

Parrák, Radovan

Working Paper

The economic valuation of variance forecasts: An artificial option market approach

IES Working Paper, No. 09/2013

Provided in Cooperation with:

Charles University, Institute of Economic Studies (IES)

Suggested Citation: Parrák, Radovan (2013) : The economic valuation of variance forecasts: An artificial option market approach, IES Working Paper, No. 09/2013, Charles University in Prague, Institute of Economic Studies (IES), Prague

This Version is available at:

<https://hdl.handle.net/10419/83303>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Institute of Economic Studies, Faculty of Social Sciences
Charles University in Prague

The Economic Valuation of Variance Forecasts: An Artificial Option Market Approach

Radovan Parrák

IES Working Paper: 09/2013



Institute of Economic Studies,
Faculty of Social Sciences,
Charles University in Prague

[UK FSV – IES]

Opletalova 26
CZ-110 00, Prague
E-mail : ies@fsv.cuni.cz
<http://ies.fsv.cuni.cz>

Institut ekonomických studií
Fakulta sociálních věd
Univerzita Karlova v Praze

Opletalova 26
110 00 Praha 1

E-mail : ies@fsv.cuni.cz
<http://ies.fsv.cuni.cz>

Disclaimer: The IES Working Papers is an online paper series for works by the faculty and students of the Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague, Czech Republic. The papers are peer reviewed, but they are *not* edited or formatted by the editors. The views expressed in documents served by this site do not reflect the views of the IES or any other Charles University Department. They are the sole property of the respective authors. Additional info at: ies@fsv.cuni.cz

Copyright Notice: Although all documents published by the IES are provided without charge, they are licensed for personal, academic or educational use. All rights are reserved by the authors.

Citations: All references to documents served by this site must be appropriately cited.

Bibliographic information:

Parrák, R. (2013). “The Economic Valuation of Variance Forecasts:
An Artificial Option Market Approach” IES Working Paper 09/2013. IES FSV. Charles University.

This paper can be downloaded at: <http://ies.fsv.cuni.cz>

The Economic Valuation of Variance Forecasts: An Artificial Option Market Approach

Radovan Parrák^a

^aIES, Charles University Prague
E-mail: rado.parrak@gmail.com

August 2013

Abstract:

In this paper we compared two distinct volatility forecasting approaches. GARCH models were contrasted to the models which modelled proxies of volatility directly. More precisely, focus was put on the economic valuation of forecasting accuracy of one-day-ahead volatility forecasts. Profits from trading of one-day at-the-money straddles on the hypothetical (artificial) market were used for assessing the relative volatility forecasting accuracy. Our contribution lies in developing a novel approach to the economic valuation of the volatility forecasts - *the artificial option market with a single market price* – and its comparison with the established approaches. Further on, we compared the relative intra- and inter-group volatility forecasting accuracy of the competing model families. Finally, we measured the economic value of richer information provided by high-frequency data. To preview the results, we show that the economic valuation of volatility forecasts can bring a meaningful and robust ranking. Additionally, we show that this ranking is similar to the ranking implied by established statistical methods. Moreover, it was shown that modelling of volatility directly is strongly dependent on the volatility proxy in place. It was

also shown, as a corollary, that the use of high frequency data to predict a future volatility is of considerable economic value.

Keywords: GARCH, Realized volatility, economic loss function, volatility forecasting

JEL: C58

Acknowledgements:

The author acknowledges the support of the Grant Agency of Charles University (GAUK), under the grant project no. 820213.

The findings, interpretations and conclusions expressed in this paper are entirely those of the author and do not represent the views of any of the above-mentioned institutions.

The author would like to express his gratitude towards Dr. Barunik for his invaluable comments during the process of writing this paper.

I. Introduction

This paper aims to compare two distinct volatility forecasting approaches, namely GARCH models and models internalizing volatility proxies directly, using an economic criterion. To do so, two economic ranking criteria are used. First, a common type of artificial option market, where option portfolios are being traded pair-wise anonymously, is utilized. Second, a framework in which a single market price emerges is developed. In both market settings, individual traders are equipped with distinct volatility forecasting techniques which are subsequently used in pricing their volatility option portfolios. Finally, traders meet at the market where pair-wise trades are conducted. Experiment is run for the course of one year, after which traders (i.e. volatility models) are ranked according to their cumulative profits.

Precise volatility forecasts are of crucial importance in many financial areas. For instance, volatility serves as an input when pricing options, hedging, in utility maximisation, in constructing the optimum portfolio of financial assets, or when computing Value-at-Risk.

However, contrary to the academic dispute over the predictability of financial returns, the issue of the returns' volatility predictability was settled more than two decades ago. The conspicuous persistence in the conditional variance of financial asset returns gave rise to voluminous literature on the topic of volatility modelling and forecasting. Many ARCH-type specifications became both academic and financial industry workhorse models, due mainly to their flexibility and relative parsimony. Similarly, a so-called direct modelling and forecasting of volatility became attractive. Once volatility became 'observable', it could be directly used as a model input in AR(FI)MA-like specifications. Rather naturally, the question of which forecasting model from the set of existing models was the best one arose.

This question was thoroughly explored within the boundaries of the individual strands of literature available (e.g. ARCH models, models internalizing volatility proxies, etc.), yet the comparison was majorly limited to using statistical ranking methodologies. Hansen and Lunde (2005), along with many others, compared the forecasting accuracy of univariate GARCH volatility models, using the superior predictive ability (SPA) test. In a similar fashion, though this time in a multivariate setting, Laurent, Rombouts and Violante (2011) compared 125 multivariate GARCH specifications. With regard to the forecasting accuracy of models using the volatility proxies, Corsi (2009) compared the forecasting ability of his HAR-RV model to the ARFIMA model of Andersen, Bollerslev, Diebold and Labys (2003), using root mean squared error loss functions.

When it comes to the cross-family comparison of volatility forecasting approaches, however, the literature is limited. Koopman, Jungbacker and Hol (2005) compared the predictive ability of forecasting models internalizing a volatility proxy vis-à-vis GARCH and stochastic volatility models, using statistical ranking methodology (SPA framework). Arguably, the main factor which limits this strand of literature from burgeoning is the inherent latency of the true volatility process. When comparing volatility forecast accuracy in statistical terms one, typically, contrasts the statistical fit of the predicted volatilities to a predefined volatility benchmark, hence becoming the major driver of the final ranking.

Going one step further, i.e. comparing the forecasting accuracy across different volatility model families using an economic ranking criterion, one finds a blank spot in the literature. Henceforth, the added value of this paper lies in the comparison of two volatility forecasting strategies, namely a) a "*direct*" volatility forecasting, internalizing volatility proxies (e.g. the

Realized Volatility) and b) an ARCH-like forecasting, treating the volatility as a latent process. Yet, unlike the statistical forecasting accuracy evaluation, we evaluate the accuracy of the volatility forecasts using observable quantities (i.e. daily returns) and an economic criterion, i.e. profits from trading in the artificial option market. More precisely, the one-day-ahead volatility forecasts from various volatility models serves as an input parameter to the Black & Scholes option pricing formula, which is subsequently used to price the portfolios of options (straddles). We create a market where each trader is assigned one volatility forecasting model with which he prices the stock option and, given the prices quoted by the rest of the traders, he sets up a buy/sell strategy. Arguably, a trader who can systematically make profit at the expense of other traders *ceteris paribus*, uses superior volatility forecasts.

There is reason to believe that comparing volatility forecasts in terms of dollar profits/losses is of more interest to a real-world investor than comparing them in terms of dimensionless quantities, i.e. accumulated point-to-point deviations from the ex-post benchmark. The eminent contribution of this paper lies in developing a novel approach to the economic valuation of volatility forecasts - *the artificial option market with a single market price* – as opposed to the rest of the literature oriented on economic valuation of the volatility forecasts (e.g. Engle, Hong and Kane (1990), Bandi et al. (2008)).

Finally, many authors have shown that volatility forecasts based on an “information-rich” variance proxy are highly accurate in statistical terms (e.g. Andersen et al. (2003), Koopman, Jungbacker and Hol (2005), Corsi (2009)). Therefore, the additional aim of this work is to verify this finding on our dataset and measure its value in economic terms.

To preview the results, we show that the economic valuation of volatility forecasts can produce a meaningful and robust ranking. In addition, we show that this ranking is similar to the ranking implied by established statistical methods. Moreover, it is shown that the modelling of volatility directly is strongly dependent on the volatility proxy in place. Hence, as a corollary, it is not surprising that the use of high frequency data to predict a future volatility is of considerable economic value.

The rest of the paper is organized as follows: Chapter 2 outlines the methodology, both the volatility models used as well as the working mechanism of the artificial market. Chapter 3 is devoted to the data description. In Chapter 4 we present the results. The following chapter outlines the robustness checks performed and, finally, in Chapter 6 we present our conclusion.

II. Methodology

In this chapter, we present the cornerstone building blocks of our experiment: namely the conditional volatility models, the option pricing framework and, finally, the artificial market settings used.

However, before outlining the conditional volatility model, we define the very fundamentals of our market settings. The daily return for day t is defined as the difference of logs of prices for day t and $t-1$ and is denoted by r_t . Similarly, intraday returns for day t computed over an intraday interval Δ are defined as: $r_{t-i\Delta} = \ln(P_{t-i\Delta}) - \ln(P_{t-(i+1)\Delta}), \forall i \in \{0, 1, 2, \dots, j\}$, where P denotes the price of an asset and j is the number of intraday intervals. Next, the unconditional volatility (of returns) is defined as a square root of variance of asset returns. The variance of return is the second unconditional moment of the return process, denoted by σ^2 . The conditional variance h_t is defined as: $\sigma_{t|t-1}^2 = h_t = E\left(\left(r_t - \mu_{t|t-1}\right)^2 \mid I_{t-1}\right)$, where I_{t-1} is the information set, containing all the relevant information about the return process up to time $t-1$. Similar to its unconditional counterpart, through the conditional volatility we understand the square root of conditional variance.

1. Volatility models

We use two distinct approaches for the modelling of a one-day-ahead conditional volatility. As already sketched in the previous chapter, the one-day-ahead conditional volatility forecasts will be used as an input for the pricing of options, which are, in turn, traded on the artificial market. The more detailed working mechanism of the option market will be provided in the subsequent chapters.

The volatility models used in this paper are divided into two families – a) GARCH models and b) volatility models that internalize volatility proxies. We will further on in this paper refer to the latter as “RV models”, since the Realized Volatility (RV) is used as a volatility proxy. Each family contains a baseline model and its variations. These variations account for particular properties of volatility, such as the leverage effect of the returns on the conditional volatility and/or its persistence. We have chosen the models for each group in a homogeneous way, such that they are inter- and intra-group comparable.

1.1 GARCH MODELS

Several GARCH models are discussed in this subsection. We focused on the most common and parsimonious models from the GARCH universe. Since GARCH models model the conditional volatility parametrically with respect to a return process, this process must be defined beforehand. Various specifications for the mean equation of the GARCH process are used throughout the literature, spanning from simple zero conditional mean assumption (e.g. Andersen et al. (1988)) to autoregressive processes (Laurent & Giot (2004)). We assumed the GARCH-in-mean type of conditional mean equations (similar to, for example, Engle et al. (1987)), where returns are specified as, $r_t = \mu + \lambda \cdot \sqrt{h_t} + \varepsilon_t$ with $\varepsilon_t = \sqrt{h_t} \cdot z_t$, where μ stands for unconditional mean, h_t is the conditional variance process which is to be further specified, $z_t \sim \mathcal{N}(0,1)$ and λ corresponds to the risk premium. This specification allows us to indirectly capture the risk premium, if any, of the asset returns by an additional parameter in the return equation. Nowadays, its variations are widely used in contingent claim pricing

literature (e.g. Duan(1995), Hardle & Hafner (2000), Christoffersen & Jacobs (2004), Christoffersen et al. (2008)).

As for the modelling of the conditional volatility, we hereby opted for a number of distinct specifications from the GARCH family. The formalized definition of all volatility models used in this paper is provided in Table 5, Table 6 and Table 7 of Appendix A. In the following two paragraphs we rationalize our model selection.

The GARCH(1,1) of Bollerslev (1986) (A.1) is a widely used tool in financial econometrics. The conditional variance in this specification is a linear function of past shocks (ε_{t-1}^2) and past conditional variances (σ_{t-1}^2). Based on something akin to an “adaptive learning mechanism” for the variance process, the GARCH(1,1) is able to model a relatively long memory without the cumbersome estimation of heavily lagged models¹. A modification to the simple GARCH(1,1) model was proposed by Taylor (1986) and Schwert (1989). The TS-GARCH (A.4) models the standard deviation as a function of absolute past innovations and past standard deviations. This is to extenuate the effect of large shocks. Thus, unlike GARCH, the variance is a non-linear function of the squared innovations. The non-linearity also captures, to some extent, the long memory feature of the conditional variance². Another important feature of the variance of stock returns is the so-called leverage effect, i.e. an asymmetric effect of positive and negative shocks on the future conditional volatility. Neither GARCH nor TS-GARCH is able to capture this effect. Therefore, we also included leverage GARCH models, namely: GJR of (Glosten et al. (1993)), Threshold ARCH (TARCH) of Zakoian (1994), Asymmetric GARCH (AGARCH) of Engle (1990), Non-linear Asymmetric GARCH (NAGARCH) and VGARCH of Engle and Ng (1993)³. All of the five above-mentioned models incorporate the leverage effect by the additional parameter(s) in the variance equation, yet the leverage effect is modelled in a rather distinct way, as reported, for example, by Engle and Ng (1993).

After enlarging our set of models able to account for asymmetric effects, we added models capable of tackling another stylized volatility feature, i.e. the long memory (Ding et al. (1993)). The long memory can be broadly defined as a “non-decaying” persistence in the conditional volatility process. There are several possible approaches approximating the long memory of the volatility process. One approach would be to focus on extending the aforementioned GARCH models for higher orders of lagged conditional volatility. However this solution would soon make the estimation technically unfeasible. Another possibility would be to consider fractionally integrated models, i.e. FIGARCH of Baillie, Bollerslev and Mikkelsen (1996) and/or FIEGARCH of Bollerslev and Mikkelsen (1996). Yet, although, as argued by Corsi (2009), the concept of fractional integration is an elegant theoretical approach, it remains difficult to interpret economically. Moreover, as also reasoned by Christoffersen et al. (2008), option pricing with fractionally integrated models is clumsy, mainly due to problematic parameter estimation. Hence, we decided to model long memory using a more parsimonious approach of the two-component models (Ding et al. (1996)). This is also more relevant in the option pricing context (Christoffersen et. al (2008)). Ding et al. (1996) showed that a volatility of financial returns usually has longer memory than the one

¹ Bollerslev (1986), p. 309.

² An extended version of this model, which models the variance to the power of δ as a function of the absolute value of the innovations and powers of past variances, was proposed by Higgins and Berra (1992). We excluded this specification due to lack of convergence when estimating parameters within particular days of our data sample.

³ We started with a broader set of models, but Alt-GARCH, EGARCH, APARCH and NARCH were excluded due to estimation problems.

implied by the GARCH model. To correct this deficiency, a model which combines a short-term conditional volatility component (s_t) with a long-term component (I_t - capable of internalizing the long-memory) is set up. For the long-term component, the Integrated GARCH of Engle & Bollerslev (1986) (IGARCH) was chosen⁴. Apart from the GARCH(1,1)-IGARCH(1,1) specification, we included other two-component models, differing in their short-term components, in order to account for the aforementioned leverage effect. Hence, we were able to construct several long memory models with asymmetric news impact curves, yet simple enough to estimate. All of the two-component models are reported in Table 6.

1.2. REALIZED VOLATILITY MODELS

The second type of models used in this paper internalizes the volatility proxies (measures), namely the *Realized Volatility* (RV) in this paper. To rationalize, if the volatility could be measured, we would be able to use it as an input to the ordinary autoregressive models (i.e. ARIMA models), typically used for the forecasting of levels of financial time series. Following Andersen et al. (2001), daily realized volatility (RV) can be defined as the sum of squared intraday returns, more formally:

$$RV(t, \Delta) = \sum_{i=0}^{1/\Delta} r_{t-i\Delta}^2. \quad (2.1)$$

In the seminal paper of Andersen and Bollerslev (1998), it was shown that RV is an unbiased and asymptotically consistent volatility estimator, i.e. the difference between RV and integrated volatility – the “true” volatility - converges to zero in probability as the interval between intraday observations decreases:

$$\text{plim}_{\Delta \rightarrow 0} \left(\int_0^1 \sigma_{t+\tau}^2 dt - \sum_{i=0}^{1/\Delta} r_{t-i\Delta}^2 \right) = 0. \quad (2.2)$$

Arising from eq. (2.1), the true volatility becomes “measurable” and hence it can be modelled directly using autoregressive models. In the following chapter, we describe and provide rationale for the RV models used in this paper.

The RV is internalized in two ways, i.e. within parametric and non-parametric models. Starting with the parametric models, we hereby opted for ARMA on RV. In the context of RV, the ARMA(1,1) model can be, to some extent, thought of as a parallel to the GARCH(1,1). The difference between equation (A.14) and (A.1) lies in the different “MA” parts of the process⁵. While the news term in the ARMA model is connected directly to the lagged variance proxy, the squared news term in the GARCH model refers to the innovations in the mean equation (returns). We included this model in the scope owing to its lag structure being similar to that of GARCH(1,1).

The non-parametric group consists of several specifications to match the characteristics of the GARCH models in the previous section. We considered the non-parametric Heterogeneous AR Realized Volatility models (HAR-RV) of Corsi (2009) and Corsi and Reno (2009). Corsi (2009) reported that his parsimonious model (A.15) performs surprisingly well in the forecasting exercise. This AR-type model is capable of mimicking the long memory properties of the volatility when using RVs recorded at different time intervals (typically daily,

⁴ Note that since the coefficients of the long-term component sum up to one: a) the long-memory is implied by construction, since the conditional variance is fully determined by past squared innovations and past conditional variances, and b) the variance process is not weakly stationary (the unconditional variance does not exist). Yet, Ding et al. (1996) showed that the two component GARCH(1,1)-IGARCH(1,1) model is covariance stationary.

⁵ Of course, also the AR parts are different, but if we assume RV to be a good proxy for the latent variance, then this difference is negligible.

weekly and monthly). In other words, an infinitely lagged ARCH, which could be used to mimic the long memory, is truncated here to three parameters, namely $\beta_0, \beta_1, \beta_2$ ⁶. By imposing suitable constraints on the parameters, a GARCH(1,1) can also be transformed into an ARCH(∞)⁷. If we assume zero daily mean return and RV to be equal to the daily squared return (i.e. the RV(low)), then a *tight linkage* between the GARCH(1,1) and HAR-RV(low) models becomes apparent. We also included a simplified version of the HAR-RV models – an AR(1) model on RV with a leverage effect, (A.17), for the sake of comparison with leverage models from the GARCH family. In the class of RV models that are able to approximate the long memory, we consider the HAR-RV (A.15) as a non-parametric counterpart to the two-component GARCH-IGARCH. We also included a version of the model generalized for the presence of asymmetric effects – the Leverage HAR-RV (A.16) of Corsi & Reno (2009), representative of a counterpart to the asymmetric two-component models.

2. Estimation

All GARCH model parameters were estimated by Gaussian maximum likelihood⁸. We assumed error terms to be conditionally normal with mean zero and variance h_t , thus the following log-likelihood function was maximised with respect to θ :

$$\ln L(\theta) = n \cdot \left(\frac{\ln(2\pi)}{2} \right) - \sum_{t=1}^n \frac{\ln(h_t(\theta))}{2} - \frac{1}{2} \sum_{t=1}^n \frac{(r_t - \mu(\theta))^2}{h_t(\theta)}, \quad (2.3)$$

where $h_t(\theta)$ represents the particular GARCH model and $\mu_t(\theta) = \mu + \lambda \cdot \sqrt{h_t(\theta)}$. For $t=1$, we set h_t equal to the sample variance. During the estimation, constraints for covariance stationarity were imposed for each model⁹. Once we had the estimated parameters at hand, we forecasted the one-step-ahead variance by recursive fitting of the model, using the rolling estimation window of 1000 days. The recursive estimation allows us to capture the parameter shifts, if any. In particular, we are able to indirectly capture the time varying risk premium.

Regarding the RV models, a simple OLS method was applied to arrive at model parameter estimates. To measure efficiency gains from the additional information when constructing the volatility proxy using data sampled at finer frequencies, we estimated all the models using RV(low), RV(mid), RV(high), totalling 11 models¹⁰. Similar to Laurent and Giot (2004), we modelled the logs of the realized volatility ($\ln(RV)$). This was based on the premise of by Andersen (1998), who found the logarithmic RV to be approximately Gaussian. This feature is important for the forecasting exercise, where it secures that volatility forecasts are not systematically biased (as might be the case for RV, the distribution of which is typically skewed to the right (Thomakos & Wang (2003))). Finally, once we had estimated the models and forecasted the one-step-ahead *logarithmic* RV, we needed to convert these quantities back into the original scale in order to be able to use them in the option pricing. We assumed

⁶ ARCH(∞) is defined as: $h_t = \alpha + \sum_{i=1}^{\infty} \delta_i \cdot \varepsilon_{t-i}^2$

⁷ For details, consult Bollerslev (1986).

⁸ Interior-point algorithm was used for finding a maximum of the corresponding likelihood function.

⁹ Stationarity conditions enforced within the parameter estimation are available upon request.

¹⁰ We denote RV(high), RV(mid) and RV(low) the realized volatility computed using 5 minute intraday returns, 30 minute intraday returns and daily returns, respectively. We did not include ARMA(low) due to estimation problems.

errors from (2.4) to be normally distributed, i.e. $\eta_t \sim N(0, \sigma_\eta^2)$. Then, following Granger and Newbold (1986), Laurent and Giot (2004) and Lutkepohl and Xu (2009), we transformed our logarithmic forecasts accordingly:

$$\hat{RV}_{t+1|t} = \exp\left(\ln \hat{RV}_{t+1} + \frac{1}{2} \hat{\sigma}_\eta^2\right), \quad (2.4)$$

where $\hat{\sigma}_\eta^2$ is the estimated variance of the error terms of (A.14-17).

3. Option pricing formula

The volatility forecast accuracy of the outlined models is tested in the artificial option market framework, the working mechanism of which will be outlined in the next chapter. We used the well-known Black & Scholes (B/S) option pricing framework to transform volatility forecasts into option prices.¹¹ We note that the B/S pricing formula implicitly assumes a constant variance of the returns over the life of the option (σ^2), which is in contrast to option pricing in the context of the time-varying conditional variance. However in our artificial option market exercise, agents forecast *one-day-ahead* conditional variances, which are, in turn, used as an input to the B/S formula to price options with *one day* to maturity. Here, we assume that at the time of the option price determination, each trader perceives his volatility forecast as the correct one and *constant* over the course of the next day.

It is noteworthy that inefficiencies in the pricing formula (i.e. pricing errors) can lead to biases in the absolute profits of traders, but as long as all traders use the same pricing formula, the relative ranking of the volatility forecasts should stay intact. There is, hence, a reason to trust the ability of this kind of exercise to assess the variance forecast precision and to deliver a reliable economic ranking of volatility forecasts.

By the same token, we add that there is nothing which could restrict us from assuming that all the underlying B/S framework assumptions are fulfilled in our artificial markets, due to their relative simplicity. In addition, we rather arbitrarily, albeit in accordance with the related academic literature and the model assumptions, set the risk-free interest rate equal to 0 (Engle et al. (1990), Bandi et al. (2008)). Also in our settings, similar to Engle et al. (1990) and Bandi et al. (2008), the agents trade at-the-money options (i.e. $K = 1$) on a \$1 share of the underlying asset.

The original closed-form solution for pricing a call (C) and a put (X) option in the B/S world is of the form:

$$C(S, t) = N(d_1) \cdot S - N(d_2) \cdot K e^{-r(T-t)}, \quad (2.5)$$

$$X(S, t) = N(-d_2) \cdot K e^{-r(T-t)} - N(-d_1) \cdot S, \quad (2.6)$$

where:

$$d_1 = \frac{\ln\left(\frac{S}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}}; d_2 = d_1 - \sigma\sqrt{T-t} \quad (2.7)$$

¹¹ For more details consult Black and Scholes (1973)

and $S, K, r, \sigma, T, t, N(\cdot)$ refers to the spot price of the underlying asset, the strike price of the option, the risk-free interest rate, the volatility of returns of the underlying asset, the maturity date, the current date and the normal cumulative distribution, respectively.

In our setting, the Black & Scholes option price formula is simplified to:

$$C(S, t) = N(d_1) - N(d_2), \quad (2.8)$$

where:

$$d_1 = \frac{\ln(1) + \left(\frac{h_t}{2}\right)}{\sqrt{h_t}} = \frac{h_t}{2}; d_2 = d_1 - h_t = -\frac{h_t}{2}, \quad (2.9)$$

with h_t being the (conditional) volatility¹² at time t implied by one of the competing volatility models from the previous chapter.

Thus:

$$C = X = N\left(\frac{h_t}{2}\right) - N\left(-\frac{h_t}{2}\right) = 2N\left(\frac{h_t}{2}\right) - 1, \quad (2.10)$$

which follows from the put-call parity.

4. Artificial Option Markets

Volatility forecasts accuracy has been predominantly studied in statistical terms (e.g. Andersen and Bollerslev (1998), Hansen and Lunde (2005), Laurent, Rombouts and Violante (2009)). In short, let us use n competing volatility models. As we obtain the volatility forecasts implied by the models, we are able to contrast them with the ex-post volatility estimates. Deviations from the ex-post volatility are in turn evaluated using a pre-specified loss function (i.e. by imposing a particular structure on the penalties for deviations from the benchmark). Then, the n models are ranked in ascending order according to their average losses¹³. In the final stage, the validity of the ranking is tested using a predefined level of confidence. This is commonly executed either by using the test for Superior Predictive Ability (SPA) of Hansen (2005) or the Model Confidence Set (MCS) of Hansen, Lunde and Nason (2003).

On the other hand, it is a rather modest area of empirical finance focused on comparison which is based on non-statistical loss functions. Being based on the economic valuation criterion, economic loss functions do not rely on the knowledge of the “true” variance. Usually, portfolio profits/standard deviations or option prices/implied volatilities are used for the ranking of the volatility models. In this paper, we use an artificial (hypothetical) option market as proposed by Engle et al. (1990).

The artificial option market consists of several hypothetical option traders. Each trader uses a different volatility model, from the set described in section 1 of this chapter, to forecast the one-day-ahead variance. The forecasted variance is then used as an input for pricing at-the-money European put and call options on the \$1 share of the underlying stock, with the

¹² This is since $S_t = 1, K = 1$ and $r = 0$.

¹³ The model with the lowest accumulated loss is ranked as the first.

maturity of one day¹⁴. For the pricing of options, the B/S formula was used. After pricing the put and call options, each agent creates a straddle, which he then trades on the artificial market with the rest of the traders.

A straddle is a portfolio of options (contingent claim/investment strategy) that *allows the holder to profit from the volatility of the underlying asset*. A long/short straddle is a long/short position in both a call and a put option on the underlying asset¹⁵.

To illustrate our trading algorithm, let us assume that two traders (*A* and *B*) at time t , priced European call options (C) on the same underlying asset with maturity at time $t+1$, given their individual volatility forecasts. Let us further assume, without the loss of generality, that $C_A > C_B$. The price of the European put option (X) is equal to the price of the call (C), owing to our assumptions on the risk-free interest rate ($r=0$), the strike price ($K=\$1$) from Section 3 and the put-call parity. The fact that the straddle can be perceived as a portfolio of put and call options in turn implies that $P_A > P_B$ ¹⁶. In other words, *A* perceives *B*'s straddle as under-priced, or equivalently, *B* perceives *A*'s price as overpriced¹⁷. Consequently, a trade takes place whenever *A* and *B* meet on the market. Note that the only factor making the prices of the straddles different (in our setting) is the forecasted volatility. Hence, if trader *A* systematically realizes the profit over trader *B*, then he is arguably using better volatility forecasts. In the subsequent sections, we extend this mechanism to the market of more than two traders.

Two distinct market settings are used. In the first setting, straddles are traded anonymously, which implies a market without a single market price. This has been of common practice in the related literature (Engle et al. (1990), Bandi et al. (2008), Chan et. al (2009)). Conversely, the second setting is based on the assumption of a single market price. To our best knowledge, volatility forecasts have not yet been compared under this type of economic valuation.

Note that both types of volatility forecast comparisons do not rely on the knowledge of the “true” volatility. Unlike the volatility forecast ranking based on the statistical criteria, we rank the models by using an ex-post observable quantity, namely daily returns of the underlying asset.

4.1 The artificial market with anonymous trading

Our artificial market comprises 27 traders. The first 24 traders were assigned 24 models discussed in Chapter 2 (11 RV models + 13 GARCH models). We also included three “indicator” traders to control for the possible pricing biases of the whole market. We constructed the *Mean* trader, whose volatility forecast was an average forecast of the 24 model based traders, on a particular day. Similarly, the *Min* and *Max* traders correspond to the minimum and maximum forecast from the set of the 24 models. For instance, high profitability of the *Max* trader would indicate that the 24 competing models were on average underestimating the variance, i.e. under-pricing the straddles.

We run the experiment on the out-of-sample period of 250 days. Let us now describe the routine for a particular day in a stepwise fashion:

In time t :

¹⁴ For the option to be at-the-money we set the strike price, $K = \$1$.

¹⁵ The two options in the straddle have the same strike price and maturity.

¹⁶ This result would hold also if assumptions on the risk-free rate and strike price are relaxed. As long as *A, B* are pricing options on the same underlying asset with different variances (but, with the same strike price and risk-free rate) the result would not change.

¹⁷ Given the individual beliefs of *A* and *B* about the future level of variance of the underlying asset.

Every trader forecasts the one-day-ahead variance with respect to the model he uses¹⁸. Traders price the one-day European put and call options according to eq. (2.10). A straddle price (P) for each trader is determined.

The traders meet at the market, where pair-wise trades take place. Trades are settled at the mid price of the ask (higher) and bid (lower) straddle price. This implies that every pair-wise trade is (possibly) executed at a different market price¹⁹. More formally, for any trader i and j :

$$P_{i,j}^{Trade} = \frac{P_i + P_j}{2}, \quad \forall i, j \in \{1, 2, \dots, 27\}; i \neq j, \quad (2.11)$$

yielding 26 trades per trader and totalling 351 pair-wise trades each day²⁰.

After all trades have been executed, each trader delta-hedges himself. When delta-hedging a plain vanilla option position, a trader enters into a counter-weighting position (ω) in the underlying asset. By doing so, he creates a risk-neutral portfolio (of an option and a underlying asset) with respect to small changes in the price of the underlying asset. For a portfolio consisting of one plain vanilla option, the weight " ω " equals the "delta of the option" (δ) - a standard finance textbook result.

In our settings (revolving around eq. 2.10), however, the delta is simplified to:

$$\hat{\delta} = \Phi\left(\frac{1}{2}h_t\right), \quad (2.12)$$

for the long position in the call option.

This result can be derived by combining the delta with the plain vanilla call option:

$$\delta = e^{-rt}\Phi(d_1), \quad (2.13)$$

with eq. (2.9) and recalling the zero interest rate assumption. Hence, the trader should go short by this quantity of stocks in order to make the portfolio delta neutral. Similarly, for the trader who goes long (buys) in the put option, we have:

$$\hat{\delta} = 1 - \Phi\left(\frac{1}{2}h_t\right), \quad (2.14)$$

i.e. the trader should go long by this quantity of stocks to attain a delta neutral position. Finally, since the straddle is a portfolio of put and call options, its hedge ratio is the difference between (2.14) and (2.13), yielding:

$$\delta_{straddle} = 1 - 2\Phi\left(\frac{1}{2}h_t\right). \quad (2.15)$$

¹⁸ Mean, Max and Min traders form their forecasts once the forecasts of the rest of the traders are known.

¹⁹ If we assume that the volatility forecast of any pair of traders is not identical. This, of course, does not hold in every pair-wise trade for indicator traders, by construction. In cases where the bid and ask price are equal, indicator traders execute the trade for this price.

²⁰ i.e. $27 \times 26 / 2$.

We note that every agent uses his variance forecast to delta-hedge himself. This means that possible profits (or losses) made by this strategy are riskless and purely determined by a speculation, i.e. richer (or weaker) information about the future volatility. This would not hold in a market consisting of traders equipped with perfect information and foresight due to no-arbitrage reasons²¹. In fact, our exercise implicitly assumes imperfect information on the market, arising from various beliefs about future volatility implied by the volatility models.

In time $t+1$:

Profits/losses for each individual trader (i) are computed and stored. The daily pair-wise profit of the trader who bought the straddle is:

$$\text{Profit}_{i,t+1} = |r_{t+1}| - P_{i,j,t}^{\text{trade}} + r_{t+1} \cdot \left(1 - 2\Phi\left(\frac{1}{2}h_{i,t+1|t}\right) \right). \quad (2.16)$$

Similarly, the daily profit of the straddle seller i is:

$$\text{Profit}_{i,t+1} = P_{i,j,t}^{\text{trade}} - |r_{t+1}| - r_{t+1} \cdot \left(1 - 2\Phi\left(\frac{1}{2}h_{i,t+1|t}\right) \right). \quad (2.17)$$

The total daily profit/loss of the individual agent is the sum of all pair-wise profits/losses from the individual trades, divided by the number of pair-wise trades executed during the day²². The routine is repeated for every day of our out-of-sample period and the daily profits are accumulated. Arising from the market assumptions of section 3, we recall that transaction costs are not deducted from the profits. This, however, should not introduce any ranking biases, since all traders would be burdened by the same transaction costs, should these be introduced.

Finally, we report two characteristics of this market setting. First, as the price of each trade is set as a mid-price of the ask and bid price, it is very likely that an individual trader is selling the same straddle at different prices within a particular day, depending on the straddle prices of the counterpart traders (i.e. their future volatility beliefs). This would be possible only in the case of *anonymous* trading, otherwise a single market price would emerge. This issue is addressed in our second market setting.

Further, as reported by Bandi et al. (2008), traders whose prices are clustered around the median price have a higher potential to be profitable, even if their volatility forecasts are not accurate enough. This is due to the fact that they are selling and buying roughly equal numbers of straddles from both sides of the market, i.e. from straddle sellers and straddle buyers. We discuss the issue in more detail in the second market setting.

4.1 The artificial market with a single market price

The second market setting of our experiment is built on seemingly more realistic assumptions. Trades are not executed anonymously; rather, all traders reveal their straddle prices and a single market price is consequently formed. For simplicity, we assume that the market price (P_{market}) equals the median price from the set of straddle prices for a particular day. This can be, to some extent, related to the market price formed by demand and supply forces²³.

²¹ In the case of perfect information and foresight, all of the agents would unify their beliefs about the future variance, implying a single market price for a contingent claim.

²² i.e. the total daily profits are scaled by the number of trades to prevent having artificially large absolute returns caused by the size of the market.

²³ When one realizes that the demand and supply side of the market consist of the same agents.

This modification of the market set-up introduces a new feature to the trading mechanism – decision-making. To illustrate this feature, let us consider two traders (A and B) in the market of n traders ($n > 2$). Without loss of generality, we assume that the market price of the straddle was formed as a median price (P_{market}) and $P_A > P_B > P_{market}$, i.e. both agents are on the same side of the market. Now, if A and B traded according to the set-up of the previous section, i.e. A buys the straddle from B (B sells the straddle to A), then B would enter the trade with a negative expected profit. This is due to the fact that both A and B would like to *buy* straddles for P_{market} in this case, perceiving them to be under-priced. Hence, we allow traders to choose whether or not to trade. Thus by construction, two traders will trade if, and only if, they are from two distinct sides of the market, i.e.:

$$P_i \geq P_{market} > P_j, \forall i, j \in \{1, 2, \dots, 27\}, i \neq j \quad (2.18)$$

Again, we provide a stepwise description of the daily trading routine for this market setting:

In day t :

The first and second steps are identical to the market setting outlined in the previous section.

The third step is similar to the previous set-up whereby traders meet at the market to buy/sell straddles, but a single market price is formed:

$$P_{market} = median(P_1, P_2, \dots, P_{27}). \quad (2.19)$$

Thus, every trade is executed for a market price. Moreover, the pair-wise trade is executed if, and only if, both traders come from distinct sides of the market. Step 4 is also identical to the previous set-up.

In $t+1$:

Step 5: Profits/losses of trader i from pair-wise trades are then computed as:

$$\text{Profit}_{i,t+1} = |r_{t+1}| - P_{market} + r_{t+1} \cdot \left(1 - 2\Phi\left(\frac{1}{2}h_{i,t+1|t}\right) \right), \quad (2.20)$$

for the straddle buyer, and

$$\text{Profit}_{i,t+1} = P_{market} - |r_{t+1}| - r_{t+1} \cdot \left(1 - 2\Phi\left(\frac{1}{2}h_{i,t+1|t}\right) \right), \quad (2.21)$$

for the straddle seller.

The rest of the simulation runs along the lines of the previous setting.

This modification of the trading setup has two advantages. Firstly, the assumption of a single price on the market is economically more meaningful. Secondly, it diminishes the potential profitability of the traders, whose straddle prices are clustered around the median price²⁴. Here, traders that are, for instance, slightly above the median price are not able/willing to sell their straddles to the traders with even higher straddle prices and, by not doing so, compensate for the possible losses arising from buying the straddles from the other side of the market.

²⁴ Also in this set-up, the trader, whose straddle price strictly equals the median price, is potentially profitable despite possible inaccuracy of his volatility forecasts.

III. Data

The dataset that we consider consists of five-minute tick prices of British Petroleum plc (BP). The BP series is a heavily traded stock of the NYSE index. This should minimise the adverse effects of the microstructure noise, which can in turn bias the volatility proxy. The data set spans from August 2003 to August 2008 (1250 days). Weekends, holidays and early closing days are excluded from the data sample. For each day, there are 78 five-minute returns implying a total of 97 500 five-minute returns for the given stock.

We split this period into two sub-samples (as indicated by the red vertical line in Figure 1), i.e. the starting estimation period, spanning from August 2003 to August 2007, and the out-of-sample evaluation period, spanning from August 2007 to August 2008.

The starting estimation period serves as an initial estimation window (information set) of 1000 observations for the parameter estimation and consequent out-of-sample volatility forecasting. For each following out-of-sample volatility forecast, a rolling window of 1000 observations is used to estimate the volatility model parameters. We note that also the estimation window used for the last out-of-sample volatility forecast comprises mainly return observations from the calm (estimation) period. This part of our data corresponds to a period of relatively low volatility (the lower panel in Figure 1) and positive average daily returns (Table 1).

The artificial market trading is run on the out-of-sample period (i.e. the data point to the right of the red vertical line) for one year – from August 2007 until August 2008. Figure also shows that the two periods differ not only in the average daily returns but also in the average volatility. The volatility and returns measured in our sample throughout the time, i.e. the RV, is depicted in the lower and upper panels respectively in Figure 1. The out-of-sample period, i.e. are the area to the right of the red line, has a significantly higher volatility and features the tumbling market (the average daily return changes from positive to negative as we move from the “starting estimation” period to the “out-of-sample” period) as evidenced by both the lower panel of Figure 1 and Table 1. These changes in the market environment were caused mainly by the impact of 2008 financial crisis. Since the out-of-sample period was more volatile, we also implicitly tested the relative flexibility of the volatility models and their speed of adjustment to the environment.

Figure 1: Data overview




Table 1: Descriptive statistics of the BP daily returns

	<i>Mean</i>	<i>Minimum</i>	<i>Maximum</i>	<i>Std. Dev.</i>	<i>Skewness</i>	<i>Ex. Kurt.</i>
<i>Entire period (1251 obs.)</i>	-0.011	-3.691	5.76	0.922	0.082	1.759
<i>"In" sample (1001 obs.)</i>	0.013	-2.535	3.4982	0.827	0.013	0.166
<i>"Out" sample (250 obs.)</i>	-0.108	-3.69	5.76	1.2236	0.288	2.063

IV. Results

Hereby, we report the results of our artificial market exercise. First, we provide the results of several ranking methods based on the statistical metric to provide a benchmark for our economic evaluation. Next, we report both the rankings resulting from the anonymous trading on the artificial option market as well as the trading on the market with a single price.

1. Ranking based on statistical loss functions

We used three consistent statistical loss functions for the volatility forecast comparison. We ranked the models in ascending order with respect to their average losses over the out-of-

sample period. The average loss (L) of the i -th's model forecast over the sample period is defined as:

$$L^i(l^i, t) = \frac{\sum_{t=1}^n l^i(\hat{\sigma}_t^2, h_t)}{n}, \forall t \in \{1, 2, \dots, n\}, i \in \{1, 2, \dots, k\}. \quad (4.1)$$

Where n denotes the number of days of the period over which the losses are accumulated, $\hat{\sigma}_t^2$ denotes the benchmark measure (i.e. the volatility proxy) and $l^i(\hat{\sigma}_t^2, h_t)$ is the form of the loss function which needs to be specified. We use the following consistent loss functions of Patton (2011):

Mean Squared Error (MSE):

$$l(\hat{\sigma}_t^2, h_t) = (\hat{\sigma}_t^2 - h_t)^2 \quad (4.2)$$

Asymmetric loss function with a heavier penalty for the under-prediction (QLIKE):

$$l(\hat{\sigma}_t^2, h_t) = \log h_t + \frac{\hat{\sigma}_t^2}{h_t} \quad (4.3)$$

Asymmetric loss function with a heavier penalty for the over-prediction:

$$l(\hat{\sigma}_t^2, h_t, b) = \frac{1}{(1+b)(2+b)} (\hat{\sigma}_t^{2b+4} - h_t^{b+2}) - \frac{1}{b+1} h_t^{b+1} (\hat{\sigma}_t^2 - h_t), \quad (4.4)$$

where b is a scalar parameter influencing shape of the loss function ($b \in \mathbb{R} - \{-1, -2\}$). We rather arbitrarily set $b=1$.

The RV constructed from five-minute returns is used as an ex-post benchmark volatility measure. In keeping with common academic practice, we compute the average loss for an individual model using the symmetric loss functions of a Mean Squared Error form (4.2). Additionally, we rank the models using asymmetric, albeit consistent, loss functions, which impose higher penalties for volatility over- and under-prediction, (4.4) and (4.3) respectively. The asymmetric loss functions are used to control for the possible systematic penalization scheme of a particular artificial option market setup. In other words, if we found that the economic ranking is close enough to that of (4.3), we would conclude, given our data, that our economic ranking criterion favours (on average) upward-biased volatility forecasts/models. If, on the other hand, we observed the economic ranking to be close to the MSE loss function, it would support the argument of the symmetric ranking of our economic volatility ranking approach.

The rankings of the volatility models are reported in Table 2. The ranking is reported in ascending order, i.e. 1 denotes the best model.

Table 2: Statistical rankings of the volatility models

	4.2	4.3	4.4
GARCH traders	MSE	Q-LIKE	
GARCH	14	15	6
GJR	12	12	14
TARCH	17	20	18
TS-GARCH	18	21	17

NAGARCH		15	14	7
VGARCH		19	19	19
AGARCH		4	11	2
2C-GARCH-IGARCH		11	13	5
2C-ALTGARCH-IGARCH		20	18	20
2C-GJR-IGARCH		13	8	16
2C-TARCH-IGARCH		8	9	8
2C-VGARCH-IGARCH		21	17	21
2C-TSGARCH-IGARCH		16	16	11
RV traders				
HAR-RV	high	1	2	1
	mid	6	6	9
	low	25	25	25
Leverage HAR-RV	high	2	1	4
	mid	5	5	10
	low	26	26	26
Short memory HAR-RV	high	10	10	15
	mid	22	22	22
	low	24	24	24
ARMA	high	3	3	3
	mid	9	7	13
Indicator traders				
Mean		7	4	12
Min		23	23	23
Max		27	27	27

We denote traders according to the volatility model they used (e.g. GARCH trader). The two component models are denoted by an acronym “2C”.

Inspecting Table 2, the high-frequency RV models rank highest, i.e. are the most accurate, throughout the spectrum of the loss functions, while the low-frequency RV models deliver rather inaccurate forecasts²⁵. These findings were also proven by the MCS test, at a 95% level of confidence²⁶.

Akin to Corsi (2009), we observed the high frequency HAR-RV model to be the most precise within the class of RV models. Moreover, we confirmed the results of Andersen et al. (2003) on our dataset when we found that even a simple AR(1)/ARMA(1,1) model on high-frequency data is on average more accurate than GARCH models on daily frequencies. The value arising from the additional information provided by high-frequency data can also be demonstrated by controlling for the particular RV model from Table 2 and varying the RV frequency. For instance, when the frequency of the HAR-RV model is changed from “high” to “mid”, its ranking drops from 1st to 6th²⁷. The leverage HAR-RV models are also highly accurate, yet their leverage feature does not make them superior to the symmetric HAR-RV. On the other hand, the long memory-mimicking feature of the HAR-RV models seems to be essential. Both HAR-RV specifications outperformed the short memory-like RV specifications (on the high and mid frequencies).

When focusing solely on the GARCH models, we find that the parsimonious GARCH(1,1) ranks lower than some of the more flexible models in this class. More concretely, the Asymmetric GARCH ranked 4th (under MSE) among all models, and 1st in the GARCH group²⁸. Also, the GJR model ranked higher than GARCH. This may suggest the presence of

²⁵ The corresponding average losses of the individual models are available upon request.

²⁶ MCS test results are not provided due to space limitation, however they are available upon request.

²⁷ Under the MSE loss function.

²⁸ Note that AGARCH is overall second when we use a loss function which penalizes for over-prediction.

asymmetric effects of shocks on the conditional volatility, described in Chapter 2. Finally, we note that both the Model Confidence Set Test and the Test for Superior Predictive Ability justified AGARCH as statistically the most accurate predictor within our class of GARCH models, at a 95% level of confidence.

Our observation that the leverage effects matter for GARCH models but appear not to be pivotal for RV models may seem surprising. However, note first that the leverage effect is modelled differently within the two model families. In the RV family, the leverage parameters $(\gamma_1, \gamma_2, \gamma_3)$ are determined *directly* from the observed values of daily, weekly and monthly (negative) returns respectively. On the other hand, in GARCH models, the value of the leverage parameter (γ) is *indirectly* dependent on the parameters of the return specification (μ, λ) ²⁹. Furthermore, it seems that the intraday frequency, at which RV is constructed, also makes a difference, as argued in the subsequent text.

2. Ranking based on economic valuation

Hereby, we present the major contribution of the paper, i.e. the results of our artificial market exercise. Table 3 presents the rankings of the volatility models from both market settings. Volatility models are ranked with respect to the accumulated profits in descending order.

2.1 ANONYMOUS TRADING

In general, Table 3 outlines a similar picture of the ranking as given by the MSE loss functions of Table 2, i.e. high and mid frequency RV models clearly outperform GARCH models. This also suggests that the ranking given by our artificial option market is, to a large extent, consistent with the MSE ranking. Being similar to that of MSE, the ranking based on anonymous straddle trading arguably does not systematically penalize traders either for under- or over-prediction of future volatility. Due to the relatively high rank of the Max trader (7), we infer that the whole market was, on average, under-estimating the variance.

In line with the statistical ranking, traders who construct their volatility forecasts using high-frequency data are ranked the highest. In this context, it is implied that traders endowed with richer information about future volatility, profit at the expense of traders with less information. This is clear when comparing the same RV models on the different intraday sampling frequencies. As we move from high to low frequency, performance deteriorates. It must be noted, however, that the higher accumulated profits of the RV traders result solely from the richer information available and not as a consequence of the higher accuracy of the RV models as such. This is demonstrated by comparing the two model families, when the intraday sampling frequency is fixed on a daily level. This corresponds to the comparison of RV(low) models with GARCH models, as advocated in Chapter 2. Almost all GARCH models outperformed the RV models at this level of aggregation of the data. Therefore, we conclude that using high-frequency data when forecasting volatility provides considerable economic value for the trader³⁰.

The “winner” of the artificial market exercise with anonymous trading is the ARMA(high) model. Using the information from the high-frequency data and the adaptive learning mechanism, this specification excels in the valuation exercise, especially since we consider short-term volatility forecasts and a turbulent economic environment. We note that also in this type of economic evaluation, as remarked by Fleming et al. (2001), *the higher statistical accuracy of the volatility forecasts does not necessarily bring higher economic profits*. Under the MSE loss function HAR-RV(high) ranked highest, but in economic terms it is

²⁹ This is given by the fact that: $\varepsilon_t = r_t - \mu - \lambda \cdot \sqrt{h_t}$

³⁰ For instance, the best “low-frequency” trader (GJR-IGARCH) earned only 34% of the profits of the best trader (ARMA(high)).

outperformed by ARMA(high). The reason for this seemingly surprising result lies in the different types of loss functions used. While MSE compares the volatility forecast to the ex-post benchmark volatility proxy in terms of the absolute point-to-point deviations, the economic loss function is defined by multiple pair-wise comparisons of competing models (i.e. relative accuracy). Nevertheless, we believe that for a real-world investor, profits earned from competition-based trading (expressed in familiar units such as dollars) may be more appealing than the statistical fit expressed in dimensionless quantities.

Finally, we note that this result can be, to some extent, driven by the length of the forecasting horizon. Had this exercise been adjusted for trading of longer maturity straddles, i.e. the volatility models providing dynamic h-step-ahead forecasts, the short-term “adaptability” of ARMA might not have been sufficient to outweigh the long memory property of the HAR-RV models. However, this issue is left for future research.

When comparing the models on the basis of characteristics, some rather interesting findings are evidenced. The long memory of the variance is approximated both by the HAR-RV and the two-component GARCH models³¹. However, it seems that the HAR-RV(high) and HAR-RV(mid) models are able to use this information more effectively. While, the long memory HAR-RV model outperforms its short memory counterpart (on high and mid frequency), in the case of the GARCH family, the result is reversed³². Here, except for the GJR-IGARCH model, all of the short memory GARCH models rank higher than the two-component models. Moreover, the long memory HAR-RV(low) is also relatively inferior to the short memory HAR-RV(low) model. One possible reason for this lies in the fact that the long memory of a variance becomes more apparent as the return sampling frequency increases. Therefore, when we fit the long memory models to the volatility as a function of low-frequency data, the long memory seems not to be the prevailing feature. At the same time, GARCH models are, by construction, unable to react to the rapid volatility changes, which are present in our sample period³³. This inability is reflected even more heavily in the two-component GARCH models (if the weight on the long-term component is not zero), which may also relatively distort the model’s performance in a volatile environment. In conclusion, the short-term characteristics (e.g. leverage effects) of the GARCH models outweighed the long memory characteristics of the two-component GARCH models in our artificial market setting.

The leverage GARCH models, in general, outperformed symmetric models, which is in line with the statistical ranking. This result is much clearer for the single component than for the two-component models. However, the leverage effect does not seem to be pivotal in the long memory HAR-RV models. A trader using the leverage HAR-RV(high/mid/low) accumulated slightly lower profits than its symmetric counterpart. By the same token, we add that these rather small profit differences may not be statistically significant.

To sum up, in our dataset we found that when volatility is estimated on the low frequencies, leverage effects seem to play a crucial role in volatility forecasting (GARCH and RV(low) models), while at mid and high frequencies these become less important. As the return sampling frequency used to compute RV increases: a) RV traders become more profitable (i.e. the volatility forecasts are more accurate with respect to the traders using low-frequency data), and b) capturing the long memory feature has a higher economic value than leverage effects.

Table 3: Economic rankings of the volatility models

Anonymous trading		Single market price	
Rank	Acc. Profit	Rank	Acc. Profit

³¹ For instance, the weights of the long-term component in the GARCH-IGARCH, TARCH-IGARCH and GJR-IGARCH models are 0.46, 0.33 and 0.30 respectively.

³² Note that a similar pattern can be found in the statistical rankings in **Table 2**.

³³ This inability to follow rather “jumpy” volatility paths is given by GARCH parameter restrictions.

GARCH traders					
GARCH	15	-5.9	22	-45	
GJR	10	10.1	11	15.43	
TARCH	9	12.6	12	9.32	
TS-GARCH	21	-9.3	17	-34	
NAGARCH	14	-5.4	24	-52.9	
VGARCH	23	-12.5	20	-41	
AGARCH	11	9.68	13	-0.4	
2C-GARCH-IGARCH	17	-8.3	19	-36.4	
2C-ALTGARCH-IGARCH	19	-8.7	23	-46.6	
2C-GJR-IGARCH	8	12.8	10	18.13	
2C-TARCH-IGARCH	18	-8.65	15	-12.45	
2C-VGARCH-IGARCH	16	-7.2	16	-29.8	
2C-TSGARCH-IGARCH	22	-12	21	-45.5	
RV traders					
HAR-RV	high	2	32.4	2	52.22
	mid	3	31.3	3	48.45
	low	24	-27.4	-	-
Leverage HAR-RV	high	5	29.1	6	35
	mid	4	29.3	4	40.6
	low	25	-28.8	-	-
Short memory HAR-RV	high	12	8.65	9	20.7
	mid	13	-4.9	14	-9.46
	low	20	-9	-	-
ARMA	high	1	37.4	1	62.3
	mid	6	28.3	7	33.4
Indicator traders					
Mean	26	-32.5	8	31.9	
Min	27	-43.28	18	-34.66	
Max	7	25.9	5	36.9	

2.2 THE MARKET WITH A SINGLE PRICE

The previous market setting compared volatility forecasts on the basis of pair-wise trades conducted at mid prices. In this setting, trades are conducted at a single market price. Introducing this modification makes trading more realistic: when trades are executed only between the pairs of traders from different sides of the market. This measure prevents us from obtaining artificially good rankings for traders whose straddle prices, though possibly inaccurate, are on average in the centre of the price distribution. In fact, we can perceive this artificial market variation as another economic “loss” function, but with a different system of penalisation. The second column of Table 3 reports the results of the ranking based on trading with a single market price. Note that we do not remark upon any change on the highest ranks, compared to the anonymous trading setting. Also, the general results of the previous section hold, i.e. traders with better information are more profitable, leverage GARCH models on average outperform long memory models (except GJR-IGARCH) and the long memory feature prevails when using intraday information.

To this extent, the ranking is consistent with that of anonymous trading and thus comparable to the statistical ranking of MSE of Table 2. Moreover, the Max trader is even more profitable in this setting, justifying the general under-estimation of the volatility.

We now focus on the differences regarding the previous set-up. The biggest fall in the ranking with respect to the anonymous trading setting was experienced by the GARCH

trader (from 15 to 22). When inspecting the usual position of the GARCH trader in the market, we observed that his straddle prices are on average underpriced, yet close to the median price. On 50% of the days, the straddle price of the GARCH trader is below, but not more than 5 positions away from the median price. While in the anonymous trading set-up the GARCH trader partly compensated for the losses from under-pricing the straddles by buying even more under-priced ones (e.g. from the Min trader), in the current set-up this is no longer possible. Thus, in this case, the penalising scheme does seem to do a good job.

On the other hand, the biggest improvement in the ranking was experienced by the Mean trader (from 26 to 8). This result may have been effected by the interaction of the following two factors: first, after we dropped the seriously upward biased RV(low) traders, the mean price on the market decreased to an arguably more realistic level, implying that the Mean trader is more accurate³⁴; second, we are reminded that this type of evaluation favours traders able to choose the “correct” side of the market³⁵. Hence, it might be the case that the improvement was mainly caused by the ability of the Mean trader to choose, on average, the “correct” side of the market. To disentangle the two effects, we compared the ranking with and without the RV(low) traders. We found that the rank of the Mean trader remains practically unchanged whether the RV(low) traders are present or not. Therefore, we conclude that the prevailing effect on this result is the relatively correct position of the Mean trader. More precisely, the Mean trader was, on average, on the upside of the market, i.e. buying the under-priced straddles on which he made a profit.

It must be noted that the aforementioned feature of the evaluation based on the option market with a single market price can, to some extent, distort the volatility ranking (especially on the lower ranks). This is mainly true if we are interested in the volatility forecast evaluation based on the symmetric *economic* “loss” function, i.e. the parallel to MSE. In this case, it is advisable to use the anonymous trading approach. If we want to penalise traders for the volatility over- and/or under-prediction, the artificial market with a single price may be of interest to us. Unlike the case of asymmetric statistical loss functions, here the penalising rule for over- or under-prediction is more flexible. More precisely, it repeatedly switches during the sample period, from penalising under-prediction to penalising over-prediction depending on the variability of the underlying asset returns, i.e. the under-prediction / over-prediction is penalised more heavily in times of high (low) variability.

To sum up, in our dataset we observed that both economic loss functions provided rankings of volatility forecasts which appear coherent with the outcome obtained when using statistical loss functions. As in Fleming et al. (2001), we found that a better statistical fit does not, on a one-to-one basis, imply higher economic value (e.g. HAR-RV(high) vs. ARMA(high)). On the other hand, we showed that the additional information gained when using high-frequency data is of considerable economic value (RV(high) vs. GARCH). Finally, we remarked that on our dataset the prevailing properties of the daily variance differ with respect to the return frequency which it is constructed on. Leverage effects were found to prevail when the daily variance was recorded on the lowest (daily) return frequency, whereas long memory seemed to prevail when estimating daily variance by the sum of squared five-minute returns (i.e. on high sampling frequency of the returns). When using the new economic valuation approach (the option market with a single market price) we found it able to deliver a reasonable volatility ranking under reasonable market restrictions.

³⁴ The reason for this will be addressed in the following paragraph.

³⁵ We say that the trader was situated on the correct side of the market, if his belief about higher/lower future volatility (straddle price) with respect to the median volatility forecast (median price) earned him a profit.

V. Robustness checks

In this section we address the issue of transitivity of rankings. By the transitivity of rankings, we mean a degree of ranking robustness with respect to the changing number of competing models. To test for this type of robustness in our economic valuation, we start with the set of all 27 competing traders from the previous chapter. After obtaining the ranking for all traders/volatility models, we exclude the four worst performing traders/models and repeat the ranking exercise. Along this line, we continue until we obtain the ranking of the smallest set consisting of the last seven models.

Obviously when the ranking of volatility models is based on the statistical loss function, there will be no change in the models ranking as some inferior models are excluded if the benchmark volatility measure stays intact. However, our economic valuation is built on the pair-wise interaction between traders, rather than point-wise comparing each model to the fixed benchmark. As noted by Engle et al. (1990), “*there is no reason to expect that the relative profit relation across forecasts is transitive*”. In other words, it might well be the case that some of the traders earn considerably less (or more) if the counter trader on whom they profited (or lost) the most is excluded. We report the results of this exercise in Table 4.

Table 4: Transitivity of the economic rankings

		Anonymous trading						Single Market Price						
		Number of Models						Number of Models						
		27	23	19	15	11	7	27	23	19	15	11	7	
GARCH traders														
GARCH		15	15	17				22	23					
GJR		10	9	10	9	9		11	12	12	12			
TARCH		9	8	9	10	10		12	11	10	9	8		
TS-GARCH		21	21					17	18	18				
NAGARCH		14	14	16				24						
VGARCH		23	22					20	20					
AGARCH		11	10	11	11	11		13	13	14	15			
2C-GARCH-IGARCH		17	20					19	21					
2C-ALTGARCH-IGARCH		19	19	19				23	22					
2C-GJR-IGARCH		8	7	7	7	8		10	9	11	11	10		
2C-TARCH-IGARCH		18	12	12	15			15	15	19				
2C-VGARCH-IGARCH		16	16	18				16	16	17				
2C-TSGARCH-IGARCH		22	23					21	19	16				
RV traders														
HAR-RV		High	2	2	2	2	3	5	2	2	2	2	5	6
		Mid	3	3	4	5	6	6	3	3	7	7	7	5
		Low	24											
Leverage		High	5	4	3	3	2	2	6	5	4	5	6	3
HAR-RV		Mid	4	5	5	4	4	4	4	4	3	4	4	7
		Low	25											
Short memory		High	12	11	8	8	7	7	9	10	8	8	9	
HAR-RV		Mid	13	13	13	12			14	14	13	10	11	
		Low	20	18	15	14								

ARMA	High	1	1	1	1	1	1	1	1	1	1	1	2
	Mid	6	6	6	6	5	3	7	7	6	6	3	4
Indicator traders													
Mean		26						8	8	9	13		
Min		27						18	17	15	14		
Max		7	17	14	13			5	6	5	3	2	1

Rather surprisingly, both market settings appear to be, in general, robust to changes in the size of the market. However, the results of the anonymous market set-up are more convincing. Here, the transitivity of the ordering is most apparent in the case of ARMA(high), which is the best model independent of the market size. We note that the remaining models are also ranked quite robustly as their ranking changes, in the majority of cases, by not more than two places as the number of competing models decreases.

The biggest shifts in the ranking positions are naturally experienced by the indicator traders, whose forecasts become more and more precise as the set of traders becomes smaller.

It is also noteworthy that, in our dataset, the robustness of ordering is the strongest in the large sets (27-15 models) and weakens as we move to the smaller sets of traders. More precisely, we observed the most dramatic change in the ranking as we moved from a set of eleven traders to a set of seven for both market set-ups. For instance, in the case of the single market price ranking, when considering the smallest market (with only the last seven traders), the ranking of leverage HAR-RV(high) and leverage HAR-RV(mid) is reversed in comparison to the larger sets.

The illustrated robustness of both ranking approaches makes us more confident about the results of our exercise and the validity of the economic ranking approaches used. Yet, there are more robustness issues which might be interesting to investigate. For example, an identical exercise can be run on the numerous datasets, to control for data mining issues. Additionally, trading of straddles with a longer maturity can be examined in a dynamic multi-step volatility forecasts framework. One can relax the assumption of trading only at-the-money options. Also, one might be interested in the robustness of this ranking procedure with respect to the various option pricing formulas (i.e. the economic loss functions). Finally, another type of market setting might stimulate interest for future research. The next chapter presents our final conclusions.

VI. Conclusions

In this work, we compared several volatility forecasting approaches using economic valuation criteria. Latent volatility modelling, represented by (G)ARCH models, was contrasted to the direct modelling of proxies of volatility, represented by RV models. In addition to comparing two modelling approaches, we also examined models with respect to their ability to mimic some of the well-known volatility features, i.e. long memory and/or leverage effects. When constructing the daily volatility proxy, we used the five-minute intraday stock returns of British Petroleum. The relative accuracy of the volatility forecasts was assessed by pair-wise trading of straddles in the artificial option market, where every trader was assigned a particular volatility model. We used two distinct market settings – *the artificial market with anonymous trading* and *the artificial market with a single market price*.

In our dataset, we observed that in this type of forecasting exercise, additional information used while constructing the volatility proxy is of considerable economic value (e.g. RV

traders). We have also shown that direct volatility modelling is only as good as the volatility proxy used. In particular, direct modelling of volatility delivers highly accurate volatility forecasts if the volatility proxy is informative enough. On the other hand, when using lower frequency returns to construct the volatility proxy, the accuracy of the RV models deteriorated rather quickly. More precisely, when the daily squared returns (i.e. our lowest sampling frequency) were used to construct the variance proxy, RV models were by far the most inaccurate in our model universe, being systematically outperformed by GARCH models.

We introduced and tested a distinct economic valuation technique – *the artificial option market with a single market price*, which delivered a reasonable and rather robust ranking. Additionally, we observed that the economic ranking from both market settings was similar to the statistical ranking using the consistent loss functions.

We are aware, however, that the profitability of the particular volatility forecasting algorithm advocated on trading of one-day options may be of little practical interest. Therefore, future research in this area may focus on the economic valuation of the variance forecasts in the dynamic multi-step framework. Additionally, it may be interesting to run similar exercises in a different option pricing framework.

VII. References

- Andersen, T. G. & Bollerslev, T., 1998. Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts. *Journal of Empirical Finance*, p. 115–158.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P., 2001. The Distribution of Realized Exchange Rate Volatility. *Journal of The American Statistical Association*.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P., 2003. Modeling and Forecasting Realized Volatility. *Econometrica* 71, pp. 579-625.
- Baillie, R. T., Bollerslev, T. & Mikkelsen, H. O., 1996. Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics* 74, p. 3–30.
- Bandi, F. M., Russell, R. J. & Yang, C., 2008. Realized Volatility Forecasting and Option Pricing. *Journal of Econometrics*, pp. 34-46.
- Black, F., 1976. Studies in Stock Price Volatility Changes. *Proceedings of the 1976 Business Meeting of the Business and Economics Statistical Section, American Statistical Association*, pp. 177-181.
- Black, F. & Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy*, pp. 637-654.
- Bollerslev, T., 1986. Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, pp. 307-327.
- Bollerslev, T. & Mikkelsen, H. O., 1996. Modelling and Pricing Long-Memory in Stock Market Volatility. *Journal of Econometrics* 73, pp. 151-184.
- Bollerslev, T. & Wright, J. H., 2001. High-frequency Data, Frequency Domain Inference and Volatility Forecasting. *Review of Economics and Statistics*.
- Clements, A., Doolan, M., Hurn, S. & Becker, R., 2009. Evaluating Multivariate Volatility Forecasts. *NCER Working Paper*.
- Corsi, F., 2009. A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics*, pp. 174-196.
- Corsi, F. & Reno, R., 2009. HAR Volatility Modelling with Heterogeneous Leverage and Jumps. *Working paper*.
- Ding, Z., Granger, J. W. & Engle, R. F., 1993. A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance*, pp. 83-106.
- Ding, Z. & Granger, W. J., 1996. Modelling Volatility Persistence of Speculative Returns: A New Approach. *Journal of Econometrics* 73, pp. 185-215.
- Duan, J., 1995. The GARCH Option Pricing Model. *Mathematical Finance* 5, pp. 13-32.
- Engle, R. F., 1990. Discussion: Stock Market Volatility and the Crash of '87. *Review of Financial Studies*, pp. 103-106.
- Engle, R. F. & Bollerslev, T., 1986. Modeling the Persistence of Conditional Variances. *Econometric Reviews* 5, pp. 1-50.
- Engle, R. F., Hong, C. & Kane, A., 1990. Valuation of Variance Forecast with Simulated Option Market. *NBER Working paper nr. 3350*.
- Engle, R. F., Lilién, D. & Robins, R. P., 1987. Estimating time varying risk premia in the term structure: The ARCH-M model. *Econometrica* 55, pp. 391-407.
- Engle, R. F. & Ng, V. K., 1993. Measuring and Testing the Impact of News on Volatility. *The Journal of Finance*, Vol 48, No. 5, pp. 1749-1778.

- Fleming, J., Kirby, C. & Ostdiek, B., 2001. The Economic Value of Volatility Timing. *The Journal of Finance*, pp. 329-352.
- Glosten, L. R., Jagannathan, R. & Runkle, D. E., 1993. On the Relation Between Expected Value and the Volatility of the Nominal Excess Return on Stocks. *Journal of Finance*, p. 1779–1801.
- Granger, C. W. & Newbold, P., 1976. Forecasting Transformed Series. *Journal of the Royal Statistical Society*, pp. 189-203.
- Hansen, P. R., 2005. A Test for Superior Predictive Ability. *Journal of Business & Economic Statistics*, Volume 23(4).
- Hansen, P. R. & Lunde, A., 2005. A Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)? *Journal of Applied Econometrics*, pp. 873-889.
- Hansen, P. R., Lunde, A. & Nason, M., 2003. Choosing the Best Volatility Models: The Model Confidence Set Approach. *Oxford Bulletin of Economics and Statistics*, Volume 65, pp. 839-861.
- Hardle, W. & Hafner, C. M., 2000. Discrete Time Option Pricing with Flexible Volatility Estimation. *Finance and Stochastics*, pp. 187-207.
- Higgins, M. L. & Bera, A. K., 1992. A Class of Non-linear ARCH Models. *International Economic Review* 33, pp. 137-158.
- Chan, W. H., Jha, R. & Kalimipalli, M., 2009. The Economic Value of Realized Volatility in Forecasting Future Implied Volatility. *The Journal of Financial Research*, pp. 231-259.
- Christoffersen, P. & Jacobs, K., 2004. Which GARCH model for Option Valuation?. *Management Science* 50, pp. 1204-1221.
- Christoffersen, P., Jacobs, K. & Wang, Y., 2008. Option Valuation with Long-run and Short-run Volatility Components. *Journal of Financial Economics*, pp. 272-297.
- Koopman, S. J., Jungbacker, B. & Hol, E., 2005. Forecasting Daily Variability of the S&P 100 Stock Index Using Historical, Realised and Implied Volatility Measurements. *Journal of Empirical Finance*, Volume 12(3), pp. 445-475.
- Lauren, S., Rombouts, J. V. & Violante, F., 2011. On the Forecasting Accuracy of Multivariate GARCH models. *Journal of Applied Econometrics*.
- Laurent, S. & Giot, P., 2004. Modelling daily Value-At-Risk Using Realized Volatility and ARCH type models. *Journal of Empirical Finance*, pp. 379-398.
- Laurent, S., Rombouts, J. V. & Violante, F., 2009. On Loss Functions and Ranking Forecasting Performances of Multivariate Volatility Models. *ECORE discussion paper*.
- Lutkepohl, H. & Xu, F., 2009. The Role of Log Transformation in Forecasting Economic Variables. *Empirical Economics*.
- Patton, A. J., 2011. Volatility Forecast Comparison Using Imperfect Volatility Proxies. *Journal of Econometrics*, p. 160.
- Schwert, G. W., 1989. Why Does Stock Market Volatility Change Over Time?. *Journal of Finance* 44, pp. 1115-1153.
- Taylor, S. J., 1986. *Modelling Financial Time Series*. s.l.:Wiley.
- Thomakos, D. D. & Wang, T., 2003. Realized volatility in the futures markets. *Journal of Empirical Finance*, pp. 321-353.
- Zakoian, J. M., 1994. Threshold Heteroskedastic Models. *Journal of Economic Dynamics and Control* 18, pp. 931-955.

VIII. Appendix A

Table 5: Simple component GARCH model specifications

GARCH models		
GARCH	$h_t = \alpha + \beta_0 \varepsilon_{t-1}^2 + \beta_1 h_{t-1}$	(A.1)
GJR	$h_t = \alpha + \beta_0 \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 \mathbf{1}_{\varepsilon_{t-1} < 0} + \beta_1 h_{t-1}$	(A.2)
TARCH	$h_t = \alpha + \beta_0 \varepsilon_{t-1}^- - \gamma \varepsilon_{t-1}^+ + \beta_1 h_{t-1}$ $\varepsilon^- = \min(\varepsilon, 0); \varepsilon^+ = \max(0, \varepsilon)$	(A.3)
TS-GARCH	$h_t^{1/2} = \alpha + \beta_0 \varepsilon_{t-1} + \beta_1 h_{t-1}^{1/2}$	(A.4)
AGARCH	$h_t = \alpha + \beta_0 (\varepsilon_{t-1} + \gamma)^2 + \beta_1 h_{t-1}$	(A.5)
NAGARCH	$h_t = \alpha + \beta_0 (\varepsilon_{t-1} - \gamma h_{t-1}^{1/2})^2 + \beta_1 h_{t-1}$	(A.6)
VGARCH	$h_t = \alpha + \beta_0 \left(\frac{\varepsilon_{t-1}}{\sqrt{h_{t-1}}} - \gamma \right)^2 + \beta_1 h_{t-1}$	(A.7)

Table 6: Two component GARCH model specifications

Two component GARCH models:		
GARCH-IGARCH	$s_t = \alpha + \beta_0 \varepsilon_{t-1}^2 + \beta_1 s_{t-1}$ $l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$ $h_t = \rho \cdot (s_t) + (1 - \rho)(l_t)$	(A.8)
ALTGARCH_IGARCH	$s_t = \alpha + \beta_0 \varepsilon_{t-1}^2 / s_{t-1} + \beta_1 s_{t-1}$ $l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$ $h_t = \rho \cdot (s_t) + (1 - \rho)(l_t)$	(A.9)
GJR-IGARCH	$s_t = \alpha + \beta_0 \varepsilon_{t-1}^2 + \gamma \varepsilon_{t-1}^2 \mathbf{1}_{\varepsilon_{t-1} < 0} + \beta_1 s_{t-1}$ $l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$ $h_t = \rho \cdot (s_t) + (1 - \rho)(l_t)$	(A.10)
TARCH-IGARCH	$s_t = \alpha + \beta_0 \varepsilon_{t-1}^- - \gamma \varepsilon_{t-1}^+ + \beta_1 s_{t-1}$ $l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$ $h_t = \rho \cdot (s_t) + (1 - \rho)(l_t)$	(A.11)
TS-GARCH-IGARCH	$s_t^{1/2} = \alpha + \beta_0 \varepsilon_{t-1} + \beta_1 s_{t-1}^{1/2}$ $l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$ $h_t = \rho \cdot (s_t) + (1 - \rho)(l_t)$	(A.12)

VGARCH-IGARCH

$$s_t = \alpha + \beta_0 \left(\frac{\varepsilon_{t-1}}{\sqrt{s_{t-1}}} - \gamma \right)^2 + \beta_1 s_{t-1} \quad (\text{A.13})$$

$$l_t = \tau \cdot \varepsilon_{t-1}^2 + (1 - \tau) l_{t-1}$$

$$h_t = \rho \cdot (s_t) + (1 - \rho) (l_t)$$

Table 7: RV models

RV models

ARMA $\ln(RV_t) = \alpha + \phi \ln(RV_{t-1}) + \varphi \cdot \eta_{t-1} + \eta_t$ (A.14)

HAR-RV $\ln(RV_t^{(d)}) = \alpha + \beta_0 \ln(RV_{t-1}^{(d)}) + \beta_1 \ln(RV_{t-1}^{(w)}) + \beta_2 \ln(RV_{t-1}^{(m)}) + \eta_t$ (A.15)

LHAR-RV $\ln(RV_t^{(d)}) = \alpha + \beta_0 \ln(RV_{t-1}^{(d)}) + \beta_1 \ln(RV_{t-1}^{(w)}) + \beta_2 \ln(RV_{t-1}^{(m)}) + \gamma_1 \cdot r^{-(d)} + \gamma_2 \cdot r^{-(w)} + \gamma_3 \cdot r^{-(m)} + \eta_t$ (A.16)

Short-memory LHAR-RV $\ln(RV_t^{(d)}) = \alpha + \beta_0 \ln(RV_{t-1}^{(d)}) + \gamma_1 \cdot r^{-(d)} + \eta_t$ (A.17)

IES Working Paper Series

2013

1. Jiří Witzany: *A Note on the Vasicek's Model with the Logistic Distribution*
2. Tomáš Havránek, Ondřej Kokeš: *Income Elasticity of Gasoline Demand: A Meta-Analysis*
3. Jiří Witzany: *Estimating Default and Recovery Rate Correlations*
4. Petra Valíčková, Tomáš Havránek, Roman Horváth: *Financial Development and Economic Growth: A Meta-Analysis*
5. Jan Soudek, Jiří Skuhrovec: *Public Procurement of Homogeneous Goods: the Czech Republic Case Study*
6. Martin Gregor: *Electoral competition for the 2+1 electoral rule and the close alternatives*
7. Aleš Čornanič, Jiří Novák: *Signaling by Underpricing the Initial Public Offerings of Primary Listings in an Emerging Market*
8. Adrian Babin: *Bank Efficiency during the Current Economic Crisis: An International Comparison*
9. Radovan Parrák: *The Economic Valuation of Variance Forecasts: An Artificial Option Market Approach*

All papers can be downloaded at: <http://ies.fsv.cuni.cz>.

