Nijmegen School of Management
Department of Economics and Business Economics
Master's Thesis Economics (MAN-MTHEC)

Forecasting implied volatility

By Tomas Gerssen (s1064123)

Nijmegen, 30 June 2024

Program: Master's Program in Economics
Specialisation: Financial Economics
Supervisor: Dr J. Qiu

Radboud Universiteit

Abstract

This study tests methods of forecasting implied volatility. Standard Black-Scholes option pricing assumes that volatility is constant over time, but numerous empirical studies have verified that the conditional volatility of stock prices is clustered and noisy. Market efficiency arguments lead to the expectation that investors would include these volatility patterns into option prices if they could, more sophisticated option valuation models such as Heston (1993) allow them to do exactly that. GARCH models (as proposed by Bollerslev (1986)) allow for the modelling and forecasting of conditional volatility which is potentially clustered. This leads to the question whether future implied volatility can be forecasted using volatility forecasts made with GARCH models.

Results show that GARCH forecasts are statistically significant predictors of implied volatility, but that these do not outperform ARMA forecasting. This leads to the conclusion that GARCH is actually not the best way of forecasting implied volatility. Moreover, it seems that the market also considers additional factors not investigated here.

Contents

1 Introduction

The value of an option contract is highly dependent on the volatility of the returns of the underlying asset because it determines the probability that the option will move in or out of the money. Intuitively, suppose volatility would be zero. This implies that the price of the underlying is perfectly constant over time. In this case, the price of an option that is out of the money is exactly zero since it has no chance to be exercised. On the other hand, if volatility of returns of the underlying would be extremely high, the probability to exercise the option would also increase since there is more chance for prices to move beyond the strike price. The price of the option would adjust accordingly.

Conventionally, in finance, the volatility of stock returns is assumed to be constant over time. However, it has been shown empirically that the variance of stock returns is actually predictably heteroskedastic. More specifically, periods of relatively high volatility are likely followed by periods of high volatility, and vice versa for low volatility. This phenomenon is commonly referred to as *clustered volatility* and can be modelled by ARCH or GARCH models as proposed by Eagle (1982) and Bollerslev (1986), respectively.

Option valuation is a notoriously difficult problem in financial economics. The Black-Scholes-Merton framework (Black & Scholes, 1973; Merton, 1973) provides a closed form solution on the assumption of constant and known volatility, but this is incompatible with the observation of clustered volatility. Heston (1993) has developed the solution for the case of stochastic and mean-reverting volatility on the basis of Hull and White (1987). This is nearly the same as volatility clustering because this mean-reversion describes the autoregressive properties of conditional volatility. This means that, in theory, investors are able to correctly evaluate options while assuming volatility to be non-constant and clustered.

Being able to forecast implied volatility allows for better understanding of how derivative markets react to shocks. If this forecasting can be done with great accuracy, then this information is useful for option traders and fund managers. Market efficiency implies that investors would price the most realistic volatility dynamics into options that they are able to. Assuming that investors do use stochastic option pricing, then a volatility forecasting model with similar underlying assumptions should be a great predictor of implied volatility. This paper addresses the following research question: Can GARCH models improve implied volatility forecasting?

The evidence of this study shows that this is actually not the case. Although GARCH forecasts are a statistically significant predictor, this is not found to be the best method of forecasting implied volatility. Unexpectedly, investors seem to take long term, unconditional, volatility estimates into account in addition to short run forecasts.

This thesis will be structured as follows: First, a theoretical overview of the relationship between volatility and option prices will be laid out, during which some difficulties with the assumptions of the Black-Scholes model are discussed. Moreover, the procedure of modelling volatility clustering with GARCH-like models is described, including a literature review of research into this phenomenon. Volatility forecasting is then related to implied volatility. Furthermore, the methodological section describes the data used and the technical procedure of this research. The results section presents the findings. After which the concluding section concludes.

2 Theoretical framework and literature overview

The value of option contracts depends on the variance of the underlying product. However, there is no genuine agreement on what this relationship should be. The seminal methods of Black and Scholes (1973) and Merton (1973) form the foundation for any modern discussion on option pricing. Their model can be summarised as follows. Stock prices are assumed to follow a geometric Brownian motion process

$$dS_t = (r - \delta)S_t dt + \sigma S_t dW_t \,,$$

where S_t is the stock price in period t, W_t is a standard Brownian motion, σ is the annualised volatility such that σ^2 is the variance of returns. The interest rate is r and δ is the continuous dividend rate. Finally, dx is the instantaneous change of x in continuous time. The familiar pricing equation of a *European* style call option on the underlying stock eventually becomes:

$$C(S,t) = Se^{-\delta(T-t)}\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

with

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + (r - \delta + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}\,,$$

$$d_2 = d_1 - \sigma\sqrt{T-t}\,.$$

In this famous equation, S is the spot price of the underlying. K is the strike price for which the option may be exercised at maturity. The current period is t and T is the date of maturity such that $T - t$ is the time to maturity. Finally, $\Phi(\cdot)$ is a cumulative density function for a standard normal distribution.

In this equation, there is a one-to-one relationship between the option price and the volatility of the underlying. This means, if the price of the option is known, the market estimate of the volatility can be inferred, this is then called the *implied volatility*.

In professional settings, the implied volatility of options is often quoted as if it is the price rather than its actual price in currency amount. This is because the price of an option can only be interpreted as high or low relative to the spot price of the underlying. For example, imagine two options with different underlying products but the same implied volatility. These options will trade at different prices if the spot prices of the underlying are different. However, in this case the options are assumed to have the same risk so one should not be called more or less expensive than the other. In a delta-hedged portfolio, when the spot price of the underlying is irrelevant, the implied volatility of the option is a more direct metric of how expensive an option is.

2.1 Volatility smile and term structure

For any given underlying, there are usually option contracts being traded for various strike prices and expiration dates. According to the Black-Scholes-Merton framework, all of these options should have the exact same implied volatility. In reality this is almost never the case. Firstly, implied volatility varies for moneyness, i.e. the difference between the current spot price of the underlying and the strike price of the option. Call options are *in-the-money* (*out-of-the-money)* when the spot price is higher (lower) then the strike price. Implied volatility observed in the market tends to be the lowest for options that are *at-the-money* (i.e. the spot price roughly equals the strike price) and increases for both higher and lower strike prices. See figure 1 for an exemplary graph. This phenomenon is commonly called the *volatility smile* and this investor behaviour cannot be explained by the Black-Scholes framework.

Macbeth and Merville (1979) make the early observation that options are predictably over- and undervalued based on moneyness. Rubinstein (1985) analyses option prices relative to strike prices and maturities. He specifically finds that out-of-the-money calls have significantly higher implied volatility. Sheikh (1991) finds more volatility smiles and evidence against the validity of the Black-Scholes framework in favour of a model that allows for stochastic volatility. Heynen et al. (1994) hypothesises that the volatility smile is the result of market imperfections after reviewing some stochastic volatility models and not finding good support for them. More empirical evidence and a way to quantify the smile by Zhang and Xiang (2008).

Figure 1: Implied volatility of options expiring in January 2022 as of 01-01-2022

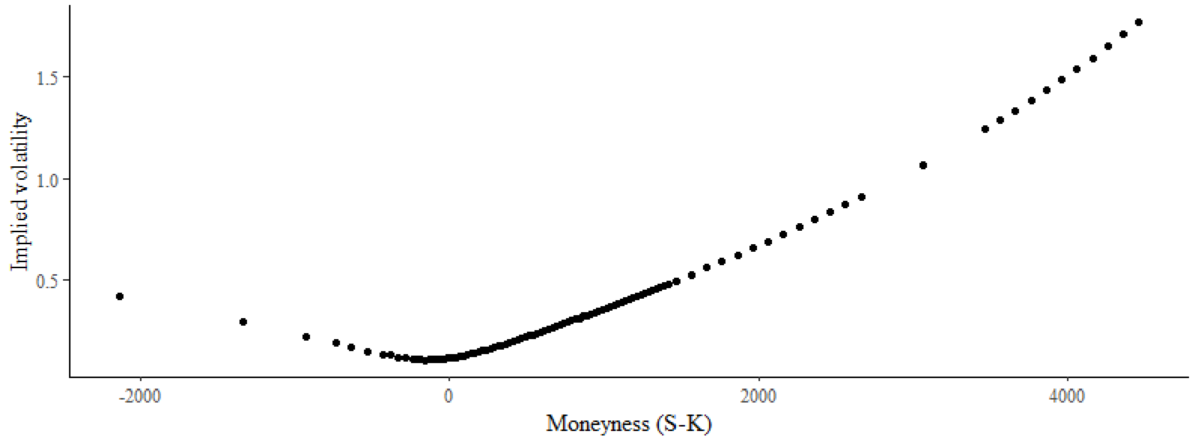

Notes: Call options on the FTSE 100 index, computations by author.

Furthermore, investors also use a different implied volatility for options with different expiration dates. This indicates that the market expectation is that volatility will change over time. In general, annualised implied volatility tends to be low for short times to expiration, but is increasing with longer maturity. See figure 2 for an exemplary graph. This *term structure* of volatility is also inconsistent with the Black-Scholes model. Xu and Taylor (1994) construct a method to quantify the term structure and use it to find empirical evidence of its persistent existence.

Figure 2: Term structure of at-the-money implied volatility as of 01-01-2022

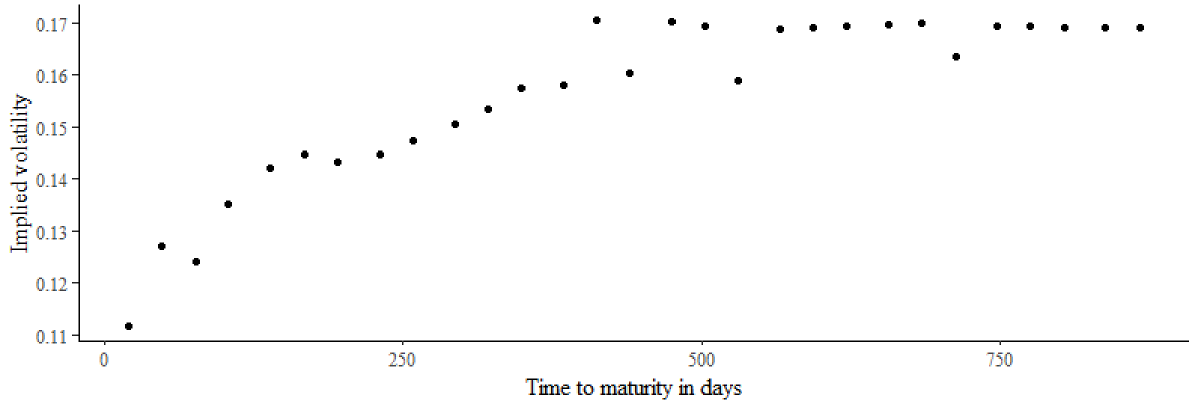

Notes: Call options on the FTSE 100 index, computations by author.

2.2 Clustered volatility and GARCH models

An crucial feature of the Black-Scholes model is that it assumes that the variance of the underlying product is known and constant over time. It is likely that the market exhibits behaviour that is inconsistent with this model because, in reality, volatility is definitely not constant over time. More specifically, there seems to be a pattern of clustering. This subsection discusses how this can be modelled, after which the empirical evidence is reviewed.

Periods of high variance are often followed by more periods of high variance, indications of this have been noted for a long time. Mandelbrot (1963) is the first to report on the observation that large price changes are often followed by more large price changes in commodity prices. This phenomenon would later be called clustered volatility. Fama (1965) finds evidence of clustered volatility in stock price returns and theorises that it is the result of slow incorporation of information into prices.

A common and powerful statistical method of modelling conditional variance in stationary time series is the Generalised Auto-Regressive Conditional Heteroskedasticity (GARCH(q, p)) model (Bollerslev, 1986), which takes the following form. Let R_t be the return in period t such that

$$R_t = \mu + \varepsilon_t \, .$$

Where μ is the expected return, modelled in some way. E.g. zero, a constant or an ARMA process. ε_t is a stochastic term with non-constant variance. More specifically:

$$\varepsilon_t \sim N(0, \sigma_t^2) \, ,$$

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2 \, .$$

The *conditional variance* σ_t^2 in any given period is determined by a constant ω, the squared shock to returns in the previous q periods and the conditional variance in the previous p periods. In this model, it is only the magnitude of returns that matter, not the sign. Moreover, the coefficients in GARCH abide to the constraint $\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i < 1$, which implies that volatility is mean-reverting. The error term here is assumed to be distributed normally, but could also be modelled to follow an alternative distribution. For example, it might be assumed to follow a student's t-distribution to allow for excess kurtosis.

For $p = 0$, the variance equation simplifies to ARCH(q) (Eagle, 1982). ARCH models assume that the range of likely shocks in this period is determined by the size of previous shocks. If in the current period the shock that materialises is not particularly large, the variance in the next period will be largely back to normal (dependent on the amount of lags q). GARCH, however, allows for the persistence of volatility by incorporating lagged conditional variance. This autoregressive property is why it is generally superior to ARCH for financial data.

Given the GARCH model, the *unconditional variance* σ^2, which can be interpreted as the long term variance, is

$$\sigma^2 = \sigma_\infty^2 = \frac{\omega}{1 - (\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i)} .$$

Theoretically, GARCH models could be long memory models for high values of p and q. However, most commonly the marginal value of additional lags is very low. It seems that in practice GARCH(1,1) is often the most useful. For interpretation purposes this is convenient, because it simplifies the conditional variance equation to

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \beta\sigma_{t-1}^2 .$$

As proposed by Nelson (1991), the exponential GARCH (EGARCH(q,p)) model is a variation on standard GARCH. Since the model is specified in log variance, coefficients are not required to be positive as they are with standard GARCH. Additionally, EGARCH allows for asymmetric effects of positive- and negative return shocks. The conditional variance equation of EGARCH is

$$\ln(\sigma_t^2) = \omega + \sum_{i=1}^{q} (\alpha_i \left|\frac{\varepsilon_{t-i}}{\sigma_{t-i}}\right| + \gamma_i \frac{\varepsilon_{t-i}}{\sigma_{t-i}}) + \sum_{i=1}^{p} \beta_i \ln(\sigma_{t-i}^2) .$$

Which is relatively similar to standard GARCH. However, it includes the additional coefficient γ. This addition allows for an added dimension to the effect of shocks on the variance in subsequent periods. Effectively, the coefficient of a positive return shock in $t-1$ on variance in t is $\alpha_1 + \gamma_1$, while the effective coefficient of a negative shock is $\alpha_1 - \gamma_1$. So, if $\alpha_1 - \gamma_1 > \alpha_1 + \gamma_1$, then negative shocks have a greater effect on variance then positive ones. This is true only for $\gamma_1 < 0$. Analogously, if gamma is positive this means that positive return shocks have a greater effect on variance in subsequent periods. In this framework, the unconditional variance (C. He et al., 2002) is

$$ln(\sigma^2) = \ln(\sigma_\infty^2) = \frac{\omega}{(1 - \sum_{i=1}^{p} \beta_i)} .$$

GARCH models can be used to quantitatively analyse patterns in the variance of exchange rates and returns on equity, among others. Importantly for the present paper, the autoregressive structure allow GARCH models to forecast volatility. Akgiray (1989) finds clustered volatility in daily and weekly stock returns using (G)ARCH models and shows that it these methods outperform other historical measures in forecasting future volatility. Jacobsen and Dannenburg (2003) find volatility clustering for monthly stock returns using GARCH on long time series. Pagan (1996, Chapter 3) is a literature summary of (G)ARCH models and their many extensions.

Theoretically, volatility clustering is thought to be a feature of heterogeneous agent markets where participants switch between trading strategies. Brock and Hommes (1998), Youssefmir and Huberman (1997), He et al. (2016), Gaunersdorfer and Hommes (2007), Gaunersdorfer et al. (2008), Kaizoji et al. (2002) and Yamamoto (2010) propose various kinds of asset pricing models with two investor strategies: fundamentalists and chartists. Following the concept of adaptive rational equilibrium (Brock & Hommes, 1997), these models allow boundedly rational agents to switch between trading strategies based on previous experiences. Under their conditions, volatility clustering is found to be the result of this switching mechanism, i.e. it is an endogenous feature of these models. The intuition can be summarised like this: "Economically, with strong trading activities of either the fundamental investors or the trend followers, market price fluctuates around either the fundamental value with low volatility or a cyclical price movement with high volatility depending on market condition. With the fundamental noise and noise traders, this triggers an irregular shifting between two volatility regimes and therefore leads to volatility clustering." (X. He et al., 2016, pp. 30–31).

2.3 Stochastic option pricing models

Evidently, volatility is not actually constant. Not only that, but there is a predictable pattern of clustering. This violates an important assumption of the Black-Scholes framework, which means it is invalid. Relaxing the assumption of constant variance in favour of stochastic variance allows for more realistic option pricing models. Stochastic variance means that the variance of returns has a random component itself, such that there are two levels of variance. Hull and White (1987) introduce a model where stock prices are modelled as

$$dS_t = (r - \delta)S_t dt + \sqrt{v_t} S_t dW_t^S.$$

Which is very similar to Black-Scholes but notice that $\sqrt{v_t}$ is the volatility of the stock returns such that v_t is the variance, which has subscript t to indicate that it varies from period to period. This *conditional* variance is a function of t and follows the stochastic process

$$dv_t = \mu dt + \xi v_t dW_t^v \,.$$

Where μ is a drift term over time and ξ is the volatility of conditional variance v_t. W_t^v is another standard Brownian motion that is the same as-, but independent of W_t^S.

In order to use this model to evaluate option contracts, one needs assumptions about the process of conditional variance. This also means that once an option price is observed, this cannot be traced back to a single value of expected variance because this process explicitly relies on multiple inputs. At best, the mean implied volatility can be induced.

Heston (1993) describes a particular variation on the Hull-White model where volatility of stock prices is stochastic and mean-reverting. Return variance is modelled as

$$dv_t = \kappa(\theta - v_t)dt + \xi v_t dW_t^v .$$

Where θ is the long run variance such that $\theta - v_t$ is the deviation in period t and κ is the degree of mean reversion of the stock return volatility to its long run level.

Although this option pricing model is closed form, it does not return a pretty analytical solution like the Black-Scholes formula. The option prices can be calculated numerically using the Fourier transform as demonstrated by Carr and Madan (1999). Analogously, for any given option price, it's possible to infer the implied parameters along which the volatility moves stochastically.

Stochastic option pricing models incorporate the empirical finding that return variance is not strictly constant and seems to change randomly. Among other things, these models explain why options deep in- and out-of-the-money are often overvalued in the market compared to Black-Scholes prices, the volatility smile. These higher implied volatilities make sense if it is assumed that return variance is random. The volatility smile shows bets being made on relatively high volatility.

Using models of stochastic volatility to infer an implied volatility from an observed option price is difficult and requires assumptions that introduce new issues. Using Black-Scholes implied volatilities instead is objectively not the same thing and introduces specification error. Day and Lewis (1992) and Hull and White (1987) show that, because stochastic volatility models are nearly identical to Black-Scholes in mean variance, the specification error is negligible for at-the-money options. This justifies the use of Black-Scholes implied volatilities regardless of the associated flaws. From now on, only at-the-money options will be considered for this very reason.

2.4 Forecasting implied volatility

If the assumptions underlying Black-Scholes would be valid, and volatility really is constant over time, then that would suggest that the implied volatility of options is constant as well. In this scenario, the notion of forecasting implied volatility is a little bit ridiculous. However, it could be done by simply extrapolating the implied volatility of the current period to the future. If volatility and thus implied volatility was constant over time, this would perfectly forecast implied volatility every time.

Under the more realistic assumption that implied volatility varies over time, one might use a simple ARMA model for forecasting, as is common for many types of financial time series data. These models have been found to be relatively efficient at forecasting implied volatility by both Ahoniemi (2006) and Hwang & Satchell (1998). ARMA(q, p) models conditional volatility in the following way:

$$\sigma_t^2 = \sum_{i=1}^{q} \varphi_i \sigma_{t-i}^2 + \sum_{i=1}^{p} \theta_i \varepsilon_{t-i} + \varepsilon_t$$

Where σ_t^2 is the conditional variance in period t and ε is an error term. φ and θ are coefficients. When applied to volatility, ARMA is very similar to the conditional volatility equation of GARCH because it does also partially model volatility clustering by incorporating the autoregressive structure of conditional volatility. The crucial difference is that in ARMA, ε is the error term of the process itself, while in GARCH it is the error term of the return function. This means that ARMA is missing the *ARCH* part of GARCH, which has been shown to be a very important predictor of volatility. Moreover, there is no moving average in GARCH models as there is in ARMA.

Forecasting implied volatility with ARMA might be done using the time series of implied volatility itself, or by defining volatility as the underlying measure of realised volatility. This is possible because in theory, implied volatility is a market estimate of the volatility that will realise in the time until expiry of the option. Alternatively, GARCH models can be used to forecast implied volatility. This is a more comprehensive way of modelling volatility clustering because, as discussed, GARCH models both size and shape effects of volatility clustering. In the case of EGARCH, even sign effects of returns are considered.

Ultimately, methods such as ARMA and GARCH can forecast volatility using historical return information, while implied volatility can be interpreted as a market expectation of volatility. Actually, quite a substantial body of research has investigated the relative information content of historical returns and option prices about the volatility of exchange rates, bond prices and stock prices. A reoccurring issue in the literature is that there is no clear benchmark to compare these methods to, because *true* volatility is difficult to define. Therefore, they have to be compared to each other, or an ex post estimate of true volatility. Fung and Hsieh (1991) find that implied volatility has little additional forecasting power for the S&P500 index.

Which would mean that they do not contain more information than historical volatility. Day and Lewis (1992) add implied volatility as an independent variable in GARCH and EGARCH regressions using data on the S&P100 index. If the explanatory power of implied volatility diminishes that of the GARCH parameters, then volatility clustering is said to be largely priced into options. Results suggest that implied volatility is not a significant factor in these regressions, which means that either options are not priced efficiently, or that the Black-Scholes formula is specified incorrectly with regard to volatility. However, Xu and Taylor (1995) find that GARCH has no additional predictive power over implied volatility using similar methods on currency options. Choi and Wohar (1992) also find that implied volatility is consistent with conditional variance and that it is thus priced in for S&P500 and NYSE index options. On the other hand, Noh et al. (1994) find that trading on GARCH forecasts of volatility gives higher abnormal returns than using methods based on implied volatility. This would mean that options are not priced efficiently. In summary, a unambiguous conclusion about the relative information content of implied volatility and GARCH models has not been achieved. It is unclear whether GARCH forecasts are priced into options.

Relating GARCH and ARMA forecasts to implied volatility introduces an issue called the *maturity mismatch*. The implied volatility is an expectation of the mean variance from now until the moment of expiration. A forecast based on a GARCH/ARMA process is valid for periods of discrete time. Depending on what kind of data is used for the estimation, this may be months, weeks, days or even intraday periods. A first order GARCH/ARMA process can only purely forecast one period ahead, after which it can go further using the first forecasted value as inputs. In most cases, the time to maturity of options does not perfectly match the forecasting period. In this case, the implied volatility and the forecast are not directly comparable because they measure variance over different time periods. It is known that variance varies over time and thus the maturity mismatch would make this comparison invalid. To solve this problem, the following analysis only concerns options that expire in the first forecasting period.

2.5 Hypothesis

In summary: Through Black-Scholes it is possible to price options if volatility is constant, but overwhelming evidence suggests that volatility is actually not constant. We also know how to evaluate options if volatility is stochastic using the Hull-White model. However, there is structure to the noise of volatility as in reality volatility is clustered. This autoregressive nature can be incorporated into option prices by the Heston model, which allows for stochastic and mean reverting volatility.

If the market really does incorporate volatility clustering into option prices, then a forecast based on a GARCH model should be a better predictor of implied volatility than anything else because GARCH is the most comprehensive way of modelling this phenomenon. Simply observing a significant correlation between the two is not good enough, because GARCH forecasts and implied volatility are expected to correlate anyway (even if investors don't use GARCH or stochastic option pricing models) because they are both estimates of the same thing: *true* volatility.

This leads to the following formal hypothesis: *Volatility forecasts made using (E)GARCH models are a better predictor of implied volatility than forecasts made with ARMA processes.*

3 Methodology

3.1 Data

The data on options is accessed through LSEG datastream. The dataset consists of monthly time series for 6891 European style call options on the FTSE 100 index in the period 2015-2023. All options expire in this timeframe. Variables used are the Black-Scholes implied volatility, closing market price and trading volume. In addition, static variables such as expiration date and strike price are available. Only observations that have a trading volume of more than 100 contracts are included in the sample, to minimise the risk of measurement error due to low trading activity. As discussed before, only at-the-money options that expire in the next period are considered. The core option data therefore consist of 60 monthly observations in the period 2019-2023. This time series is the dependent variable in all analyses.

Continuously compounded returns are calculated using the monthly total return index of the FTSE100 index through LSEG datastream for the period 1986-2023. The total return index includes both price movements and dividend yields. Daily returns are also acquired and used to calculate monthly realised volatility.

Spot prices for the FTSE100 index are acquired from Yahoo finance through the quantmod R package. Monthly prices are used to calculate the moneyness of options.

3.2 Forecasting

GARCH-like models are autoregressive. With information at t, we can estimate the value at $t+1$. This allows us to forecast volatility. For the intents of this research, a rolling one-step-ahead forecast is done as follows: There are four GARCH specifications varying with regard to two aspects. The model is either GARCH(1,1) or EGARCH(1,1), and the distribution of the conditional returns is assumed to be either normal or following a student's t-distribution. This leads to four GARCH models. All models have the expected return structure ARMA(1,1). At any given period t, all return data up to t is used to fit the four GARCH models. The estimated coefficients of these models are then used to make forecasts for period $t+1$. This procedure is then repeated for all periods in the sample such that for every period t, the forecasts that were made about t in $t-1$ are available for all four models.

Additionally, at every period, the estimated GARCH coefficients are used to calculate the unconditional volatility according to all four models for that information set.

Similarly, rolling volatility forecasts are made using ARMA models. While GARCH models use raw return data, ARMA for volatility forecasting requires a volatility time series as input. For this, volatility is defined in two ways. Firstly, the time series of implied volatility itself is used. Secondly, volatility is forecasted using realised monthly volatility. ARMA models of various orders are tested to see what produces the best forecasts of implied volatility.

3.3 Comparing model performance

The aforementioned methods result in various forecasts of implied volatility in every period. Each of these forecasts is the independent variable in an OLS regression to assess its explanatory power of implied volatility. As described, the explanatory power of GARCH models is only meaningful to the research question in comparison to the explanatory power of alternative methods of forecasting implied volatility, the alternatives used are ARMA models. This means that it is not good enough to simply observe significant correlations between GARCH forecasts and implied volatility. It must be determined whether GARCH forecasts are not only a good predictor of implied volatility, but whether they are the best. To make this comparison properly, models are compared by both R-squared and log-likelihood.

Additionally, to test and conclude whether GARCH forecasts contain more information than any other forecasting method, the best performing GARCH forecasts are added into the model of the best performing alternative forecasts. If the inclusion of GARCH forecasts diminishes the significance of coefficients associated with the alternatives, then there is more information about implied volatility in GARCH forecasts. This would lead to the conclusion that it is indeed a superior forecasting method. For this analysis, performance is measured by both adjusted R-squared and the Akaike information criterion (AIC). The AIC is a log-likelihood based measure that is adjusted for the amount of independent variables in the regression. These methods are similar to those applied by Day and Lewis (1992).

One issue with the latter analysis is the possibility of collinearity among explanatory variables. To assess whether this is a problem, the variance inflation factor (VIF) is calculated for all independent variables. To calculate the VIF, an OLS regression is computed to estimate every explanatory variable by using all other explanatory variables as regressors. The explanatory power (R^2) of this model determines the VIF, which equals $1/(1-R^2)$. The threshold for problematic collinearity is set at five.

4 Results

4.1 ARMA forecasting

In this subsection, the best ARMA model for forecasting implied volatility determined. As described, two kinds of data and various ARMA specifications are tested. For all variable names of forecasts, subscript "$t, t-1$" means that it is the forecast of period t made with information up to and including period $t-1$.

See table 1 for the explanatory power of forecasts made with ARMA processes on historical implied volatility itself. Table 2 shows the results for forecasts made with ARMA models on historical monthly realised volatility.

Table 1: Explaining current implied volatility with ARMA forecasts based on historical implied volatility

	Dependent variable:		
	$Implied\ volatility_t$		
	(1)	(2)	(3)
$ARMA(1,0)^{IV}_{t,t-1}$	0.857***		
	(0.153)		
$ARMA(1,1)^{IV}_{t,t-1}$		0.803***	
		(0.157)	
$ARMA(3,3)^{IV}_{t,t-1}$			0.720***
			(0.145)
Constant	0.024	0.032	0.041
	(0.025)	(0.026)	(0.025)
Observations	56	56	56
R-squared	0.368	0.327	0.314
Adjusted R-squared	0.357	0.315	0.302
Residual SE	0.056	0.058	0.059
F statistic	31.499***	26.277***	24.771***
Log Likelihood	82.723	80.958	80.428
AIC	-159.446	-155.917	-154.856

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. Implied volatility of FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month).*

Table 1 shows that ARMA processes can significantly forecast future implied volatility and explain a little over 30 percent of variance. It seems that this relationship is quite straightforward when using implied volatility to forecasts. The model in column one indicates that implied volatility is simply autoregressive. Higher order models or incorporating any moving average at all does not seem improve the ability to forecasts future implied volatility.

However, the estimation results in table 2 show that the same is definitely not true for realised volatility. It seems that implied volatility forecasts made with the higher order ARMA(3,3) model on monthly realised volatility actually is a significant improvement relative to simpler models. $ARMA(3,3)^{RV}_{t,t-1}$ has the best R-squared, log-likelihood, AIC and residual standard error. Therefore, this is the model that GARCH forecasts are compared to in order to assess the hypothesis.

Table 2: Explaining current implied volatility with ARMA forecasts based on historical realised volatility

	Dependent variable:		
	$Implied\ volatility_t$		
	(1)	(2)	(3)
$ARMA(1,0)^{RV}_{t,t-1}$	0.829***		
	(0.116)		
$ARMA(1,1)^{RV}_{t,t-1}$		0.853***	
		(0.115)	
$ARMA(3,3)^{RV}_{t,t-1}$			0.909***
			(0.113)
Constant	0.032*	0.029	0.020
	(0.019)	(0.019)	(0.018)
Observations	59	59	59
R-squared	0.473	0.493	0.532
Adjusted R-squared	0.463	0.484	0.524
Residual SE	0.051	0.050	0.048
F statistic	51.103***	55.316***	64.722***
Log Likelihood	93.187	94.315	96.687
AIC	-180.374	-182.63	-187.375

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. Implied volatility of FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month). Realised volatility calculated on monthly basis using continuously compounded daily returns of FTSE 100 index.*

The implied volatility forecasts made with ARMA(3,3) on realised volatility are plotted over time in figure 3. Here, the monthly realised volatility and implied volatility are plotted as well to show the dynamics of these variables over time. To put this graph into perspective, the simple annualised standard deviation of monthly returns over this period is 0.148.

Figure 3: Volatility and ARMA(3,3) volatility forecast over time, with implied volatility (dotted)

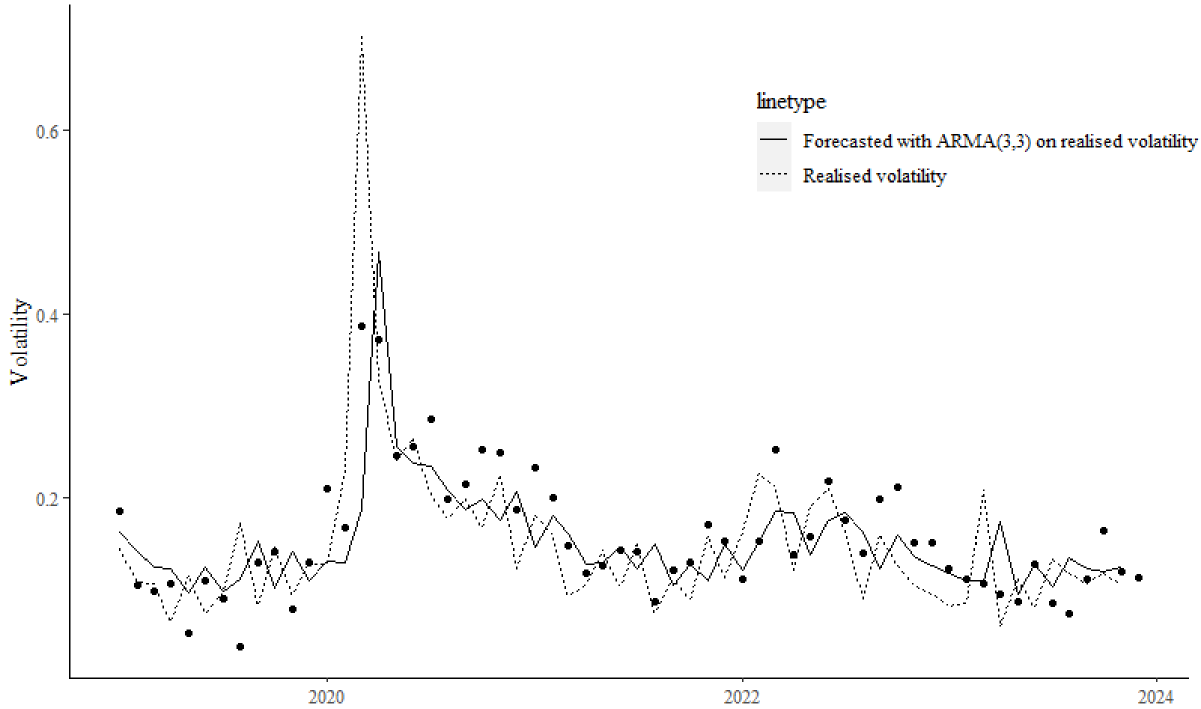

Notes: Computations by author. Volatility forecast made with an ARMA(3,3) model on realised monthly volatility. Dots are the implied volatility for the at-the-money call options on the FTSE 100 index expiring in the given month.

4.2 Volatility clustering

See table 3 for the (E)GARCH estimations using the entire sample, this means that this is not a rolling estimation, but ex post. This is done to formally assess conditional volatility dynamics such as volatility clustering. Columns contain estimations from models varying in GARCH specification and assumption about conditional return distribution. For all models, expected returns are modelled as an ARMA(1,1) process.

Table 3: Ex post (E)GARCH fits

	Model:			
	GARCH(1,1)		EGARCH(1,1)	
	(1)	(2)	(3)	(4)
α	0.267***	0.383***	0.400***	0.364***
	(0.084)	(0.118)	(0.093)	(0.116)
β	0.679***	0.284*	0.885***	0.656***
	(0.115)	(0.147)	(0.070)	(0.127)
γ			-0.033	-0.211***
			(0.038)	(0.077)
Observations	455	455	455	455
Conditional return distribution	Normal	Student's t	Normal	Student's t
Log Likelihood	784.043	807.233	787.060	811.968
AIC	-3.420	-3.518	-3.429	-3.534

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviation in brackets. Data consists of monthly continuously compounded returns of the FTSE 100 index in the period 1986-2023. Expected returns are modelled as an ARMA(1,1) process.*

Among other things, these results show that there is volatility clustering found in this sample by all model specifications because the betas are significantly positive. Furthermore, the negative value of gamma in the EGARCH estimations suggests that negative return shocks do have a higher effect on volatility than positive shocks indeed. Although, gamma is not found to be significant in the model in column three. However, when assuming a student's t-distribution of returns, it is found that gamma is significantly negative. This suggests that the effective coefficient associated with a positive return shock $(\alpha + \gamma)$ is smaller than that of negative return shocks $(\alpha - \gamma)$. Given that the model in the fourth column has the higher log likelihood, it is concluded that there is asymmetry found in this sample.

4.3 GARCH forecasting

Using the four (E)GARCH models described in table 4, volatility forecasts are made one step ahead in a rolling manner. This way, the volatility in period t is forecasted using information up to and including period $t-1$. These forecasts are then used to predict implied volatility in an OLS regression, just like was done with ARMA forecasts before. Here, the unconditional volatility level as calculated in every period is also included as an independent variable to test whether this is relevant for implied volatility forecasting. See appendix 1 and 2 for the regression results for all four GARCH model specifications. These results show that the most effective forecast of implied volatility is produced by the exponential GARCH(1,1) model under the assumption that conditional returns follow a student's t-distribution. In table 4, only the results of this model are highlighted, as well as the ARMA model that was previously identified as the best forecasting method of implied volatility excluding GARCH models. This means that column one in table 4 has the same results as column three from table 2, so that it is easier to compare these to the EGARCH forecasts.

Table 4: Explaining current implied volatility with EGARCH forecasts

	Dependent variable:			
	$Implied\ volatility_t$			
	(1)	(2)	(3)	(4)
$ARMA(3,3)^{RV}_{t,t-1}$	0.909***			
	(0.113)			
$EGARCH(1,1)_{t,t-1}$		0.980***		0.739***
		(0.148)		(0.138)
$Unconditional\ volatility_{t-1}$			31.401***	21.732***
			(5.337)	(4.733)
Constant	0.020	0.014	-4.310***	-3.044***
	(0.018)	(0.023)	(0.760)	(0.666)
Observations	59	59	59	59
Conditional return distribution		Student's t	Student's t	Student's t
GARCH specification		Exponential	Exponential	Exponential
R-squared	0.532	0.434	0.378	0.589
Adjusted R-squared	0.524	0.424	0.367	0.574
Residual SE	0.048	0.053	0.055	0.045
F statistic	64.722***	43.750***	34.615***	40.122***
Log Likelihood	96.687	91.109	88.305	100.535
AIC	-187.375	-176.218	-170.61	-193.069

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. Implied volatility of FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month). Realised volatility calculated on monthly basis using continuously compounded daily returns of FTSE 100 index.*

These results show that EGARCH forecasts alone (column two) are actually not a better predictor of implied volatility than $ARMA(3,3)^{RV}_{t,t-1}$ as indicated by a lower R-squared and log-likelihood. Perhaps surprisingly, it appears that unconditional volatility is a highly significant predictor of implied volatility as well. Ultimately, column four shows that combining conditional and unconditional volatility forecasts actually creates the superior model for explaining implied volatility, even better than the ARMA model.

Figure 4 shows the estimated unconditional volatility and forecasted conditional volatility plotted over time together with monthly realised volatility and implied volatility.

Figure 4: Volatility and EGARCH(1,1) volatility forecast over time, with implied volatility (dotted)

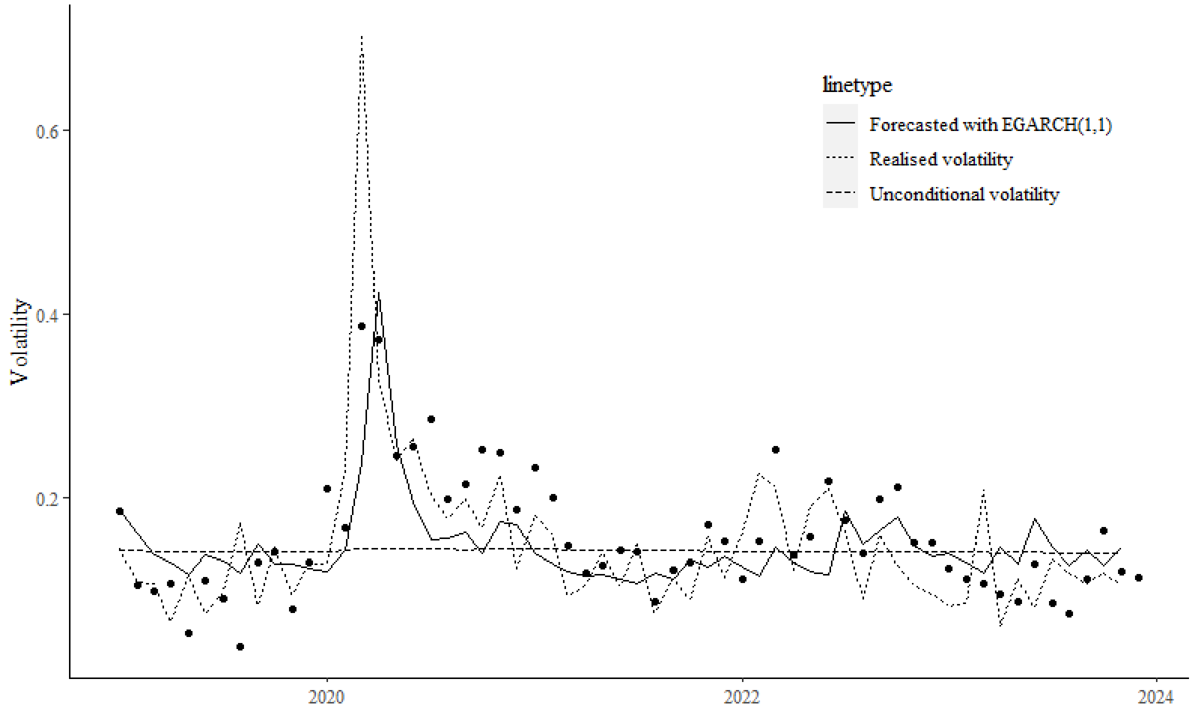

Notes: Computations by author. EGARCH(1,1) model with expected return modelled as ARMA(1,1) assuming conditional returns to follow a student's t-distribution used for forecasting. Unconditional volatility computed based on EGARCH(1,1) estimation results. Realised volatility calculated on monthly basis using continuously compounded daily returns of FTSE 100 index. *Dots are the implied volatility for the at-the-money call options on the FTSE 100 index expiring in the given month.*

In this graph it seems that unconditional volatility is constant over time. Although the scale makes it nearly invisible, unconditional volatility does truly vary over time as GARCH parameters are updated with new information. Not only that, this small variance is actually highly correlated with implied volatility as was noted earlier.

4.4 Comparing the best models

The previous analysis surprisingly suggests that (E)GARCH might actually not be a better method of forecasting implied volatility than ARMA on realised volatility. To confirm whether EGARCH forecasts have more or less information about implied volatility in them than forecasts made based on historical realised volatility, these forecasts are put together in regressions. As was described in section three, this is done to assess whether the significance of the coefficient associated with either forecasting method is diminished by the inclusion of the other. If one forecast would lose its significance, that would be an indication that it contains inferior information about future implied volatility. See table 5 for the estimation results and explanatory power of these models.

Table 5: Combining forecasts to assess relative information content

	Dependent variable:		
	$Implied\ volatility_t$		
	(1)	(2)	(3)
$ARMA(3,3)^{RV}_{t,t-1}$	0.750***	0.699***	0.409*
	(0.211)	(0.128)	(0.221)
$EGARCH(1,1)_{t,t-1}$	0.224		0.378
	(0.252)		(0.238)
$Unconditional\ volatility_{t-1}$		15.475***	17.139***
		(5.226)	(5.260)
Constant	0.011	-2.150***	-2.399***
	(0.021)	(0.733)	(0.740)
Observations	59	59	59
Conditional return distribution	Student's t	Student's t	Student's t
GARCH specification	Exponential	Exponential	Exponential
R-squared	0.538	0.595	0.613
Adjusted R-squared	0.522	0.581	0.592
Residual SE	0.048	0.045	0.044
F statistic	32.640***	41.157***	29.035***
Log Likelihood	97.102	100.979	102.308
AIC	-186.205	-193.958	-194.617

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. Implied volatility of FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month). Realised volatility calculated on monthly basis using continuously compounded daily returns of FTSE 100 index.*

These results show that the significance of EGARCH forecasts is completely wiped out by controlling for ARMA(3,3) forecasts based on realised volatility. This means that, quantitatively, GARCH forecasts have no added value. Interestingly, the unconditional volatility does retain its significance. Again, this might suggest that investors calculate the volatility based on both long- and short-term dynamics.

For these models, the VIF is calculated for all independent variables. (Multi)collinearity is a concern because the forecasts are naturally correlated with each other and this may compromise the accuracy of estimation results. Fortunately, the VIF is always below the threshold of 5. The highest VIF is 4.485 for $ARMA(3,3)^{RV}_{t,t-1}$ in the model in column three, which is on the higher end. This might cause the standard deviations to be inflated a bit in this model. Even so, ARMA is still significant on the 10% level.

Additionally, see figure 5 for scatterplots relating implied volatility to the two forecasting methods that are considered. In these plots a diagonal line is drawn with slope one and intercept zero. If these forecasting methods would perfectly explain implied volatility, all observations would be on this line. The plots show that these forecasting methods are very similar, but that observations are a little bit closer to the line for the ARMA forecasts.

Figure 5: Forecasted and implied volatility, scatterplots.

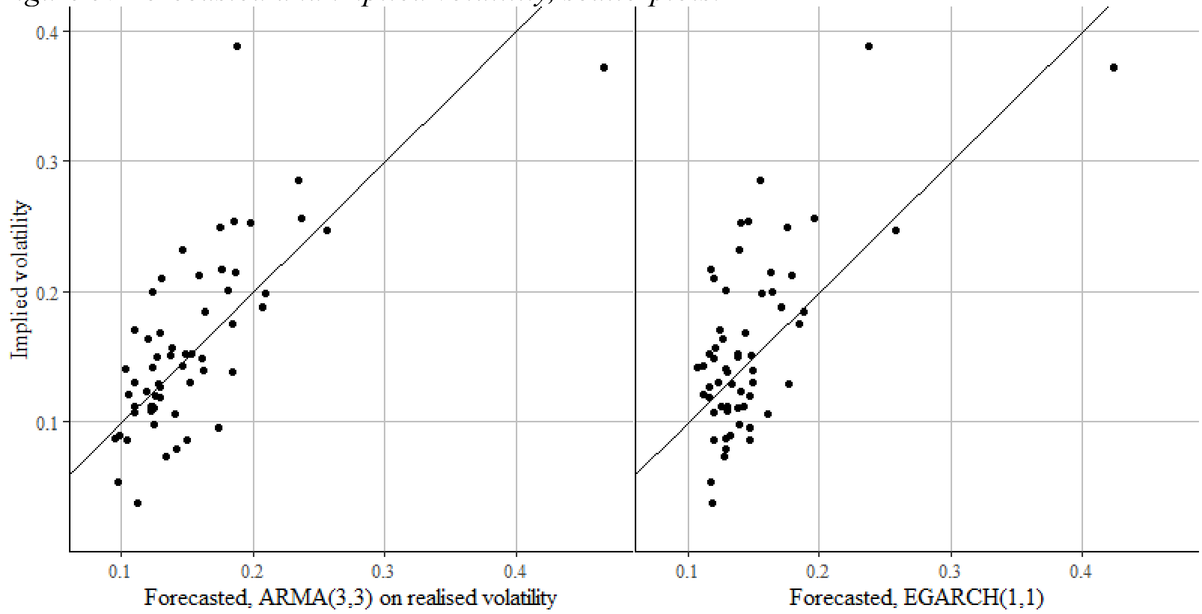

Notes: Computations by author. Implied volatility of at the money, close to maturity options on FTSE 100 index. Realised volatility calculated on monthly basis using daily returns. Line is where $x = y$, not a fitted line.

5 Conclusion

Option valuation is dependent on what one assumes about the variance of the underlying product. This is difficult, because variance cannot be neutrally observed, even after the fact. In the very definition of variance that one uses are assumptions nested that reveal how financial markets are thought to function. It is also possible to infer what investors assume about volatility given the value that they attach to options. Observations such as the volatility smile and the term structure of volatility reveal that the market does not assume volatility to be constant as the Black-Scholes framework assumes. This is not surprising, because it is well documented that volatility of equity prices is actually conditionally heteroskedastic. This volatility clustering can be modelled and used for forecasting by GARCH models. Given these features of financial markets, the research question is whether GARCH models can forecast implied volatility well.

If market participants really do take volatility clustering into account when evaluating options, then volatility forecasts based on GARCH models should correlate highly with implied volatility. However, even if there would be no causality, GARCH volatility forecasts and implied volatility would still be correlated. Using monthly data on FTSE 100 index options in the period 2019-2023, it is tested whether GARCH forecasts have better explanatory power of implied volatility than alternative forecasting methods.

The alternatives forecasting methods used are rolling ARMA models using two different kinds of data and various parameter specifications. First of all, using ARMA forecasting on the time series of implied volatility itself produces relatively poor forecasts of future implied volatility. However, using monthly realised volatility to forecast future implied volatility works very well. More specifically, ARMA(3,3) on realised volatility is selected as the best method for forecasting future implied volatility in the absence of GARCH models.

Using estimation results from various GARCH specifications, the existence of volatility clustering is verified once more. The exponential GARCH model reveals that negative return shocks have a greater effect on volatility in subsequent periods than positives do, the so-called *sign effect*. Furthermore, GARCH models based on non-normally distributed conditional returns (specifically the student's t-distribution) fit the data very well, incorporating the excess kurtosis observed in the market. All GARCH specifications are used to make rolling volatility forecasts. Using these, the EGARCH(1,1) model with non-normal conditional returns is selected as the best GARCH specification for forecasting implied volatility. Although this model is the best GARCH specification, it does not lead to better forecasts of implied volatility than ARMA. This is investigated further by combining EGARCH and ARMA forecasts in the same model, which diminishes the significance of EGARCH forecasts. This suggests that GARCH is actually not the best way to forecast implied volatility.

Another result is that unconditional volatility is a great predictor of future implied volatility as well. A possible interpretation of this is that investors do take conditional forecasts into consideration, but they do not completely disregard long-term dynamics. The deviation from the unconditional volatility is sticky.

In conclusion, volatility forecasts made using GARCH models do not contain more information about future implied volatility than historical realised volatility and thus it is not a better predictor. Therefore, the hypothesis has to be refuted. However, it should be noted that the implied volatility observed in the market is likely the result of more than just volatility forecasts of any kind because a lot of noise remains unexplained. Whether that is because investors use a mix of short term forecasts and long term volatility levels or because different market participants use various conflicting methods, is unknown.

6 References

Ahoniemi, K. (2006). Modeling and forecasting implied volatility: an econometric analysis of the VIX index. In *Helsinki Center for Economic Research* (Vol. 17, Issue 129).

Akgiray, V. (1989). Conditional Heteroscedasticity in Time Series of Stock Returns: Evidence and Forecasts. *The Journal of Business*, *62*(1), 55. https://doi.org/10.1086/296451

Black, F., & Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *Journal of Political Economy*, *81*(3), 637–654. https://doi.org/10.1086/260062

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Brock, W. A., & Hommes, C. H. (1997). A Rational Route to Randomness. *Econometrica*, *65*(5), 1059. https://doi.org/10.2307/2171879

Brock, W. A., & Hommes, C. H. (1998). Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control*, *22*(8–9), 1235–1274. https://doi.org/10.1016/s0165-1889(98)00011-6

Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *The Journal of Computational Finance*, *2*(4), 61–73. https://doi.org/10.21314/jcf.1999.043

Choi, S., & Wohar, M. E. (1992). Implied Volatility in Options Markets and Conditional Heteroscedasticity in Stock Markets. *Financial Review*, *27*(4), 503–530. https://doi.org/10.1111/j.1540-6288.1992.tb01329.x

Day, T. E., & Lewis, C. M. (1992). Stock market volatility and the information content of stock index options. *Journal of Econometrics*, *52*(1–2), 267–287. https://doi.org/10.1016/0304-4076(92)90073-Z

Eagle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation Author (s): Robert F . Engle Published by : The Econometric Society Stable URL : http://www.jstor.org/stable/1912773 REFERENCES Linked references ar. *Econometrica*, *50*(4), 987–1007.

Fama, E. F. (1965). The Behavior of Stock-Market Prices. *The Journal of Business*, *38*(1), 34. https://doi.org/10.1086/294743

Fung, William, K. ., & Hsieh, David, A. (1991). Empirical Analysis of Implied Volatility: Stocks, Bonds and Currencies. In *Manuscript, Duke Uniersity* (Vol. 37).

Gaunersdorfer, A., & Hommes, C. (2007). A nonlinear structural model for volatility clustering. *Long Memory in Economics*, 265–288. https://doi.org/10.1007/978-3-540-34625-8_9

Gaunersdorfer, A., Hommes, C. H., & Wagener, F. O. O. (2008). Bifurcation routes to volatility clustering under evolutionary learning. *Journal of Economic Behavior & Organization*, *67*(1), 27–47. https://doi.org/10.1016/j.jebo.2007.07.004

He, C., Teräsvirta, T., & Malmsten, H. (2002). MOMENT STRUCTURE OF A FAMILY OF FIRST-ORDER EXPONENTIAL GARCH MODELS. *Econometric Theory*, *18*(4), 868–885. https://doi.org/10.1017/S0266466602184039

He, X., Li, K., & Wang, C. (2016). Volatility clustering: A nonlinear theoretical approach.

Journal of Economic Behavior & Organization, *130*, 274–297. https://doi.org/10.1016/j.jebo.2016.07.020

Heston, S. L. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies*, *6*(2), 327–343. https://doi.org/10.1093/rfs/6.2.327

Heynen, R., Kemna, A., & Vorst, T. (1994). Analysis of the Term Structure of Implied Volatilities. *The Journal of Financial and Quantitative Analysis*, *29*(1), 31. https://doi.org/10.2307/2331189

Hull, J., & White, A. (1987). The Pricing of Options on Assets with Stochastic Volatilities. *The Journal of Finance*, *42*(2), 281–300. https://doi.org/10.1111/j.1540-6261.1987.tb02568.x

Hwang, S., & Satchell, S. E. (1998). Implied Volatility Forecasting: A Comparison of Different Procedures Including Fractionally Integrated Models with Applications to UK Equity Options Soosung. In J. Knight & S. Satchell (Eds.), *Forecasting Volatility in the Financial Markets*. Butterworth-Heinemann.

Jacobsen, B., & Dannenburg, D. (2003). Volatility clustering in monthly stock returns. *Journal of Empirical Finance*, *10*(4), 479–503. https://doi.org/10.1016/S0927-5398(02)00071-3

Kaizoji, T., Bornholdt, S., & Fujiwara, Y. (2002). Dynamics of price and trading volume in a spin model of stock markets with heterogeneous agents. *Physica A: Statistical Mechanics and Its Applications*, *316*(1–4), 441–452. https://doi.org/10.1016/S0378-4371(02)01216-5

Macbeth, J. D., & Merville, L. J. (1979). An Empirical Examination of the Black-Scholes Call Option Pricing Model. *The Journal of Finance*, *34*(5), 1173. https://doi.org/10.2307/2327242

Mandelbrot, B. (1963). The Variation of Certain Speculative Prices. *The Journal of Business*, *36*(4), 394. https://doi.org/10.1086/294632

Merton, R. C. (1973). Theory of Rational Option Pricing. *The Bell Journal of Economics and Management Science*, *4*(1), 141. https://doi.org/10.2307/3003143

Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica*, *59*(2), 347. https://doi.org/10.2307/2938260

Noh, J., Engle, R. F., & Kane, A. (1994). Forecasting Volatility and Option Prices of the S&P 500 Index. *The Journal of Derivatives*, *2*(1), 17–30. https://doi.org/10.3905/jod.1994.407901

Pagan, A. (1996). The econometrics of financial markets. *Journal of Empirical Finance*, *48*(1–3), 15–102. https://doi.org/10.1016/0927-5398(95)00020-8

Rubinstein, M. (1985). Nonparametric Tests of Alternative Option Pricing Models Using All Reported Trades and Quotes on the 30 Most Active CBOE Option Classes from August 23, 1976 Through August 31, 1978. *The Journal of Finance*, *40*(2), 455. https://doi.org/10.2307/2327895

Sheikh, A. M. (1991). Transaction Data Tests of S&P 100 Call Option Pricing. *The Journal of Financial and Quantitative Analysis*, *26*(4), 459. https://doi.org/10.2307/2331406

Xu, X., & Taylor, S. J. (1994). The Term Structure of Volatility Implied by Foreign Exchange

Options. *The Journal of Financial and Quantitative Analysis*, *29*(1), 57. https://doi.org/10.2307/2331190

Xu, X., & Taylor, S. J. (1995). Conditional volatility and the informational efficiency of the PHLX currency options market. *Journal of Banking and Finance*, *19*(5), 803–821. https://doi.org/10.1016/0378-4266(95)00086-V

Yamamoto, R. (2010). Asymmetric volatility, volatility clustering, and herding agents with a borrowing constraint. *Physica A: Statistical Mechanics and Its Applications*, *389*(6), 1208–1214. https://doi.org/10.1016/j.physa.2009.11.038

Youssefmir, M., & Huberman, B. A. (1997). Clustered volatility in multiagent dynamics. *Journal of Economic Behavior and Organization*, *32*(1), 101–118. https://doi.org/10.1016/s0167-2681(96)00021-2

Zhang, J. E., & Xiang, Y. (2008). The implied volatility smirk. *Quantitative Finance*, *8*(3), 263–284. https://doi.org/10.1080/14697680601173444

7 Appendix

Appendix 1: Explaining current implied volatility with GARCH volatility forecasts.

	Dependent variable:					
	$Implied\ volatility_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Conditional\ volatility\ forecast_{t-1}$	0.763***		0.514***	0.735***		0.529***
	(0.106)		(0.151)	(0.119)		(0.114)
$Unconditional\ volatility_{t-1}$		1.641***	0.776**		14.339***	9.974***
		(0.255)	(0.345)		(2.422)	(2.280)
Constant	0.041**	-0.182***	-0.082	0.049**	-2.278***	-1.615***
	(0.018)	(0.053)	(0.057)	(0.019)	(0.411)	(0.381)
Observations	59	59	59	59	59	59
Conditional distribution	Normal	Normal	Normal	Student's t	Student's t	Student's t
GARCH specification	Standard	Standard	Standard	Standard	Standard	Standard
R-squared	0.477	0.421	0.520	0.400	0.381	0.553
Adjusted R-squared	0.468	0.411	0.503	0.389	0.370	0.537
Residual SE	0.051	0.053	0.049	0.054	0.055	0.047
F statistic	51.965***	41.437***	30.356***	37.953***	35.059***	34.577***
Log Likelihood	93.421	90.424	95.97	89.361	88.448	98.03
AIC	-180.843	-174.848	-183.939	-172.722	-170.896	-188.06

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month). Forecasts made using a GARCH(1,1) model, expected returns modelled as an ARMA(1,1) process.*

Appendix 2: Explaining current implied volatility with EGARCH volatility forecasts.

	Dependent variable:					
	$Implied\ volatility_t$					
	(1)	(2)	(3)	(4)	(5)	(6)
$Conditional\ volatility\ forecast_{t-1}$	0.910***		0.896***	0.980***		0.739***
	(0.139)		(0.147)	(0.148)		(0.138)
$Unconditional\ volatility_{t-1}$		2.643*	0.384		31.401***	21.732***
		(1.512)	(1.240)		(5.337)	(4.733)
Constant	0.024	-0.277	-0.037	0.014	-4.310***	-3.044***
	(0.022)	(0.249)	(0.199)	(0.023)	(0.760)	(0.666)
Observations	59	59	59	59	59	59
Conditional distribution	Normal	Normal	Normal	Student's t	Student's t	Student's t
GARCH specification	Exponential	Exponential	Exponential	Exponential	Exponential	Exponential
R-squared	0.428	0.051	0.429	0.434	0.378	0.589
Adjusted R-squared	0.418	0.034	0.409	0.424	0.367	0.574
Residual SE	0.053	0.068	0.053	0.053	0.055	0.045
F statistic	42.636***	3.056*	21.028***	43.750***	34.615***	40.122***
Log Likelihood	90.781	75.847	90.832	91.109	88.305	100.535
AIC	-175.562	-145.694	-173.664	-176.218	-170.61	-193.069

*Note: *p<0.1; **p<0.05; ***p<0.01. Standard deviations in brackets. FTSE 100 index options from 2019-2023. Sample consists of only at-the-money and nearest maturity options for every period (month). Forecasts made using an EGARCH(1,1) model, expected returns modelled as an ARMA(1,1) process.*