AN ANALYSIS OF STOCK INDEX DISTRIBUTIONS OF SELECTED EMERGING MARKETS

*Silvio John Camilleri**

Abstract. Stock market data often display distinct characteristics commonly known as "stylised facts." These include non-stationarity of price levels as well as peak-shaped, fat-tailed and heteroskedastic log returns. This paper presents empirical evidence of these characteristics for emerging market indices spanning over different geographic regions. The results do not disclose asymmetry in the tails of log return distributions in any particular direction. In addition, the frequently observed characteristic that high volatility follows large negative returns does not show up in the data. Yet, when the results for the latter characteristic are grouped by geographic regions, some similarities become apparent.

Introduction

The analysis of the return distributions of financial assets is a vital topic in the finance discipline, not only on account of the academic research undertaken in this area, but also due to its relevance for practitioners when making portfolio choices and in risk management processes. When examining stock market price series, the data is typically non-stationary and deviates from the normal distribution. Indeed, stock market data often displays distinct characteristics which are commonly referred to as "stylised facts".

The aim of this paper is to glean empirical evidence of such characteristics, in respect of various emerging market stock indices. It is structured as follows. The basic principles of stock market data and relevant literature

* Silvio John Camilleri lectures at the Banking and Finance Department of the University of Malta. He is currently following a PhD programme at Loughborough University, UK. The author would like to thank the Stock Exchanges of Buenos Aires, Cairo and Alexandria, Bombay, Jamaica, Lithuania, Ljubljana, Malta, Mauritius and Taiwan for providing the data. The useful comments of Prof. Philip Hans Franses (Erasmus University, Rotterdam) are thankfully acknowledged. Any errors remain the sole responsibility of the author.

are reviewed in the next section. The third section includes a description of the methodologies which are typically used to model the stylised facts of stock market data. The analysis proceeds with a note on the data-set and its limitations. The last sections discusses the empirical results and presents the conclusions.

Statistical Principles and a Brief Literature Review

Most econometric tools assume stationary characteristics of the data set being analysed. A time series may be defined as stationary if its properties such as the mean and variance are unchanged over different sub-samples of the data set. Financial time series tend to deviate from stationarity and they often exhibit a time-changing mean and variance, as outlined by Mills (1999:37). Researchers thus avoid the application of econometric techniques to non-stationary data given that this could lead to flawed conclusions such as spurious regression results as shown by Granger and Newbold (1974). Non-stationary series can be transformed to stationary ones; for example researchers may difference the series or may analyse the logarithms of the observed time series. Working with logarithms presents distinct advantages including:
- using logarithms one may transform a non-linear relationship into a linear one;
- when using linear regressions on logarithmic series, the estimated coefficients have an immediate interpretation as elasticities; and
- when applying a log transformation to the data, the series is "compressed" often resulting in a constant variance for the transformed series.

Thus, one approach to modelling a time series of stock prices is to use the continuously compounded return (or log return r_t), which in the case of a non-dividend paying asset is equal to:

$$r_t = \log (P_t/P_{t-1}) = p_t - p_{t-1}$$

where P_t is the price level and $p_t = \log P_t$

The distribution of log returns often deviates from the normal distribution. Various authors such as Fama (1965) presented empirical evidence

34

which exposed the drawbacks of using a normal distribution to model logarithmic returns. Financial log return distributions tend to be peak-shaped (leptokurtic) and fat tailed. This characteristic of financial log returns was explained by patterns in the arrival of information, as well as patterns in traders' reactions to news, as discussed in Peters (1991). Dacorogna *et al*. (2001: 133), in an empirical investigation of USD exchange rate returns, showed that as the frequency of the data increases (say from weekly to hourly) the tails of the distribution become fatter.

Stock markets tend to be characterised by periods of substantial volatility interspersed with other periods of lower volatility. This implies a time-changing variance of returns as reviewed in Bollerslev, Chou, and Kroner (1992). Jacobsen and Dannenburg (2003) used stock market data from various developed countries and showed that this characteristic is not only present in high frequency data, but also in time series of lower frequencies such as monthly data.

Franses and van Dijk (2000:13-19) used stock index data to present empirical evidence of two further characteristics of log returns as follows:
• Large negative returns are more common than large positive returns. This feature was not confirmed by Longin (1996) in an empirical analysis of US stock market data, and by Jondeau and Rockinger (2003) who studied different stock market indices. The latter authors suggested that the common "perception" that left tails are thicker than the right ones might have been cultivated by the presence of data outliers.
• Franses and van Dijk (2000) also noted that high volatility often follows large negative returns. The authors also showed that the above two features are not as clearly evident in exchange rate data. Further empirical evidence of asymmetric volatility responses in relation to positive and negative returns is found in Koutmos (1999), who used stock price indices from G-7 countries. Yet, DeGennaro and Zhao (1998) found mixed evidence on the relationship between returns and volatility for US stock market data and concluded that this relationship is either "weak or variable".

The aim of this paper is to investigate the degree to which the above characteristics are evident in the index data of selected emerging stock markets.

Methodology

This study focuses on the degree to which the following properties are evident in selected indices:
• original price series are non-stationary, but transforming the series to logarithmic returns induces stationarity;
• logarithmic returns are not normally distributed–they are peak-shaped and fat tailed;
• logarithmic returns exhibit a time-changing variance;
• the left tail of the distribution is fatter than the right tail which implies that large negative returns are more common than large positive returns; and
• high volatility follows large negative returns.

Stationarity of the Original Series and Logarithmic Returns

One preliminary method through which stationarity of a data set may be inferred is to inquire whether the plot of the data discloses a changing mean and variance for different sub-samples of the series.

Dickey and Fuller (1979) proposed a procedure which may be used to infer whether a data set is stationary or otherwise and this is known as the augmented Dickey-Fuller (ADF) test. The differenced time series is expressed as a function of a constant, an optional trend, a lag of the levels, as well as *n* lags of the first difference. Thus:

$$\Delta x_t = f \text{ (constant, trend, } x_{t-1}, \Delta x_{t-1}, \dots , \Delta x_{t-n}).$$

A test for a unit root may by formulated by comparing the coefficient on x_{t-1} with its standard error. The null hypothesis is that the series contains a unit root and is therefore non-stationary; whilst the alternative hypothesis is that the series does not have a unit root. The critical values used in this hypothesis test are those of the augmented Dickey-Fuller statistic.

Distribution of Logarithmic Returns

A frequently used test of normality was proposed by Jarque and Bera (1980). This test jointly considers the skewness and kurtosis of the distribution as follows:

36

$$JB_y = n\left[\frac{\hat{SK}_y^2}{3!} + \frac{(\hat{K}_y - 3)^2}{4!}\right] \xrightarrow{\ d\ } \chi_2^2$$

where $\hat{SK}_y$ and $\hat{SK}_y$ are the skewness and kurtosis of the distribution respectively, whilst n is the sample size. The null hypothesis of a normal distribution is accepted or rejected by comparing the Jarque-Bera statistic to the value of the one-tailed χ^2 distribution statistic with two degrees of freedom.

Another way in which one may inquire whether a data set is peak-shaped or otherwise is to look at the location of the central percentile, say, the mid 20% observations starting from the end of the 0.4 percentile to the beginning of 0.6 percentile of the standardised returns. If the mid 20% observations of the data set lie in a narrower range of standardised values as compared to the normal distribution, then the distribution is likely to be peak-shaped.

Heteroskedasticity of Logarithmic Returns

The plot of the data may reveal traces of heteroskedasticity in terms of whether the variance tends to change over different sub-samples in the set.

Yet, a more formal Lagrange Multiplier (LM) test may also be applied. The data set is regressed on a constant, a lag and an error term as follows:

$$r_t = \rho_0 + \rho_1\, r_{t-1} + u_t$$

The LM statistic is then used to test whether there are autoregressive conditional heteroskedasticity (ARCH) effects in the error term u_t, as proposed by Engle (1982). The squared residuals u^2_t are regressed on q lags as follows:

$$u^2_t = \rho_0 + \rho_1 u^2_{t-1} + \rho_2 u^2_{t-2} + \ldots + \rho_q u^2_{t-q}$$

The null hypothesis of no ARCH effects, i.e. $\rho_1 = \rho_2 = \ldots = \rho_q = 0$, is tested against the alternative hypothesis that $\rho_1 \neq 0,\ \rho_2 \neq 0,\ \ldots \rho_q \neq 0$.

Symmetry of the Tails of the Logarithmic Distribution

The empirical investigation in the next section also inquires whether there is any general trend for a fatter right or left tail as compared to the other. This is done by comparing the location of the extreme percentiles of the distributions. If these percentiles, say the left 1% and the right 1% of the data lie within approximately the same area, the tails are likely to be symmetric. Yet if one of the percentiles is "squeezed" into a narrower range of standardized values as compared to the other one, then the former tail is likely to be fatter.

High Volatility Follows Large Negative Returns

The methodology used by Franses and van Dijk (2000) for inquiring whether high volatility follows large negative returns, was to estimate the correlation between the squared return at day *t* and the return at day *t-1*. A negative correlation coefficient indicates that the larger returns were preceded by a negative return.[1]

Data and Limitations

In the next section, the tests outlined above are used to inquire whether the former statistical properties are present in stock index data from different emerging markets. The data set shows daily closing values of nine emerging markets indices: BOLSA (Argentina), CASE 30 (Egypt), BSE 500 (India), JSE Index (Jamaica), LITIN (Lithuania), SBI 20 (Slovenia), MSE Index (Malta), SEMDEX (Mauritius), and TSEC 50 (Total Return) (Taiwan).

The data was obtained from the respective exchanges i.e. Buenos Aires Stock Exchange, Cairo and Alexandria Stock Exchange, Bombay Stock

---

[1] An alternative methodology to infer the level of asymmetry in volatility was proposed by Engle and Ng (1993). The squared error term of the first order autoregressive process [AR(1)] for log returns series is regressed over a constant and a dummy variable of the lagged sign of the AR(1) error term. The dummy variable takes a value of 1 when the lagged error term of the AR(1) process is negative, whilst it takes a value of zero otherwise. A significantly positive coefficient for the dummy variable is an indicator that high volatility follows negative returns. The results obtained from this methodology were broadly in line with those following the method proposed by Franses and van Dijk (2000).

Exchange, Jamaica Stock Exchange, National Stock Exchange of Lithuania, Ljubljana Stock Exchange, Malta Stock Exchange, Mauritius Stock Exchange, and Taiwan Stock Exchange Corporation. The indices were compiled by the exchanges, with the exception of TSEC 50 (Total Return) Index which was compiled by FTSE International Limited.

The particular indices were selected in order to achieve a comprehensive cross section of emerging markets within different geographic regions: Africa, Asia, Europe and Latin America. The selection of the actual time span of the data was done in order to minimize non-trading periods and/ or missing observations which exceeded five calendar days. In those cases where missing observations or non-trading periods exceeding five calendar days remained in the sample, the index did not show any major fluctuation during the particular period. Preliminary plots of the time series did not reveal any outlier observations.

The empirical study shown below is subject to the limitations inherent in analysing security price data:

• Stock prices are discrete prices; for example price changes have to be in one-sixteenth of a dollar, or multiples thereof. Possible effects of price discreteness include price clustering. Such effects might still be present to some degree in price series where trading is decimalised, given that in such cases prices still have to be quoted in cents and therefore they are still not continuous. According to Campbell, Lo and MacKinlay (1997: 110-112), the impacts of price discreteness become more evident as the sampling period shortens.

• When analysing stock market data which spans over long periods of time, one should be aware that the conditions which underlie the pricing process are likely to change. For example, a long sample period is likely to include changes in the composition of stock indices and changes in company structure due to merger and takeover activity. At times, changes in the trading procedures and changes in the trading hours might also be present. Dacorogna et. al. (2001: 5) referred to these effects as the "breakdown of the permanence hypothesis" and the authors also questioned whether researchers can actually claim that they are analysing the same market when working with a long time-series.

Empirical Results

We now turn to the empirical results with respect to the statistical properties described above.

Stationarity of the Original and Logarithmic Series

The plots of the original price series give a preliminary indication that the series are non-stationary due to a time changing mean and / or variance. The time series plots for CASE 30 and JSE index are shown in Figures 1 and 2 as examples. The other index plots are not being reproduced for the sake of brevity.

Figure 1
CASE 30 Index (Levels Data)
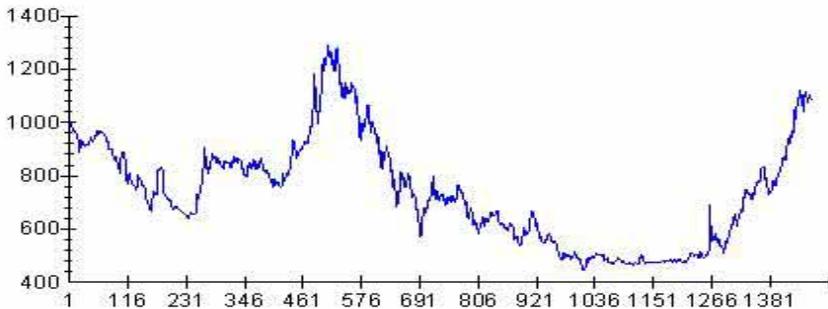

Figure 2
**JSE Index (Levels Data)**

40

The slowly declining autocorrelation coefficients (not reproduced here) constituted another indication of non-stationarity. The autocorrelation coefficients remain significant (at the 5% level) till around lag 30, indicating that it is not advisable to analyse the original price levels due to non-stationarity.

As noted above, log returns are more suitable for the application of econometric techniques given that they tend to be closer to stationarity. The plots of the log returns of BOLSA and TSEC 50 are shown in Figure 3 and 4. The other plots are not being shown again for the sake of conciseness.

Figure 3
BOLSA Log Returns
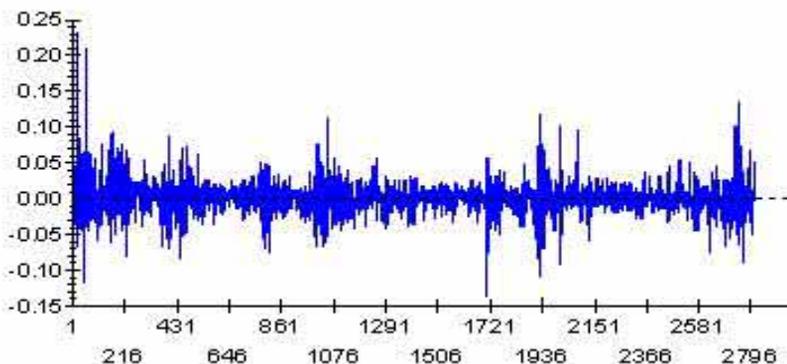
Figure 4
TSEC 50 Log Returns

41
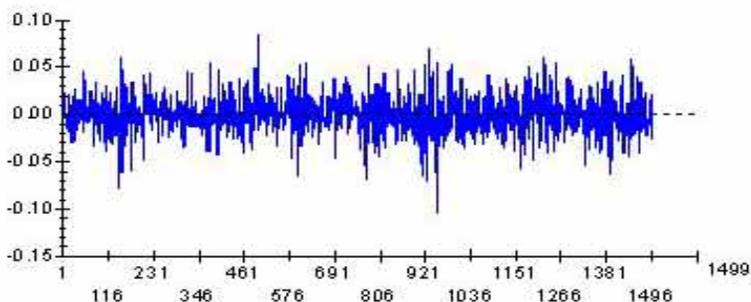
The plots visually demonstrate that log returns have a mean of approximately zero, or perhaps slightly positive. In addition, most of the plots such as BOLSA disclose a time changing variance, where periods of a relatively low variance alternated with others of higher variance. The time series plot which was visually closest to a constant variance is TSEC 50, yet even in this case a time-changing variance is plausible.

The autocorrelation test confirms that log returns are better candidates for analysis purposes than the original series since taking the log returns, reduced the serial correlation. Yet, given that some of the autocorrelation coefficients were still significant, Augmented Dickey-Fuller tests were used to test for stationarity of the log returns.

In applying this test to the data, specifications without a trend were selected given that the plots of the log returns suggest that it is unlikely that these series include a trend. The ADF test results are shown in Table 1. The values of the test statistic as compared to the 95% critical value of the ADF statistic permit rejection of the null hypothesis of a unit root for all the nine indices.

This indicates that the log returns series are stationary. Overall, the above tests indicate that it is reasonable to analyse the log returns series.

Table 1
Augmented Dickey-Fuller Test Results

| Index (Log Returns)	Order Selection (AICC)	Test Statistic	95% Critical Value
BOLSA	2	-28.345	-2.863
CASE 30	1	-25.603	-2.864
BSE 500	0	-31.012	-2.864
JSE Index	2	-25.164	-2.863
LITIN	2	-11.237	-2.868
SBI 20	5	-18.063	-2.863
MSE Index	0	-22.281	-2.865
SEMDEX	3	-16.568	-2.864
TSEC 50	4	-17.065	-2.864

Distribution of Logarithmic Returns

The histograms of the log returns of the indices indicate that the data is peak-shaped and perhaps fat-tailed. The two histograms which most prominently displayed these characteristics are shown in Figures 5 and 6. Normal distributions are superimposed on the histograms for ease of comparison.

Figure 5
Histogram of MSE Index Log Returns and
Normal Distribution Curve
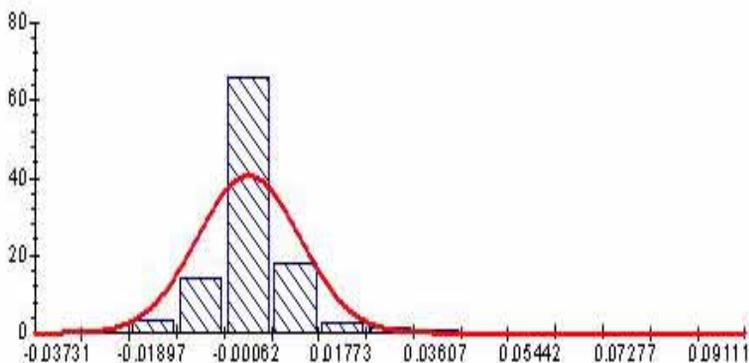
Figure 6
Histogram of TSEC 50 Index Log Returns and
Normal Distribution Curve

43
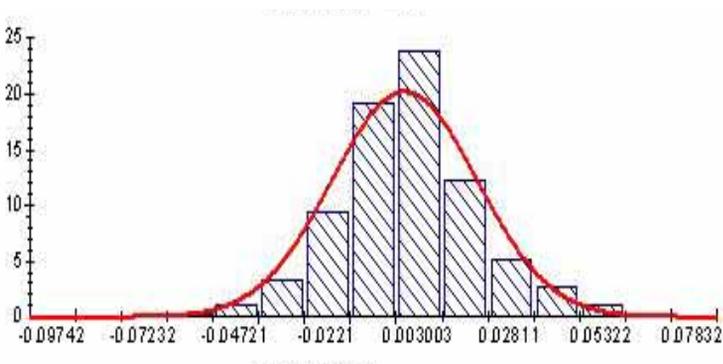
The Jarque-Bera statistic enables the rejection of the null hypothesis of a normal distribution at the 99% level of confidence for all indices. The excess kurtosis values show that in all cases, log return distributions are peak-shaped. An alternative way in which one may infer whether a distribution is peak-shaped is by looking at the location of the standardised values of the mid-observations. The mid-20% observations in a normal distribution lie between the standardised values of +/- 0.251. In comparison, the mid-20% (standardised) observations for all indices lie within a narrower range of standardised values, as shown in Table 2. This is an alternative indication that these distributions are peak-shaped.

In inquiring whether the distributions of the index log returns are fat-tailed, the location of the "extreme percentiles" may be compared to that of the normal distribution. The 0.01 and 0.99 percentiles in a normal distribution occur at -/+ 2.326. With the exception of the CASE 30 index, the 0.01 percentile occurs "earlier" than expected, indicating fat left tails for eight of the indices being analysed. The location of the 0.99 percentile for the distributions of all indices indicates that the right tail is fatter than the normal one, given that we enter this percentile "later" than expected. This confirms that the tails of the log returns are fatter than normal. Yet, the distributions become narrower than the normal distribution as we move further towards the centre, say when the location of the 0.1 and 0.9 percentiles is considered. The distributions then become "fat" again in the centre, given that they are peak-shaped as discussed above. Seven of the indices are positively skewed, whilst BSE 500 and LITIN have a negative skewness.

Heteroskedasticity of Logarithmic Returns

The plots of logarithmic returns for the indices show that it is plausible that the series feature a time-changing variance. The Lagrange Multiplier statistics for the indices are shown in Table 2, both for an order 1 test and an order 12 test. These statistics are compared to the 95% critical value of the χ^2 distribution at the respective degrees of freedom. The values of the LM statistics (for all indices and for tests of both orders) are high enough to permit rejection of the null hypothesis of no ARCH effects. This is a sign of heteroskedastic time series where clusters of large returns are interspersed with clusters of smaller returns.

44

Table 2

Basic Characteristics of Index Log Returns

Country Index	Argentina BOLSA	Egypt CASE 30	India BSE 500	Jamaica JSE	Lithuania LITIN	Slovenia SBI 20	Malta MSE	Mauritius SEMDEX	Taiwan TSEC 50
Initial Observation	2 Jan91	1 Jan98	1 Feb99	2 Jan91	2 Jan02	7 Jan93	18 May98	25 Nov97	1 Apr97
No. of Observations	2807	1460	1207	2712	485	2686	1166	1501	1498
Mean	0.0008	0.0001	0.0005	0.0012	0.0007	0.0006	0.0005	0.0002	-0.0002
Std Deviation	0.0231	0.0168	0.0171	0.0137	0.0119	0.0141	0.0098	0.0046	0.0197
Skewness	0.9669	0.8811	-0.4209	0.9422	-0.1715	0.4201	2.6681	0.7222	0.0697
Excess Kurtosis	9.8529	13.4197	2.0853	12.9245	5.1688	21.3521	20.2589	10.6578	1.4655
Minimum	-0.1367	-0.1098	-0.0735	-0.1245	-0.0738	-0.1161	-0.0419	-0.0282	-0.1037
Maximum	0.2321	0.1837	0.0693	0.1085	0.0585	0.1893	0.0957	0.0382	0.0846
Jarque-Bera Test	11,791	11,144	254	19,277	542	51,103	21,323	7,234	135
Coeff. of Variation	28.4726	302.9185	32.2536	11.4520	16.0366	22.1661	19.9678	20.3734	95.3220
0.01Percentile (of standardised returns)	-2.5069	-2.2938	-2.8627	-2.5989	-2.5461	-2.6888	-2.3255	-2.7730	-2.4954
0.1 Percentile	-1.0356	-1.0433	-1.1746	-0.8620	-1.0046	-0.8872	-0.8547	-0.9787	-1.1508
0.4 Percentile	-0.1780	-0.2101	-0.1246	-0.1750	-0.1879	-0.1371	-0.1519	-0.1788	-0.2136
0.6 Percentile	0.1381	0.1215	0.2441	0.0296	0.1167	0.1176	0.0416	0.1022	0.1180
0.9 Percentile	0.9740	1.1324	1.1220	0.9643	1.0741	0.9067	0.7699	0.9670	1.2478
0.99 Percentile	3.0915	2.4371	2.6199	3.4539	2.5713	2.9567	3.5933	3.0423	2.6342
LM(1)	212.8591	291.2590	126.0123	161.9786	11.9594	314.7868	71.8324	119.5308	25.8294
LM(12)	331.6720	315.7366	179.2452	233.0858	26.2908	341.7391	87.7094	156.0175	82.9735
Correlation (r^2_t, r_{t-1})	0.1174	0.2003	-0.2101	0.1657	-0.0049	-0.1216	0.2313	0.1076	-0.0434

45

Symmetry of the Tails of the Logarithmic Distribution

The histograms of the indices do not visually indicate that the left tail is fatter than the right one, as suggested by Franses and van Dijk (2000). Indeed, comparing the location of the 0.01 percentile with that of the 0.99 percentile (which should be equidistant from zero in a symmetric distribution) indicates a fatter right tail for all indices except BSE 500 (Table 2). When comparing the location of the 0.1 and the 0.9 percentiles, the evidence in favour of fatter right tails declines, given that BOLSA, BSE 500, MSE and SEMDEX indicate a fatter left tail. Therefore, the empirical results for these indices are in line with the suggestions of Longin (1996) and Jondeau and Rockinger (2003), that asymmetry in the tails in any one direction is not a general characteristic of stock market returns.

High Volatility Follows Large Negative Returns

Table 2 reports the correlation coefficients for the squared return r^2_t with the lagged return r_{t-1}. The correlation is negative only in case of four of the nine indices being analysed, and overall this does not confirm the observations of Koutmos (1999) and Franses and van Dijk (2000) that high volatility follows large negative returns.[2] This is more in line with the mixed evidence presented by DeGennaro and Zhao (1998).

The result that high volatility does not seem to follow large negative returns in emerging markets may be due to the absence of any relationship between these variables in the first place, or because high volatility tends to be a more common feature in emerging stock markets, and therefore it tends to follow both negative and positive large returns. The latter hypothesis may be explained by the notion that stock market volatility tends to be interconnected with macroeconomic volatility.[3]

Another particular feature of the results obtained with respect to this characteristic is that if the countries are grouped by geographic regions,

---

[2] The correlation r^2_t, r_{t-3} was also estimated and the signs of the coefficients were unchanged. When the correlation r^2_t, r_{t-5} was considered, the signs of the coefficients were confirmed again, with the exception of SBI 20 and SEMDEX.

[3] For instance Morelli (2002) used UK data to present empirical evidence of the interrelationship between the conditional volatility of the stock market and that of various macroeconomic variables.

some patterns emerge. For instance Argentina and Jamaica show a positive correlation between volatility and lagged returns. The same applies for African countries (Egypt and Mauritius). Asian countries (India and Taiwan) reveal a negative correlation between volatility and lagged returns. The results are somewhat mixed in case of European countries where this relationship is positive in the case of Malta, negative for Slovenia and (slightly) negative for Lithuania. Factors which contribute to these differences may include the shorter sample period for Lithuania, and the relatively low market activity on the Malta Stock Exchange as compared to the other exchanges. Overall, sub-dividing the countries in geographic regions reveals certain trends, which may be taken as an indication of the interdependence of "proximate" markets.

Conclusion

This paper presented empirical evidence of the "stylized facts" of stock market data from selected emerging market indices. Following a brief exposition of the characteristics and the relevant literature relating to stock market time series, the methodology, data set and limitations were subsequently discussed. The empirical results confirmed that stock price levels are often non-stationary and that it is more reasonable to transform the data into logarithms. It was also confirmed that the latter deviate from normality, and they tend to be peak-shaped, fat-tailed and heteroskedastic. The empirical results did not confirm the observations of other authors regarding the asymmetry of the distribution tails in any one direction and that high volatility tends to follow large negative returns. Yet some patterns for the latter characteristic were found over different geographic regions.

In interpreting the above results, one should keep in mind that they might be sensitive to differing sampling intervals, as found for instance by Balaban, Ouenniche and Politou (2005). The use of higher frequency data rather than daily price series might result in even more pronounced deviations from normality.

The deviations of stock return distributions from normality are relevant for portfolio selection and risk management decisions. In modelling the price risk of financial assets, particular attention should be devoted to

the tails of the distributions since these constitute the largest price fluctuations and are thus highly relevant for the risk management function. The modelling of these extreme fluctuations may require more focussed econometric models, and this issue provides an interesting avenue for further research.

References

BALABAN, E., OUENNICHE, J., and POLITOU, D. (2005) "A Note on Return Distribution of UK Stock Indices," *Applied Economics Letters*, Vol.12: 573-576.

BOLLERSLEV, T., CHOU, R.Y. and KRONER, K.F. (1992) "ARCH Modelling in Finance: A Review of the Theory and Empirical Evidence," *Journal of Econometrics*, Vol. 52: 5-59.

CAMPBELL, J.Y., LO, A.W. and MacKINLAY, A.C. (1997) *The Econometrics of Financial Markets*, Princeton, NJ: Princeton University Press.

DACOROGNA, M., GENÇAY, R., MÜLLER, U., OLSEN, R., and PICTET, O. (2001) *An Introduction to High-Frequency Finance,* San Diego, CA: Academic Press.

DeGENNARO, R.P. and ZHAO, Y.L. (1998) "Stock Returns and Volatility: Another Look," *Journal of Economics and Finance*, Vol. 22(1): 5-18.

DICKEY, D.A. and FULLER, W.A. (1979) "Distribution of the Estimators for Autoregressive Time Series with a Unit Root," *Journal of the American Statistical Association*, Vol.74: 427-431.

ENGLE, R.F. (1982) "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of U.K. Inflation," *Econometrica*, Vol. 50: 987-1008.

ENGLE, R.F. and Ng, V.K. (1993) "Measuring and Testing the Impact of News on Volatility," *The Journal of Finance*, Vol. 48(5):1749-1778.

FAMA, E.F. (1965) "The Behaviour of Stock Market Prices," *Journal of Business*, Vol. 38: 34-105.

FRANSES, P.H. and VAN DIJK, D. (2000) *Non-Linear Time Series Models in Empirical Finance*, Cambridge and New York: Cambridge University Press.

GRANGER, C.W.J. and NEWBOLD, P. (1974) "Spurious Regressions in Econometrics," *Journal of Econometrics,* Vol.2:111-120.

JACOBSEN, B. and DANNENBURG, D. (2003) "Volatility Clustering in Monthly Stock Returns," *Journal of Empirical Finance*, 10:479-503.

JONDEAU, E. and ROCKINGER, M. (2003) "Testing for Differences in the Tails of Stock-Market Returns," *Journal of Empirical Finance*, Vol.10: 559-581.

JARQUE, C.M. and BERA, A.K. (1980) "Efficient Tests for Normality, Homoscedasticity, and Serial Independence of Regression Residuals," *Economics Letters,* Vol.6: 255-259.

KOUTMOS, G. (1999) "Asymmetric Index Stock Returns: Evidence from the G-7," *Applied Economics Letters,* Vol.6: 817-820.

LONGIN, F.M., (1996) "The Asymptotic Distribution Of Extreme Stock Market Returns," *Journal of Business*, 69: 383-408.

MILLS, T.C. (1990) *Time Series Techniques for Economists*. Cambridge and New York: Cambridge University Press.

MORELLI, D. (2002) "The Relationship Between Conditional Stock Market Volatility and Conditional Macroeconomic Volatility: Empirical Evidence Based on UK Data," *International Review of Financial Analysis,* Vol.11: 101-110.

PETERS, E. (1991) *Chaos and Order in the Capital Markets. A New View of Cycles, Prices, and Market Volatility*. New York: John *Wiley* and Sons.

49