NON-STATIONARY VOLATILITY WITH HIGHLY ANTI-PERSISTENT INCREMENTS: EVIDENCE FROM RANGE-BASED VOLATILITY

Ladislav KRISTOUFEK*

Abstract. *We analyze range-based volatility of three highly capitalized companies to show that the volatility process is non-stationary and its logarithmic transformation together with logarithmic increments is approximately normally distributed. Further, the increments are shown to be highly anti-persistent. Together with the assertion that logarithmic returns are normally distributed, and uncorrelated with time-varying volatility, we propose a new returns-generating process. The whole procedure is based on empirical observations without any limiting assumptions. We reconstruct returns series which remarkably mimic the real-world series and posses the standard stylized facts – uncorrelated returns with heavy tails, strongly auto correlated absolute returns and volatility clustering.*

Keywords*: volatility modelling, anti-persistence, non-stationarity.*

1. Introduction

Accurate modelling and forecasting of volatility is one of the biggest challenges in financial economics and financial econometrics. Historically, there are four major groups of volatility forecasting approaches – historical volatility, conditional heteroskedasticity models, stochastic volatility models, and implied volatility models. Historical volatility models contain many different approaches such as random walks, historical averages, HAR models, autoregressive moving averages and their fractional generalizations, exponential smoothing approaches and others. Conditional heteroskedasticity group contains a wide portfolio of (G)ARCH models and their specifications such as EGARCH, IGARCH, FIGARCH and many others. Stochastic volatility models add a stochastic term into the conditional variance equation contrary to the previous group. Implied volatility models are based on the Black-Scholes option pricing formula

* Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Institute of Economic Studies, Charles University, kristouf@utia.cas.cz

[2, 27] and its various specifications and generalizations. These approaches are nicely reviewed and compared in two studies of Poon & Granger [30, 31]. In their second study [31], the authors argue that the implied volatility models outperform the others, followed by the historical volatility models, in volatility forecasting. This is a rather interesting, and disturbing, finding since the Black-Scholes formula is known to be based on highly unrealistic assumptions of the returns process. This might imply that the other approaches are actually not optimal and there is some other approach closer to reality.

As nicely summarized by Morana & Beltratti [28], there are two basic approaches to volatility modeling – structural breaks (and regime switching) and long-range dependence. In this paper, we propose an alternative approach to the volatility modelling. Based on a simple statistical analysis, we show that volatility can be effectively modelled as a non-stationary process with highly anti-persistent logarithmic increments, which are, moreover, normally distributed. By doing so, we are able to mimic the basic stylized facts of the financial returns – no autocorrelation, highly persistent absolute returns, non-normality, fat tails, and volatility clustering [7]. Note that mainly the ability to produce fat tails is a great achievement as the leptokurtosis is usually imposed to the simulated processes by various fat-tailed distributions of the innovations – Student's *t*, generalized error distributions, and others. We approach the problem from the opposite direction, starting from the real-world data. By analyzing the series of three American highly capitalized stocks, we are able to state eight basic Claims. Based on these Claims, we argue that the logarithmic returns are uncorrelated and normally distributed with approximately zero mean and time-varying standard deviation (volatility). The logarithm of the standard deviation is non-stationary and approximately normally distributed with approximately normally distributed increments, which are stationary and highly anti-persistent. Based on these findings, we are able to reconstruct the series of returns, which resemble the actual financial returns and the stylized facts very closely. Note that the construction of returns is based solely on the normal distributions. These results might have a crucial effect on volatility modeling and connected topics – mainly portfolio analysis and (conditional) Value-at- Risk.

The paper is structured as follows. In Section 2, we present the basic methodology needed for the statistical analysis conducted in the paper. In Section 3, we describe the dataset and present the crucial findings about the process of volatility. In Section 4, we show whether the findings can be used for simulation of series, which would resemble the real-world financial series. Section 5 concludes and states directions for further research in this area.

2. Methodology

In this section, we describe the basic methodology used in the rest of the paper. Hence, we describe the notions of long-range dependence, tests for stationarity (and unit roots) and tests for long-range dependence (specifically anti-persistence in our case).

2.1. Long-range dependence and stationarity

Long-range dependence is highly connected to the Hurst effect, i.e. a situation characteristic with long periods when the series is above the mean which are followed by long periods when the series is below the mean of the series while the series still remains stationary, and so also Hurst exponent H. A critical value of Hurst exponent is 0.5 and suggests two possible processes. H being equal to 0.5 implies either an uncorrelated or a short-term dependent process [22, 1]. Uncorrelated process has zero auto-covariances at all non-zero lags and short-term dependent process shows non-zero auto-covariances at low lags vanishing exponentially to zero for higher lags. For $H > 0.5$, auto-covariances of the process are positive at all lags so that the process is called long-range dependent with positive correlations or persistent. Auto-covariances of such process are hyperbolically decaying and non-summable [1]. On the other hand, for $H < 0.5$, the process is said to be long-range dependent with negative correlations or anti-persistent. Similarly to the previous case, auto-covariances of such process are slowly decaying but summable. The persistent process is visibly locally trending while the anti-persistent process switches signs more frequently than a random process would [34, 25, 13, 1].

Hurst exponent H is connected to parameter d of fractional integration so that $H = d + 0.5$. Long-range dependent processes are frequently defined in two domains – time and frequency:

- *(time domain)* A stationary process with an autocorrelation function $\rho(k)$ decaying as $\rho(k) \propto k^{2H-2}$ for $k \rightarrow +\infty$ is called long-range dependent with Hurst exponent H.
- *(frequency domain)* A stationary process with a spectral density $f(\lambda)$ following $f(\lambda) \propto |\lambda|^{1-2H}$ for $\lambda \rightarrow 0$ is called long-range dependent with Hurst exponent H.

Note that a notion of long-range dependence is tightly connected to stationarity of the process, which is obvious from both presented definitions. Without stationarity, there can be no standard long-range dependence. Going back to the two given definitions, we can construct neither autocorrelation function nor spectrum for a non-stationarity process due to infinite variance; hence we cannot talk about standard long-range dependence. To check for stationarity, we apply two standard tests – ADF [10] and KPSS [19]. ADF has a null hypothesis of a unit root and can take lags of the differenced series into consideration and thus controlling for the memory effects. KPSS, on contrary, has a stationarity null hypothesis and it also can control for memory effects with a use of auto-covariance adjusted variance with Barlett weights [14].

2.2. Anti-persistence tests

As we are interested primarily in long-range dependence and not necessarily in a specific value of H or d, we use (modified) rescaled range and rescaled variance analyses for testing the presence of long-range dependence in the studied process.

2.2.1. Classical rescaled range analysis

Rescaled range analysis (*R/S*) is the traditional Hurst exponent estimation method proposed by Hurst [16] and further adjusted by Mandelbrot & Wallis [26]. In the procedure, the series of returns of length T is divided into N adjacent sub-periods of length υ so that $N\upsilon = T$. For each sub-period, a rescaled range of a profile X_t (cumulative deviations from an arithmetic mean) is calculated as R/S, where $R = max(X_t) - min(X_t)$ is a range of the corresponding profile and S is a standard deviation of the corresponding returns. Rescaled ranges are calculated for each sub-period of length υ and an average rescaled range is calculated [29]. The rescaled ranges scale as:

$$(R/S)_\upsilon \propto \upsilon^H.$$

V statistic, which is also used for a cycles detection, a stability testing of the Hurst exponent or a change in scaling behaviour (crossover) detection, is defined as:

$$V_\upsilon = (R/S)_\upsilon/\sqrt{\upsilon}$$

and converges to the distribution with a rather complicated distribution function F_V for independent processes (see Refs. [24, 16, 29] for details and critical values). For our purposes, we set $\upsilon = T$ to test for the presence of long-range dependence in the underlying process by the means of V statistic.

2.2.2. Modified rescaled range analysis

Due to the bias of *R/S* analysis for short-term memory processes, Lo [24] proposed the modified rescaled range analysis (*M−R/S*) which utilizes a modified standard deviation S_M. S_M is defined with a use of auto-covariances of the original increment series γ_j up to lag ξ as:

$$S^M = \sqrt{S^2 + 2\sum_{j=1}^{\xi} \gamma_j \left(1 - \frac{j}{\xi + 1}\right)} \ .$$

R/S is thus a special case of $M - R/S$ with $\xi = 0$. The distribution of the modified V statistic converges to F_V not only for independent processes but also for short-range dependent ones. Choice of the correct lag ξ is critical for the estimation of the modified rescaled ranges [35, 33]. Lo [24] suggested optimal lag based on the sample first-order autocorrelation coefficient of the original series $\rho(1)$. However, we set some fixed values of ξ (defined later) since the automatic lag assumes that the underlying process is $AR(1)$, which is rather strict.

2.2.3. Rescaled variance analysis

Rescaled variance analysis (V/S) was proposed by Giraitis and co-authors [14] as a modified version of KPSS statistic [19], which is usually used for testing of stationarity but was also shown to have good power for series with long-term memory [20, 21]. The procedure is very similar to the modified rescaled range analysis and differs in a use of a sample variance of the profile of the series instead of the range. As an alternative to the V statistic, the M statistic is defined as:

$$M_\upsilon = \frac{\mathrm{var}(X_t)}{\upsilon (S^M)^2}.$$

Note that the modified standard deviation is used so that the method is robust to short-range dependence as well. Variance of the M statistic is much lower than the one of R/S and $M - R/S$ so that the confidence intervals are much narrower and the method is thus more efficient. Similarly to R/S and $M - R/S$, one needs to set $\upsilon = T$ and construct M statistic if testing for the presence of long-range dependence in the process.

2.2.4. Moving block bootstrap

The bootstrap method [11] has been proposed to deal with the statistical properties of small samples. The basic notion behind the procedure is resampling with replacement from the original series and repeated estimation of a specific parameter. By shuffling, a distribution of the original series remains unchanged while possible dependencies are distorted [8]. Hypothesis can be then tested with a use of p-values based on the bootstrapped estimates. For our purposes and for the time series analysis in general, the simple bootstrap is not enough as the shuffling rids us not only from the long-range dependence but the short-range dependencies as well. Srinivas & Srinivasan [32] proposed a modified method which retains the short-term dependence characteristics but lacks the long ones – the moving block bootstrap with pre-whitening and post-blackening.

In the procedure, the original time series $\{x_t\}^T_{t=1}$ =1 is firstly pre-whitened (filtered) by a specific process – usually $AR(p)$ or $MA(q)$ – and residuals ε_t are obtained. Residuals are centered so that the centered residuals $\{\varepsilon_{t,c}\}^T_{t=1}$ are just the demeaned residuals. The series $\{\varepsilon_{t,c}{}^T_{t=1}$ is divided into m blocks of length ζ while $m\zeta = T$. The blocks are reshuffled and post-blackened with the use of the model from the pre-whitening part and residuals $\varepsilon_{t,c}$ to form the new bootstrapped time series $\{x_{t,b}\}^T_{t=1}$. Such time series retains the short-range dependence, potential heteroskedasticity and short-term trends as well as the distribution of the original time series. Importantly for the long-range dependence testing, for small enough ζ, the long-term correlations are torn. The parameter of interest is then estimated on the new time series. Bootstrapping is repeated B times so that distributions of M_υ and V_υ are obtained and used to get p-values for the null hypothesis of no long-range dependence in the series.

For our purposes, we use $AR(1)$ process for pre-whitening and post-blackening, which is standard in finance literature [12, 17], and $\zeta = 20$.

Such choice of ζ should be sufficient for ridding of the potential long-range dependence while the other properties remain similar to the original process. The procedure is repeated for lags $\xi = 0,1,2,5,10,15,30,50,75,100$ for both modified rescaled range analysis and rescaled variance analysis with $B = 1000$ bootstrap repetitions. As will be visible in the following sections, the analyzed series are on the edge between stationarity and non-stationarity, and we will eventually analyze the first differences of the series. However, for such boundary cases, there is a high risk of over-differencing which would inflict MA(1) process in the series. To control for this, we also apply the moving block bootstrap with the same parameters as noted before but with MA(1) for pre-whitening and post-blackening. This way, we can be more confident about our findings while controlling for the most problematic cases.

2.3. Garman-Klass daily variance estimator

There are various estimators of daily volatility (or variance) ranging from very simple absolute and squared returns through model-based estimators (e.g. GARCH or implied volatility based) to range-based estimators and realized variance [6]. From many possibilities, we choose Garman-Klass estimator:

$$\widehat{\sigma^2_{GK,t}} = \frac{(\log(H_t / L_t))^2}{2} - (2\log 2 - 1)(\log(C_t / O_t))^2,$$

where H_t and L_t are daily highs and lows, respectively, and C_t and O_t are daily closing and opening prices, respectively. This estimator does not take overnight volatility into consideration but is very simple and efficient (much more efficient than absolute and various power-returns, comparable with other range-based measures and less efficient than the realized variance) [6]. Even though the realized variance would be a more efficient choice, it is not easily obtainable for all assets while for the Garman-Klass estimator, all necessary variables are available freely for practically all financial assets. Moreover, the Garman-Klass estimator does not impose any specific model on the daily variance process, compared to (*G*)ARCH and stochastic volatility models.

3. Data and statistical analysis

We analyze series of three stocks with some of the highest capitalizations in the US markets – AAPL (Apple Inc.), IBM (International Business Machines Corporation) and MSFT (Microsoft Corporation) – between 3.1.2000 and 29.2.2012 (3059 observations). Normalized prices are shown in figure 1. We can see that even though all three companies are technological, they underwent very different dynamics during the analyzed period. Apple, as a favorite of the most recent days, has grown remarkably while IBM and Microsoft have been rather stagnant. Estimates of daily volatility are shown in figure 2. IBM and MSFT experienced very similar dynamics of volatility as well as its levels. AAPL, on contrary, shows markedly higher average values of volatility with more extreme values than the other two. Nevertheless, we do not observe any drastic jumps in the volatility levels and we rather find smooth transitions from lower to higher levels or vice versa. Moreover, all three series seem non-stationary at least visually as we notice phases with varying means. This observation is checked on statistical basis later in the text.

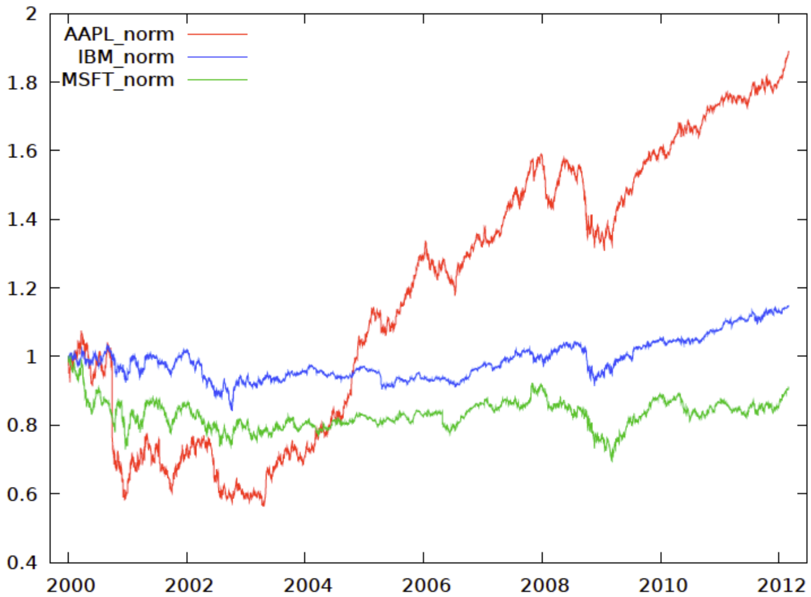

Figure 1. Normalized stock indices prices. Logarithmic prices divided by the first observation of the respective series for a better comparison.

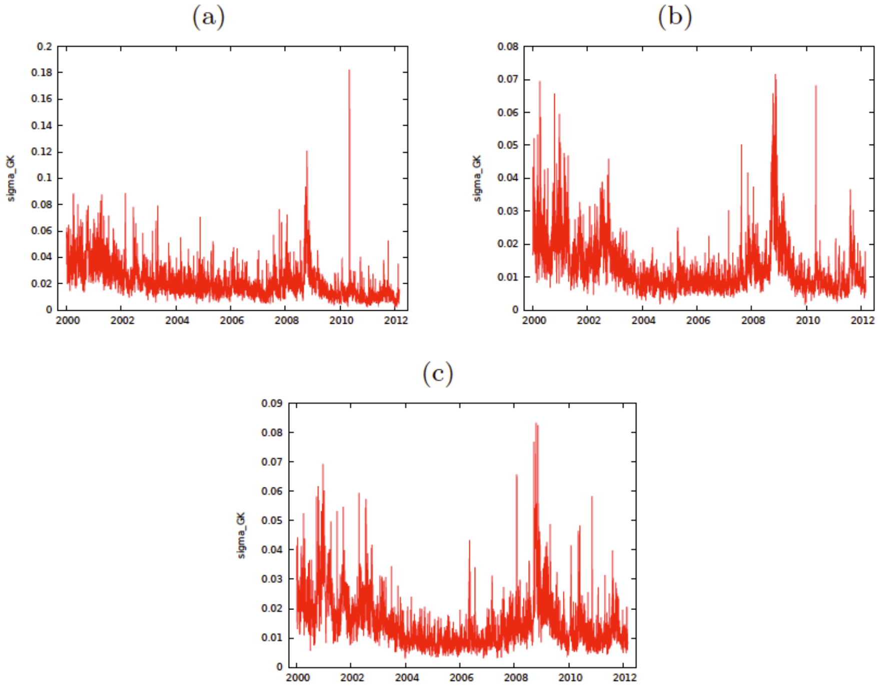

Figure 2. Garman-Klass estimates $\widehat{\sigma_{GK}}$ of volatility for (a) AAPL, (b) IBM and (c) MSFT.

As the Garman-Klass estimator does not take overnight dynamics into consideration, we analyze daily logarithmic returns defined as $r_t = \log(C_t) - \log(O_t)$. Descriptive statistics of r_t are summarized in table 1. The table also includes the descriptive statistics of the standardized returns (returns standardized by the square root of the estimated daily variance). We can see that the raw returns are fat-tailed and positively skewed while the standardized returns are very close to having normal tails. This result is supported by the Jarque-Bera test [18] – the raw returns are not normally distributed but the standardized returns are very close to being normally distributed. This is also supported by the quantile-quantile (*QQ*) plots in figure 3. Based on Ljung-Box test [23], we find no significant autocorrelations in the first thirty lags for the standardized returns. Based on these findings, we propose the first Claim[1]:

Claim 1. *Logarithmic open-close returns are uncorrelated and normally distributed with time varying volatility* $\approx N(0, \sigma_t)$.

[1] All ”Claims” presented in this paper should be taken as approximate results. Nevertheless, we show later in the text that these ”Claims” can be used to construct the series which strongly resemble the basic stylized facts of the financial returns.

Table 1. *Descriptive statistic of returns and returns standardized by* $\widehat{\sigma_{GK}}$.

	mean	SD	skew.	ex. kurt.	JB	p-val.	Q(30)	p-val.
r_{AAPL}	0.0001	0.0247	0.0432	2.4504	766	0.0000	65	0.0000
r_{IBM}	0.0008	0.0150	0.0836	4.7032	2823	0.0000	69	0.0000
r_{MSFT}	-0.0001	0.0173	0.2105	3.4390	1530	0.0000	46	0.0300
$r_{st,AAPL}$	0.0336	1.0374	0.0510	-0.3711	19	0.0000	40	0.1120
$r_{st,IBM}$	0.1059	1.0111	0.0551	-0.1963	6	0.0396	47	0.0250
$r_{st,MSFT}$	0.0027	0.9880	0.0552	-0.2216	8	0.0201	36	0.2180

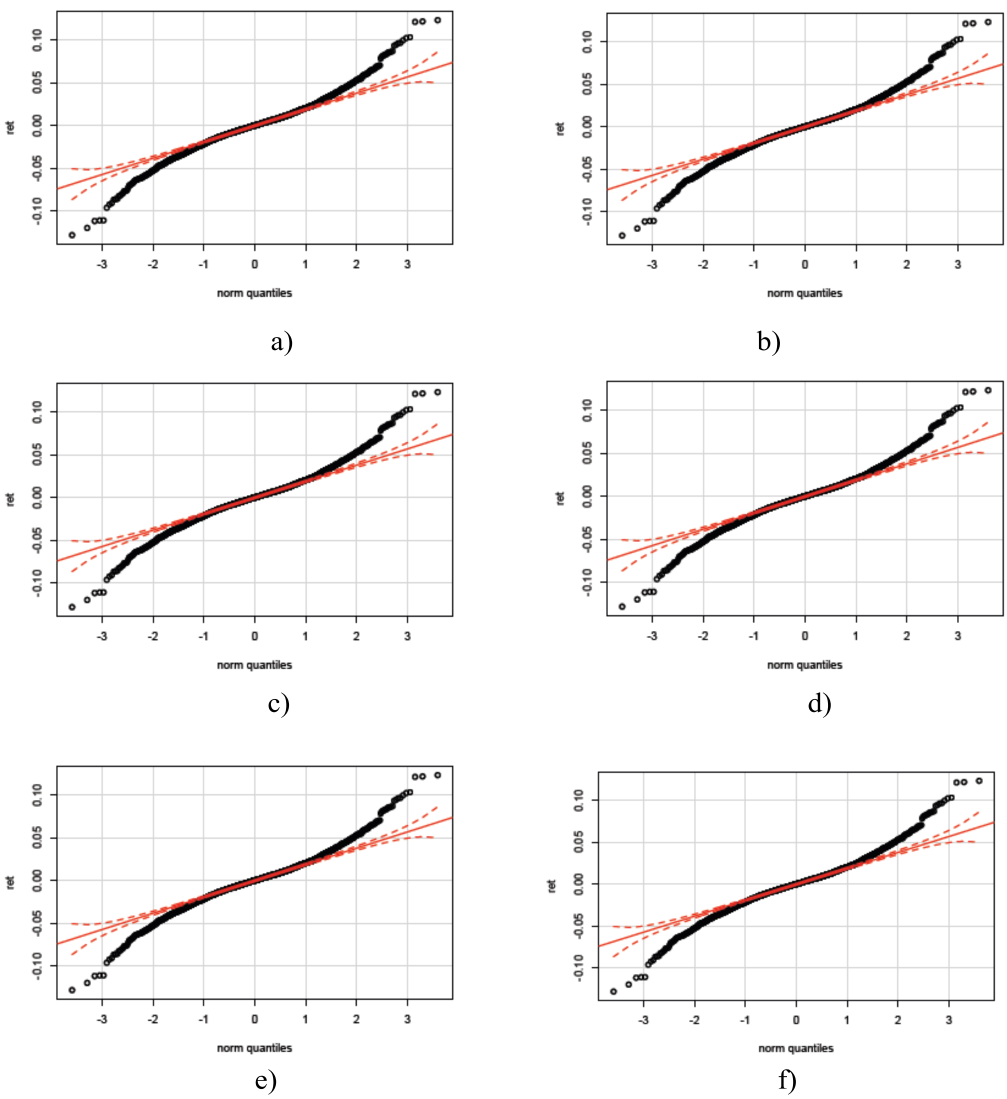

Figure 3. *QQ plots for raw (first row) and standardized (second row) returns*. AAPL (first column), IBM (second column) and MSFT (third column). 95% confidence intervals for normal distribution are marked by red dashed lines.

Let us now focus on the approximate distribution of the volatility process. Figure 4 presents the *QQ* plots against the normal distribution and uncovers that the process of volatility is far from being normally distributed whereas its logarithmic transformation is approximately normally distributed and evidently, so are the increments of logarithmic volatility. Note that these findings also indicate that there are no obvious structural breaks which would either cause the logarithmic increments to be fat-tailed (if there were some sudden jumps in the volatility levels) and/or the logarithmic volatility to be bi- or multi-modal (neither the *QQ* plots nor the histograms, not shown here, indicate such a characteristic). Therefore, we propose two other Claims:

Claim 2. *Logarithmic volatility is close to being normally distributed with a mean* $\mu_{\log\sigma}$.

Claim 3. *Increments of logarithmic volatility are close to being normally distributed* $\approx N(0,\sigma_{\Delta})$.

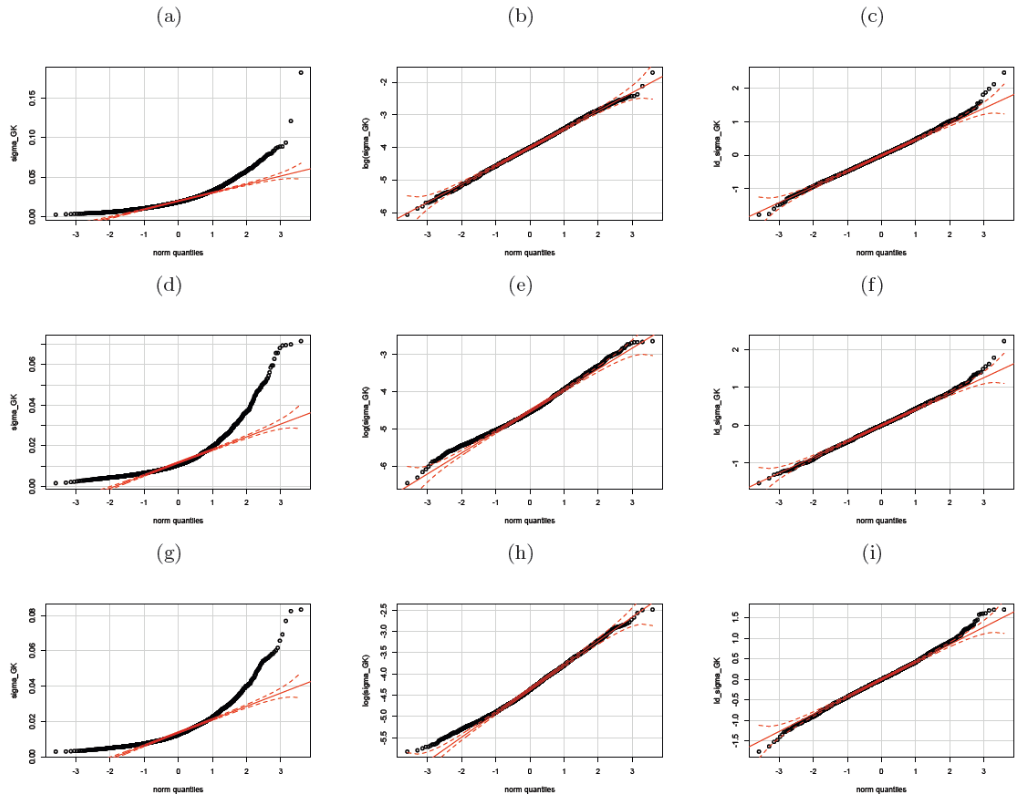

Figure 4. *QQ plots for* $\widehat{\sigma_{GK}}$ *(first column),* $\log\widehat{\sigma_{GK}}$ *(second column) and* $\Delta\log\widehat{\sigma_{GK}}$ *(third column).* AAPL (first row), IBM (second row) and MSFT (third row). 95% confidence intervals for normal distribution are marked by red dashed lines.

Table 2. Stationarity tests for $\widehat{\sigma_{GK}}$. KPSS with the null hypothesis of stationarity and ADF with the null hypothesis of a unit root. *, ** and *** for significance at 10%, 5% and 1% significance level, respectively.

	KPSS(1)	KPSS(10)	KPSS(100)	ADF(1)	ADF(10)	ADF(100)
$AAPL$	43.8283***	11.2969***	1.7515***	−18.8717***	−6.3622***	−2.2830
IBM	20.5034***	4.7213***	0.7410***	−15.0817***	−5.7168***	−3.1221**
$MSFT$	21.2672***	5.0439***	0.7780***	−16.1605***	−5.9202***	−2.6700*
$\log AAPL$	57.5655***	13.8110***	1.9881***	−16.5976***	−5.6643***	−1.7975
$\log IBM$	26.2098***	5.9606***	0.8626***	−14.8127***	−5.4914***	−2.6609*
$\log MSFT$	24.8671***	5.7543***	0.8176***	−15.9073***	−5.2080***	−2.2931
$\Delta \log AAPL$	0.0012	0.0055	0.0391	−62.8961***	−24.8072***	−7.7096***
$\Delta \log IBM$	0.0009	0.0038	0.0241	−63.5708***	−23.2364***	−6.7480***
$\Delta \log MSFTL$	0.0015	0.0072	0.0431	−64.9717***	−24.9563***	−7.1949***

Now, we focus on an essential question of stationarity of the series. The results for ADF and KPSS tests for various lags are summarized in Table 3. We use lags 1, 10 and 100 to control for practically no memory, short memory and long-term memory, respectively. The results are quite straightforward. Firstly, volatility process is neither stationary not an evident unit root process. For long lags, we are not able to reject the unit root null for AAPL and MSFT. Secondly, the same is true for the logarithmic transformation of volatility but for all three stocks. Moreover, the results are much stronger here as the series are very close to being normally distributed which is assumed for the tests. Here again, after controlling for long-range dependence (100 lags), we cannot reject the unit root of the series. This indicates that after controlling for long-range dependence in the increments, we cannot reject unit root for the logarithmic volatility. Thirdly, increments of the logarithmic volatility are asymptotically stationary even after controlling for long-term memory. These indicate that the first differences of the logarithmic volatility are likely to be long-range dependent. Based on these findings, we propose four other Claims:

Claim 4. *Volatility and logarithmic volatility are non-stationary.*

Claim 5. *Logarithmic volatility is close to a unit root process after controlling for the long-term memory.*

Claim 6. *Increments of logarithmic volatility are asymptotically stationary.*

Claim 7. *Increments of logarithmic volatility are likely to be long-range dependent.*

Table 3. Anti-persistence tests – bootstrapped p-values for the null of “no anti-persistence” controlling for $AR(1)$ process. ξ stands for the number of lags taken into consideration for standard deviation S^M for V and M statistic.

ξ	V_{AAPL}	M_{AAPL}	V_{IBM}	M_{IBM}	V_{MSFT}	M_{MSFT}
0	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
1	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
2	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
5	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
10	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
15	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000
30	0.0090	0.0000	0.0040	0.0000	0.0000	0.0000
50	0.0710	0.0000	0.0150	0.0000	0.0060	0.0030
75	0.3930	0.0150	0.0040	0.0000	0.0280	0.0180
100	0.7700	0.0680	0.0090	0.0000	0.1610	0.0830

Therefore, it is needed to analyze the increments of logarithmic volatility and its correlation structure. Figure 5 presents the autocorrelation and partial autocorrelation functions. We see a common pattern for all three analyzed series – strongly negative autocorrelation at the first lag for ACF which vanishes for further lags, and negative partial autocorrelations which decay relatively slowly to zero for PACF. This is indicative for two possible processes – a strong MA(1) process [4] or a strongly anti-persistent ARFIMA(0,d,0) process [15, 3, 9]. Debowski [9] actually shows that ARFIMA processes can be generalized so that we obtain stationary and invertible processes even for anti-persistent processes with $d \in (-1, 0)$, i.e. $H \in (-0.5, 0.5)$. To test for anti-persistent processes while still controlling for short-term memory as well as potential over-differencing[2], we use the modified rescaled range analysis and rescaled variance analysis with moving block bootstrap p-values for the null hypothesis of no anti-persistence with AR(1) and MA(1) processes in pre-whitening and post-blackening procedures. The results for ξ = 0, 1, 2, 5, 10, 15, 30, 50, 75, 100 are summarized in tables 3 and 4. We observe that the results for both AR(1) and MA(1) controls are practically the same – for short to medium lags, we reject ”no anti-persistence” null hypothesis, while for long lags, we do not. However, it is hard to distinguish between short and long-term memory for such long lags so that we treat the series as anti-persistent. For further discussion of the issue, see Beran [1]. Based on these results, we propose the last Claim:

[2] Since the unit root tests have low power when too many lags are taken into consideration, it is possible that taking the first differences of potentially spuriously detected unit root process imposes a strong moving average process, MA(1).

Claim 8. *Increments of logarithmic volatility are strongly anti-persistent with* $d \in (-1, 0)$, *i.e.* $H \in (-0.5, 0.5)$.

Table 4. Anti-persistence tests – bootstrapped *p*-values for the null of "no anti-persistence" controlling for *MA*(1) process. ξ stands for the number of lags taken into consideration for standard deviation S^M for *V* and *M* statistic.

ξ	V_{AAPL}	M_{AAPL}	V_{IBM}	M_{IBM}	V_{MSFT}	M_{MSFT}
0	0.0080	0.0000	0.0050	0.0000	0.0000	0.0000
1	0.0070	0.0000	0.0040	0.0010	0.0000	0.0030
2	0.0040	0.0000	0.0040	0.0000	0.0000	0.0010
5	0.0040	0.0000	0.0020	0.0000	0.0000	0.0030
10	0.0100	0.0000	0.0030	0.0000	0.0010	0.0020
15	0.0030	0.0000	0.0030	0.0000	0.0000	0.0020
30	0.0210	0.0000	0.0070	0.0000	0.0000	0.0050
50	0.1030	0.0010	0.0220	0.0000	0.0090	0.0220
75	0.4050	0.0500	0.0030	0.0000	0.0430	0.0480
100	0.7140	0.1930	0.0040	0.0000	0.1550	0.1390

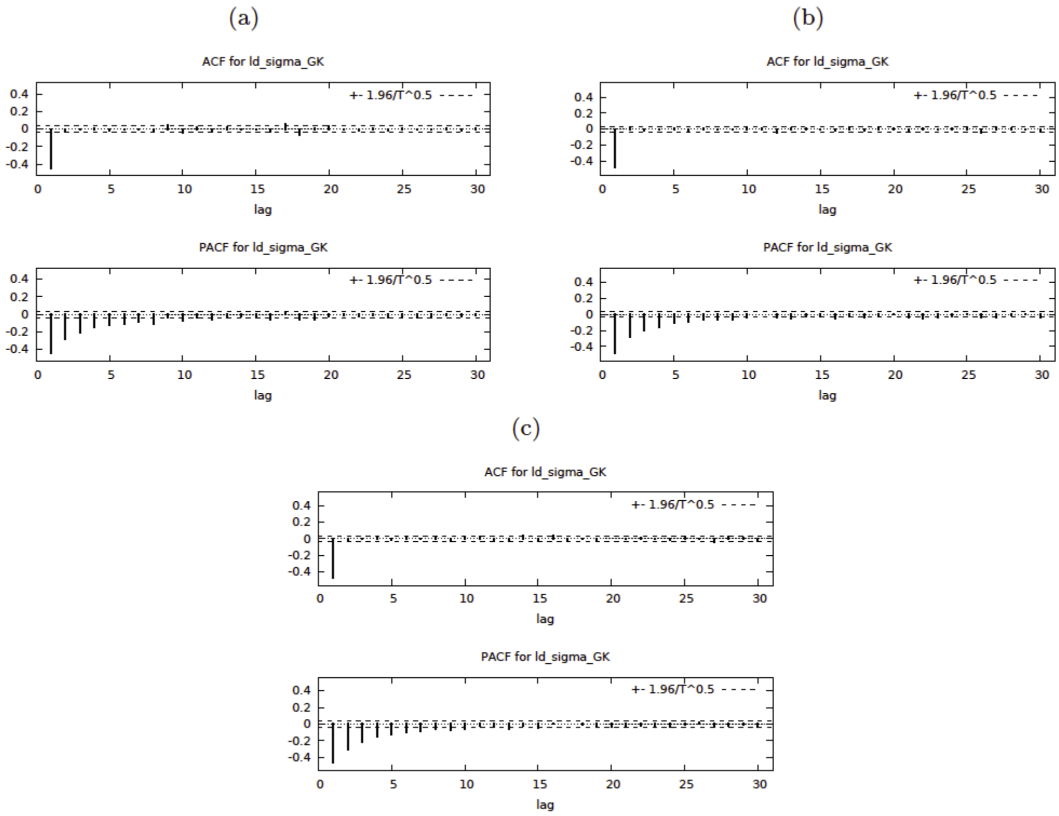

Figure 5. Autocorrelation and partial autocorrelation functions of logarithmic differences Of $\widehat{\sigma_{GK}}$ (a) AAPL, (b) IBM and (c) MSFT.

4. Simulations

Based on all the Claims we made, we now try to reconstruct the series with the observed statistical properties and compare whether these are in hand with the real financial series. To do so, we need to estimate three parameters – $\mu_{log\sigma}$, σ_Δ and *d*. We then simulate the series of the logarithmic returns in the following way. First, we simulate ARFIMA (0,*d*,0) process[3] for the increments of the logarithmic volatility with a standard deviation of σ_Δ. Second, we integrate the series and adjust it so that the average value of the integrated series is $\mu_{log\sigma}$. Third, we take the exponential of the integrated series to get the series of the time-dependent standard deviation. In the last step, we use the standard deviation for uncorrelated normally distributed series with zero mean.

The average logarithmic volatilities $\mu_{log\sigma}$ are –3.9962, –4.4959 and –4.3439 for AAPL, IBM and MSFT, respectively. The standard deviations of the increments of the logarithmic volatility σ_Δ are 0.4817, 0.4388 and 0.4463 for AAPL, IBM and MSFT, respectively. The biggest issue is the estimation *d* because majority of the *d* and *H* estimators are constructed primarily for the persistent processes with $d > 0$, i.e. $H > 0.5$, and their finite sample performance for anti-persistent processes has not been seriously discussed in the literature yet. To overcome this issue, we present the simulations for $-0.9 \leq d \leq -0.1$ with a step of 0.1. The other two parameters are set to $\mu_{log\sigma} = -4.5$ and $\sigma_\Delta = 0.45$.

In figure 6, we present the simulated series of the standardized logarithmic returns for various values of *d*. We can see that for $d = -0.1$, the series is degenerate and shows some extreme values and does not resemble the empirical series of logarithmic returns. The lower the *d* parameter gets the less extreme the simulated values are and also the less extreme the volatility clustering is. The series which resemble the real world data the most are the ones with $d = -0.4, -0.5, -0.6$. For the lower values of *d*, the volatility clustering becomes very negligible. The best fit to the real-world series is found for $d = -0.5$ where we can observe very different phases of market behavior – calm periods are followed by highly volatile and vice versa. It is needed to note that the returns for $-0.9 \leq d \leq -0.4$ are uncorrelated (the ACF and PACF are not shown here).

[3] We choose ARFIMA(0,d,0) because we need an anti-persistent Gaussian process which allows for strong anti- persistence. ARFIMA is an obvious and logical choice.

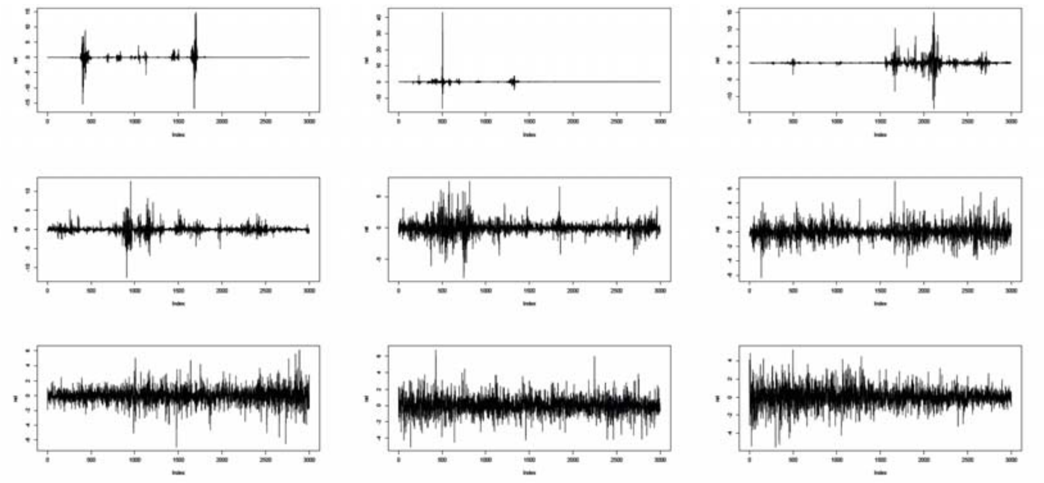

Figure 6. Simulated standardized returns. Starting from $d = -0.1$ (top left) and $d = -0.2$ (top middle) to $d = -0.9$ (down right).

Simulated time-dependent standard deviations are summarized in figure 7. Here again, we observe that for d parameter close to zero, the standard deviation process is degenerate with several extreme spikes. The time-dependent volatility is the most realistic for $d = -0.4$ and $d = -0.5$ and it is actually very hard to distinguish between these two and the standard deviations in figure 2 for the three analyzed stocks. For lower d, the standard deviation series become too volatile and do not produce phases with high volatility switching with phases of low volatility. Volatility persistence is further illustrated by autocorrelation functions of absolute returns in figure 8. Persistence of absolute returns is considered to be a stylized fact of financial markets. Even though the question of stationarity of absolute returns themselves remains, we do not deal with it and just check whether the stylized fact is observed for the simulated series as well. We observe that for the processes with d between –0.1 and –0.5, the autocorrelations decay slowly and remain positive for high lags. Again for $d = -0.5$, the autocorrelations remain slightly above 0.2 for all analyzed lags, which is most closely in hand with what is reported for the real-world time series [7].

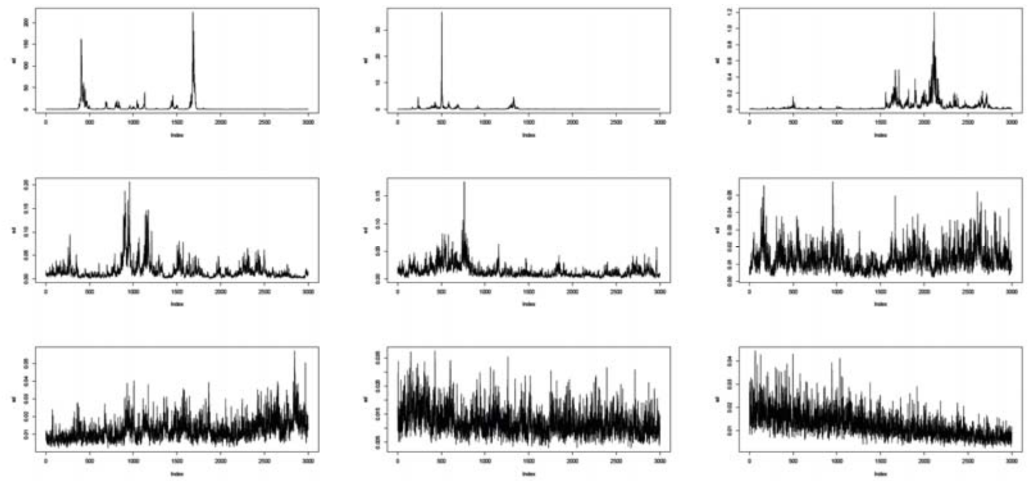

Figure 7. Simulated standard deviation processes. Starting from $d = -0.1$ (top left) and $d = -0.2$ (top middle) to $d = -0.9$ (down right).

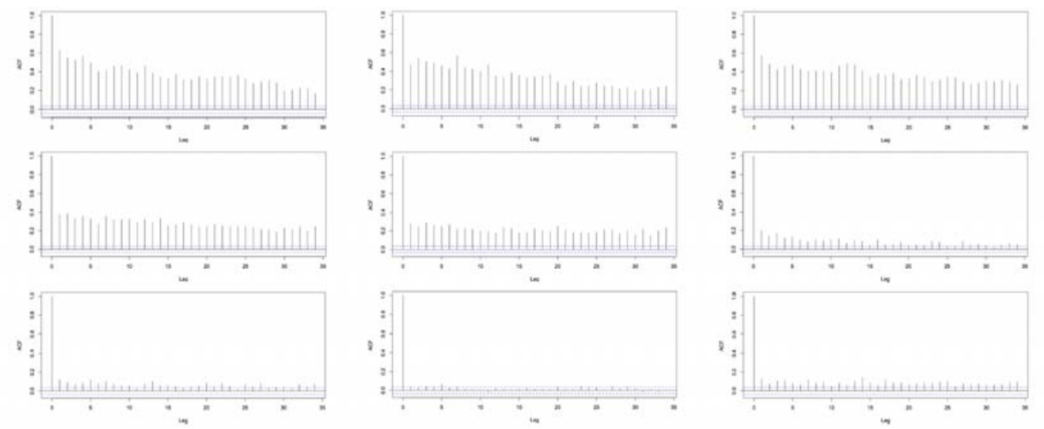

Figure 8. Autocorrelation function for absolute returns of the simulated processes. Starting from $d = -0.1$ (top left) and $d = -0.2$ (top middle) to $d = -0.9$ (down right).

5. Conclusions and discussion

We introduced the alternative paradigm to volatility modelling of financial securities. On the example of three stocks of highly capitalized companies, we showed that volatility process is non-stationary and its logarithmic transformation together with logarithmic increments is approximately normally distributed. Further, the increments have been shown to be highly anti-persistent. Together with the assertion that logarithmic returns are normally distributed, and uncorrelated with time-varying volatility, we proposed the new returns-generating process. Note that the whole procedure is based on empirical observations without any limiting assumptions. We are able to construct the returns series which remarkably mimic the real-world series and posses the standard stylized facts – uncorrelated returns with heavy tails, strongly auto correlated absolute returns and volatility clustering. Therefore, the proposed methodology opens a rather unexplored field in research of financial volatility. As this paper rather introduces the framework, there are many possibilities for further research in the field.

First of all, the crucial distinguishing between stationary and non-stationary series may raise questions. However, all standard stationarity tests have relatively low power and size against alternatives presented in this paper. Therefore, the correctness of this approach should be decided based on its ability to forecast volatility. This leads us to other future directions. Second, the proposed methodology should be compared to standardly used methods on basis of their forecasting accuracy. Third, the Value-at-Risk usefulness of the methodology is to be checked. Fourth, the correct specification of the logarithmic differences of the volatility process should be more thoroughly analyzed. ARFIMA(0,d,0) is only the initial proposition and there are more possibilities such as general ARFIMA(p,d,q) specifications or multifractal generalizations. Fifth, consistent and efficient estimators of d and H for highly anti-persistent processes should be discussed. And finally, the proposed methodology needs to be checked on other estimators of volatility (mainly various estimators of realized volatility).

Acknowledgments

The support from the Grant Agency of Charles University (GAUK) under project 118310, Grant Agency of the Czech Republic (GACR) under project 402/09/0965 and project SVV 265 504 are gratefully acknowledged.

REFERENCES

[1] Beran, J. (1994), *Statistics for Long-Memory Processes,* Volume **61** of Monographs on Statistics and Applied Probability. New York: Chapman and Hall.

[2] Black, F. and M. Scholes (1973), *The pricing of options and corporate liabilities*. Journal of Political Economy **81**(3), 637-654.

[3] Bondon, P. and W. Palma (2006), *A class of antipersistent processes*. Journal of Time Series Analysis **28**(2), 261-273.

[4] Box, G. and G. Jenkins (1970), *Time Series Analysis: Forecasting and Control. Holden Day*, San Francisco.

[5] Cajueiro, D. and B. Tabak (2005), *The rescaled variance statistic and the determination of the hurst exponent. Mathematics and Computers in Simualtion* **70**(3), 172-179.

[6] Chou, R.-Y., H.-C. Chou, and N. Liu (2010), *Handbook of Quantitative Finance and Risk Management, Chapter Range Volatility Models and Their Appliations in Finance,* pp. 1273-1281. Springer.
[7] Cont, R. (2001), *Empirical properties of asset returns: stylized facts and statistical issues*. Quantitative Finance **1**(2), 223-236.
[8] Davison, A. and D. Hinkley (1997), *Bootstrap Methds and Their Application*. Cambridge University Press.
[9] Debowski, L. (2011), *On processes with hyperbolically decaying autocorrelations*. Journal of Time Series Analysis **32,** 580-584.
[10] Dickey, D. and W. Fuller (1979), *Estimators of autoregressive time series with a unit root.* Journal of the American Statistical Association **74,** 427-431.
[11] Efron, B. (1979), *Bootstrap methods: Another look at the jacknife*. Annals of Statistics **7,** 1-26.
[12] Efron, B., R. Tibshirani, and R. Tibshirani (1993), *An introduction to the bootstrap*. Chapman & Hall.
[13] Embrechts, P. and M. Maejima (2002), *Selfsimilar Processes*. Princeton University Press. [14] Giraitis, L., P. Kokoszka, R. Leipus, and G. Teyssiere (2003). Rescaled variance and related tests for long memory in volatility and levels. Journal of Econometrics **112,** 265-294.
[15] Hosking, J. (1981), *Fractional differencing*. Biometrika **68**(1), 165-176.
[16] Hurst, H. (1951), *Long term storage capacity of reservoirs*. Transactions of the American Society of Engineers **116,** 770-808.
[17] Izzeldin, M. and A. Murphy (2000), *Bootstrapping the small sample critical values of the rescaled range statistic*. The Economic and Social Review **31**(4), 351-359.
[18] Jarque, C. and A. Bera (1981), *Efficient tests for normality, homoscedasticity and serial independence of regression residuals: Monte carlo evidence*. Economics Letters **7**(4), 313-318.
[19] Kwiatkowski, D., P. Phillips, P. Schmidt, and Y. Shin (1992), *Testing the null of stationarity against alternative of a unit root: How sure are we that the economic time series have a unit root?* Journal of Econometrics **54,** 159-178.
[20] Lee, D. and P. Schmidt (1996). *On the power of the kpss test of stationarity against fractionally-integrated alternatives*. Journal of Econometrics **73,** 285-302.
[21] Lee, H. and C. Amsler (1997), *Consistency of the kpss unit root against fractionally integrated alternative.* Economics Letters **55,** 151-160.
[22] Lillo, F. and J. Farmer (2004). *The long memory of the efficient market. Studies in Nonlinear Dynamics and Econometrics* **8**(3).
[23] Ljung, G. and G. Box (1978). *On a measure of a lack of fit in time series models*. Biometrika **65**(2), 297-303.
[24] Lo, A. (1991), *Long-term memory in stock market prices*. Econometrica **59**(5), 1279-1313.
[25] Mandelbrot, B. and J. van Ness (1968), *Fractional brownian motions, fractional noises and applications*. SIAM Review **10**(422), 422-437.
[26] Mandelbrot, B. and J. Wallis (1968), *Joah, joseph and operational hydrology*. Water Resources Research **4,** 909-918.
[27] Merton, R. (1973), *Theory of rational option pricing*. Bell Journal of Economics and Management Science **4**(1), 141-183.

[28] Morana, C. and A. Beltratti (2004), *Structural change and long-range dependence in volatility of exchange rates: either, neither or both?* Journal of Empirical Finance **11,** 629-658.
[29] Peters, E. (1994), *Fractal Market Analysis – Applying Chaos Theory to Investment and Analysis.* New York: John Wiley & Sons, Inc.
[30] Poon, S.-H. and C. Granger (2003), *Forecasting volatility in financial markets: A review*. Journal of Economic Literature **41,** 478-539.
[31] Poon, S.-H. and C. Granger (2005), *Practical issues in forecasting volatility.* Financial Analysts Journal **61**(1), 45-56.
[32] Srinivas, V. and K. Srinivasan (2000), *Post-blackening approach for modeling dependent annual streamflows*. Journal of Hydrology **230,** 86-126.
[33] Teverosky, V., M. Taqqu, and W. Willinger (1999), *A critical look at lo's modified r/s statistic.* Journal of Statistical Planning and Inference **80**(1-2), 211-227.
[34] Vandewalle, N., M. Ausloos, and P. Boveroux (1997), *Detrended fluctuation analysis of the foreign exchange market.* In Econophysic Workshop, Budapest, Hungary.
[35] Wang, W., P. van Gelder, J. Vrijling, and X. Chen (2006), *Detecting long-memory: Monte Carlo simulations and application to daily streamflow processes*. Hydrology and Earth System Sciences Discussions **3,** 851-862.