LUISS

Dipartimento
Di Economia e Finanza

Cattedra di Asset Pricing

Modeling and forecasting EUR/USD volatility with GARCH models

Nicola Borri
RELATORE

Pietro Reichlin
CORRELATORE

Gianmarco Fabbri 699751
CANDIDATO

Anno Accademico 2019/2020

Table of contents

Introduction

The EUR/USD exchange rate is an extremely relevant economic and financial variable; In this research, GARCH models are used to model its conditional volatility. The paper is structured as follows: In section 1 the exchange rate and its significance as an economic variable are presented, together with GARCH models; I go on to present the news impact curve and consider the opportunity of introducing asymmetry in the form of GJR models; In section 2, daily data from 2002 to 2014 are used to estimate GARCH(p,q) and GJR(p,q) models with p and q up to 2 lags, and estimated models are used to produce a one-day-ahead forecast series in the year 2014. Then 5-minute intraday data for the year 2014 are used to compute the Realized Volatility, and finally the forecast series is compared with the realized series using the same six loss functions in Hansen and Lunde (2005). The model with the best forecast fit, according to five functions out of six, is found to be the asymmetric GJR(2,1).

Section 1

The exchange rate

Issues concerning the exchange rate are of great interest to researchers in modern economic theory. Foreign exchange rate, defined as the value of a currency in relation to another currency, comes in two formulations: direct quote, when it is expressed as price of one unit of foreign currency expressed in terms of the domestic currency; indirect quote when it is expressed as the price of one unit of the domestic currency in terms of the foreign currency. From the perspective of an individual or institution residing in the European Union, for example, the exchange rate at the time of writing is about 1.19 $ for 1€ (indirect quote) or 0.84 € for 1 $ (direct quote). In this research, I will always refer to the indirect quote and I will label it €/$, EUR/USD, or EURUSD. The exchange rate is a variable of primary importance which is taken into account by international economic operators, be it investment firms devising their strategies or central banks for policymaking purposes. Thus, the exchange rate and its volatility heavily influence the flow of capitals across borders for investment decisions and import/export purposes alike, and so extensive research has sparked over the years, with the intent of modeling - and forecasting it. The ARCH and GARCH type of models, introduced by Engle (1982) and Bollerslev (1986) respectively, together with the copious amount of derivations devised in research subsequent to Engle's seminal paper, have been applied to a large extent in an attempt at modeling the volatility of the exchange rate.

Generalized Autoregressive Conditional Heteroskedasticity

As the empirical evidence suggests, the classic assumption of homoskedasticity in a regression model is typically not verified. Assuming a time varying volatility can help to explain the empirical evidence of volatility clustering, that is, the phenomenon in which periods of overall low volatility alternate with periods of high volatility. By and large, volatility is not constant over time, and during periods of financial turmoil it tends to rise, possibly in response to behavioral dynamics (i.e. the VIX index, sometimes called the "fear indicator"). Assuming a time varying volatility is also useful to explain those which have become known as stylized facts of asset returns: non-normal distribution, little to no correlation between returns of different days (and thus little return predictability), and positive and statistically significant correlations between magnitudes of returns on nearby days. As far as the distribution is concerned, it has been observed that it tends to be slightly negatively skewed (third moment), in the sense that large negative returns happen more often than large positive ones, and is leptokuritc, i.e. it has higher kurtosis (fourth moment) with respect to the normal distribution, such that "extreme" or "tail" events occur more often.

The Autoregressive Conditional Heteroskedasticity model was introduced by Engle in 1982 and it explicitly recognize the difference between the conditional and the unconditional variance, allowing the former to change over time as a function of past errors. Bollerslev (1986) assumed that the error variance followed an autoregressive moving average model, and gave the conditional variance a linear functional form that allowed it to depend not only on past errors, but also on its own past values. Here is Bollerslev's GARCH(p,q) model:

$$\varepsilon_t|\Psi_{t-1}\sim N(0,\sigma_t)$$

$$\sigma\sigma_t^2{}_t = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma^2{}_{t-i}$$

with

$P > 0$ $q \geq 0$ $\omega > 0$

$\alpha_i > 0$ For $i = 1 \dots p$

$\beta_i > 0$ For $i = 1 \dots q$

Where ε_t denotes a real valued, discrete time stochastic process, Ψ_{t-1} denotes the information set (σ-field) of all information through time *t*, and stationarity is guaranteed if:

$$\sum_{i=1}^{q} \alpha_i + \sum_{i=1}^{p} \beta_i < 1$$

For *p=0* the process reduces to Engle's ARCH(q) process, and for p=q=0, ε_t degenerates to a white noise.

Remarkably, it can be shown that the GARCH(1,1) process presents excess unconditional kurtosis:

$$E(\varepsilon_t^2) = \frac{\omega}{1-\alpha-\beta}$$

And

$$E(\varepsilon_t^4) = 3\omega^2(1+\alpha+\beta)[(1-\alpha-\beta)(1-\beta^2-2\alpha\beta-3\alpha^2)]^{-1}$$

Such that the coefficient of kurtosis is

$$\mathrm{k} = (\mathrm{E}(\varepsilon_t^4) - 3\mathrm{E}(\varepsilon_t^2)^2)\mathrm{E}(\varepsilon_t^2)^{-2}$$

which is greater than zero by assumption, hence the GARCH(1,1) process is leptokurtic. For estimation purposes, various specifications have been proposed for the conditional mean of the return series: ARMA(p,q), mean zero, unconditional mean, and GARCH-in-mean ($m_t = \mu_0 + \mu_1 \sigma_{t-1}^2$). Though it cannot be ruled out *ex-ante* that a more sophisticated specification such as the GARCH-in-mean could lead to better estimation and forecasts of volatility, Hansen and Lunde (2005) find that the performance is almost identical across the three mean specifications.

Engle (1982) also proposed a useful Lagrange-Multiplier type of test to assess the presence of ARCH effects in the return series. It uses two auxiliary regressions: one estimates the best fitting AR(q) model, and the other is a regression of the squares of the errors from the previous, on a constant and their lagged valued up to *q*.

$$r_t = \alpha_0 + \sum_{i=1}^{q} \alpha_i r_{t-i} + \varepsilon_t$$

$$\hat{\varepsilon}_t^2 = \beta_0 + \sum_{i=1}^{q} \beta_1 \hat{\varepsilon}_{t-1}^2$$

Under the null hypothesis, we have $\alpha_i = 0$ for all $i = 1, \dots, q$, that is, absence of ARCH effects, while under the alternative hypothesis at least one of the estimated coefficients is significant. The test statistic has the form $T \times \mathrm{R}^2$, where T is equal to the number of observations minus q and R^2 is the r-squared from the second auxiliary regression, and it is distributed as a χ^2 with q degrees of freedom.

News impact curve and asymmetry

The GARCH, as it was initially proposed by Bollerslev, was a symmetric model, that is, it did not allow to take into account the so-called leverage effect, which corresponds to a negative correlation between past returns and future volatility. Engle and Ng (1993) introduce the news impact curve, which measures how new information is incorporated into volatility estimates. In their findings, they emphasize the asymmetry of the volatility response to news. A number of extensions to the original GARCH model were proposed to capture this effect, including the EGARCH (Nelson, 1991) and the GJR-GARCH (Glosten, Jagannathan, and Runkle, 1993). The latter was found to be the best parametric model to explain asymmetry in Japanese stock returns by Engle and Ng (1993). Its functional form is similar to the baseline GARCH, but it includes additional coefficients which are only activated by an indicator function when the respective past innovation is negative:

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} (\alpha_i \varepsilon_{t-i}^2 + \gamma_i S_{t-i}^- \varepsilon_{t-i}^2) + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2$$

$$S_t^- = \begin{cases} 1 & if \quad \varepsilon_t < 0 \\ 0 & if \quad \varepsilon_t \geq 0 \end{cases}$$

The equation of the news impact curve for a GARCH(1,1) model is

$$\sigma_t^2 = A + \alpha \varepsilon_{t-1}^2$$

$$A = \omega + \beta \sigma^2$$

Where σ_t^2 is the conditional variance at time *t*, ε_{t-1} is the innovation on the return at time *t*-1, σ is the unconditional return standard derivation, ω, α and β are the parameters of the model.

The equation of the news impact curve for a GJR(1,1) model would be

$$\sigma_t^2 = \begin{cases} A + \alpha \varepsilon_{t-1}^2 & for\ \varepsilon_{t-1} > 0 \\ A + (\alpha + \gamma) \varepsilon_{t-1}^2 & for\ \varepsilon_{t-1} \leq 0 \end{cases}$$

$$A = \omega + \beta \sigma^2$$

Where σ is the unconditional return standard deviation, ε_{t-1} is the innovation on the return at time *t*-1, ω, α, γ and β are the parameters of the model. Figure 1 shows that the GJR news impact curve captures the asymmetry in the effect of news on volatility because it has a steeper slope in its negative side than on its positive side, whereas the GARCH news impact curve is just a symmetric parabola centered in $\varepsilon_{t-1} = 0$.

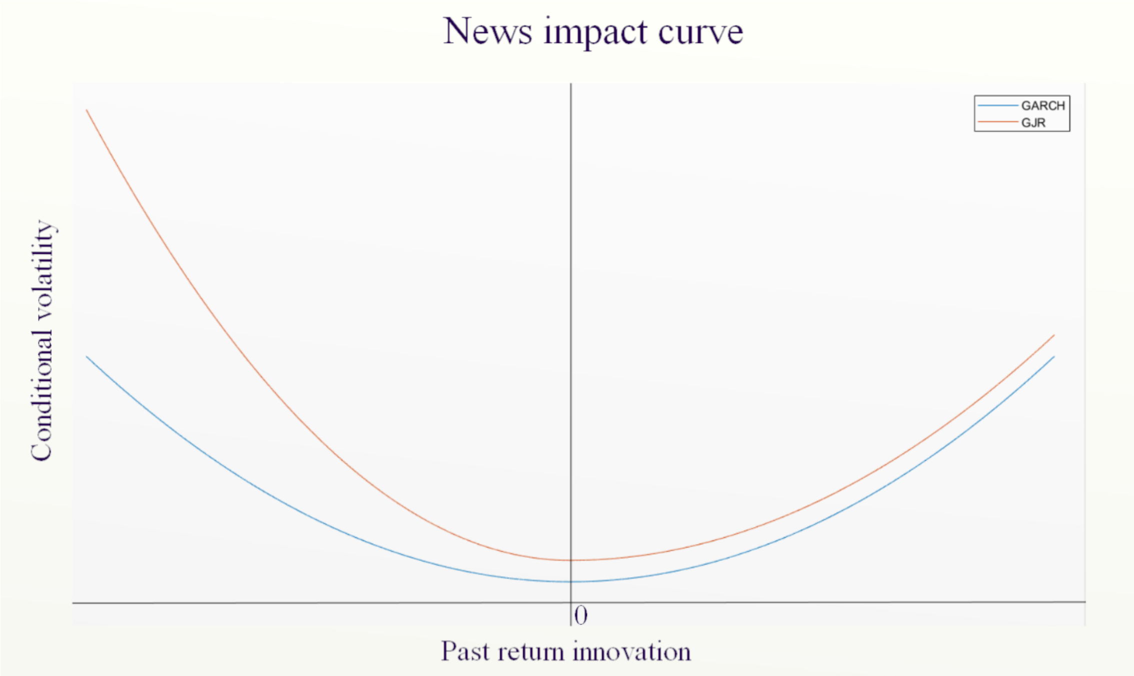

Figure 1: The news impact curve of the GARCH(1,1) and the GJR(1,1). The image is traced for equal values of α and β. GJR is shifted upwards (i.e. higher ω) for visualization purposes only, as its positive portion would overlap with the GARCH. Though they both have a minimum at $\varepsilon_{t-1} = 0$, it's abundantly clear how the GJR model predicts higher variance than the GARCH when the past return innovation ε_{t-1} is negative, due to its steeper slope.

The empirical evidence for the leverage effect was first documented by Black (1976), who attributed it to the effects of financial leverage: as a company's stock price declines, it becomes more highly leveraged for a given fixed level of debt outstanding, and this increase in leverage induces a higher equity-return volatility. Subsequently, it has been suggested that the phenomenon might not be driven by financial leverage: Hasanhodzic and Lo (2011) found that, in a sample of all-equity-financed companies, the leverage effect is just as strong or even stronger. However, their results are consistent with a behavioral interpretation: individuals' perceptions of risk might be altered by negative returns, giving rise to changes in their demand of risky assets. Research in experimental psychology has suggested that most people tend to overreact to unexpected and dramatic events; DeBondt and Thaler (1985) found that such behavioral bias matters at a market level. But does this bias apply also to the currency exchange rate? Hanse and Lunde (2005) found that volatility models which allow to account for asymmetry performed better with data for the IBM stock returns, but as long as their analysis of the DM/$ exchange rate is concerned, they found no evidence that the GARCH(1,1) is outperformed by more sophisticated models. This might be due to some differences between the stock market and the foreign exchange market: the former is a regulated market, while currencies are traded over the counter; Stocks are priced in an “absolute” way, whereas each currency is traded in pairs with every other currency; Stocks and currencies may also be traded by different type of actors and for different purposes: stocks are traded by investment firms, banks, and retail investors for investment and speculation purposes. Currencies and currency derivatives are mostly traded by institutional investors such as banks, multination corporations and hedge funds for hedging purposes, and also notably by central banks for monetary policy purposes. Being among the biggest markets in terms of liquidity and volumes traded, the foreign exchange may be less likely to be influenced by behavior-induced inefficiencies and biases.

Section 2

Data and methodology

This research aims to model the conditional volatility of the €/$ exchange rate with GARCH models, and evaluate the forecasts they can provide. The data used span from February 1st, 2002 to February 1st, 2015 and are divided as follow: An estimation window consisting in daily observations which span from the beginning up until February 1st, 2014, and a forecasting window consisting of 5-minute intraday returns which span from February 3rd, 2014 until February 2nd, 2015, which are used to compute the daily Realized variance. The source for the data in the estimation period is the Bank of Italy[1], while the source for the intraday data is BacktestMarket[2].

The estimation is carried out as follows: daily exchange rate data is firstly transformed into log-returns. The log-return series is used for estimation of GARCH and GJR models with up to two lags for each parameter. This choice mirrors Hansen and Lunde (2005) and is also motivated by the fact that typically more parametrized models fit increasingly well in-sample, but tend to perform poorly out-of-sample. The estimation is performed with MATLAB using a fairly standard maximum-likelihood approach with gaussian likelihood function. The log-likelihood for the *i*-th observation is

$$\mathcal{L}_t(\theta; r) = -\frac{1}{2}(\log(2\pi) + \log(\sigma_i^2) + \frac{\varepsilon_i^2}{\sigma_i^2})$$

Where θ is the vector of parameters which includes ω, α_i for i=0, 1, 2; β_i for i=1, 2; In estimation of the GJR models θ also contains the parameters γ_i for i=0, 1, 2. The conditional variance σ_i^2 is specified as stated above, and the residuals ε_i are computed as difference between the observation and the sample average.

[1] https://tassidicambio.bancaditalia.it/timeSeries
[2] https://www.backtestmarket.com/

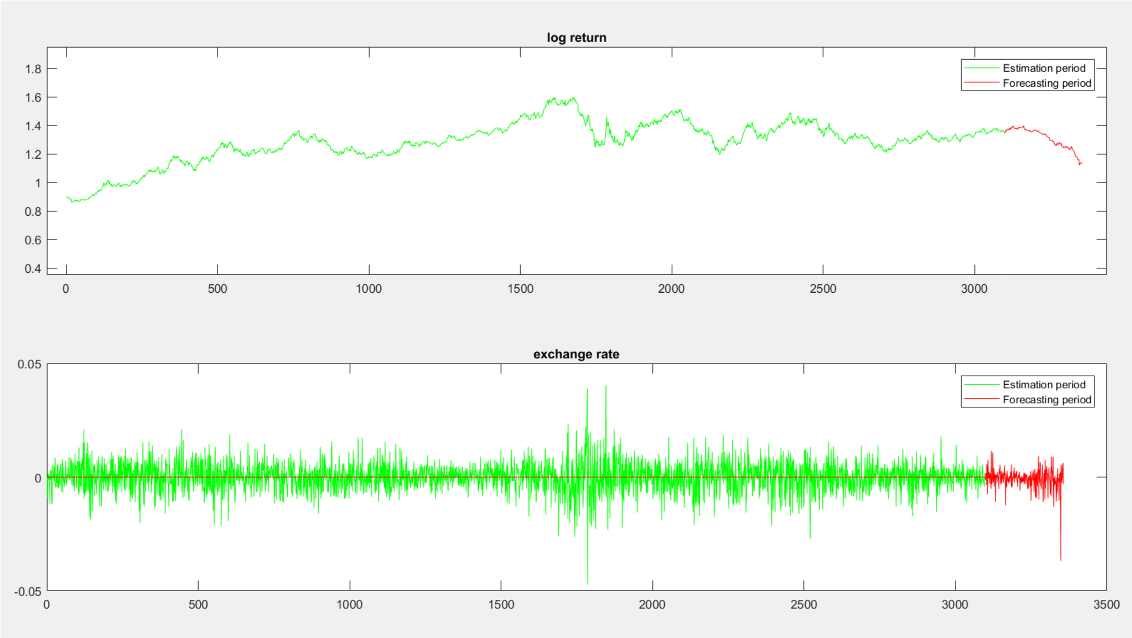

Figure 2: The exchange rate and log returns. Estimation period consists of 3098 daily observations from Febuary 1st, 2002 until February 1st, 2014; forecasting period consists of 258 daily observations from February 3rd, 2014 until February 2nd, 2015.

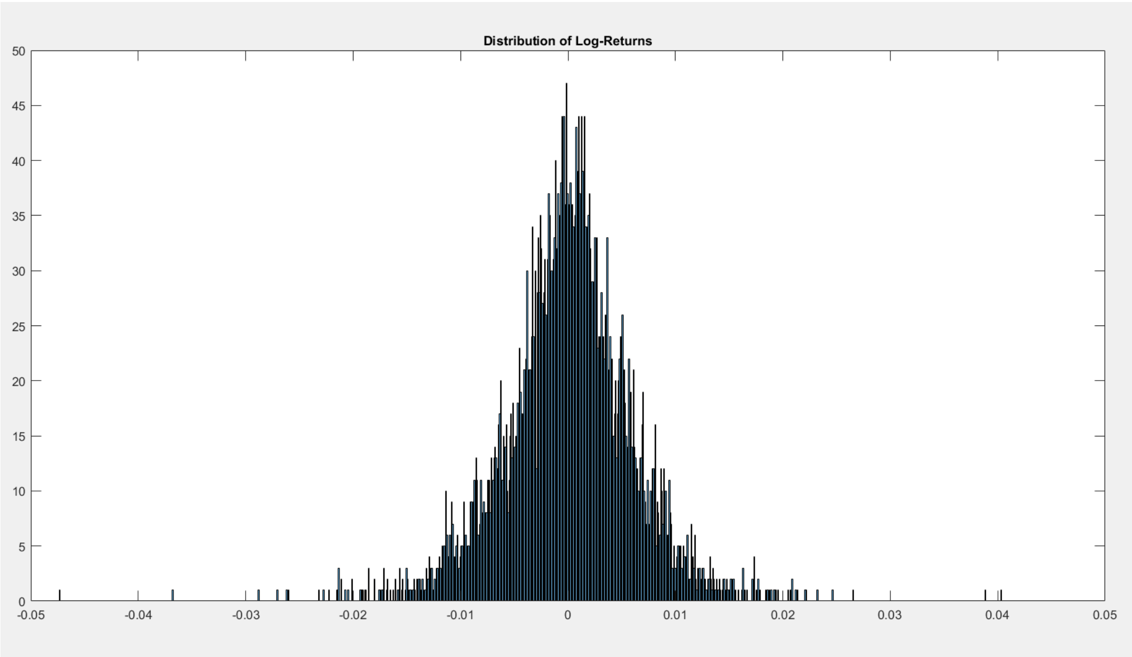

Figure 3: The histogram of the log-return series highlights a somewhat asymmetric distribution, roughly centered in zero, with some outliers far away both in positive and negative territory.

Judging by Figure 2, the log-return series seems fairly stationary. Shocks are not really persistent, and we can observe periods in which the magnitude of the returns is fairly contained, as opposed to periods in which returns display a higher volatility (the phenomenon known as volatility clustering). Figure 4 represents the histogram of the og-return series, and highlights at a first glance a somewhat symmetric distribution, roughly centered in zero, with a notable amount of extreme observations. Remarkably, the furthest observation away from zero is in negative territory. The Jarque-Bera test rejects the null hypothesis of the returns coming from a normal distribution at a 1% significance level.

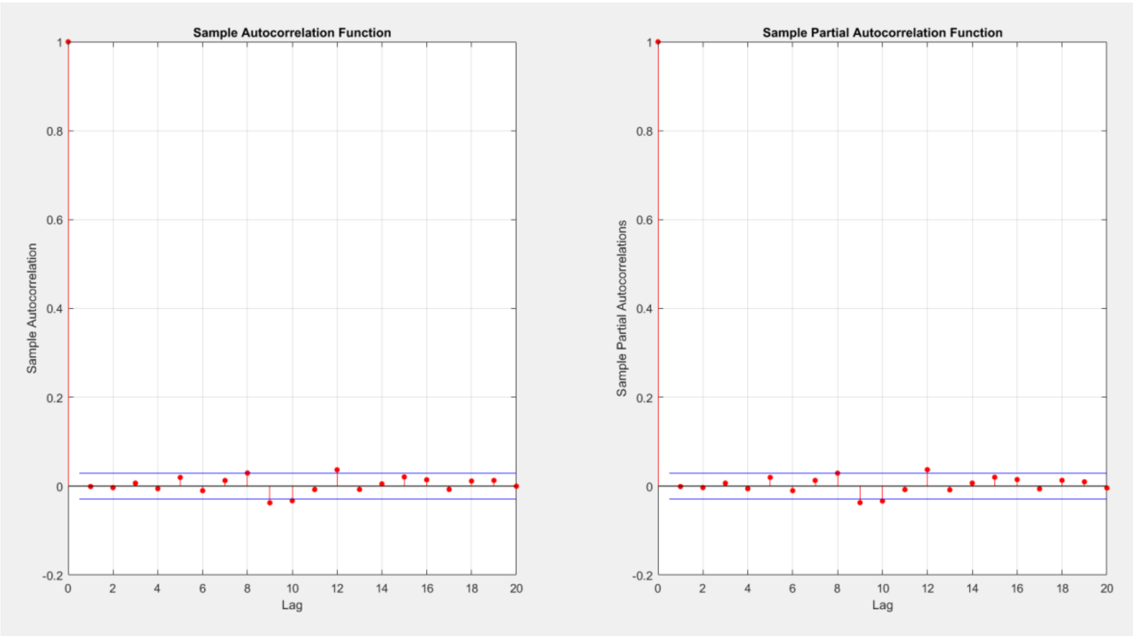

Figure 4: The autocorrelograms do not highlight significative intertemporal correlations across the log-returns. This is consistent with the "little return predictability" stylized fact.

As expected, and somewhat anticipated by the log-returns plot, the autocorrelograms in Figure 4 do not highlight significant intertemporal correlation. Little correlation between returns in subsequent days is consistent with the “little or no predictability of returns” stylized fact, and it does not justify, in my opinion, the use of complicated specifications for the conditional mean. For this reason, and also because it is extremely close to zero (although slightly positive), I chose to use the sample average as specification for the mean, in estimation of the models. Finally, as a conclusion to the preliminary analysis, Engle's ARCH test confirms the presence of significative conditional heteroskedasticity effects in the residuals' series.

Table 1: Results of estimation of GARCH models

ARCH(1)

	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	29.8328	0.0000
β	0.1520	0.0325	4.6740	0.0000

ARCH(2)

	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	26.0012	0.0000
β_1	0.1335	0.0333	4.0052	0.0001
β_2	0.0657	0.0268	2.4493	0.0144

GARCH(1,1)

	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.8128	0.0700
α_1	0.0345	0.0043	8.0870	0.0000
β_1	0.9623	0.0047	203.2033	0.0000

GARCH(1,2)

	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.8027	0.0715
α_1	0.0453	0.0097	4.6699	0.0000
β_1	0.4694	0.3062	1.5332	0.1253
β_2	0.4800	0.2970	1.6159	0.1062

GARCH(2,1)

	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.8725	0.0612
α_1	0.0246	0.0119	2.0734	0.0382
α_2	0.0088	0.0138	0.6340	0.5261
β_2	0.9636	0.0048	200.1227	0.0000

GARCH(2,2)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	0.7116	0.4768
α_1	0.0326	0.0076	4.2849	0.0000
α_2	0.0337	0.0098	3.4514	0.0006
β_1	0.0296	0.3584	0.0827	0.9341
β_2	0.9031	0.3469	2.6033	0.0093

Table 2: Results of estimation of GJR models

GJR(1,0)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	29.7208	0.0000
α_1	0.1482	0.0417	3.5562	0.0004
γ_1	0.0107	0.0581	0.1842	0.8539

GJR(2,0)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	25.9651	0.0000
α_1	0.1312	0.0423	3.1031	0.0019
α_2	0.0636	0.0401	1.5865	0.1127
γ_1	0.0072	0.0573	0.1260	0.8997
γ_2	0.0033	0.0465	0.0710	0.9434

GJR(1,1)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.5422	0.1231
α_1	0.0291	0.0036	8.0165	0.0000
γ_1	0.0018	0.0025	0.7336	0.4632
β_1	0.9675	0.0040	243.7243	0.0000

GJR(1,2)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.3429	0.1794
α_1	0.0418	0.0113	3.7105	0.0002
γ_1	0.0035	0.0032	1.0784	0.2809
β_1	0.5567	0.3772	1.4761	0.1400
β_2	0.3969	0.3662	1.0840	0.2784

GJR(2,1)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0000	1.8150	0.0696
α_1	0.0313	0.0043	7.3543	0.0000
α_2	0.0010	0.0036	0.2938	0.7690
γ_1	0.0011	0.0234	0.0450	0.9641
γ_2	0.0010	0.0231	0.0453	0.9639
β_1	0.9634	0.0047	206.5108	0.0000

GJR(2,2)				
	Parameter	St. err.	Student t	p-value
ω	0.0000	0.0001	0.2071	0.8360
α_1	0.0893	9.4487	0.0094	0.9925
α_2	0.1432	8.2645	0.0173	0.9862
γ_1	0.0465	13.2189	0.0035	0.9972
γ_2	0.0791	0.2629	0.3009	0.7635
β_1	0.1192	62.7918	0.0019	0.9985
β_2	0.2662	52.2262	0.0051	0.9959

Forecasting and Realized Variance

The forecasting exercise was carried out as follows: coefficients found in estimation, together with daily data from February 2014 to February 2015, are used to compute the fitted values of the conditional variance $\hat{\sigma}_t^2$ in the GARCH and GJR models respectively:

$$\hat{\sigma}_t^2 = \hat{\omega} + \sum_{i=1}^{q} \hat{a}_i \varepsilon_{t-i}^2 + \sum_{i=1}^{p} \hat{\beta}_i \hat{\sigma}^2{}_{t-i}$$

And

$$\hat{\sigma}_t^2 = \hat{\omega} + \sum_{i=1}^{q} (\hat{\alpha}_i \varepsilon_{t-i}^2 + \hat{\gamma}_i S_{t-i}^- \varepsilon_{t-i}^2) + \sum_{i=1}^{p} \hat{\beta}_i \hat{\sigma}^2{}_{t-i}$$

Where the residuals ε_t are computed as difference between the return in the *t*-th day and the sample average of the returns in the forecasting window. The forecast values are one-day-ahead fitted values $\hat{\sigma}_t^2$, given the information set Ψ_{t-1}, which contains past returns up to time *t*-1 and past values of σ_t^2 itself up to time *t*-1, and are used to measure the forecasting ability of our models against the Realized Volatility in each day.

The realized volatility is a non-parametric estimator of the variance and is defined as the second uncentered sample moment of the return process over a fixed interval of length *h*, scaled by the number of observations *n*, which corresponds to the sampling frequency 1/*n*, so that it provides a volatility measure calibrated to the *h*-period measurement interval. The realized volatility over the period $[t-h, t]$ is given by:

$$v^2(t,h;n) \equiv \sum_{i=1}^{n} r\,(t - h + \left(\frac{i}{n}\right) \cdot h, \frac{h}{n})^2$$

In their empirical analysis, Hansen and Lunde (2005) use the realized variance as substitute for the latent σ_t^2:

$$RV_t^{(m)} \equiv \sum_{i=1}^{m} r_{t,i,m}^2$$

Where m is the sampling frequency and

$$r_{t_1 i,m} \equiv p_{t-\frac{(i-1)}{m}} - p_{t-\frac{i}{m}}$$

is the return over a time interval with length $1/m$ on day t. It can often be shown that $E[RV_t^{(m)} - \sigma_t^2]$ is decreasing in m, such that the realized variance is an increasingly more precise estimator of σ_t^2 as the sampling frequency increases. However, if the sampling frequency is too high, market microstructure noise can emerge (irregular trading, missing values, bid-ask spread bounce).

I computed the realized volatility over a sample of 5-minute returns for 256 trading days. There are, however, some imperfections in the data, namely on average there are 299 intraday observations per day instead of 288. This is due to the presence of extra observations before the opening on some Mondays and after the closing on some Fridays. For some other weeks, the data for Mondays and Fridays is not even complete (i.e. less than 288 observations). As I could not clean the data in an automated way, let alone manually, I solved this issue by dividing the sample in 256 days, containing 299 observations each, and using only the first 288 of those observations to compute the realized variance in each day. I suspect this approach led to a spillover of volatility between consecutive days. Had I used all of the observations, the realized volatility would have resulted in an upward-biased estimator of the variance on Fridays and Mondays and potentially on every other weekday as well, due to the presence of basically one extra trading hour per day on average. I also used the last 288 observations as well as the middle 288 observations and found no significative change in the estimation of the realized volatility – perhaps with higher quality data the realized volatility could be more accurate and the overall result might vary[3].

[3] Although very blunt, this is the best method I could devise to address the issue. While it was easy to find and institutional source for the daily data, this was not the case for the intraday data. Since the foreign exchange is not a regulated market and currencies are traded over the counter, I struggled to find good quality data from an institutional source, and finally settled on purchasing them here.

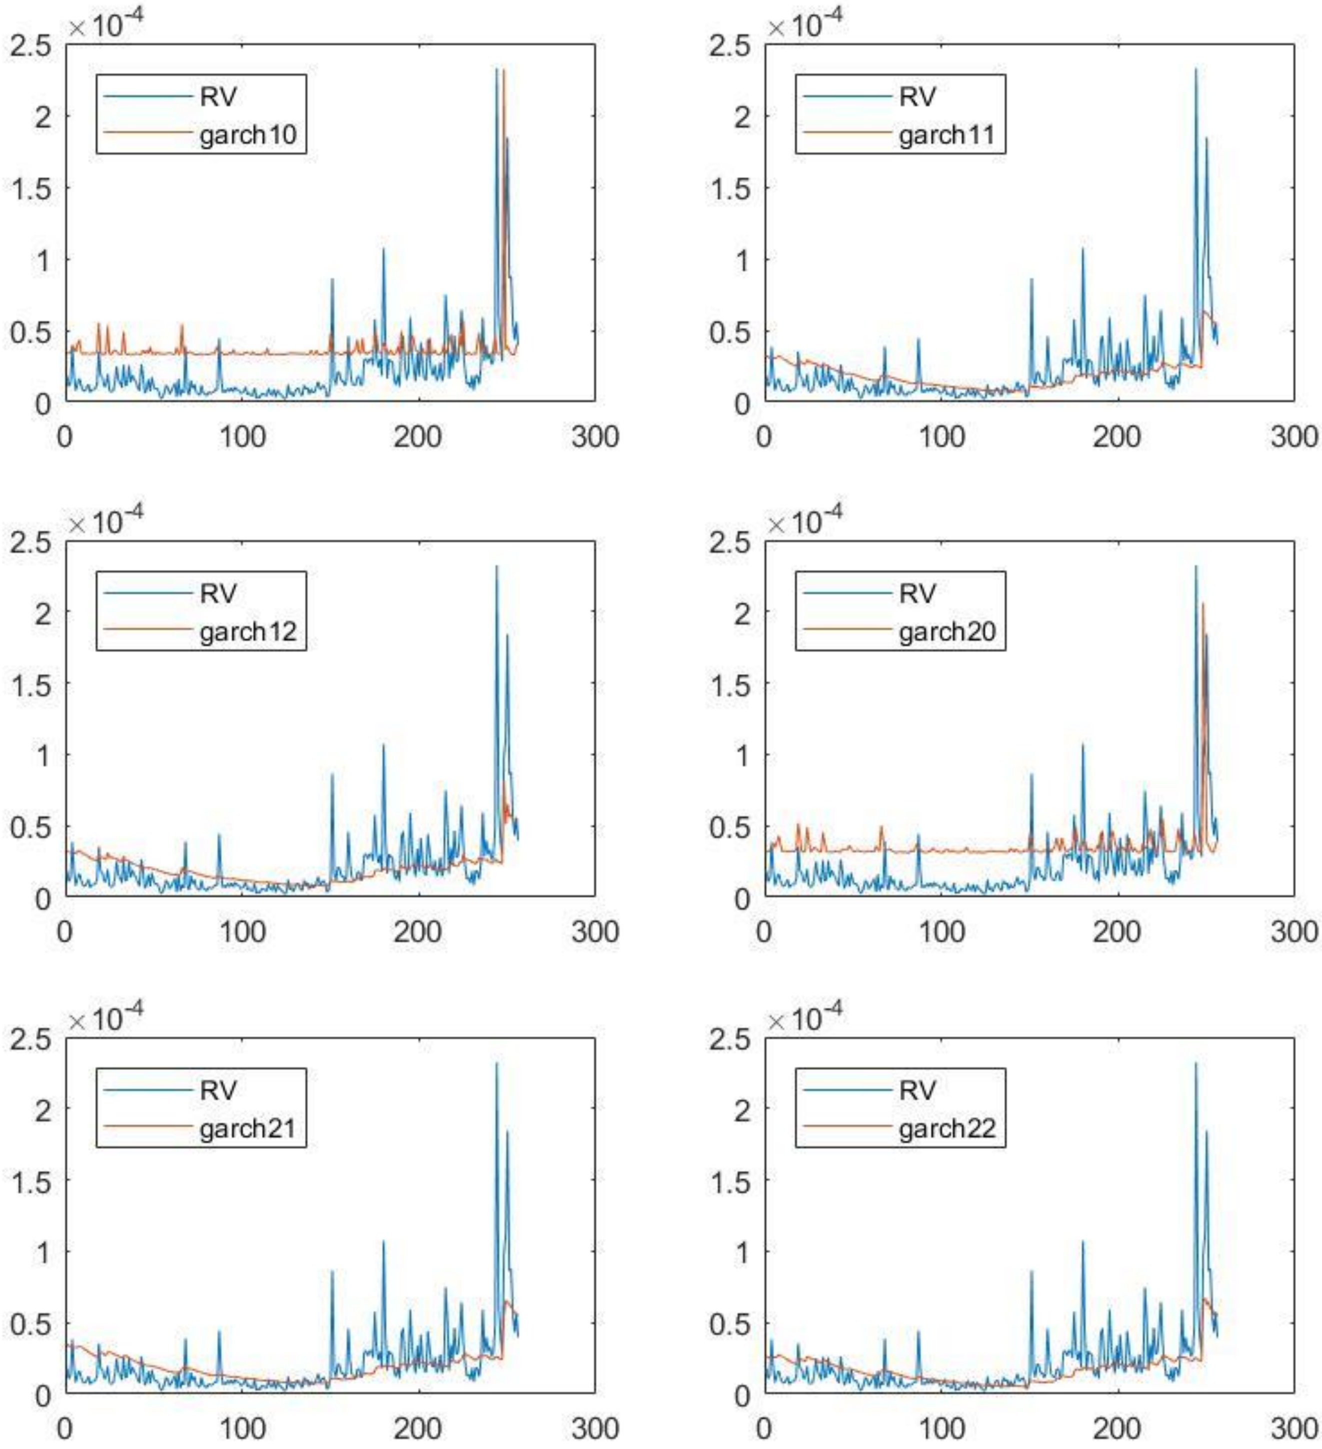

Figure 5: conditional volatility for different GARCH models and realized volatility.

Figure 6: conditional volatility for different GJR models and realized volatility.

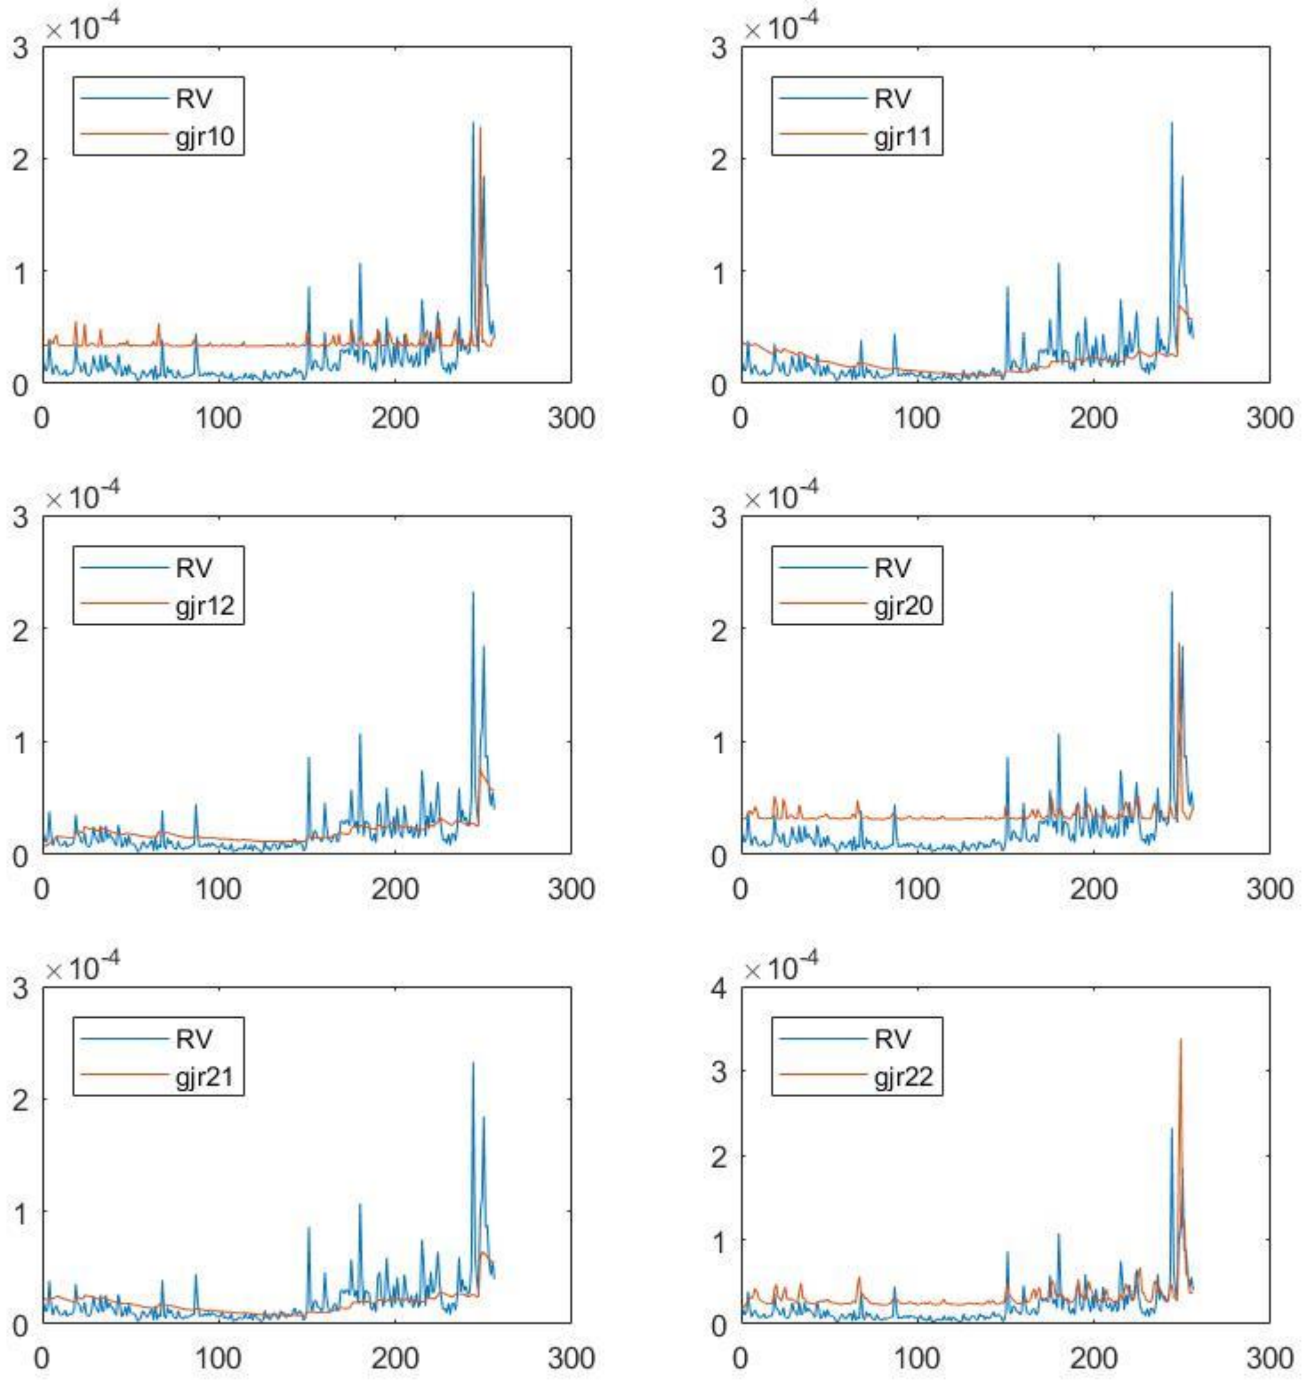

Figure 6: conditional volatility for different GJR models and realized volatility.

Evaluation via loss functions

Forecast are benchmarked against the realized volatility using the same six loss function proposed by Hansen and Lunde (2005):

$$MSE_1 = n^{-1} \sum_{t=1}^{n} (\sigma_t - h_t)^2 \qquad MSE_2 = n^{-1} \sum_{t=1}^{n} ({\sigma^2}_t - h_t^2)^2$$

$$MAE_1 = n^{-1} \sum_{i=1}^{n} |\sigma_t - h_t| \qquad MAE_2 = n^{-1} \sum_{i=1}^{n} |\sigma_t^2 - h_t^2|$$

$$R^2LOG = n^{-1} \sum_{t=1}^{n} [\log(\sigma_t^2 h_t^{-2})]^2 \qquad QLIKE = n^{-1} \sum_{i=1}^{n} (\log(h_t^2) + \sigma_t^2 h_t^{-2})$$

Where h_t^2 are the forecast values produced by the model and the latent process ${\sigma^2}_t$ is substituted with the realized variance. Table 3 reports the value of the proposed loss functions for each model. As expected, ARCH models (both symmetric and asymmetric) are outperformed by their generalized counterparts. The GJR(2,1) achieves the lowest value among the models across all error metrics except for the $QLIKE$, where GARCH(1,2) prevails instead. Interestingly, GJR(2,1) estimated γ_1 and γ_2 coefficients are very close to zero and not very significant, judging on the high p-values. This may suggest that, in fact, leverage effects are not really significant as far as EUR/USD is concerned (and possibly other foreign exchange pairs). Indeed, if leverage actually is the driver of the effect, it would only be observed in stocks. If the effect has a behavioral origin, the bias could be mitigated by the fact that possibly relatively less operators have speculative positions on the F/X market.

Table 3: value of the loss function for each model

	MSE_1	MSE_2	MAE_1	MAE_2	R^2LOG	$QLIKE$
ARCH(1)	6.9797e-06	8.5909e-10	0.0023	2.3004e-05	1.4843	1.5661e+04
ARCH(2)	6.3766e-06	7.6047e-10	0.0022	2.1605e-05	1.3969	1.6111e+04
GARCH(1,1)	2.5950e-06	4.4570e-10	0.0012	1.1231e-05	0.4974	6.8997e+04
GARCH(1,2)	2.6201e-06	4.4534e-10	0.0012	1.1293e-05	0.5108	**6.3992e+04**
GARCH(2,1)	2.6310e-06	4.4713e-10	0.0012	1.1370e-05	0.5041	6.7705e+04
GARCH(2,2)	2.4171e-06	4.3343e-10	**0.0011**	1.0509e-05	0.4419	9.8050e+04
GJR(1,0)	6.9894e-06	8.6913e-10	0.0023	2.3030e-05	1.4834	1.5690e+04
GJR(2,0)	6.3856e-06	7.6667e-10	0.0022	2.1640e-05	1.3968	1.6117e+04
GJR(1,1)	2.4274e-06	4.3153e-10	**0.0011**	1.0751e-05	0.4649	6.8846e+04
GJR(1,2)	2.3710e-06	4.2322e-10	**0.0011**	1.0535e-05	0.4532	6.9821e+04
GJR(2,1)	**2.2188e-06**	**4.1395e-10**	**0.0011**	**1.0086e-05**	**0.4120**	7.3134e+04
GJR(2,2)	4.5540e-06	7.2107e-10	0.0018	1.7945e-05	1.0298	1.9561e+04

Conclusions and caveats

In this research I have shown how GARCH models can be used to characterize the variance of the EUR/USD returns. The graphs in Figures 5 and 6 highlight that in general, the one-day-ahead forecast obtained via GARCH models is not too far away from the realized variance. The use of loss functions to compare the model's predictive ability leads to affirm that the best performing model is the GJR(2,1). Interestingly though, estimated gammas are not statistically different from zero. However, loss functions aside, it can be seen that, at least qualitatively, the models' prediction are somewhat consistent with the observed realized volatility.

One thing worth noting would be that the EUR/USD volatility is heavily influenced by the release of macroeconomic data: in my experience, as an example, a spike in volatility could typically be observed during a press release speech by (former) president of the BCE Mario Draghi; another example would be the publishing of Non-Farm Payroll, an important piece of economic information released by the Bureau Of Labor Statistics, which measures the variation in employment in the secondary and tertiary sectors in the United States. Hence, it should be kept in mind that the proposed models are just econometric tools which do not take into account anything that is not contained in the information set at time t-1 (i.e., unexpected variation in occupation, unexpected inflation, policy decisions, etc.).

Another issue worth mentioning in my opinion, is that econometrics is not really an exact science. Maximum likelihood is a fairly standard and well-established estimation methodology, however MATLAB's optimization tools are entirely numerical, iterative algorithms, and as such, for some kind of functions and/or data sets, it cannot be totally excluded that they suffer from issues like optimal points depending on starting values, tolerances, iterations etc. Software able to perform computation, which would take a lifetime, in a matter of seconds is an extremely useful resource, but it does not come without its drawbacks.

Finally, what is most relevant is not the exact value of the parameter or the loss function, but the fact that at least qualitatively, the theory behind the models is able to explain the empirical evidence to a certain extent. Looking at the graph of the predicted volatilities in figure 5 and 6, it can be noticed how the overall “shape” is fairly consistent across almost all of the proposed models: which one has the “best fit” is just a matter of extremely small values, as one can gather from the value of the loss functions being so similar. But then again, the GJR(2,1) might very well be surpassed by other models in forecasting volatility of another year, let alone another asset or currency pair, so there does not exist the ultimate, true model. Models are just models, and as such they arise as direct consequences to the assumptions underlying; whether or not those assumptions are consistent with the empirical evidence of the real world is often debatable. I like to quote George Box: “All models are wrong, but some are useful”.

References

Hansen P, Lunde A. 2005. A forecast comparison of volatility models: does anything beat a GARCH(1,1)? *Journal of Applied Econometrics* 20: 873–889.

Andersen TG, Bollerslev T, Diebold FX. 2003. Parametric and nonparametric volatility measurement. In *Handbook of Financial Econometrics*, Vol. I, Ait-Sahalia Y, Hansen LP (eds). Elsevier-North Holland: Amsterdam.

Andersen TG, Bollerslev T. 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance 4: 115–158.*

Black, F. (1976) Studies of Stock Price Volatility Changes. In: *Proceedings of the 1976 Meeting of the Business and Economic Statistics Section, American Statistical Association, Washington DC, 177-181.*

Bollerslev T, Engle RF, Nelson D. 1994. ARCH models. In: *Handbook of Econometrics*, Vol. IV, Engle RF, McFadden DL (eds). Elsevier Science B.V.: Amsterdam; 2961–3038.

Bollerslev T. 1986. Generalized autoregressive heteroskedasticity. *Journal of Econometrics 31: 307–327.*

Engle RF. 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. *Econometrica 45: 987–1007.*

Engle RF, Ng V. 1993. Measuring and testing the impact of news on volatility. *Journal of Finance 48: 1747–1778.*

Baillie RT, Bollerslev T. 1989. The message in daily exchange rates: a conditional variance tale. *Journal of Business & Economic Statistics 7(4): 297–305.*

Cont R. 2001. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance. 1. 223-236. 10.1080/713665670.*

De Bondt, Thaler. 1985. Does the Stock Market Overreact? *The Journal of Finance Vol. 40, No. 3, Papers and Proceedings of the Forty-Third Annual Meeting American Finance Association*, Dallas, Texas, December 28-30, 1984 (Jul., 1985), pp. 793-805.

Engle R. and Patton A. 2001.What Good is a Volatility Model? NYU Working Paper No. S-DRP-01-03.

Glostein L. Jagannathan R. Runkle D. 1993. On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *The journal of finance vol. XLVIII, no. 5, pp. 1779-1801*.

Hasanhodzic J. Lo A. 2011. Black's Leverage Effect is not Due to Leverage. *SSRN Electronic Journal. 10.2139/ssrn.1762363.*

Ait-Sahalia Y. Fan J. Li Y. 2011. The Leverage Effect Puzzle: Disentangling Sources of Bias at High Frequency. NBER Working Paper No. w17592.

Andersen, T. G., & Bollerslev, T. 1998. Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, *39*(4), 885-905.