.22 (0.16)	0.66 (0.011)	-5.13 (2.670)
Including VIX	-0.71 (0.039)	-	0.99 (0.003)	0.06 (0.026)	-0.53 (0.052)	0.38 (0.015)	-0.16 (0.016)	0.12 (0.007)	2.84 (0.094)	-0.17 (0.016)	0.01 (0.001)	0.66 (0.009)	-0.02 (0.123)
GARCH													
Returns	4.47e−6 (0.571e−6)	0.22 (0.021)	0.75 (0.020)	0.13 (0.018)	-	-	-	-	-	-	-	-	-
Including VIX	6.59e−6 (0.226e−6)	0.18 (0.005)	0.78 (0.006)	0.29 (0.015)	-	-	-	-	-	-	-	-	-

Table 3: Estimated coefficients for the (Realized) GARCH models for the period January 3th 2011 to November 24th 2021. These are used to calculate the implied VIX9D. Standard errors are given in parentheses.

Using the estimated coefficients and the implied VIX formulas, the comparison with the CBOE VIX9D can be made. The reported results are given in Table 4. Looking at the first row, it can be seen that the Realized GARCH implied VIX9D without the inclusion of the CBOE VIX9D has the worst performance by a wide margin. With a mean error or 9.89, this model heavily understates the CBOE VIX9D. While its skewness and kurtosis are accurate compared to the CBOE VIX9D, the variance is considerably lower.

In the case where the CBOE VIX9D is included in estimation, the Realized GARCH implied VIX9D performs better. The mean error is just 0.10, and all other error measures outperform the other models. Moreover, the skewness and kurtosis come quite close to those of the CBOE VIX9D, while the autocorrelations and variance are slightly inaccurate.

Comparing these results to those of the benchmark univariate GARCH model, it is noted that the GARCH model does not outperform the Realized GARCH model. For the GARCH model using only returns, the implied VI9D understates the CBOE VIX9D, and it has inaccurate higher moments (except for the variance). When the CBOE VIX9D is included in estimation, the fit with the CBOE VIX9D is better, but still worse than that of the Realized GARCH model (a mean error of 0.26 versus 0.10). However, the skewness and kurtosis are still substantially higher compared to the CBOE VIX9D.

Overall, the Realized GARCH implied VIX9D (including VIX) performs best in both modelling errors and higher moments (except the variance). Moreover, the equity risk premium λ is distorted to fit the VIX9D in the case of the GARCH model including the VIX, because its value is too high compared to empirical studies. The Realized GARCH model has a more realistic equity risk premium. However, to provide some nuance, there is an argument to be made about distorted coefficients for the Realized GARCH model as well, because of the big differences in certain values between the exclusion and inclusion of the CBOE VIX9D in estimation.

	ME	SD	MAE	MSE	P-value	AR1	AR10	Var	Skew	Kurt
RealGARCH										
No VIX	9.89	5.46	9.89	127.62	0.0000	0.97	0.69	14.84	3.70	25.38
Including VIX	0.10	3.70	2.54	13.71	0.6647	0.98	0.76	63.96	3.88	28.04
GARCH										
Returns	2.03	4.04	3.11	20.45	0.0000	0.96	0.64	73.21	5.43	48.36
Including VIX	0.26	3.93	2.66	15.53	0.2666	0.97	0.66	68.93	5.70	51.39
CBOE VIX9D						0.94	0.69	76.18	3.63	25.31

Table 4: The fit of the (Realized) GARCH implied VIX9D on the CBOE VIX9D and various statistical properties. These results are for the period January 3th 2011 to November 24th 2021. The meaning of the goodness of fit measures and statistical properties are identical to those in Table 2.

4.3 Realized GARCH implied VIX1D

Moving on to the Realized GARCH implied VIX1D results, the estimated coefficients are given in Table 5. These results are for the period May 16th 2022 to May 31st 2024. Note that this is a relatively short period, but necessary because of the trading history of the CBOE VIX1D.

Looking at the Realized GARCH model, the results of two different variants are shown. Firstly, there is the case where just the returns and realized variance are used in estimation (i.e. $\ln\ L_R$ and $\ln\ L_X$ in Equation 38). Secondly, we report the case where the CBOE VIX is included in estimation. There are no major differences in estimation results between the two variants, but we do note that the standard errors are substantially lower for the case where VIX is included. This implies that the estimations are more efficient when the VIX is used in estimation.

For the univariate GARCH model that serves as a benchmark, it is noted that there is again a unrealistically high equity risk premium λ of 0.29 when VIX is included in estimation. This is consistent with the results in Section 4.1 for the VIX30D. However, compared to the case with just returns, there is a decrease in β when VIX is used, which is inconsistent with the results in Section 4.1. While this could be because of the fact that a different time period is used, it is highly probable that a lower VIX horizon is the cause of this. As the VIX1D is substantially more volatile and less persistent compared to the VIX30D, it is not surprising to see a lower β (representing variance persistence) when the CBOE VIX1D is used in estimation.

	α_0	α_1	β	λ	τ_1	τ_2	δ_1	δ_2	γ	ω	ϕ	σ	ζ
RealGARCH													
No VIX	-0.68 (0.686)	-	0.99 (0.072)	0.05 (0.057)	-0.13 (0.014)	0.05 (0.009)	-0.14 (0.021)	0.04 (0.012)	0.21 (0.010)	-2.56 (3.519)	0.05 (0.362)	0.46 (0.014)	-5.00 (5.242)
Including VIX	-0.76 (0.004)	-	0.99 (0.001)	0.05 (0.001)	-0.14 (0.001)	0.06 (0.000)	-0.15 (0.001)	0.05 (0.001)	0.22 (0.000)	-2.63 (0.016)	0.07 (0.002)	0.46 (0.001)	-5.03 (1.250)
GARCH													
Returns	4.31e-7 (4.386e-7)	0.04 (0.014)	0.95 (0.015)	0.07 (0.044)	-	-	-	-	-	-	-	-	-
Including VIX	2.42e-6 (0.234e-6)	0.11 (0.004)	0.87 (0.004)	0.29 (0.036)	-	-	-	-	-	-	-	-	-

Table 5: Estimated coefficients for the (Realized) GARCH models for the period May 16th 2022 to May 31st 2024. These are used to calculate the implied VIX1D. Standard errors are given in parentheses.

In Table 6, the fit of the implied VIX1D on the CBOE VIX1D is reported. Comparing the two cases of the Realized GARCH implied VIX, there is a clear improvement in mean error for the inclusion of the VIX. This is not surprising, as in estimation this difference is minimized when VIX is included. Other measures such as the MAE and MSE show small relative enhancements in terms of fitting.

Another interesting comparison would be that between the Realized GARCH with just returns and realized variance, and the univariate GARCH with just returns. This essentially shows the added value of realized measures under the physical measure P. While the mean error is roughly the same, the MAE and MSE show notable improvements, implying that the realized variance has some added value in pricing the VIX using GARCH models. A last comparison is that between both models when VIX is included. However, the Realized GARCH model implied VIX1D does not show any gains in fitting compared to the regular GARCH model.

Looking at the statistical properties of all models, it can be seen that autocorrelations are substantially overstated compared to the CBOE VIX1D. This can be explained by the fact that persistent volatility (volatility clustering) is in the nature of GARCH models, implying relatively high autocorrelations. However, in practise the VIX1D is highly volatile on a day-to-day basis, which is confirmed by the relatively lower autocorrelations. Regarding the higher moments (variance, skewness and kurtosis), all models understate them quite heavily, with the GARCH model using only returns is the worst performer.

All in all, the Realized GARCH implied VIX1D does not improve the performane of the univariate GARCH model in terms of fitting the CBOE VIX1D. However, the coefficients of the GARCH model are again distorted to fit the CBOE VIX1D, because the equity risk premium is unrealistically high. This makes the Realized GARCH model more adequate in pricing the CBOE VIX1D, but it still cannot fit its statistical properties.

	ME	SD	MAE	MSE	P-value	AR1	AR10	Var	Skew	Kurt
RealGARCH										
No VIX	1.27	3.87	2.64	16.57	0.0023	0.98	0.81	34.00	0.94	3.09
Including VIX	0.18	3.88	2.62	15.04	0.6736	0.98	0.81	37.32	0.93	3.08
GARCH										
Returns	1.21	4.57	3.27	22.29	0.0018	0.99	0.94	26.07	0.62	1.95
Including VIX	0.13	4.05	2.76	16.43	0.7575	0.97	0.80	34.71	1.03	3.29
CBOE VIX1D						0.79	0.57	50.02	1.20	4.61

Table 6: The fit of the (Realized) GARCH implied VIX1D on the CBOE VIX1D and various statistical properties. These resultsa are for the period May 16th 2022 to May 31st 2024. The meaning of the goodness of fit measures and statistical properties are identical to those in Table 2.

In Figure 1, the CBOE VIX1D and Realized GARCH implied VIX1D are displayed together with the daily realized volatility. As all variables are daily measures of volatility, it could be interesting to see its relation visually. It is seen that, most of the time, the realized volatility understates both the CBOE and implied VIX1D. The difference between the realized volatility and VIX1D can be interpreted as the volatility premium. Another observation in this graph is that the implied VIX1D is not able to capture the heavy spikes of the CBOE VIX1D, which is consistent with the statistical properties in Table 6.

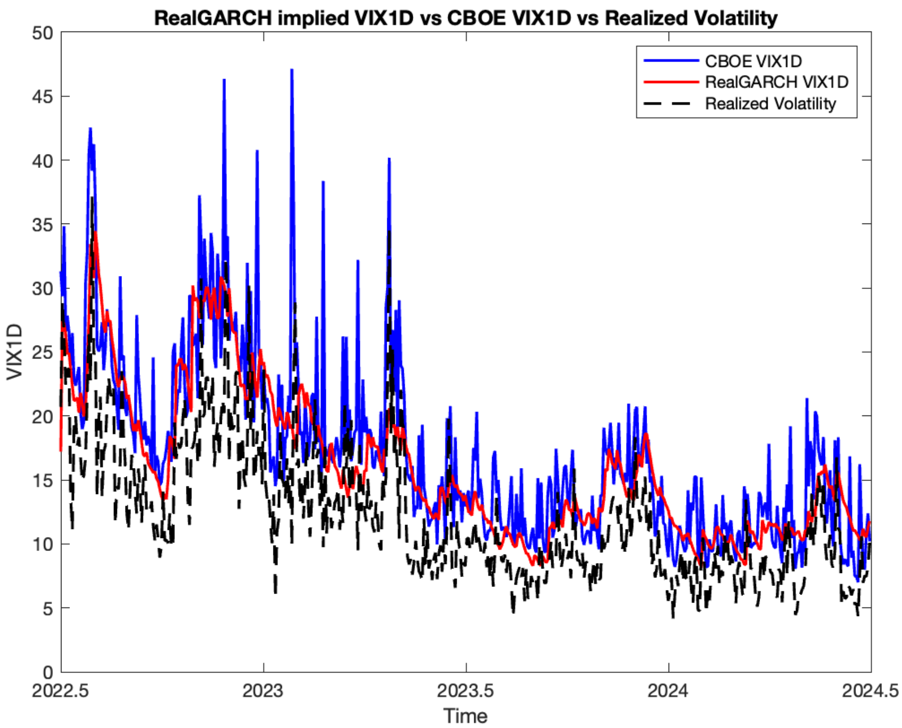

Figure 1: This graph shows the time series for the CBOE VIX1D, the Realized GARCH implied VIX1D (including VIX in estimation) and the daily annualized realized volatility.

5 Conclusion

In this paper, multiple (Realized) GARCH option pricing models are used to calculate and price the VIX at different time horizons. Specifically, the expected volatility for the next 30, 9 and 1 days are investigated. This is done under the LRNVR of Duan (1995) for the GARCH, TGARCH and AGARCH models, while an exponentially affine SDF is implemented in the case of the Realized GARCH model. The coefficients of all models were estimated and used to calculate specific implied VIX formulas.

For the GARCH, TGARCH and AGARCH models, the research entails a replication of Hao and Zhang (2013). Similar results are obtained, where these models understate the CBOE VIX when only index returns are used in estimation, implying that this model is unable to capture the variance risk premium. When the CBOE VIX is included in estimation, the coefficients are distorted to fit the CBOE VIX, resulting in an unrealistic equity risk premium.

This research is extended using a more sophisticated model that implements realized measures: the Realized GARCH model. With this model, the implied VIX for the next 9 (VIX9D) and 1 (VIX1D) days are calculated and compared to the univariate GARCH model. For the VIX9D, the Realized GARCH implied VIX9D including the VIX9D in estimation outperforms the GARCH implied VIX9D in terms of both goodness of fit and certain statistical properties. Regarding the VIX1D, the Realized GARCH implied VIX1D does not outperform the GARCH implied VIX1D. However, the coefficients of the GARCH model are distorted to fit the VIX1D, while this is not the case for the Realized GARCH model.

Suggestions for further research includes further exploration of the VIX1D, as this index is relatively new and spans a short time frame. The VIX1D only contains the expected volatility for the next day, such that methods like high-frequency data (in other ways besides the Realized GARCH) or machine learning could be implemented. Furthermore, one could also look at non-Gaussian GARCH models and at the VIX measurement error in estimation that is not independent and identically distributed (as is in Equation 20).

References

Aharon, D. Y. & Dimpfl, T. (2022). Future portfolio returns and the vix term structure. *Journal of Risk*.

Albers, S. & Kestner, L. N. (2024). The daily rise and fall of the vix1d: causes and solutions of its overnight bias. *Finance Research Letters*, 105186.

Bakshi, G. & Kapadia, N. (2003). Delta-hedged gains and the negative market volatility risk premium. *The Review of Financial Studies*, *16*(2), 527–566.

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

Bollerslev, T., Gibson, M. & Zhou, H. (2011). Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities. *Journal of econometrics*, *160*(1), 235–245.

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

Carr, P. & Wu, L. (2016). Analyzing volatility risk and risk premium in option contracts: A new theory. *Journal of Financial Economics*, *120*(1), 1–20.

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

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

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

Engle, R. F. & Ng, V. K. (1993). Measuring and testing the impact of news on volatility. *The journal of finance*, *48*(5), 1749–1778.

Glosten, L. R., Jagannathan, R. & 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.

Hansen, P. R. & Huang, Z. (2016). Exponential garch modeling with realized measures of volatility. *Journal of Business & Economic Statistics*, *34*(2), 269–287.

Hansen, P. R., Huang, Z. & Shek, H. H. (2012). Realized garch: a joint model for returns and realized measures of volatility. *Journal of Applied Econometrics*, *27*(6), 877–906.

Hansen, P. R., Huang, Z., Tong, C. & Wang, T. (2024). Realized garch, cboe vix, and the volatility risk premium. *Journal of Financial Econometrics*, *22*(1), 187–223.

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

Huang, Z., Wang, T. & Hansen, P. R. (2017). Option pricing with the realized garch model: An analytical approximation approach. *Journal of Futures Markets*, *37*(4), 328–358.

Jiang, Y. & Lazar, E. (2022). Forecasting vix using filtered historical simulation. *Journal of Financial Econometrics*, *20*(4), 655–680.

Johnson, T. L. (2017). Risk premia and the vix term structure. *Journal of Financial and Quantitative Analysis*, *52*(6), 2461–2490.

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

Prokopczuk, M. & Simen, C. W. (2014). The importance of the volatility risk premium for volatility forecasting. *Journal of Banking & Finance*, *40*, 303–320.

Todorov, V. (2010). Variance risk-premium dynamics: The role of jumps. *The Review of Financial Studies*, *23*(1), 345–383.

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

Appendix A: Implied VIX Formulas

Rewriting the corresponding GARCH models into their SR-SARV(1) variants and substituting its parameters in Equation 12, gives the following implied VIX formulas (see Hao and Zhang (2013) for additional details). For all models, n is the amount of trading days and the implied VIX is calculated linearly:

$$Vix_t = A + Bh_{t+1} \tag{A1}$$

GARCH:

$$\begin{aligned} A &= \frac{\alpha_0}{1-\eta}(1-B), \\ B &= \frac{1-\eta^n}{n(1-\eta)}, \\ \eta &= \alpha_1(1+\lambda^2)+\beta. \end{aligned} \tag{A2}$$

TGARCH:

$$\begin{aligned} A &= \frac{\alpha_0}{1-\eta}(1-B), \\ B &= \frac{1-\eta^n}{n(1-\eta)}, \\ \eta &= \alpha_1(1+\lambda^2)+\beta+\theta S. \\ S &= \Big[\frac{\lambda}{\sqrt{2\pi}}e^{-\frac{\lambda^2}{2}} + (1+\lambda^2)N(\lambda)\Big], \end{aligned} \tag{A3}$$

where $N(x)$ is the cdf of the normal distribution and $\xi/\sqrt{h_t}$ is independent and identically standard normal distributed.

AGARCH:

$$\begin{aligned} A &= \frac{\alpha_0}{1-\eta}(1-B), \\ B &= \frac{1-\eta^n}{n(1-\eta)}, \\ \eta &= \alpha_1(1+(\lambda+\kappa)^2)+\beta. \end{aligned} \tag{A4}$$

Appendix B: Additional Graphs

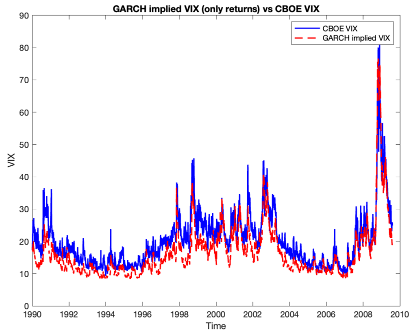

(a) GARCH implied VIX using only returns vs CBOE VIX.

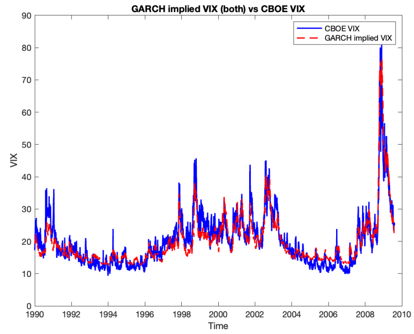

(b) GARCH implied VIX using returns and VIX vs CBOE VIX.

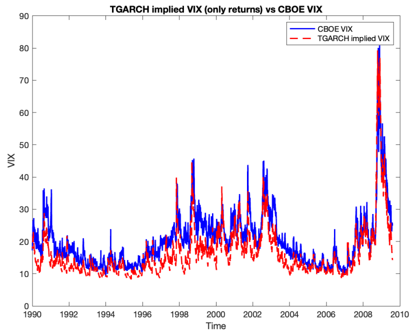

(c) TGARCH implied VIX using only returns vs CBOE VIX.

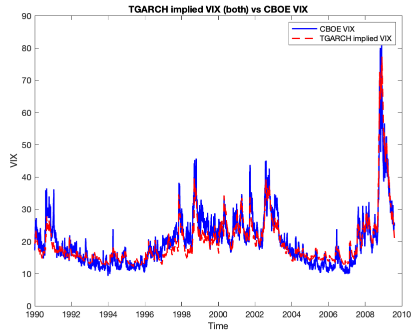

(d) TGARCH implied VIX using returns and VIX vs CBOE VIX.

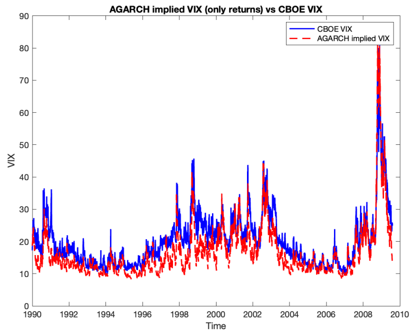

(e) AGARCH implied VIX using only returns vs CBOE VIX.

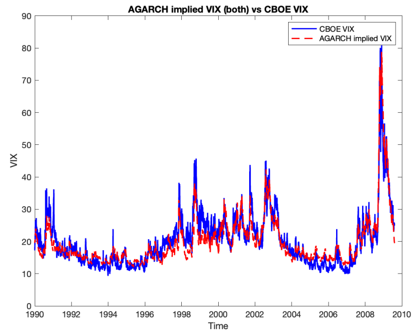

(f) AGARCH implied VIX using returns and VIX vs CBOE VIX.

Figure 2: Graphs of the CBOE VIX versus all implied VIXes by various GARCH models. In all graphs, the CBOE VIX is represented by the blue line, and the implied VIX by the dotted red line.

Appendix C: Explanation of Programming Code

As mentioned in Section 4, all programming was done through MATLAB. In this part of the Appendix, the code structure will be shortly explained. The code file consists of four folders, each representing the estimation of either the GARCH, TGARCH, AGARCH or Realized GARCH models. The files within these folders each follow a similar pattern. As the TGARCH and AGARCH files have identical structures (except for the actual model characteristics of course), only the AGARCH model is discussed. In the following paragraphs, the data, GARCH folder, AGARCH folder and Realized GARCH folder are explained.

The data consists of three different files, each representing the different time periods that are provided in Section 2. The file 'RawData' belongs to the period for the regular VIX. This file consists of the corresponding dates, S&P 500 closing prices, CBOE VIX closing prices and the daily risk free rates. The file 'RawDataFreq' represents the period for the VIX9D. This also includes the high-frequency daily realized variances. Lastly, the file 'RawDataFreq1D' belongs to the VIX1D and also contains the daily realized variances.

The GARCH folder uses all three different datasets. The file 'GARCH_Estimation' is used to run all code. All other files are used as functions and make the code structure more concise and simple. The 'GARCH_Estimation' file has comments that give a clear overview of the estimation procedures. Other noticeable files are 'getSigmas' (which returns the values of the variance process given parameter values) and all the loglikelihood files (which returns the loglikelihood values for different estimation methods). Note that the function 'fmincon' in MATLAB is used to minimize the negative loglikelihood. Lastly, the file 'constraint' consist of the stationary constraint and the file 'getStatisticalProperties' returns the values that are given in Table 2, 4 and 6.

The AGARCH folder is easier and smaller compared to the GARCH folder. The AGARCH folder only makes use of the 'RawData' dataset and the main file that runs all code is called 'AGARCH _Estimation'. All other files are again used as functions and these are similar to that of the GARCH folder.

The last folder is the Realized GARCH folder and it makes use of the datasets 'RawDataFreq' and 'RawDataFreq1D'. The main file is 'RealGARCH_Estimation'. The other files are functions, and most of them are similar to those in the GARCH folder. Exceptions are the 'getF' and 'getWholeRealVIX' functions that return the F in Equation 39 and the implied VIX9D given certain parameters, respectively. The latter is used in including the CBOE VIX9D in estimation.