6

GRZEGORZ KANOZA

Volatility Modelling with GARCH Models – Application to KGHM Returns

Opiekun naukowy: dr Anna Czarczyńska

Grzegorz Kanoza – absolwent studiów pierwszego stopnia na kierunku ekonomia – rynki zagraniczne Akademii Leona Koźmińskiego. Wcześniej studiował także ekonomikę przedsiębiorstw na Uniwersytecie Europejskim Viadrina. Obecnie pracuje dla grupy Citi. Jego główne zainteresowania to zastosowanie metod ilościowych w inwestycjach na rynku kapitałowym.

1. Introduction

One of the main empirical properties of financial time-series is time-varying conditional volatility. As long as volatility remains the fundamental risk measure both in modern portfolio theory and risk management, there is a practical need to define, model and forecast volatility. There is a number of ways to define volatility, but it is a common practice to define volatility as the standard deviation of returns. Unlike asset prices, volatility cannot be observed directly so that it has to be estimated from historical data. To model and forecast volatility many possible models have been proposed and the literature on volatility models is very abundant. This article deals with univariate ARCH (Autoregressive Conditional Heteroskedasticity) model and its extension: the generalized ARCH (GARCH) model, both of which constitute only the starting point for more complex descriptions.

An application of volatility models to historical data of a polish company's stock KGHM will be carried out with the aim to check the statistical properties of the time series and to see to how well the series might be described by the models presented. Based on the size and its history the polish stock market along with its component companies might still be regarded as emerging or regional. There are currently only 438 companies traded on the Warsaw Stock Exchange (WSE) whereas, for example, on the London Stock Exchange (LSE) 2477 companies were traded as of Nov 2012. Furthermore, by looking at the distributional properties of the stock return series and its autocorrelation functions, one can measure the extent of market efficiency. Firstly, the return distribution should be approximately normal with typical for financial time series fat tails. Secondly, in a highly efficient market no serial autocorrelation in the stock returns should be present so that any serial autocorrelation of high order might be an argument against market efficiency[1].

This article is structured as follows. Section 1 provides an overview of typical features of financial time series in order to give a glimpse of the kind of models needed to describe return series. In section 2 volatility definition is given. Section 3 presents ACH model structure and its methodology and thereafter an analysis of KGM stock returns using the framework presented earlier follows.

[1] When investigating autocorrelations of returns, one should consider the aspect of data frequency as for some data frequencies autocorrelations of a certain order are typical which might not be true for other frequencies. For example, high-frequency data (intraday of shorter) will exhibit declining autocorrelation due to trading technicalities, but daily stock data should not exhibit any serial autocoreelation. For more details see the article of Andrew Lo mentioned in the references. For more details see the article of Andrew Lo mentioned in references.

All charts and estimations were made with R software version 3.0.0. To confirm the results of statistical tests and model estimations, the same procedures were also performed in JMulTi and revealed none significant differences.

2. Characteristic Features of Financial Time Series

As shown by extensive empirical research, financial assets like bonds, stocks or foreign exchange rates, do exhibit certain similar features common to high-frequency time series. Figure 1 below shows time plots of daily logarithmic returns of USD/GBP exchange rate, daily log returns of DAX index and of weekly log returns between USD and JPY. The log returns were calculated using the following formula: $r_t = 100ln\,(P_t/P_t - 1)$, where P_t is the price of an asset at time point t.

Figure 1. (a) Daily logarithmic returns of USD/GBP exchange rate. The span begins Apr 1, 1996 and ends Apr 30, 2013 and includes 6239 observations. (b) DAX daily logarithmic returns from May 1, 1999 to Apr 30, 2013 with 3570 of observations. (c) USD/JPY weekly log returns starting at Jan 7, 1990 and ending on May 3, 2013 – 744 observations in total. In all three series trading days were used

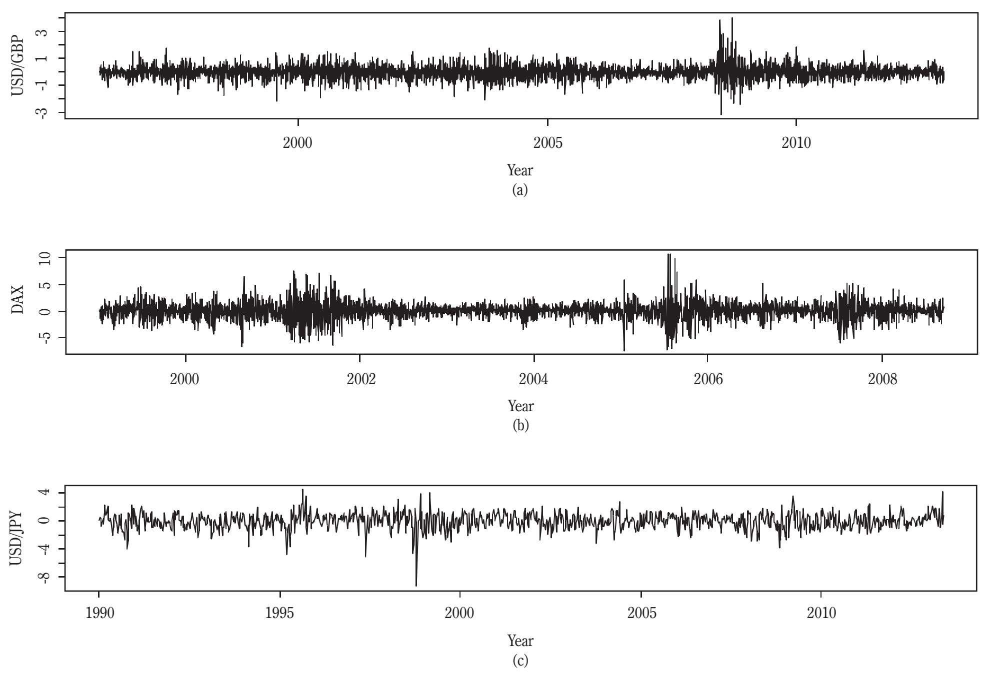

Source: The exchange rate data was taken from www. oanda.com and the data for the index comes from www. yahoo.finance.com.

The series seem to be very similar to a realization of a white noise process[2], i.e. the returns vary around a constant zero mean, however, the size of this variability changes over time. The latter feature is referred to as time-varying conditional volatility. Both series were more volatile in some periods than in other periods and periods of low (high) volatility tend to be followed by corresponding periods of low (high) volatility – a commonly encountered phenomenon in financial time series which is called volatility clustering or volatility persistence. These features give grounds for the assumption of serial correlation in the conditional volatility – the cornerstone assumption behind volatility models described in this article. Furthermore, a significant number of extreme observations (outliers) indicates that the return distribution might be non-normal. And the descriptive statistics (Table 1)[3] confirm that the probability of extreme observations is higher than if the returns were normally distributed, i.e. the distribution is leptocurtic – this is the so called heavy-tails property. Consequently, an appropriate volatility model needs to possibly well capture at least some of the mentioned stylized facts.

Table 1. Data summary of the series used in the article

Data	**N**	**Frequency**	**Beginning Date**	**End Date**
USD/GBP	6239	daily	01/04/1996	04/30/2013
DAX	3569	daily	03/05/1999	04/30/2013
USD/JPY	1217	weekly	07/01/1990	04/28/2013
KGHM	744	weekly	05/02/1999	03/05/2013
	Mean	**Standard Deviation**	**Skewness**	**Excess Kurtosis**
USD/GBP	-0.0003	0.4238	0.3227	6.0152
DAX	0.0108	1.5894	0.0247	4.0021
USD/JPY	-0.0310	1.1578	-0.6970	4.3534
KGHM	0.4385	6.3853	-0.5363	5.9950

3. Volatility Definition

Firstly, the precise definition of the term volatility used in the article shall be given. The starting point is the variable's continuously compounded return defined as a logarithmic difference of two adjacent observations. The subsequently

[2] A white noise process denotes a sequence of uncorrelated random variables with zero mean and constant variance. For details see, for example Hamilton (1994, p. 47).

[3] Because normally distributed data has excess kurtosis equal 0, the sample excess kurtosis of 2.26 of the series indicates that returns have more probability mass on the tail areas.

estimated standard deviation of the return defines volatility usually denoted in the financial literature as σ. There are other definitions of returns like geometric, harmonic, absolute etc. as well as there are other definitions of volatility like for example semi-variance. However, the definitions used in the article are the most typical in the literature. It is important to note that both the return and the standard deviation are values expressed per unit of time. For example, a series might have a return of 3.5% per year and volatility of 12.3% per year. However, in certain applications like options pricing or some VaR calculations other time units are needed so the values for standard deviation are recomputed using the following formula:

$$\sigma_{year} = \sigma_{day} \sqrt{\Gamma},$$

where σ_{year}, σ_{day} stand for the standard deviation on an annual and daily basis, respectively. And with T denoting the number of time points in the year, so if $T = 252$ the formula simply annualizes the daily standard deviation assuming 252 trading days in a year[4].

4. Univariate Conditional Heteroskedasticity Models

The widely used class of models designed to replicate the behavior of volatility was introduced by Engle (1982). His autoregressive conditional heteroskedasticity (ARCH) model and its extensions turned out to be very useful in studying the volatility of foreign exchange rates and others financial time-series.

4.1. ARCH Model Specification

To describe a log return series $\{x_t\}$, $t = 1, 2, \ldots, T$ with T observations a general model will be used:

$$x_t = \mu_t + a_t, \tag{1}$$

where μ_t denotes the conditional mean equation for x_t and a_t is referred to as an innovation, a shock or as a mean corrected return. In practice, the conditional mean equation will be often well represented by an white noise or an autoregressive model of small order (Tsay, 2002, p. 111). The shocks are assumed to be a white noise so that they have the following properties:

[4] For details see Hull (2012, p. 111).

$$E(a_t) = 0, \quad E(a_t^2) = \sigma_t^2, \quad Cov(a_t, a_{t-k}) = 0, \quad \text{for} \quad k \neq 0. \tag{2}$$

This means that the shocks oscillate around zero value[5], have time-varying variance σ_t and are serially uncorrelated. In the ARCH setting the $\{a_t\}$ series is modeled as

$$a_t = \sigma_t \varepsilon_t, \text{ with } \varepsilon_t \sim N(0,1) \text{ or } \sqrt{\frac{v}{v-2}}_t \varepsilon_t \sim t_v, v > 2.$$

The term σ_t denotes the conditional standard deviation of the process at a time t and $\{\varepsilon_t\}$ is a sequence of independent and identically distributed random variables with mean zero and variance one. In practice, $\{\varepsilon_t\}$ is often assumed to be standard normally or Student-t distributed[6]. The ARCH (1) model simply sets the conditional variance σ_t on its lagged value from the most previous period so that the model becomes

$$\sigma_t^2 = \alpha_0 + \alpha_1 a_{t-1}^2 = \alpha_0 + \alpha_1 (x_{t-1} - \mu_{t-1})^2 = Var(x_t \mid F_{t-1}),$$

where $\alpha_0 > 0$, $\alpha_1 \geq 0$ to ensure positiveness of the conditional variance and F_{t-1} denotes the information set at time $t-1$. It is also necessary to set $\alpha_0 + \alpha_1 < 1$, otherwise the model produces exploded values. To work with these models time-series need to satisfy the condition of stationarity. Put it simply, a stationary time-series has a constant unconditional variance which means that even though its conditional variance changes in time the unconditional variance remains constant. The models described in these article are not able to capture the behavior of non-stationary time-series. The equation above is called the volatility equation for x_t. It shows that the conditional variance in the period t has two components: a constant and last period's volatility.

To relate the conditional variance to more than only one lagged value of a_t an ARCH(m) model can be built:

$$\sigma_t^2 = \alpha_0 + \alpha_1 a_{t-1}^2 + \ldots + \alpha_m + a_{t-m}^2.$$

Consequently, the conditional variance σ_t^2 is thereby dependent upon the past squared shocks up to m lags. Moreover, because the shocks are squared

[5] None of the series presented above have statistically significant mean values different than zero, see Table 1.

[6] Because empirical distributions of financial data are often leptocurtic, a heavier-tailed t-distribution might fit data better than the normal disribution (Franke, 2011, p. 225).

both positive and negative values of the deviations have the same effect on the dependent variable – a feature that is rather inconsistent with empirical findings about volatility of financial data. In fact, volatility does not respond symmetrically to positive and negative returns; positive returns tend to produce smaller volatility than price drops[7]. This is an evident drawback of the model.

4.2. Generalized Autoregressive Conditional Heteroskedasticity (GARCH) Model

The application of ARCH models can be cumbersome because one often needs many parameters to adequately describe the volatility process of an asset return (Tsay, 2002, p. 93). Particularly, in financial applications involving the use of daily or weekly data the conditional variance can depend on volatilities going back a large number of periods. Because precise estimation of large number of parameters poses difficulties, an extension of the ARCH(m) model known as the generalized ARCH (or GARCH) model was introduced by Bollerslev (1986). The idea was to let the conditional variance be a function not only of squared innovations but also of its own lagged values. The GARCH (1,1) has the following form:

$$\sigma_t^2 = \alpha_0 + \alpha_1 a_{t-1}^2 + \beta_t \sigma_{t-1}^2, \text{ where } \alpha_0 > 0 \text{ and } \alpha_1,\ \beta_1 \geq 0.$$

The model can be extended to GARCH(p,q)

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{m} \alpha_1 a_{t-1}^2 + \sum_{j=1}^{s} \beta_j \sigma_{t-1}^2,\ \alpha_0 > 0,\ \alpha_j,\ \beta_j \geq 0 \text{ for } j = 1, 2 \ldots, n.$$

Unlike ARCH(m) models, GARCH models are very parsimonious so in practice it is sufficient to use only GARCH models of low orders, that is to set m and s to 1 or 2 lags maximally (Tsay, 2002, p. 134.).

4.3. Testing for ARCH Effects and Model Selection

After fitting the mean equation to the series, squared residuals can be used to check for conditional heteroskedasticity (ARCH effects). Usually, two tests are performed: Ljung-Box statistics of squared residuals and the Lagrange

[7] Tsay (2002, p. 80). One of the most important models that addresses this issue is the exponential GARCH (EGARCH) model (Nelson, 1991) where innovations are weighted (for details see Tsay, 2002, p. 102–103).

Multiplier test of Engle (1982), which is equivalent to F statistic for testing the joint hypothesis that all the regression parameters are zero (see Mills, 1999, p. 143 for details).

If ARCH effects are present, the a natural question is to how to determine the order of the model. In the case of ARCH (m) models the PACF of squared shocks a_t^2 might be used because a_t^2 is an unbiased estimate of σ_t^2 (Tsay, 2002, p. 119). When modeling financial data, it is recommended to fit models of different orders and then to choose the model that minimizes a given information criterion. In ARCH modeling the Akaike's Information Criterion (AIC) and Bayesian Information Criterion (BIC) are commonly used. The models with the smallest criterion value are chosen (for details see Brockwell, Davis, 2002, p. 355 and Mills (1999, p. 34–35).

4.4. Estimation and Diagnostic Checking

The estimation of ARCH/GARCH models is normally done using the maximum likelihood method. In the first step, the likelihood function needs to be determined and thereafter the parameters maximizing the function are searched for. Here, it is assumed that a_t are conditionally normally distributed with zero mean and variance σ_t^2. However, it is often more appropriate to assume t-distributed error terms albeit the computation becomes more complex. In order to check for the adequacy of the volatility model the usual Ljung-Box statistic on the squared standarized shocks from the estimation can be used. If p-values of statistics of the squared standarized residuals are high, then the null hypothesis of no autocorrelation cannot be rejected, residuals are white noise and it is concluded that the model adequately describes observations (for details see for example Hamilton, p. 117 and Franke p. 295 (ARCH) and p. 306 GARCH)).

5. Empricial Analysis of KGHM Returns

5.1. Checking for Autocorrelations and Finding the Conditional Mean Equation

Figure 2 displays prices and weekly log returns of polish stock ticketed 'KGHM'. Whereas price series does show strong trend and evident nonstationarity so the series in this form cannot be analyzed in the framework described above, log returns exhibit usual features of financial time series, ire observations oscillate around a constant (zero) mean and conditional volatility occurs in clusters. Additionally, around February 2009 some extreme observations are present[8]. To formally check if the log return series is indeed stationary the augmented Dickey-Fuller (ADF) unit-root test is carried out[9]. The test statistic of –5.321 with p-value of 0.01 enables to reject the null hypothesis at 5% significance level[10].

Figure 2. (a) KGHM weekly prices (in PLN). (b) KGHM weekly log returns. The data for both series spans from February 5, 1999 to May 3, 2013 and both series contain 744 observations. Datasource: www.stooq.pl

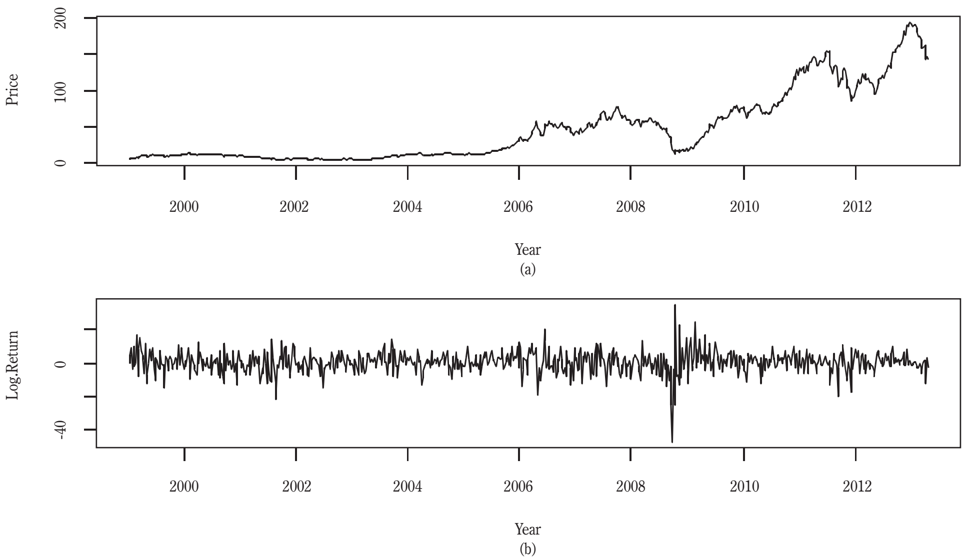

Source: own elaboration.

In order to find the conditional mean equation according to (1) sample autocorrelation and partial autocorrelation functions of log returns will be used[11]. Figure 3 shows that the series is not serially correlated at statistically significant level[12].

---

8 Outliers might be seen either as a measurement error or as a part of a data genereting process and their presence might influence standard error estimates in an OLS regression. For details on kinds of outliers and how to handle them see Montgomery (2003, p. 61).

9 A description of the test can be found in Tsay (2002, p. 77).

10 The null hypothesis is that a unit root is present in the autoregression against the alternative hypothesis of stationarity.

11 For definitions of SACF and PACF see for example Tsay (2002, p. 30). These functions are useful in determing the order of an autoregressive process.

12 Even though there seems to be a significant autocorraltion at lag 21, such a distinct lag has very a low weight and as further analysis will show the process might be well approximated by a white noise.

Figure 3. Sample autocorrelation (a) and partial autocorrelation (b) functions of KGHM weekly log returns. The two dashed lines denote the two standard error limits

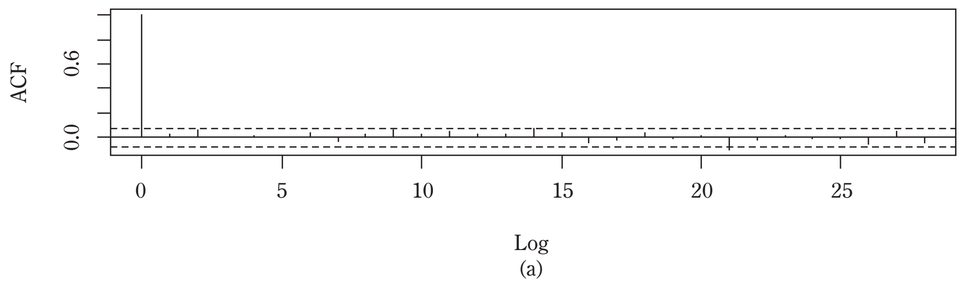

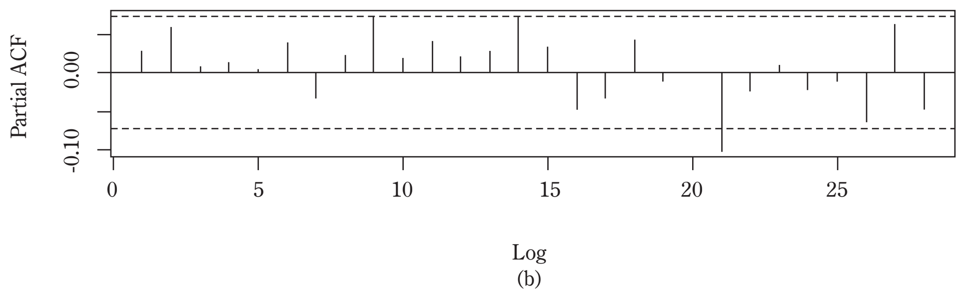

Source: own elaboration.

Additionaly, the statstical tests of Ljung-Box and F-test[13] confirm that there is none significant autocorrelations: the Ljung-Box statistic even for 21 degress of freedom is 29.2494 and its p-value is 0.1081 so that the null hypothesis of no autocorrelation cannot be rejected at standard significance levels.

As a result, the conditional mean equation will simply consist of a constant and an error term. Fitting an AR(0) model with non-zero mean[14] gives a value for the constant coefficient 0.4385 with std. error 0.2341 (and approx. p-value of 0.06143) so it might be concluded that the coefficient is statstically insignificant at 5% level and the return generating process as yet gets the following form:

$$x_t = \varepsilon_t a_t,\ \hat{\sigma}_a = 40.7717$$

To check the model adequacy the Ljung-Box statistic for the residuals might be used. If the model is adequate the standarized residuals should be similar

[13] For the description of the tests see Tsay (2002, p. 114). More results can be found in the table in the appendix.

[14] Other selection criteria like Akaike's Information Criteria also point to an autoregressive model of zero order, see Tsay (2002).

to a white noise. The statistic for 12 lags is 12.1813 with p-value of 0.4312 so that the null hypothesis of no serial autocorrelation cannot be rejected at any reasonable level[15].

5.2. Testing for ARCH effects

Having fitted the mean equation it is necessary to test for heteroskadasticity (ARCH effects)[16]. The usual way is to begin with sample autocorrelation and partial autocorrelation functions of the squared log returns. Figure 4 shows conditional volatility (as squared returns) together with sample autocorrelation and sample partial autocorrelations functions of the same series. Significant autocorrelation are present in both in the SACF and SPACF. Specifically, spikes in the latter indicate the presence of conditional heteroskedasticity. Additionaly, the test statistic from the Ljung-Box test equals 234.4170 with p-value of 0.00 which allows to reject the null hypthesis of no serial correlation in squared observations – ARCH effect are present (i.e. conditional volatility depends on its past values).

Figure 4. Conditional volatility (a) sample autocorrelation (b) and partial autocorrelation functions (c) of KGHM squared logarithmic returns

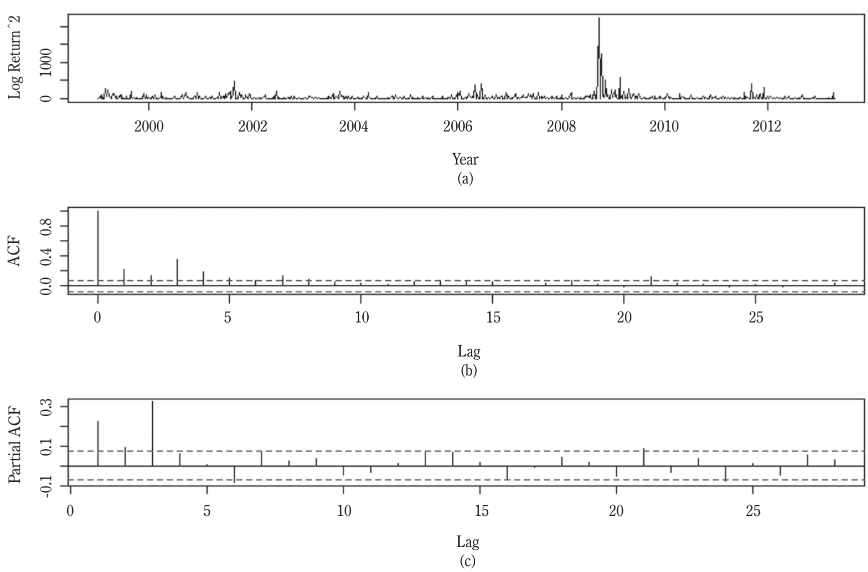

Source: own elaboration.

[15] See the table in the appendix for more results.

[16] Time-varying conditional variace distorts standard error estimates in a regression. Secondly, accounting for heteroskadasticity migh improve forecasts accuracy. See Engle (2001, p. 3–4).

5.3. Model Selection and Estimation

If ARCH effects have been confirmed, a model of the appropriate order needs to be selected. One way to initially determine the order of an ARCH model is to use the sample partial autocorraltion function of squared shocks (Tsay, 2002, p. 120). Figure 4 (b) shows that an ARCH model of order 3 might be appropriate. However, to get more precision it is recommended to fit models of different orders and then to choose the model that maximizes the log-likelihood function and/or minimizes a given information criterion. Table 2 presents a few models of different orders and its Akaike's information criteria and it can be seen that GARCH (1,1) model is the model which minimizes the AIC, so this model is selected. Consequently, the model for the whole series takes the following form:

$$\sigma_t^2 = 2.23064 + 0.10391 a_{t-1}^2 + 0.83749 \sigma_{t-1}^2$$

Table 2. Estimation results of different ARCH models

	ARCH(3,0)		**ARCH(2,0)**		**GARCH(1,1)**	
	parameter	**t-value**	**parameter**	**t-value**	**parameter**	**t-value**
ARCH(1)	0.15435	6.26988	0.19513	6.96941	0.10268	7.56289
ARCH(2)	0.06736	2.09736	0.06516	2.11339	X	X
ARCH(3)	0.08720	2.32984	X	X	X	X
GARCH(1)	x	X	X	X	0.83971	31.49698
GARCH(2)	x	X	X	X	X	X
Intercept	26.57210	15.21691	28.83570	18.49637	2.19797	3.35148
Log-Likelihood	-2384.99		-2392.79		-2370.95	
AIC	6.4401		6.4477		6.3836	

Source: own elaboration.

To check the model adequacy autocorrelation and partial autocorrelations of the residuals are investigated. As can be seen from Figure 5 showing standarized residuals and conditional volatility, the selected model quite accurately describes the data generating process.

Figure 5. (a) Residuals from the model. (b) Estimated conditional volatilty for the kghm log returns

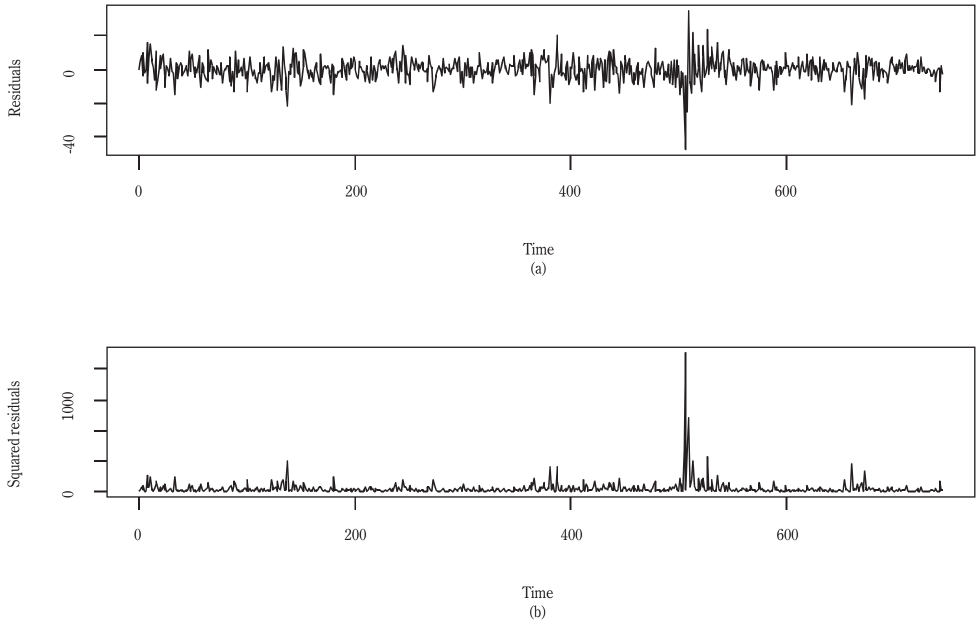

Source: the estimation made by the author.

6. Conclusion

Based on abundant research as well as examples given in this article many financial time series exhibit certain common features. Usually, returns of different asset classes oscillate around a constant mean, but observations tend to be more volatile in some periods than in other periods. This phenomenon, called volatility clustering, suggests some nonlinear (conditional on past observations) dependence in observations and became the main motivation fot the introduction of ARCH (autoregressive conditional heteroskedasticity) models by R. Engle in 1987. This model has proved to be very succesful and gave rise to many other, more complex models. This article describes only its widely used extension the generalised ARCH or GARCH model – a more parsimonious but very useful version, however these models often serve just as a building block for other versions.

The author of the article tried to find out to what extent a polish stock KGHM exhibits typical features of financial time series and how well the series might be described by the ARCH modeling framework presented in the article. Because the polish stock market and its main components, of which KGHM is by far the biggest, would be still classified as emerging, it might be expected to see some discrepancies in the features of time series in comparison to more mature markets. As it is widely recognized that mature markets like for example the Geraman stock index DAX are fairly efficient, then such serious discrepancies might be interpreted as evidence against market efficiency.

What has been shown in the article is that weekly log returns of polish stock KGHM do exhibit typical features of financial time series, that is, the series varies around zero mean and exhibits volatility clusters. Additionally, descriptive statistics show that the return distribution is non-normal as the series displays slightly negative skewness and has significant excess kurtosis. As shown by tests and the sample autocorrelation as well as partial autocorrelation functions there is none significant serial autocorrelations in returns, however, this is not the case for squared observations which show significant autocorrelation up to lag 3. Obviously, it is evident that ARCH effects are present in the series. Subsequent analysis and model fitting revealed that the best model (as selected by the log-likelihood and AIC) to fit the series, is the GARCH(1,1).

As a result, the analysis of the KGHM returns series given revealed no discrepancies with respect to features typical in other financial time series and because the series might be well described with the usual approach, there is lack of evidence of market inefficiency in comparison to mature markets.

Bibliography

Bollerslev, T. (1986). Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, *31*: 307–327.

Cryer, J., Kung-Sik Chan (2008). *Time Series Analysis With Applications in R*. Springer.

Engle, R. (Autumn, 2001), GARCH 101: The Use of ARCH/GARCH models in Applied Econometrics. *The Journal of Economic Perspectives, 18*: 157–168.

Engle, R. (1982). Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of the U.K. Inflation. *Econometrica*, *50*: 987–1008.

Franke, J., Höardle, W.F., Hafner, Ch.M. (2011) (3rd ed.). *Statistics of Financial Markets*. Springer.

Franses, P.H. and van Dijk, D. (2003). *Nonlinear Time Series Models in Empirical Finance*, Cambridge University Press (Virtual Publishing).

Green, H.W. (2003) (5th ed.). *Econometric Analysis*. Prentice Hall.

Hamilton, D.J. (1994). *Time series analysis*. Princeton: Princeton University Press.

Harris, R. and Sollis, R. (2003). *Applied Time Series Modeling and Forecasting*. Wiley.

Hull, J.C. (2012). *Risk Management and Financial Institutions*. Pearson.
Shumway, R. and Stoffer, D. (2011) (3rd ed.). *Time Series Analysis and Its Applications With R Examples*. Springer.
Tsay, R. (2002) (3rd ed.). *Analysis of Financial Time Series*. Wiley.
Zumbach, G. (2013). *Discrete Time Series, Processes, and Applications in Finance*. Springer.