Non-stationary volatility with highly anti-persistent increments: An alternative paradigm in volatility modeling?

Ladislav Krištoufek [1]

Abstract. We introduce the alternative paradigm to volatility modeling. On the example of three stocks of highly capitalized companies, we show that volatility process is non-stationary and its logarithmic transformation together with the logarithmic increments are approximately normally distributed while the latter are strongly anti-persistent. Together with the assertion that logarithmic returns are normally distributed, and uncorrelated with time-varying volatility, we propose the new returns-generating process, which is able to remarkably mimic the real-world series and the standard stylized facts – uncorrelated returns with heavy tails, strongly autocorrelated absolute returns and volatility clustering. The proposed methodology opens a wholly new field in research of financial volatility.

Keywords: volatility, anti-persistence, non-stationarity

JEL classification: C53, C58, G17

1 Introduction

Accurate modeling and forecasting of volatility is one of the biggest challenges in financial economics and econometrics. Historically, there are four major groups of volatility forecasting approaches – historical volatility, conditional heteroskedasticity models, stochastic volatility models, and implied volatility models. These approaches are nicely reviewed and compared in two studies of Poon & Granger [13, 14]. Poon & Granger [14] 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.

In this paper, we propose a new approach to volatility modeling. Based on a simple statistical analysis, we show that volatility can be effectively modeled 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 [3]. Analyzing the real-world data, we are able to state seven basic Claims and argue that the logarithmic returns are uncorrelated and normally distributed with approximately zero mean and time-varying standard deviation. 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.

The paper is structured as follows. In Section 2, we present the basic methodology. In Section 3, we describe the dataset and present the crucial findings about the process of volatility. In Section 4, we show simulations for estimated parameters. Section 5 concludes.

[1] Institute of Information Theory and Automation, Academy of Sciences of the Czech Republic, Pod Vodárenskou věži 4, CZ-182 08, Prague, Czech Republic; Institute of Economic Studies, Faculty of Social Sciences, Charles University in Prague, Opletalova 26, CZ-110 00, Prague, Czech Republic, e-mail: kristouf@utia.cas.cz

2 Methodology

Long-range dependence is highly connected to the Hurst effect, i.e. a situation characteristic by 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. For $H > 0.5$, auto-covariances of the process are positive at all lags so that the process is labeled as persistent. On the other hand, for $H < 0.5$, the process is said to be long-range dependent with negative correlations or anti-persistent. The persistent process is visibly locally trending while the anti-persistent process switches signs more frequently than a random process would [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:

- A stationary process with an autocorrelation function $\rho(k)$ decaying as $\rho(k) \propto k^{2H-2}$ for $k \to \infty$ is called long-range dependent with Hurst exponent H.
- A stationary process with a spectral density $f(\lambda)$ following $f(\lambda) \propto |\lambda|^{1-2H}$ for $\lambda \to 0$ 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 is no standard long-range dependence. To check for (non-)stationarity, we apply two standard tests – ADF [5] and KPSS [11]. 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 also can control for memory effects with a use of auto-covariance adjusted variance with Barlett weights [8].

As we are interested primary 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. For more details, see [12, 8].

Lo [12] proposed the modified rescaled range analysis (M-R/S) as a generalization of the classical rescaled range analysis [10] to control for short-term correlations. With a use of adjusted standard deviation $S^M = \sqrt{S^2 + 2\sum_{j=1}^{\xi} \gamma_j \left(1 - \frac{j}{\xi+1}\right)}$, where S^2 is the standard deviation of the returns, γ_j is the jth auto covariance of the process up to lag ξ, the V statistic is defined as $V_\upsilon = \frac{(R/S)_\upsilon}{\sqrt{\upsilon}}$, where R/S is a range of the profile of the series (cumulative deviations from the mean of the original series) standardized by the standard deviation of the returns process. 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 [12, 10].

Rescaled variance analysis (V/S) was proposed [8] as a modified version of KPSS statistic, which is usually used for testing of stationarity but was also shown to have good power for series with long-term memory. 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 $M_\upsilon = \frac{var(X_{t,I_n})}{\upsilon (S^M_{I_n})^2}$. Note that the modified standard deviation S^M 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 reliable. Similarly to R/S and $M - R/S$, one needs to follow the steps of the R/S if estimating H or set $\upsilon = T$ and construct M statistic if only testing for the presence of long-range dependence in the process, which is actually the case in our study.

The bootstrap method has been proposed to deal with the statistical properties of small samples [6]. 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. 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 [15] proposed a modified method which retains the short-term dependence characteristics but lacks the long one – the moving block bootstrap with pre-whitening and post-blackening.

For our purposes, we use AR(1) process for pre-whitening and post-blackening 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. Let us now turn to the volatility estimate choice.

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 [2]. From many possibilities, we choose Garman–Klass estimator [7]:

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

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) [2]. 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.

3 Data and statistical analysis

We analyze series of three stocks with one 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). 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. 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.

As the Garman–Klass estimator does not take overnight dynamics into consideration, we analyze daily logarithmic returns defined as $r_t = \log(C_t) - \log(O_t)$ as well as the standardized returns $r_{st,t} = \frac{r_t}{\widehat{\sigma_{GK,t}}}$. We find 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 Jarque-Bera test – the raw returns are not normally distributed but the standardized returns are very close to being normally distributed. Based on Ljung-Box test, we find no significant autocorrelations in the first thirty lags for the standardized returns. Further, we analyzed the distributional properties of volatility process and with support of the QQ plots and Jarque-Bera tests, which are not shown here, we uncover that the process of volatility is far from being normally distributed whereas its logarithmic transformation and the increments of logarithmic volatility are approximately normally distributed. Based on these findings, we propose the first three Claims[1]:

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

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

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

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 1. We use lags 1, 10 and 100 to control for practically

[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.

	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***

Table 1: *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.

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. Secondly, the same is true for the logarithmic transformation of volatility. Moreover, the results are much stronger here as the series are very close to being normally distributed which is assumed for the tests. Interestingly, 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. Based on these findings, we propose three 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 long-term memory.*

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

Therefore, it is needed to analyze the increments of logarithmic volatility and its correlation structure. The autocorrelation and partial autocorrelation functions, not shown here, share 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 slowly to zero for PACF. This is indicative for two possible processes – a strong MA(1) process or a strongly anti-persistent ARFIMA(0,d,0) process [9, 4]. Debowski [4] 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 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 $\xi = 0, 1, 2, 5, 10, 15, 30, 50, 75, 100$ are summarized in Table 2. We observe that the results for both AR(1) and MA(1) controls are practically the same – for short to medium lags, we reject the "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 [1]. Based on these results, we propose the last Claim:

Claim 7 *Increments of logarithmic volatility are strongly anti-persistent with $d \in (-1, 0)$.*

4 Simulations

Based on all the Claims we made, we now try to reconstruct the series with the observed statistical properties and observe 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 logarithmic returns in the following way. First, we simulate ARFIMA(0,d,0) process[3] for increments of 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

[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)

[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.

the time-dependent standard deviation. In the last step, we use the standard deviation for uncorrelated normally distributed series with zero mean.

ξ	V_A	M_A	V_I	M_I	V_M	M_M	V_A	M_A	V_I	M_I	V_M	M_M
0	0.000	0.000	0.000	0.000	0.000	0.000	0.008	0.000	0.005	0.000	0.000	0.000
1	0.000	0.000	0.000	0.000	0.000	0.000	0.007	0.000	0.004	0.001	0.000	0.003
2	0.000	0.000	0.000	0.000	0.000	0.000	0.004	0.000	0.004	0.000	0.000	0.001
5	0.000	0.000	0.000	0.000	0.000	0.000	0.004	0.000	0.002	0.000	0.000	0.003
10	0.000	0.000	0.000	0.000	0.000	0.000	0.010	0.000	0.003	0.000	0.001	0.002
15	0.000	0.000	0.000	0.000	0.000	0.000	0.003	0.000	0.003	0.000	0.000	0.002
30	0.009	0.000	0.004	0.000	0.000	0.000	0.021	0.000	0.007	0.000	0.000	0.005
50	0.071	0.000	0.015	0.000	0.006	0.003	0.103	0.001	0.022	0.000	0.009	0.022
75	0.393	0.015	0.004	0.000	0.028	0.018	0.405	0.050	0.003	0.000	0.043	0.048
100	0.770	0.068	0.009	0.000	0.161	0.083	0.714	0.193	0.004	0.000	0.155	0.139

Table 2: *Anti-persistence tests – bootstrapped p-values for the null of "no anti-persistence" controlling for AR(1) [left side] and MA(1) [right side] process.* ξ stands for the number of lags taken into consideration for standard deviation S^M for V and M statistic. Notation "A", "I" and "M" stand for AAPL, IBM and MSFT, respectively

The average logarithmic volatility $\mu_{\log\sigma}$ is -3.9962, -4.4959 and -4.3439 for AAPL, IBM and MSFT, respectively. The standard deviation of the increments of the logarithmic volatility σ_Δ is 0.4817, 0.4388 and 0.4463 for AAPL, IBM and MSFT, respectively. The biggest issue is the estimation d because majority of d and H estimators are constructed primarily for 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 analyzed 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 $\sigma_\Delta = 0.45$.

Due to lack of space, we only show the simulations for the process which resembles the real-world financial series the best, i.e. $d = -0.5$, in Fig. 1. In the figure, we show the simulated standardized returns, simulated standard deviation process and autocorrelation function of the absolute returns. We observe that the simulations mimic the financial stylized facts remarkably – returns are uncorrelated (not shown here), volatility clustering is apparent, heavy tails are obvious and persistence of the absolute returns is visible as well (while the level of autocorrelations is close to 0.2 for all presented lags, which is exactly what is observed for empirical series [3]).

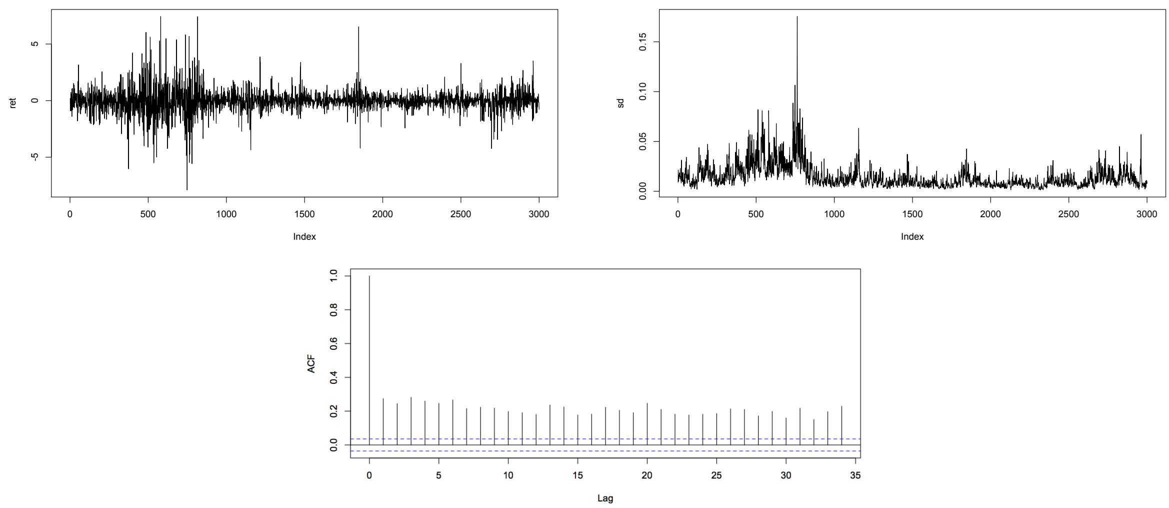

Figure 1: *Simulated process with $\mu_{\log\sigma} = -4.5$, $\sigma_\Delta = 0.45$ and $d = -0.5$. (left)* Standardized returns, *(right)* standard deviation process, and *(bottom)* autocorrelation function of absolute returns.

5 Conclusions

We introduced the alternative paradigm to volatility modeling 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 are 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 autocorrelated absolute returns and volatility clustering. Therefore, the proposed methodology opens a wholly new field in research of financial volatility. As this paper rather introduces the framework, there are many possibilities for further research in the field.

Acknowledgements

The support from GAUK 118310, GACR 402/09/0965, and SVV 265 504 are gratefully acknowledged.

References

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

[2] R.-Y. Chou, H.-C. Chou, and N. Liu. *Handbook of Quantitative Finance and Risk Management*, chapter Range Volatility Models and Their Appliations in Finance, pages 1273–1281. Springer, 2010.

[3] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[4] L. Debowski. On processes with hyperbolically decaying autocorrelations. *Journal of Time Series Analysis*, 32:580–584, 2011.

[5] D. Dickey and W. Fuller. Estimators of autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74:427–431, 1979.

[6] B. Efron. Bootstrap methods: Another look at the jacknife. *Annals of Statistics*, 7:1–26, 1979.

[7] M. Garman and M. Klass. On the estimation of security price volatilities from historical data. *Journal of Business*, 53:67–78, 1980.

[8] L. Giraitis, P. Kokoszka, R. Leipus, and G. Teyssiere. Rescaled variance and related tests for long memory in volatility and levels. *Journal of Econometrics*, 112:265–294, 2003.

[9] J. Hosking. Fractional differencing. *Biometrika*, 68(1):165–176, 1981.

[10] H.E. Hurst. Long term storage capacity of reservoirs. *Transactions of the American Society of Engineers*, 116:770–808, 1951.

[11] D. Kwiatkowski, P. Phillips, P. Schmidt, and Y. Shin. 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, 1992.

[12] A. Lo. Long-term memory in stock market prices. *Econometrica*, 59(5):1279–1313, 1991.

[13] S.-H. Poon and C. Granger. Forecasting volatility in financial markets: A review. *Journal of Economic Literature*, 41:478–539, 2003.

[14] S.-H. Poon and C. Granger. Practical issues in forecasting volatility. *Financial Analysts Journal*, 61(1):45–56, 2005.

[15] V. Srinivas and K. Srinivasan. Post-blackening approach for modeling dependent annual streamflows. *Journal of Hydrology*, 230:86–126, 2000.