Wavelet Exploratory Analysis of the FTSE ALL SHARE Index

Antonios Antoniou [a] Constantinos E. Vorlow [a,*]

[a] *Durham Business School, University of Durham, Mill Hill Lane, Durham, DH13LB, UK.*

Abstract

Wavelets by construction are able to show us "the forest as well as the trees". They are compactly supported functions that allow us to localise in frequency as well as in time whereas traditional Fourier analysis focuses only on frequency. This makes wavelets useful when examining time sequences that exhibit sharp spikes and irregularities, such financial time series. In this paper we demonstrate how a wavelet semi-parametric approach can provide useful insight on the structure and behaviour of stock index prices, returns and volatility. By using wavelets we capture salient features such as changes in trend and volatility and reveal dynamic patterns at various scales.

Key words: Wavelet transforms, scalograms, multiresolution analysis, financial time series analysis.
JEL classification: G14; C29; Z00.

* Corresponding author
Email addresses: `Antonios.Antoniou@durham.ac.uk` (Antonios Antoniou), `K.E.Vorloou@durham.ac.uk` (Constantinos E. Vorlow).
URL: `http://www.vorlow.org` (Constantinos E. Vorlow).

Preprint submitted to WSEAS NOLASC 2003 *20 November 2003*

1 Introduction

Wavelets are becoming more popular as the academic community appreciates their ability to detect localised events as well as periodic structures. Originally introduced through financial and economic applications in the beginning of the last decade, they have not yet enjoyed their deserved popularity. This is in contrast to their potentially wide applicability and great versatility. Their ability to isolate breaks and shifts in dynamics, to manipulate intermittent sequences and their usefulness in denoising and smoothing makes them an important tool for univariate time series analysis, comparable to that of ARIMA and spectral analysis. The difference is that wavelets can provide the exact locality of any changes in the dynamical patterns of the sequence whereas the other two techniques concentrate mainly on their frequency. [1] Moreover, Fourier transforms assume infinite-length signals, whereas wavelet transforms can be applied to any kind and size of time series, even when these sequences are not homogeneously sampled in time. In general, wavelet transforms can be used to explore, denoise and smoothen financial time series, aid in forecasting, contribute to other empirical analysis frameworks (efficiency tests, event studies etc. etc.) and to calibrate or improve trading models. In this paper we aim to reintroduce wavelet transforms with a more “digestible” demonstration. We show how wavelet transforms can be applied and their results interpreted through a very simple exploratory framework. This approach can be of great interest to practitioners as it is avoids the pitfalls of a parametric model and statistical hypothesis testing approach, while being fast and accurate. It also provides results in various scales and frequencies, delivering a decomposition of great detail and information.

In the following sections we provide a brief outline of existing research in the area of wavelet applications in finance and economics. Following this review, we present in section 2.1 a simple and brief introduction to the internals of wavelet transforms and describe our database. We then explain in section 3 how results of “time-scale” analysis can be interpreted and finally in section 4 how wavelet transforms could be used in a discrete framework to provide rudimentary data smoothing and noise filtering.

2 Past research

Strang (1989) and Graps (1995) provide an interesting introduction to the general subject (see also Ramsey, 2002). A more recent introduction, focused on economics, is Schleicher (2002). A very insightful and early paper is also Jensen

[1] This excludes the case of Short Time Fourier Transforms.

(1997) who discusses the general potential of wavelets in financial empirical research. One of the first applications were by Greenblatt (1996) who used wavelets for outlier detection and Ramsey and Zaslavsky (1995) who influence this paper. Jensen (1994) uses wavelets to estimate fractionally integrated processes. He shows an alternative way to estimate the fractional differencing parameter and shows the advantages of wavelets over the existing method of Geweke and Porter-Hudak (1983). Continuing, Jensen (1999a) and (1999b) explores more deeply the applicability of wavelet transforms in long memory models. Olmeda and Fernandez (2000) provide a criticism of the theory by drawing our attention to the pitfalls of using wavelet filtering for denoising and forecasting purposes. Capobianco (2001), uses wavelets for describing financial returns processes. He studies the Nikkei stock index in high frequency and shows results about modelling with GARCH when the data have been preprocessed by wavelets transforms. He demonstrates that one can obtain better volatility prediction power for one-step-ahead forecasts, implying that latent volatility features can be detected more efficiently. Capobianco (2002) uses multiresolution analysis to approximate volatility processes. He focuses on intra-day dynamics and again shows how wavelet transforms can improve our view of volatility dynamics provided by a GARCH specification on wavelet pre-processed sequences (see also Capobianco, 1999). Antoniou and Volrow (2003), in a similar approach to that of Chen (1996), use wavelets to denoise index returns from various countries and obtain evidence of deterministic nonlinearities. Los and Karuppiah (2000) apply wavelet multiresolution analysis on high-frequency Asian FX rates, in support to the *Fractal Market Hypothesis* of Mandelbrot (1966) and Peters (1994). Recently, Jamdee and Los (2003) examine the risk profile of the US term structure by the use of wavelet scalograms and multiresolution analysis, suggesting that the term structure of interest rate is segmented and that the basic assumptions of the “traditional” models are violated. In a recent extensive work, Los (2003) focuses in financial markets even more. In this extensive work wavelets and multiresolution analysis are used to measure persistence and to provide explanations on financial crises, turbulence and volatility.

In this paper we use wavelets in a continuous and discrete multiresolution framework to show their usefulness in the empirical investigation of asset price dynamics. Our work follows the research of Ramsey and Zhifeng (1994), Ramsey and Zhifeng (1995) and Ramsey and Lampart (1998). The necessary theoretical background for our approach is discussed in depth in Gencay et al. (2001) and Percival and Walden (2000). A very gentle introduction is Hubbard (1998) and Ogden (1997) is more focused on the statistical aspect of the theory. A thorough treatment of wavelet theory is Strang and Nguyen (1996) and interesting on-line starting point and resource is the Wavelet Digest: `http://www.wavelet.org`.

2.1 *Data description & methodology*

The data set consists of 8192 daily observations of the FTSE ALL SHARE index closing prices. Furthermore, we investigate the continuously compounded returns of the FTSE and the realised volatility of these returns. All daily sequences start form the 10th of July 1970 and end the 30th of November 2001. In table 1 we display the descriptive statistics for the FTSE index, its logarithmic returns and the realised volatility. From this table and the inspection of the relevant distribution histograms [2] we can deduce that the distribution of the closing prices of the FTSE index is positively skewed whereas the corresponding returns are leptokurtic and the realised volatility positively skewed as well. The Jarque-Bera test (see Bera and Jarque, 1981) for normality clearly rejects the null for all sequences, as expected. In general, the sequences analysed here follow closely the stylised facts as these are described in Cont (2001).

[insert table 1 about here]

Our aim in this section is not to provide a rigourous treatment of the mathematical background of wavelet transforms but to communicate in brief of the way wavelet transforms operate. Wavelets are compactly supported functions (usually orthogonal), i.e. defined over a compact subset of $\mathbb{R}$ (see figure 1). The positioning of wavelets on $\mathbb{R}$ can vary and using a scaling function, we can "stretch" or compress wavelets, i.e. change their length. This variability in their shape, allows us to use them for pin-pointing singularities in any sequence. Because of this property, wavelets can localise in time as well as in frequency, whereas the traditional Fourier transforms concentrate solely on the latter. The way a wavelet transform functions is similar to that of the Fourier transform. We have the continuous wavelet transforms (CWT) and their discrete counterparts (DWT). Wavelet analysis of time series originally required the sequences to be of length 2^n, where n is discrete. A modification though (see Percival and Walden, 2000) allows us to work with signals of any length. The transform involves convolutions of the stretched or compressed wavelets with the signals at various frequency levels. The obtained result is a set of wavelet coefficients. Their value is high at the point where the wavelet function approximates locally the signal's structure. In that sense, high wavelet coefficient values imply that the function mimics the local structure well or that the correlation of the wavelet with the signal at that point is high. Time-scale analysis through wavelet transforms can provide us with the "scalograms". These indicate in time (space) and frequency (scale) the way the dynamics of the sequence involve.

[2] Results available from the authors upon request.

Representing a time-series as a function $f(t)$, the CWT of this function is defined as:

$$W_\psi(a,b) = \frac{1}{\sqrt{a}} \int_{-\infty}^{\infty} f(t)\psi(\frac{t-b}{a})\mathrm{d}t \tag{1}$$

and the inverse transform is defined as:

$$f(t) = \frac{1}{C_\psi} \int_{-\infty}^{\infty} \int_{-\infty}^{\infty} \frac{1}{a^2}[W_\psi(a,b)]\frac{1}{\sqrt{a}}\psi(\frac{t-b}{a})\mathrm{d}a\mathrm{d}b, \tag{2}$$

where $\psi(t)$ is the "mother" wavelet and $a, b \in \mathbb{R}$ continuous variables with $a \neq 0$. The parameter a is called the "scale" or "dilation" parameter which determines the level of stretch (expansion) or compression of the wavelet. Parameter "b" is called the "shift" or "translation" parameter. For low scales i.e. when $|a| \ll 1$, the wavelet function is highly concentrated (shrunken-compressed) with frequency contents mostly in the higher frequency bands. Inversely when $|a| \gg 1$, the wavelet is stretched and contains mostly low frequencies (for example: a time trend). For small scales we obtain thus a more detailed view of the signal (known also as a "higher resolution") whereas for larger scales we obtain a more general view of the signal structure.

Continuous transforms provide us with a large amounts of redundant information. For this reason, there is a discrete version (DWT), generated by critically subsampling the coefficients of the CWT. This is conducted in such a way that the energy of the signal is preserved. DWT is faster and useful for denoising and providing smoothed versions of the sequences at various discrete levels of analysis. In the following sections we will demonstrate how one can obtain interesting qualitative information on the structure of financial time series by using both continuous and discrete wavelet transforms.

3 Scalograms and time-scale analysis

For the time-scale analysis that follows, we utilise two wavelet functions: the Haar and the symmlet 8 (or s8 for short). These two "mother" wavelet functions are depicted in the two top subfigures of figure 1, together with their corresponding "father" or scaling functions. These wavelets were chosen after careful experimentation. There is no general rule on how to select mother wavelets. Usually a mother wavelet should exhibit a similar structure to that of the analysed signal. Often this choice is not such a crucial issue. Any relatively small (short) wavelet should be a suitable starting point (the Haar usually excluded). Given that a wavelet should resemble the time series, the discontinuous and blocky Haar is not really appropriate for our data. It is though a very simple and economical function, with fast transforms. Moreover, as we shall demonstrate, the qualitative information we obtain from the time-scale analysis is not that different from that of the s8 wavelet. The s8 mother function seems a more appropriate candidate and produces slightly more detailed scalograms as opposed to the "chunky" Haar ones.

[insert figure 1 about here]

In figures 2 (a-c) we have produced the Haar and s8 scalograms for the first 8000 observations of the FTSE closing prices, the corresponding logarithmic returns and realised volatility. Darker regions in these scalograms correspond to higher wavelet transform coefficients. From an initial inspection, we can easily discern that the s8 wavelet provides a finer decomposition ("thinner") than the Haar for the returns (figure 2 (b)) and the realised volatility sequence (figure 2 (c)) due the discontinuous nature of the latter wavelet depicted in figure (1). Because of the smoothness of the s8 mother wavelet, the scalograms for that transform exhibit smoother transitions between low, medium and high-valued wavelet coefficients, and produce clearer bifurcations. An interesting point here from the comparison of the Haar and s8 scalograms is that for both the levels and the transformations of the FTSE series, the message they deliver is the same.

By careful examination of the *index closing prices* scalograms in figure 2 (a), we can see that both the Haar and s8 CWT coefficients change patterns after the 3000th (mainly 4000th i.e. roughly the 1st half of the history of the series) observation onwards and especially from the short "booming" period before the 1987 crash. Until that point, the prices of the wavelet transform coefficients are lower, indicative of the relative smoothness or lack of excessive volatility and of very weak positive trend. Both Haar and s8 scalograms are quite similar. For the larger scales of between 150 and 250 days though, the Haar based scalogram reports wider and fewer periods of smaller coefficients than that of the s8 wavelet. We attribute this to the structure of the mother wavelet function itself. Through both scalograms though, we are able to discern that for the lower scales there is a relative absence of trend (*i.e.* of a low frequency component) whereas for larger scales and larger time "windows", a very weak trend is more obvious for the first half of the index series. One may recognise those as the darker areas (*"tree trunk"* like formations) i.e. collections of high coefficients on the top of both scalograms in figure 2. It is obvious form the two scalograms that the wavelet coefficients are able to capture the change of the pattern after the 4000th observation. They also reveal the increase in volatility and trend of the series. The difference between the Haar and the s8 based scalogram is more evident on the right bottom half of the graphs where for the s8 wavelet, one can identify more easily the bifurcations formed by the coefficients (see also details of scalograms in figure 3). This shows mainly how small and large sequences of coefficients interchange and may imply a multifractal structure or some kind of non-periodic seasonality (see also Ramsey and Zaslavsky, 1995). See Wornel (1993 and 1995) for an extensive treatment of the relevant theoretical background and framework regarding this issue.

The patterns discussed so far, clearly change for the last half of the series. It is obvious from the time series plot of figure 2 (a), that there is a increase of the steepness and the variance of the index sequence. This follows up historically the occurrence of the 1987 stock market crash. The crash occurs in the vicinity on the 4500th x-axis coordinate, where both Haar and s8 scalograms show a concentration of high coefficients on all scales (shown as an inverted dark peak). We see that regardless of the choice of the mother wavelet, the actual timing of the crash of 1987 is detected successfully. Following that point, the volatility of the series seems to increase considerably with finer bifurcations of wavelet coefficients occurring in low, medium and large scales. It is obvious that the frequency and the intensity of the aperiodic cycle structures has changed for the last half of the series. The keen eye can also identify the rest of the famous crises as they occur after observation 7000 such as the Asian crisis, the NASDAQ and others. These and the effect of the incident of the 11th of September 2001 can be seen in figure 3 (a) where we have produced the s8 based scalogram of the whole 8192 FTSE ALL SHARE observations. In our analysis so far we choose to limit to the first 8000 observations in order to exclude the intensive fluctuations of the last part of the history of the series. Although our discrete and continuous analysis has included all 8192 observations, we choose to truncate the sample in order to avoid depicting the large valued coefficients at the end of the scalograms by excluding 192 points. We do this as we are mainly interested in the 1987 crash which seems more isolated and clear to interpret (we can examine though the scalogram of this last cluster of observations at the end of figure 3 (a)). We can thus concentrate on the oil crisis, the 1987 and Asian markets crashes and avoid the “blurring” of the results at the right edge of the series because of the increased concentration of high valued coefficients due to the clustering of well known recent events (mainly September the 11th). It would be interesting though to see in a couple of years how these scalograms would have “evolved” with the inclusion of the recent and future history of the series.

In figure 3 (a), we can clearly detect after the vertical barrier line, the change in the scalogram’s pattern. We can also locate the intense oscillations following the Asian crisis, the NASDAQ crash and the September the 11th events at the darker regions of the right edge of the scalogram. An interesting point is that the oil crisis of the 70s is not that evident from the levels of the index as in the scalograms of the returns and realised volatility. This is more clearly shown in figure 2 (d), where we show the s8 scalograms from figures 2 (a-c), side by side for comparison purposes.

For further analysis purposes we chose three sub-samples from the history of the FTSE ALL SHARE index series. We generated only the s8 scalograms in order to obtain smoother and clearer graphs. The first period covers 500 sample points, starting from the 1000th one and ending on the 1500th one. It covers the daily observations ranging from 09/05/1974 to 08/04/1976. The second and the third have both length of 1000 observations. The second starts on the 4000th one and ends on the 5000th observation. It covers the time-span 07/11/1985 - 07/09/1989. The third and last sub-sample refers to the period between the 7000th and 8000th observation i.e. the dates 08/05/1997 and 07/03/2001. The analysis of the first, second and third sub-samples is displayed in figures 3 (a), (b) and (c) respectively. In table 3 we have listed the 19 largest shocks or oscillations encountered in the history of the whole sample by date of occurrence and position in the sample for reference reasons. These have been isolated on the basis of the 19 highest wavelet trabform coefficients. In all the above mentioned figures, we display on top the corresponding realised volatility sequences which provides an adequate representation of the magnitude of the oscillations.

[insert table 2 about here]

[insert figure 2 about here]

[insert figure 3 about here]

[insert table 3 about here]

4 The DWT of the FTSE

In the previous section we used scalograms to reveal information on the structure of the FTSE prices, returns and volatility. In this section we concentrate on the application of the discrete wavelet transform (DWT) for the analysis and smoothing of the sequences used in the preceding section. We concentrate on the examination of the index and the logarithmic returns.

In the analysis that follows, we have applied the Maximum Overlap Discrete Wavelet Transform (MODWT) (see Percival and Walden, 2000). We examined the series using the Haar, daubechies 6 and 20 (d6, d20), symmlet 20 (s20) and coiflet 30 (c30) based MODWT (wavelet functions are depicted in figure 1). We found that the results qualitatively similar for the d20, s20 and c30 wavelets so we concentrate in this paper on the report of the Haar and d20 DWT and multiresolution analysis (MRA) results.

In figure 4 we display the multiresolution approximations and decompositions of the FTSE ALL SHARE index closing prices (a) and its returns respectively (b). The MRA in those sub-figures was conducted with a Haar wavelet. We used 8 levels of approximation which we found that provide adequate analysis. For each of those figures, the right part contains the detail coefficients sequences which when added to the "smooth" series S8, generate the reconstruction sequences for each of the 8 levels. For example, in order to obtain the smoothed sequence S6 of the FTSE index in figure 4 (a), one just needs to add to S8 the detail coefficient sequences D8, and D7 i.e. S6 is simply:

$$\mathrm{S6} = \mathrm{S8} + \mathrm{D8} + \mathrm{D7} \tag{3}$$

[insert figure 4 about here]

In figure 5 (a) we show how the FTSE ALL SHARE index closing prices are approximated (smoothed) by the S6 smooth level. In figure 5 (b) we show the original series overlapped by the S6 sequence for the period 1985-1988 and the level 6 residuals ϵ_6 which are computed as:

$$\epsilon_i = \mathrm{DATA} - \sum_{i=1}^{j} \mathrm{D}i \tag{4}$$

for j=6. The arbitrarily chosen MRA S6 level of smoothing follows the FTSE very close, especially during the 1987 crash. As a rule of thumb, we choose the decomposition level which delivers a relatively noise-free approximation of the original sequence. We can see in figure 4 (c) that the level 1-4 approximations for the returns, contain a strong noise signature.[3] So levels 5 or 6 may be the most likely candidates for reconstruction. The Haar DWT transform coefficients for the first 6 levels, the S6 smooth and the inverse discrete wavelet transform of the FTSE series is depicted in figure 5 (c). The DWT of the FTSE logarithmic returns sequence, computed for a maximum of 13 levels, is represented in figure 5 (d). Again we applied here a Haar wavelet. In this attempt we have accounted for all the possible discrete wavelet decomposition levels to demonstrate the 2^j order of the DWT algorithm. We should recall here that $8192 = 2^{13}$ which allows for 13 levels of discrete decomposition. In practice, one may choose to concentrate on the first 5 to 7 levels as higher decompositions provide no further information on the variability of the series.

[insert figure 5 about here]

[3] We should not oversee the fact the differencing is a high-pass filter and returns are expected to contain more noise than the corresponding level prices.

In the same figure we may see how the oil crisis and the crash of 1987 have been picked up by the $d1$ to $d4$ coefficients series, seen as negative and positive spikes on those levels (see figure 5 (c)).

Apart from the MRA, the DWT is also useful in denoising. The most popular technique is called "Waveshrink" (e.g. (see Bruce and Gao, 1996)) and is based on the elimination ("shrinking") of the small (below some threshold) wavelet transform coefficients to 0. The denoised sequence is then obtained from the inverse discrete wavelet transform. Antoniou and Volrow (2003) demonstrate how this can be applied on stock index returns and reveal evidence of deterministic nonlinearities in the denoised series.

5 Conclusions

Wavelets transforms can be used on data in order obtain information on various frequencies as well as in time. This is a clear advantage over spectral analysis which can only focus on frequencies. Financial time series exhibit volatility and sharp localised fluctuations. This makes them ideal candidates for wavelet analysis. Through continuous and discrete wavelet transforms, we can conduct outlier detection and smoothen and denoise time-series without resorting to a theoretical model.

In this paper we showed how continuous and discrete wavelet transforms can be applied on financial time series. We also showed how the results from these transforms can be used to gain an insight on the dynamics and the structure of these time series. Initially, we conducted time-scale analysis through wavelet scalograms. We revealed a wealth of structures in various scales and showed how these graphs can identify significant events that altered the structure or volatility of the sequences. More precisely, using the FTSE ALL SHARE daily time series, we were able to identify the timing of shocks such as the oil crisis of the '70s or the 1987 stock market crash. We were also able to find how these events translate in various scales. It is obvious from our brief analysis here that the ability of wavelets to localise in time as well as in frequency, makes them a very useful tool that should find its way into the mainstream time series analysis econometric toolkits.

Acknowledgements

We wish to thank Prof. Krishna Paudyal for his useful comments and suggestions.

References

Antoniou, A., Volrow, C. E., 2003. Recurrence quantification analysis of wavelet pre-filtered index returns. Tech. rep., University of Durham, Economics & Finance, Discussion Paper Series.

Bera, A., Jarque, C., 1981. Efficient Tests for Normality, Heteroskedasticity and Serial Independence of Regression Residuals: Monte-Carlo evidence. Economics Letters 7, 313–318.

Bruce, A., Gao, H.-Y., 1996. Understanding WaveShrink: Variance and bias estimation. Biometrika 83 (4).
URL `ftp://ftp.statsci.com/pub/gao/varbias.ps.Z`

Capobianco, E., 1999. Statistical Analysis of Financial Volatility by Wavelet Shrinkage. Methodology And Computing In Applied Probability 1 (4), 423–443.

Capobianco, E., 2001. Wavelet Transforms for the Statistical Analysis of Returns Generating Stochastic Processes. International Journal of Theoretical and Applied Finance 4 (3), 511–34.

Capobianco, E., 2002. Multiresolution Approximation for Volatility Process. Quantitative Finance 2 (2), 91–110.

Chen, P., 1996. A random walk or color chaos on the stock market? time-frequency analysis of S&P indexes. Studies in Nonlinear Dynamics and Econometrics 1 (2), 87–103.

Cont, R., 2001. Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance 1 (2), 223–36.

Gencay, R., Selcuk, F., Whitcher, B., 2001. An Introduction to Wavelets and Other Filtering Methods in Finance and Economics. Academic Press, San Diego.

Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time series models. Journal of Time Series Analysis 4, 221–238.

Graps, A. L., 1995. An Introduction to Wavelets. IEEE Computational Sciences and Engineering 2 (2), 50–61.

Greenblatt, S. A., 1996. Wavelets in econometrics: An application to outlier testing. In: Gilli, M. (Ed.), Computational Economic Systems: Models, Methods & Econometrics. Vol. 5 of Advances in Computational Economics. Kluwer Academic Publishers, Dordrecht, pp. 139–160.

Hubbard, B. B. A., 1998. The World According to Wavelets: The Story of a Mathematical Technique in the Making, 2nd Edition. A K Peters Limited Publisher Record, Natick.

Jamdee, S., Los, C. A., January 2003. Dynamic Risk Profile of the US Term Structure by Wavelet MRA. Working paper, Kent State University, Department of Finance.

Jensen, M. J., 1994. Wavelet Analysis of Fractionally Integrated Processes. Working paper, University of Missouri.

Jensen, M. J., 1997. Making wavelets in finance. Financial Engineering News (1), 1–10.

Jensen, M. J., 1999a. An Approximate Wavelet MLE of Short- and Long-Memory Parameters. Studies in Nonlinear Dynamics and Econometrics 3 (4), 239–53.

Jensen, M. J., 1999b. Using wavelets to obtain a consistent ordinary least squares estimator of the fractional differencing parameter. Journal of Forecasting 18, 17–32.

Los, C. A., July 2003. Financial Market Risk: Measurement & Analysis. Routledge International Studies in Money and Banking. Taylor & Francis Books Ltd, London, UK.

Los, C. A., Karuppiah, J., December 2000. Wavelet Multiresolution Analysis of High-Frequency Asian FX Rates. In: Quantitative Methods in Finance & Bernoulli Society 2000 Conference. University of Technology, Sydney, Australia, pp. 171–198.

Mandelbrot, B. B., 1966. Forecasts of Future Prices, Unbiased Markets and Martingale Models. Journal of Business 39 ((Special Supplement)), 242–255.

Ogden, R. T., 1997. Essential wavelets for statistical applications and data analysis. Birkhuser.

Olmeda, I., Fernandez, E., 2000. Filtering with wavelets may be worst than you think. In: Computing in Economics and Finance, Society for Computational Economics.

Percival, D. B., Walden, A. T., 2000. Wavelet Methods for Time Series Analysis. Cambridge University Press, New York, series in Statistical and Probabilistic Mathematics, No. 4.

Peters, E. E., 1994. Fractal market analysis : applying chaos theory to investment and economics. J. Wiley & Sons, New York.

Ramsey, J. B., October 2002. Wavelets in economics and finance: Past and future. Studies in Nonlinear Dynamics and Econometrics 6 (3), Article 1.

Ramsey, J. B., Lampart, C., 1998. The decomposition of economic relationships by time scale using wavelets: Expenditure and income. Studies in Nonlinear Dynamics and Econometrics 3 (1), 23–42.

Ramsey, J. B., Zaslavsky, G. M., 1995. An Analysis of U.S. Stock Price Behaviour Using Wavelets. Fractals 3, 377–389.

Ramsey, J. B., Zhifeng, Z., 1994. The application of wave form dictionaries to stock market index data. Tech. Rep. 94-05, STAR Discussion Paper.

Ramsey, J. B., Zhifeng, Z., 1995. The analysis of foreign exchange data using waveform dictionaries. Tech. Rep. 95-03, STAR Discussion Paper.

Schleicher, C., 2002. An Introduction to Wavelets for Economists. Tech. rep., Bank of Canada Working Paper.

Strang, G., 1989. Wavelets and dilation equations: a brief introduction. SIAM Review 31, 614–627.

Strang, G., Nguyen, T. A., 1996. Wavelets and Filter Banks. Wellesley-Cambridge Press, Wellesley.

Wornell, G., 1993. Wavelet-based representation for the 1/f family of fractal processes. Proc. IEEE (Sept. 1993).

Wornell, G., 1995. Signal Processing with Fractals: A Wavelet Based approach.

Prentice Hall.

Table 1
Descriptive statistics. Jarque-Bera p-values within parenthesis.

Statistic	index	returns	realised volatility
minimum	61.92	-0.1191000	0.000e+000
Q1	221.60	-0.0048350	5.287e-006
median	768.90	0.0003343	2.859e-005
mean	1001.00	0.0003669	9.995e-005
Q3	1514.00	0.0058370	9.368e-005
maximum	3266.00	0.0894300	1.419e-002
st.deviation	899.0817	0.009992	0.000335
skewness	0.941687	-0.332639	19.48220
kurtosis	2.741635	12.32305	608.4498
Jarque-Berra	1233.378 (0.0)	29815.85 (0.0)	1.26e+08 (0.0)

Table 2
The 3 subsamples used in figures 3, subfigures (b)-(c)

Subsample	**Dates**	**Range**	**Size**
1	09/05/1974-08/04/1976	1000-1500	500
2	07/11/1985-07/09/1989	4000-5000	1500
3	08/05/1997-07/03/2001	7000-8000	1500

Table 3
Dates and positions of th 19th largest oscillations in the FTSE series as these are identified by the 19 largest DWT wavelet coefficients.

Dates	**Observation**
06/12/1973	890
14/12/1973	896
01/03/1974	951
02/01/1975	1170
24/01/1975	1186
27/01/1975	1187
29/01/1975	1189
30/01/1975	1190
07/02/1975	1196
10/02/1975	1197
11/03/1975	1218
17/04/1975	1245
19/10/1987	4507
20/10/1987	4508
21/10/1987	4509
22/10/1987	4510
26/10/1987	4512
10/04/1992	5676
11/09/2001	8134

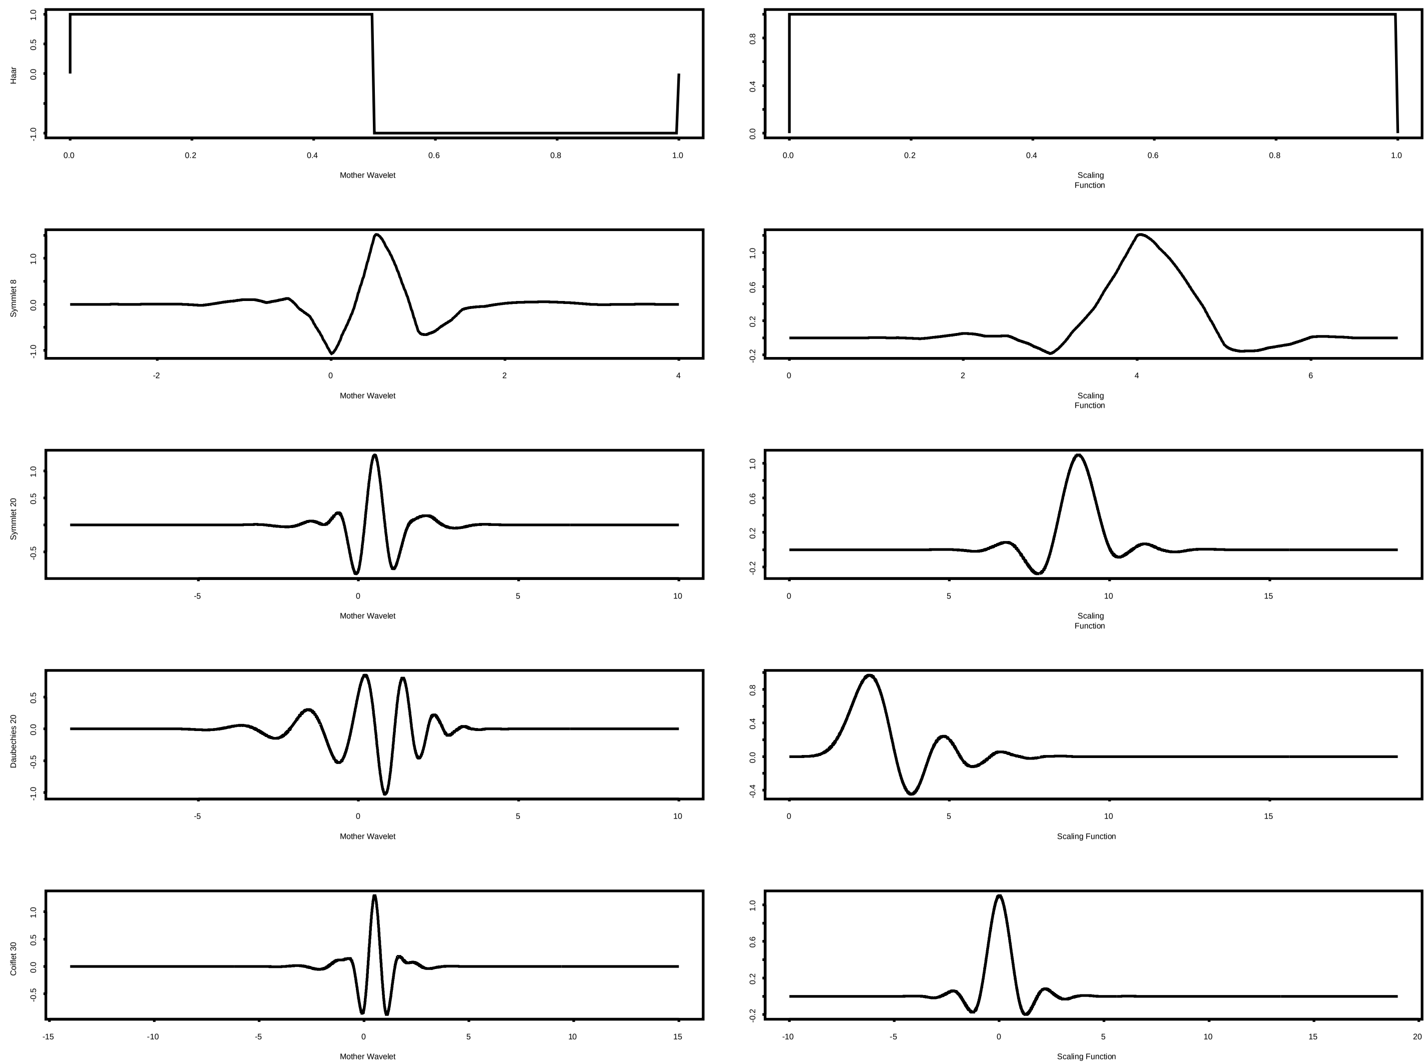

Fig. 1. Wavelets and scaling functions: Haar, Symmlet 8 and 20, Daubechies 20, Coiflet 30

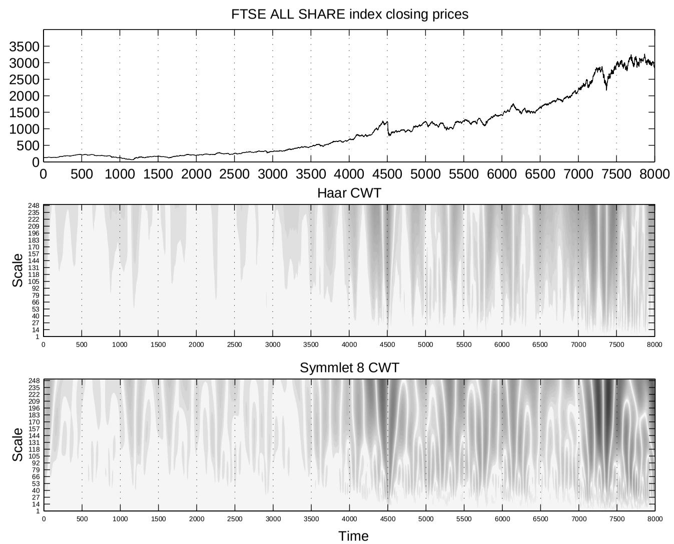

(a) Closing prices

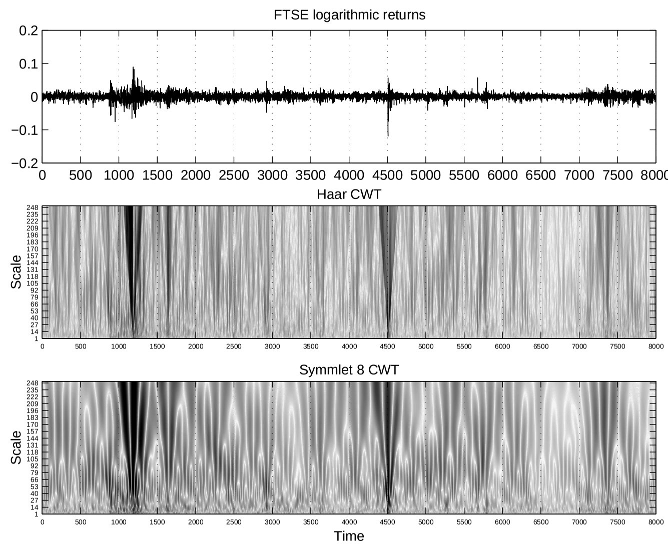

(b) Returns

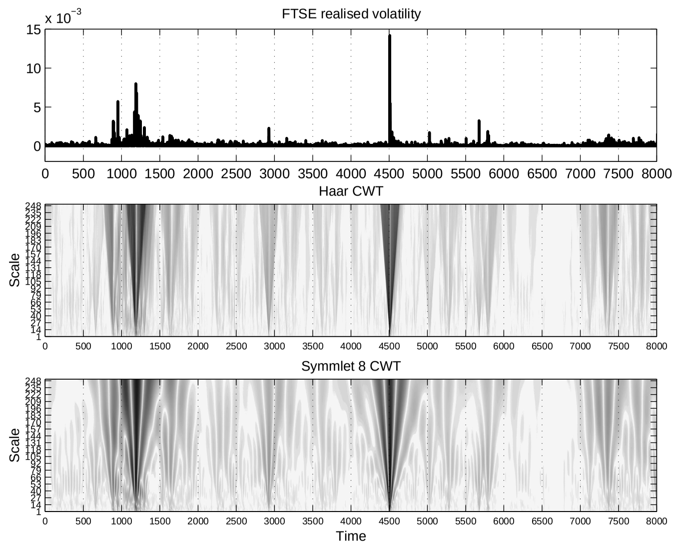

(c) Realised Volatility

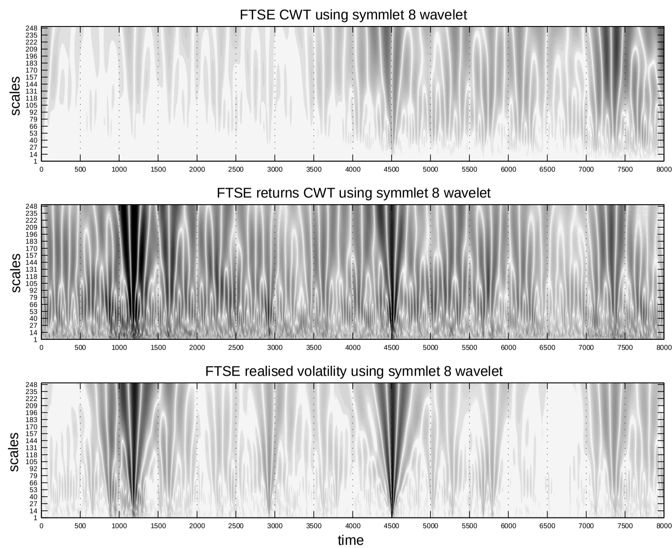

(d) Comparison of s8

Fig. 2. Time-scale analysis of the FTSE closing prices, returns and volatility.

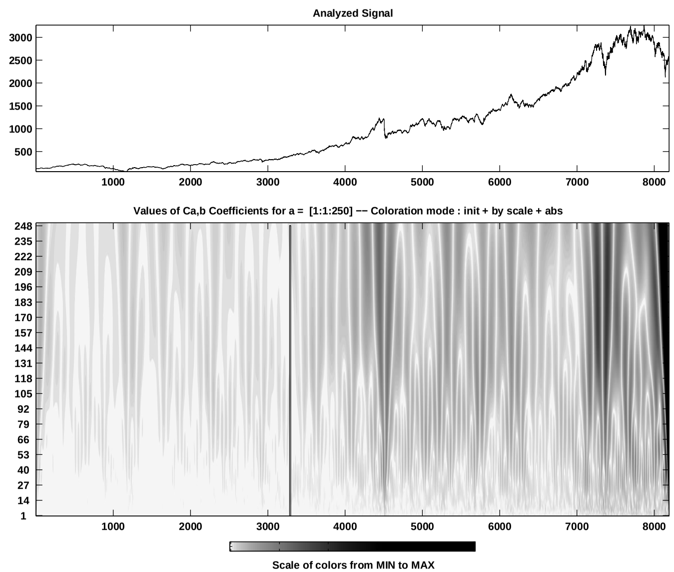

(a) 10/07/1970 till 30/11/2001

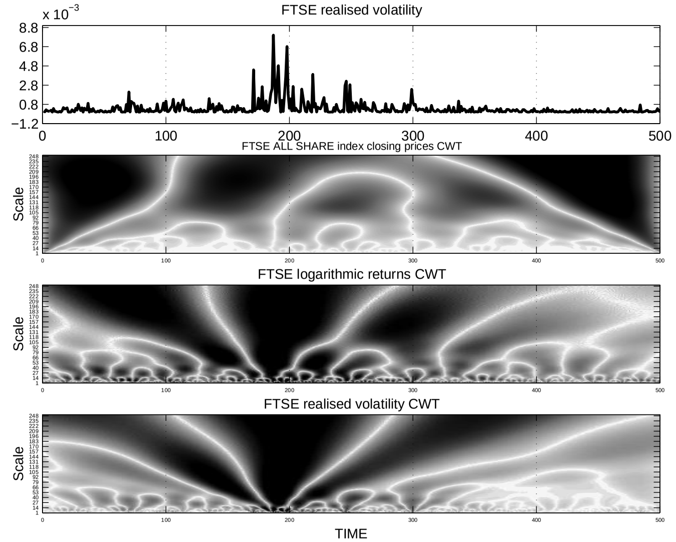

(b) 09/05/1974 till 08/04/1976

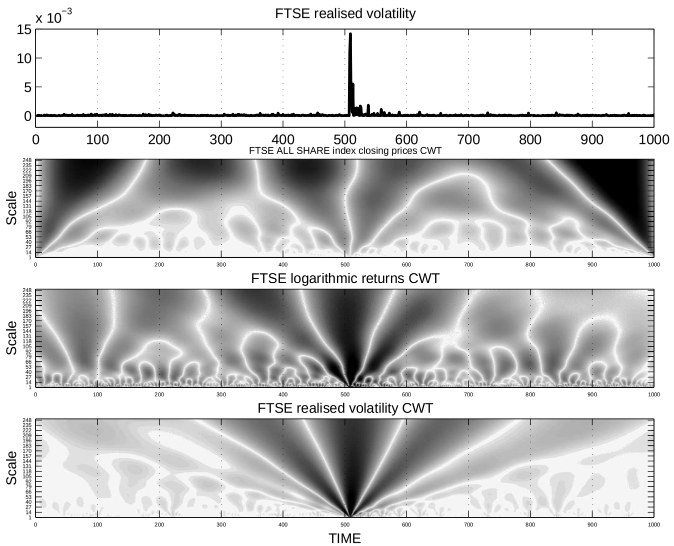

(c) 07/11/1985 till 07/09/1989

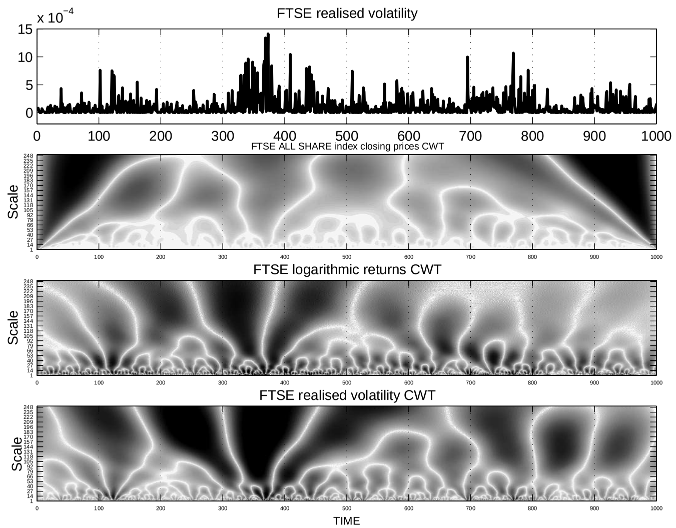

(d) 08/05/1997 till 07/03/2001

Fig. 3. Examining the details of the s8 scalograms in figure 2.

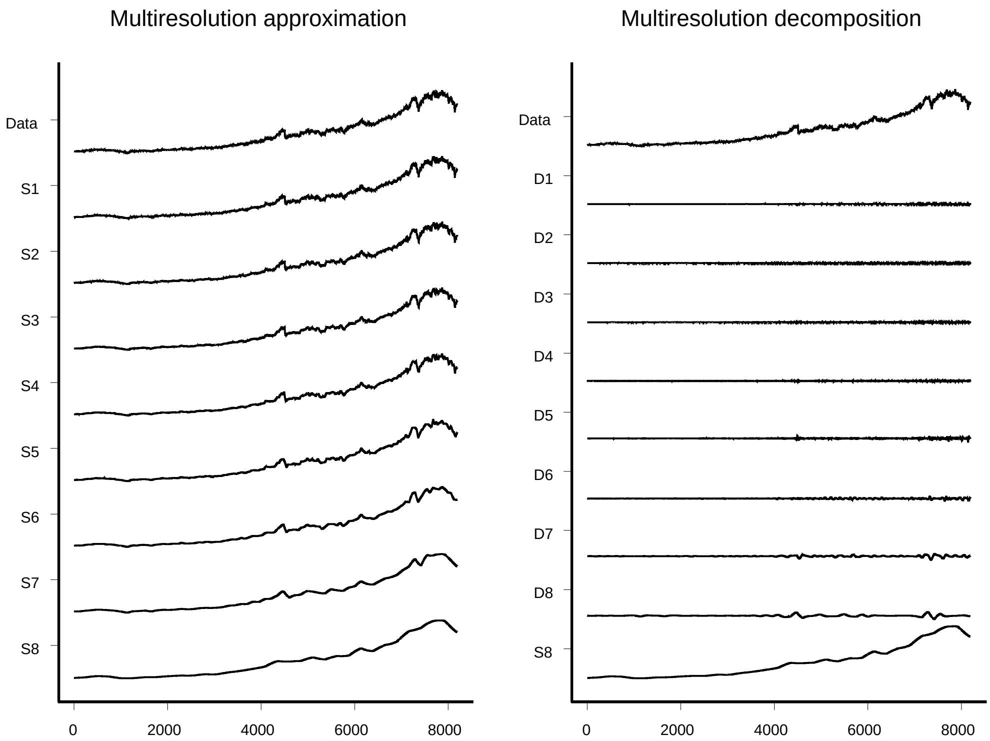

(a) FTSE closing prices

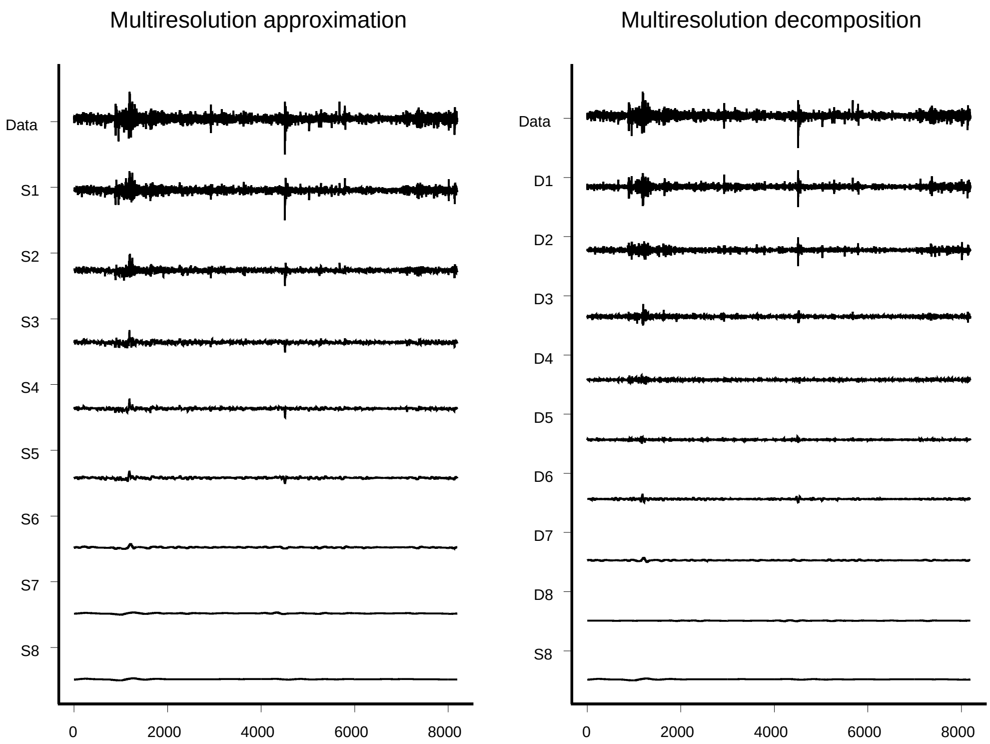

(b) FTSE returns

Fig. 4. Multiresolution approximations and decompositions

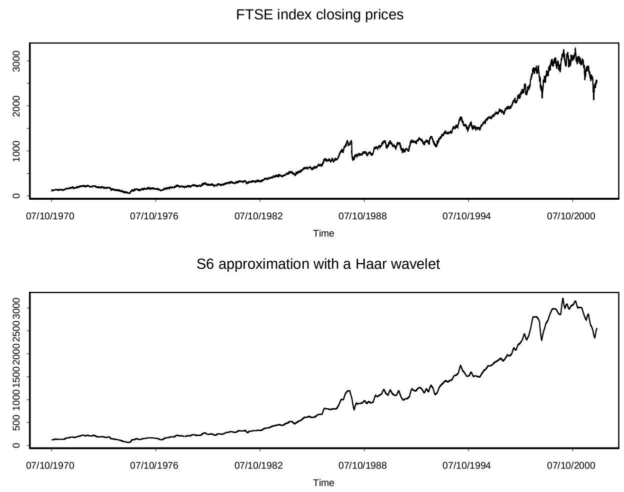

(a) FTSE and S6 smooth

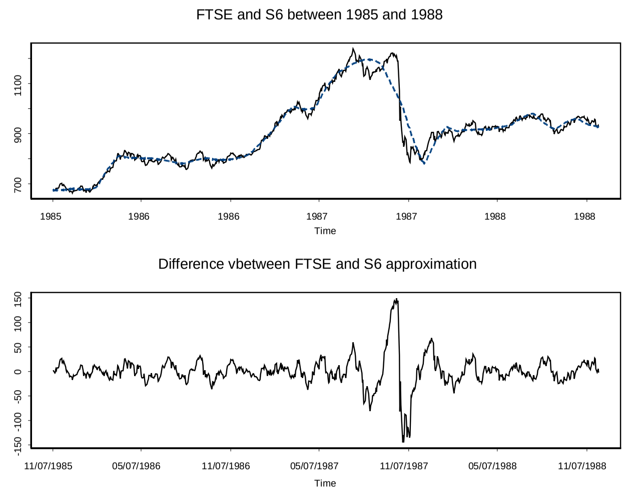

(b) A detail from subfigure (a)

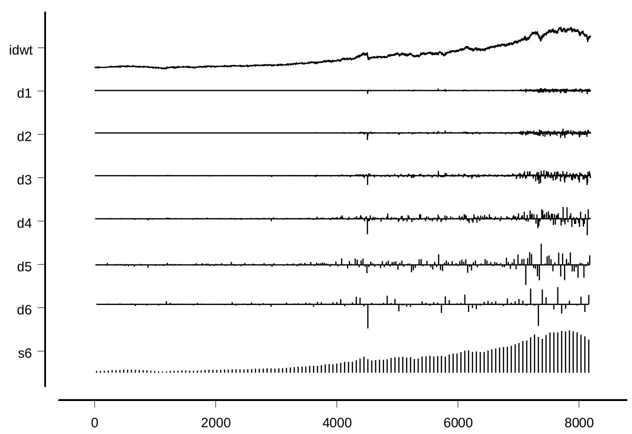

(c) Multiresolution analysis

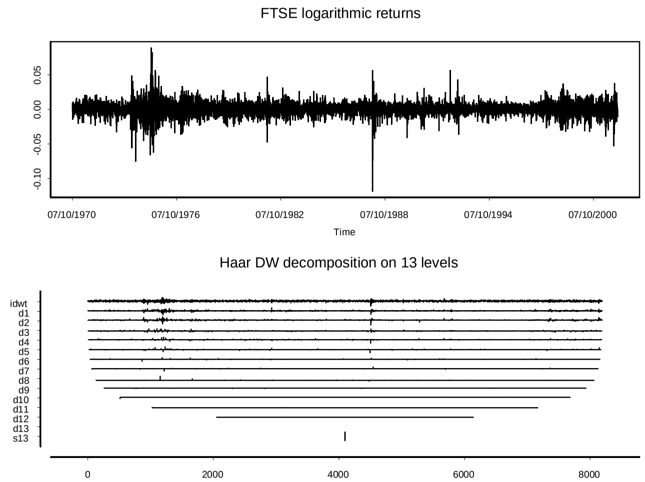

(d) The DWT of the FTE returns

Fig. 5.