

Testing for Homogeneity of Variance in Time Series: Long Memory, Wavelets and the Nile River

Brandon Whitcher

EURANDOM

Simon D. Byers

AT&T Labs Research

Peter Guttorp

University of Washington

Donald B. Percival

University of Washington & MathSoft, Inc.

Abstract

We consider the problem of testing for homogeneity of variance in a time series with long memory structure. We demonstrate that a test whose null hypothesis is designed to be white noise can in fact be applied, on a scale by scale basis, to the discrete wavelet transform of long memory processes. In particular, we show that evaluating a normalized cumulative sum of squares test statistic using critical levels for the null hypothesis of white noise yields approximately the same null hypothesis rejection rates when applied to the discrete wavelet transform of samples from a fractionally differenced process. The point at which the test statistic, using a non-decimated version of the discrete wavelet transform, achieves its maximum value can be used to estimate the time of the unknown variance change. We apply our proposed test statistic on a time series of Nile River yearly minimum water levels covering 622 to 1284 AD. The test confirms an inhomogeneity of variance at short scales and identifies the change point around 720 AD, which coincides closely with the construction of a new device around 715 AD for measuring these water levels.

Some key words: Cumulative sum of squares; Discrete wavelet transform; Fractional difference process; Variance change.

1 Introduction

Suppose we have a time series that we are considering to model as a realization of one portion $Y_0, \dots, Y_{N-1}$ of a stationary Gaussian fractionally differenced (FD) process Y_t . This process can be represented as

$$Y_t = \sum_{k=0}^{\infty} \frac{\Gamma(k+d)}{\Gamma(k+1)\Gamma(d)} \epsilon_{t-k},$$

where $-\frac{1}{2} < d < \frac{1}{2}$, and ϵ_t is a Gaussian white noise process with mean zero and variance σ_ϵ^2 . The spectral density function (SDF) for this process is given by $S(f) = \sigma_\epsilon^2 |2 \sin(\pi f)|^{-2d}$ for $|f| \leq \frac{1}{2}$, while its autocovariance sequence $s_{Y,\tau}$ can be obtained using

$$s_{Y,\tau} = s_{Y,\tau-1} \frac{\tau+d-1}{\tau-d}, \quad \tau = 1, 2, \dots, \quad \text{with } s_{Y,0} = \frac{\sigma_\epsilon^2 \Gamma(1-2d)}{\Gamma^2(1-d)} \quad (1)$$

(for $\tau < 0$, we have $s_{Y,\tau} = s_{Y,-\tau}$). When $0 < d < \frac{1}{2}$, the SDF has a pole at zero, in which case the process exhibits slowly decaying autocovariances and constitutes a simple example of a long memory process; see Granger and Joyeux (1980), Hosking (1981), and Beran (1994, Sec. 2.5). In this case, d is called the long memory parameter.

An important assumption behind any stationary process is that its variance is a constant independent of the time index t . In the context of short memory models, such as stationary autoregressive and moving average (ARMA) processes, a number of tests has been proposed for homogeneity of variance. For a time series consisting of either independent Gaussian random variables with zero mean and possibly time-dependent variances σ_t^2 or a moving average of such variables, Nuri and Herbst (1969) proposed to test the hypothesis that σ_t^2 is constant for all t by using the periodogram of the squared random variables. Wichern, Miller, and Hsu (1976) proposed a moving block procedure for detecting a single change of variance at an unknown time point in an autoregressive model of order one. Hsu (1977, 1979) looked at detecting a single change in variance at an unknown point in time in a series of independent observations. Davis (1979) studied tests for a single change in the innovations variance at a specified point in time in an autoregressive process. Abraham and Wei (1984) used a Bayesian framework to study changes in the innovation variance of an ARMA process.

Tsay (1988) looked at detecting several types of disturbances in time series – among them variance changes – by analyzing the residuals from fitting an ARMA model. Srivastava (1993) found a cumulative sum of squares procedure to perform better than an exponentially weighted moving average procedure for detecting an increase in variance in white noise sequences. Inclán and Tiao (1994) investigated the detection of multiple changes of variance in sequences of independent Gaussian random variables by recursively applying a cumulative sum of squares test to pieces of the original series. As discussed in Section 3.1, the test we propose can be regarded as an adaptation of the work of Hsu (1977, 1979) and Inclán and Tiao (1994) to handle long memory processes, but, except for the tests with a null hypothesis of white noise with constant variance, it is not an easy matter to adapt the other tests.

In this paper we demonstrate how the discrete wavelet transform (DWT) can be used to construct a test for homogeneity of variance in a time series exhibiting long memory characteristics. The DWT is a relatively new tool for time series analysis, but has already proven useful for investigating other types of nonstationary events. For example, Wang (1995) tested wavelet coefficients at fine scales to detect jumps and sharp cusps of signals embedded in Gaussian white noise, and Ogden and Parzen (1996) used wavelet coefficients to develop data-dependent thresholds for removing noise from a signal. The key property of the DWT that makes it useful for studying possible nonstationarities is that it transforms a time series into coefficients that reflect changes at various scales and at particular times. For FD and related long memory processes, the wavelet coefficients for a given scale are approximately uncorrelated; see, e.g., Tewfik and Kim (1992), McCoy and Walden (1996), Wornell (1996) and our discussion in Section 3.2. We show here that this approximation is good enough that a test designed for a null hypothesis of white noise can be used for testing homogeneity of variance in a long memory process on a scale by scale basis. An additional advantage of testing the output from the DWT is that the scale at which the inhomogeneity occurs can be identified. Using a variation of the DWT, we also investigate an auxiliary test statistic that can estimate the time at which the variance of a time series changes.

An outline of the remainder of this paper is as follows. Section 2 provides a brief description of the DWT and points out some key properties we use later on. Section 3 defines the test statistic for detecting sudden changes of variance in a sequence of independent Gaussian random variables, summarizes some simulation results for the DWT of FD processes, and applies our proposed test to a time series of Nile River minimum water levels. Section 4 introduces a procedure for determining the location of a variance change, presents simulation results for FD processes with a known variance change, and demonstrates the procedure using the Nile River data. Section 5 gives some concluding remarks.

2 Discrete wavelet transforms

Let $h_1 = (h_{1,0}, \dots, h_{1,L-1}, 0, \dots, 0)^T$ denote the wavelet filter coefficients of a Daubechies compactly supported wavelet for unit scale (Daubechies 1992, Ch. 6), zero padded to length N by defining $h_{1,l} = 0$ for $l \geq L$. Let

$$H_{1,k} = \sum_{l=0}^{N-1} h_{1,l} e^{-i2\pi lk/N}, \quad k = 0, \dots, N-1,$$

be the discrete Fourier transform (DFT) of h_1 . Let $g_1 = (g_{1,0}, \dots, g_{1,L-1}, 0, \dots, 0)^T$ be the zero padded scaling filter coefficients, defined via $g_{1,l} = (-1)^{l+1} h_{1,L-1-l}$ for $l = 0, \dots, L-1$, and let $G_{1,k}$ denote its DFT. Now define the length N wavelet filter h_j for scale $\tau_j = 2^{j-1}$ as the inverse DFT of

$$H_{j,k} = H_{1,2^{j-1}k \bmod N} \prod_{l=0}^{j-2} G_{1,2^l k \bmod N}, \quad k = 0, \dots, N-1.$$

When $N > L_j = (2^j - 1)(L - 1) + 1$, the last $N - L_j$ elements of h_j are zero, so the wavelet filter h_j has at most L_j non-zero elements.

Let $Y_0, \dots, Y_{N-1}$ be a time series of length N . For scales such that $N \geq L_j$, we can filter the time series using h_j to obtain the wavelet coefficients

$$W_{j,t} = 2^{j/2} \widetilde{W}_{j,2^j(t+1)-1}, \quad \left\lceil (L-2) \left(1 - \frac{1}{2^j}\right) \right\rceil \leq t \leq \left\lfloor \frac{N}{2^j} - 1 \right\rfloor,$$

where

$$\widetilde{W}_{j,t} = \frac{1}{2^{j/2}} \sum_{l=0}^{L_j-1} h_{j,l} Y_{t-l}, \quad t = L_j - 1, \dots, N - 1.$$

The $W_{j,t}$ coefficients are associated with changes on a scale of length τ_j and are obtained by subsampling every 2^j th value of the $\widetilde{W}_{j,t}$ coefficients, which forms a portion of one version of a ‘non-decimated’ DWT called the ‘maximal overlap’ DWT (see Percival and Guttorp (1994) and Percival and Mofjeld (1997) for details on this transform; other versions of the non-decimated DWT are discussed in Shensa (1992), Beylkin (1992), Coifman and Donoho (1995), Nason and Silverman (1995) and Bruce and Gao (1996)). In practice the DWT is implemented via a pyramid algorithm (Mallat 1989) that, starting with the data Y_t , filters a series using h_1 and g_1 , subsamples both filter outputs to half their original lengths, keeps the subsampled output from the h_1 filter as wavelet coefficients, and then repeats the above filtering operations on the subsampled output from the g_1 filter. A simple modification, namely, not subsampling the output at each scale and inserting zeros between coefficients in h_1 and g_1 , yields the algorithm for computing $\widetilde{W}_{j,t}$ described in Percival and Mofjeld (1997).

3 Testing for homogeneity of variance

If $Y_0, \dots, Y_{N-1}$ constitutes a portion of an FD process with long memory parameter $0 < d < \frac{1}{2}$, and with possibly nonzero mean, then each sequence of wavelet coefficients $W_{j,t}$ for Y_t is approximately a sample from a zero mean white noise process. This enables us to formulate our test for homogeneity of variance using wavelet coefficients for FD processes, as follows.

3.1 The test statistic

Let $X_0, \dots, X_{N-1}$ be a time series that can be regarded as a sequence of independent Gaussian (normal) random variables with zero means and variances $\sigma_0^2, \dots, \sigma_{N-1}^2$. We would like to test the hypothesis $H_0 : \sigma_0^2 = \dots = \sigma_{N-1}^2$. A test statistic that can discriminate between this null hypothesis and a variety of alternative hypotheses (such as $H_1 : \sigma_0^2 = \dots = \sigma_k^2 \neq$

$\sigma_{k+1}^2 = \dots = \sigma_{N-1}^2$, where k is an unknown change point) is the normalized cumulative sums of squares test statistic D , which has previously been investigated by, among others, Brown, Durbin, and Evans (1975), Hsu (1977) and Inclán and Tiao (1994). To define D , let

$$\mathcal{P}_k \equiv \frac{\sum_{j=0}^k X_j^2}{\sum_{j=0}^{N-1} X_j^2}, \quad D^+ \equiv \max_{0 \leq k \leq N-2} \left(\frac{k+1}{N-1} - \mathcal{P}_k \right) \quad \text{and} \quad D^- \equiv \max_{0 \leq k \leq N-2} \left(\mathcal{P}_k - \frac{k}{N-1} \right). \quad (2)$$

The desired statistic is given by $D \equiv \max(D^+, D^-)$. For $N > 2$ there is no known tractable closed form expression for the critical levels for D under the null hypothesis. Brown, Durbin, and Evans (1975) obtained critical levels by an interpolation scheme that makes use of the fact that, if N is even and if we group the squared deviates by pairs, then D reduces to the well-known cumulative periodogram test for white noise (Bartlett 1955), for which critical levels are available (Stephens 1974). Hsu (1977) used two methods, Edgeworth expansions and fitting the first three moments to a beta distribution, in order to obtain small sample critical levels for a statistic equivalent to D . The results in Inclán and Tiao (1994) indicate that, for large N and $x > 0$,

$$Pr \left\{ (N/2)^{\frac{1}{2}} D \leq x \right\} \approx Pr \left(\sup_t |B_t^0| \leq x \right) = 1 + 2 \sum_{l=1}^{\infty} (-1)^l e^{-2l^2 x^2},$$

where B_t^0 is a Brownian bridge process, and the right-hand expression is equation (11.39) of Billingsley (1968). Critical levels for D under the null hypothesis can be readily obtained through Monte Carlo simulations and are shown in Table 1 for comparison with levels determined by the Brownian bridge process.

3.2 Simulation study

Here we investigate whether in fact the DWT of an FD process is a good approximation to white noise as far as performance of the test statistic D is concerned. We first note that

$$\text{cov}\{W_{j,t}, W_{j,t+\tau}\} = \sum_{m=-(L_j-1)}^{L_j-1} s_{Y,2^j\tau+m} \sum_{l=0}^{L_j-|m|-1} h_{j,l} h_{j,l+|m|}, \quad (3)$$

where $s_{Y,\tau}$ is given in (1). Using this equation, we can compute the unit lag correlations $\text{corr}\{W_{j,t}, W_{j,t+1}\}$ given in Table 2 for the Haar, D(4) and LA(8) wavelet filters and scales 1,

2, 4 and 8; here ‘D(4)’ and ‘LA(8)’ refer to the Daubechies extremal phase filter with four nonzero coefficients and to her least asymmetric filter with eight coefficients (Daubechies 1992). Note that all unit lag correlations are negative, with departures from zero increasing somewhat as j increases. Computations indicate that $|\text{corr}\{W_{j,t}, W_{j,t+2}\}| < 0.033$ and $|\text{corr}\{W_{j,t}, W_{j,t+3}\}| < 0.009$ for all three wavelet filters, and $\text{corr}\{W_{j,t}, W_{j,t+\tau}\}$ is negligible for $\tau \geq 4$. To ascertain the effect of these small remaining correlations on D , we determined the upper 10%, 5% and 1% quantiles for the distribution of D based upon a large number of realizations of white noise for sample sizes commensurate with time series of length $N \in (128, 256, 512, 1024, 2048, 2^{15})$. Using these quantiles, we then generated a large number of realizations of length N from FD processes with $d \in (0.05, 0.25, 0.4, 0.45)$; computed wavelet coefficients for scales 1, 2, 4 and 8 using the three wavelet filters; computed the test statistic D for all four scales based on the wavelet coefficients; and compared the resulting D ’s to the white noise critical levels. We found the deviations between the actual rejection rates and the rates established for white noise to be generally around 10%, with the agreement decreasing somewhat with increasing scale (this is consistent with Table 2). We can thus conduct an approximate α level test for variance homogeneity of an FD process, on a scale by scale basis, by simply using critical levels determined under the assumption of white noise (details about this Monte Carlo study are given in Whitcher (1998)).

One may not want to perform Monte Carlo studies in order to obtain critical values for the test statistic D . The simulation study described above was run again substituting the asymptotic critical values (last column of Table 1) for the Monte Carlo critical values. For sample sizes greater than 128, the percentage of times D exceeded the asymptotic critical levels was within 10% of the theoretical quantile. The Haar wavelet filter was found to be conservative for all sample sizes; i.e., the percentage of times D exceeded the asymptotic critical levels was below the nominal percentile. Hence, wavelet coefficients of length 128 or greater using asymptotic critical values will give reasonable results if Monte Carlo critical values have not been computed.

3.3 Application to the Nile River water levels

As an example of a time series exhibiting both long memory and possible inhomogeneity in variance, we consider the Nile River minimum water level time series (Toussoun (1925); for an interesting recent analysis of this series, see Eltahir and Wang (1999)). This time series consists of $N = 663$ yearly values from 622 AD to 1284 AD and is plotted in the top panel of Figure 3. Beran (1994, p. 118) found an FD model to be fit this time series well and obtained an estimate of $d = 0.40$ using an approximate maximum likelihood approach. Visually there seems to be greater variability in the first part of this series. Beran (1994, Sec. 10.3) investigated the question of a change in the long memory parameter in this time series by partitioning the first 600 observations into two sub-series containing, respectively, the first 100 and the remaining 500 measurements. Maximum likelihood estimates of the long memory parameter d were quite different between the two sub-series, 0.04 and 0.38 respectively. This analysis suggests a change in d , a conclusion also drawn in Beran and Terrin (1996) using a procedure designed to test for change in the long memory parameter.

We can perform a similar analysis using the wavelet variance, which decomposes the variance of Y_t on a scale by scale basis. The wavelet variance ν_j^2 for scale $\tau_j = 2^{j-1}$ is defined to be the variance of $\widetilde{W}_{j,t}$ and can be estimated by

$$\tilde{\nu}_j^2 = \frac{1}{N - L_j + 1} \sum_{l=L_j-1}^{N-1} \widetilde{W}_{j,l}^2;$$

see Percival (1995) for a discussion of the statistical properties of this estimator. The estimated wavelet variances, given a partitioning scheme similar to the one used by Beran (1994), are displayed in Figure 1. The 95% confidence intervals for scales of 1 and 2 years do not overlap, suggesting that the greater variability seen in the first one hundred years might be attributable to changes in variance at just these two scales.

For an FD process we have $\nu_j^2 \propto \tau_j^{2d-1}$ approximately, so we can estimate d by regressing $\log \tilde{\nu}_j^2$ on $\log \tau_j$ and using the estimated slope $\hat{\beta}$ to form $\hat{d} = \frac{1}{2}(\hat{\beta} + 1)$. This yields estimates of $\hat{d} = 0.38, 0.42$ and -0.07 for the whole time series, the last 563 observations and the first

100 observations. These compare favorably with Beran's values of 0.40, 0.38 and 0.04, but Figure 1 says that the smaller value for $\hat{d}$ in the first 100 years is due to increased variability at scales of 2 years or less. The observed difference in $\tilde{\nu}_j^2$ at longer scales between the first and last portions of the time series is consistent with sampling variability.

Let us now apply the methodology developed in this paper to the Nile River minima. Using all $N = 663$ values in the time series, we computed our test statistic for scales of 1, 2, 4 and 8 years based, respectively, on 331, 115, 57 and 28 wavelet coefficients. The results from the test, shown in Table 3, confirm an inhomogeneity of variance at scales of 1 and 2 years, but fail to reject the null hypothesis of variance homogeneity at scales of 4 and 8 years.

4 Locating the change in variance

4.1 Auxiliary test

We shift our attention to determining the location of a variance change in the original time series. A naive choice of location can be based on the test statistic D ; i.e., on the location of the wavelet coefficient at which the cumulative sum of squares at level j achieves its maximum. Since the wavelet coefficient is a linear combination of observations from the original series, this procedure will yield a range of times trapping the variance change. The subsampling inherent in the DWT, however, causes a loss of resolution in time at each scale. We thus propose to use the non-decimated coefficients $\widetilde{W}_{j,t}$ to more accurately determine the location of a variance change after detection by the DWT.

4.2 Simulation study

A study was conducted to investigate how well the test statistic D , now using the $\widetilde{W}_{j,t}$ coefficients, locates a single variance change in a series with long memory structure. To do this, we implemented a setup motivated by the Nile River example by (i) generating a realization of length $N = 663$ from an FD process with a specified long memory parameter d

= 0.4; (ii) adding Gaussian random variables with $\sigma^2 = 2.07\delta$ to the first 100 observations of the FD process, where 2.07 is the variance $s_{Y,0}$ of an FD process with $d = 0.4$ and $\sigma_\epsilon^2 = 1$ as given by Equation (1); (iii) computing the $\widetilde{W}_{j,t}$ coefficients for $j = 1, \dots, 4$, using the Haar, D(4) and LA(8) wavelet filters; and (iv) recording the location of the wavelet coefficient from which the test statistic D attains its value, adjusting for the phase shift of the filter output $\widetilde{W}_{j,t}$ by shifting the location $\frac{1}{2}L_j$ units to the left. The above was repeated 10,000 times each for $\delta = 0.5, 1, 2$ and 3. Those estimated locations of the variance changes are displayed in Figure 2. The estimates are roughly centered around the 100th wavelet coefficient for $j = 1, 2$ with the spread narrowing as the variance ratio increases. There is a very slight difference between wavelet filters, the broader spread being associated with the longer wavelet filters. However, for variance ratios of 2:1 or greater all three wavelets appear to perform equally well. The estimates from the $j = 1$ level have a median value closer to the truth with much less spread at every combination of variance ratio and wavelet filter, when compared to the second level. We therefore recommend using the unit scale estimate when trying to locate an sudden change of variance in a time series.

4.3 Application to the Nile River water levels

We apply the above procedure to locate the variance change in the Nile River minimum water levels. Figure 3 displays the normalized cumulative sum of squares as a function of wavelet coefficient for the first two scales. We see a sudden accumulation of variance in the first 100 years and a gradual tapering off of the variance (by construction the series must begin and end at zero). The maximum is actually attained in 720 AD for the level 1 coefficients and 722 AD for level 2. The subsequent smaller peaks occurring in the ninth century are associated with large observations, as seen in the original series, not changes in the variance.

The source document for this series (Toussoun 1925) and studies by Popper (1951) and Balek (1977) all indicate the construction in 715 AD (or soon thereafter) of a ‘nilometer’ in a mosque on Roda Island in the Nile River. After its construction, the yearly minimum water

levels up to 1284 AD were measured using this device, or a reconstruction of it in 861 AD. How measurements were made prior to 715 AD is unknown, but most likely devices with less accuracy than the Roda Island nilometer were used. Our estimated change point at 720 or 722 AD coincides well with the construction of this new instrument, and it is reasonable that this new nilometer led to a reduction in variability at the very smallest scales.

5 Discussion

The discrete wavelet transform has been shown to adequately decorrelate time series with long memory structure for the purpose of evaluating a normalized cumulative sum of squares test statistic. It provides a convenient method for detecting and locating inhomogeneities of variance in such time series. The DWT produces a test statistic that can be evaluated under the assumption of white noise while the non-decimated coefficients $\widetilde{W}_{j,t}$ offer good time domain resolution for locating a variance change. This methodology should be a useful analysis tool applicable to a wide variety of physical processes.

Ogden and Parzen (1996) use a test statistic, similar to D , as a solution to a change-point problem for nonparametric regression. Their statistic involves estimating the standard deviation of the squared wavelet coefficients. We avoid this estimation by dividing the cumulative sum by the total sum of squares (see Whitcher (1998) for a discussion on the relative merits of these two approaches). Whereas we are looking for changes in the variance, they looked at changes in the mean of a process in order to determine an appropriate level-dependent wavelet threshold. The similarities between the Ogden–Parzen test statistic and D indicate that, while we have discussed D in the context of detecting changes in variance, in fact D can pick up others kinds of nonstationarities, a fact that must be taken into account before drawing any conclusion when D rejects the null hypothesis of variance homogeneity.

Beran and Terrin (1996) looked at the Nile River minimum water levels and used a test statistic to argue for a change in the long memory parameter in the time series. The results

from our analysis, in conjunction with an examination of the historical record, suggest an alternative interpretation. There is a decrease in variability at scales of 2 years and less after about 720 AD and that this decrease is due to a new measurement instrument, rather than a change in the long term characteristics of the Nile River.

Finally, we note that, while we have concentrated on FD processes in order to validate our proposed homogeneity of variance test, in fact our test is by no means limited to just these processes. The key to the methodology proposed here is the decorrelation property of the DWT, which can be verified for other processes by using Equation (3) and comparing the results with Table 2. Alternatively, because the DWT yields an octave band decomposition of the SDF, we can expect the decorrelation property to hold for level j wavelet coefficients as long as there is relatively little change in the SDF over the octave band $[1/2^{j+1}, 1/2^j]$, as is true for FD processes. Other processes for which the decorrelation property holds include first order autoregressive processes with nonnegative lag one autocorrelations, fractional Gaussian noise, stationary long memory power law processes, and certain fractionally integrated autoregressive, moving average processes (i.e., extensions to FD processes that do not exhibit rapid variations within octave bands).

References

- Abraham, B. and W. W. S. Wei (1984). Inferences about the parameters of a time series model with changing variance. *Metrika* 31, 183–194.
- Antoniadis, A. and G. Oppenheim (Eds.) (1995). *Wavelets and Statistics*, Volume 103 of *Lecture Notes in Statistics*, New York. Springer-Verlag.
- Balek, J. (1977). *Hydrology and Water Resources in Tropical Africa*, Volume 8 of *Developments in Water Science*. New York: Elsevier Scientific Pub. Co.
- Bartlett, M. S. (1955). *An Introduction to Stochastic Processes with Special Reference to Methods and Applications*. London, UK: Cambridge University Press.

- Beran, J. (1994). *Statistics for Long-Memory Processes*, Volume 61 of *Monographs on Statistics and Applied Probability*. New York: Chapman & Hall.
- Beran, J. and N. Terrin (1996). Testing for a change of the long-memory parameter. *Biometrika* 83(3), 627–638.
- Beylkin, G. (1992). On the representation of operators in bases of compactly supported wavelets. *SIAM Journal on Numerical Analysis* 29, 1716–1740.
- Billingsley, P. (1968). *Convergence of Probability Measures*. New York: John Wiley & Sons.
- Brown, R. L., J. Durbin, and J. M. Evans (1975). Techniques for testing the constancy of regression relationships over time. *Journal of the Royal Statistical Society, Series B* 37, 149–163.
- Bruce, A. and H.-Y. Gao (1996). *Applied Wavelet Analysis with S-PLUS*. New York: Springer.
- Coifman, R. R. and D. Donoho (1995). Time-invariant wavelet denoising. See Antoniadis and Oppenheim (1995), pp. 125–150.
- Daubechies, I. (1992). *Ten Lectures on Wavelets*, Volume 61 of *CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia: Society for Industrial and Applied Mathematics.
- Davis, W. W. (1979). Robust methods for detection of shifts of the innovation variance of a time series. *Technometrics* 21(3), 313–320.
- Eltahir, E. A. B. and G. Wang (1999). Nilometers, El Niño, and climate variability. *Geophysical Research Letters* 26, 489–492.
- Granger, C. W. J. and R. Joyeux (1980). An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.
- Hosking, J. R. M. (1981). Fractional differencing. *Biometrika* 68(1), 165–176.

- Hsu, D.-A. (1977). Tests for variance shift at an unknown time point. *Applied Statistics* 26(3), 279–284.
- Hsu, D.-A. (1979). Detecting shifts of parameter in gamma sequences with applications to stock price and air traffic flow analysis. *Journal of the American Statistical Association* 74, 31–40.
- Inclán, C. and G. C. Tiao (1994). Use of cumulative sums of squares for retrospective detection of changes of variance. *Journal of the American Statistical Association* 89(427), 913–923.
- Mallat, S. G. (1989). A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 11(7), 674–693.
- McCoy, E. J. and A. T. Walden (1996). Wavelet analysis and synthesis of stationary long-memory processes. *Journal of Computational and Graphical Statistics* 5(1), 26–56.
- Nason, G. P. and B. W. Silverman (1995). The stationary wavelet transform and some statistical applications. See Antoniadis and Oppenheim (1995), pp. 281–300.
- Nuri, W. A. and L. J. Herbst (1969). Fourier methods in the study of variance fluctuations in time series analysis. *Technometrics* 11(1), 103–113.
- Ogden, R. T. and E. Parzen (1996). Change-point approach to data analytic wavelet thresholding. *Statistics and Computing* 6(2), 93–99.
- Percival, D. B. (1995). On estimation of the wavelet variance. *Biometrika* 82(3), 619–631.
- Percival, D. B. and P. Guttorp (1994). Long-memory processes, the Allan variance and wavelets. In E. Foufoula-Georgiou and P. Kumar (Eds.), *Wavelets in Geophysics*, Volume 4 of *Wavelet Analysis and its Applications*, pp. 325–344. San Diego: Academic Press, Inc.
- Percival, D. B. and H. O. Mofjeld (1997). Analysis of subtidal coastal sea level fluctuations

- using wavelets. *Journal of the American Statistical Association* 92(439), 868–880.
- Popper, W. (1951). *The Cairo Nilometer*, Volume 12 of *Publications in Semitic Philology*. Berkeley: University of California Press.
- Shensa, M. J. (1992). The discrete wavelet transform: Wedding the à trous and Mallat algorithms. *IEEE Transactions on Signal Processing* 40(10), 2464–2482.
- Srivastava, M. S. (1993). Comparison of CUSUM and EWMA procedures for detecting a shift in the mean or an increase in the variance. *Journal of Applied Statistical Science* 1(4), 445–468.
- Stephens, M. A. (1974). EDF statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association* 69(347), 730–737.
- Tewfik, A. H. and M. Kim (1992). Correlation structure of the discrete wavelet coefficients of fractional Brownian motion. *IEEE Transactions on Information Theory* 38(2), 904–909.
- Toussoun, O. (1925). Mémoire sur l’histoire du nil. In *Mémoires a l’Institut d’Egypte*, Volume 18, pp. 366–404.
- Tsay, R. S. (1988). Outliers, level shifts, and variance changes in time series. *Journal of Forecasting* 7, 1–20.
- Wang, Y. (1995). Jump and sharp cusp detection by wavelets. *Biometrika* 82(2), 385–397.
- Whitcher, B. (1998). *Assessing Nonstationary Time Series Using Wavelets*. Ph. D. thesis, University of Washington.
- Wichern, D. W., R. B. Miller, and D.-A. Hsu (1976). Changes of variance in first-order autoregressive time series models – with an application. *Applied Statistics* 25(3), 248–256.
- Wornell, G. W. (1996). *Signal Processing with Fractals: A Wavelet Based Approach*. New Jersey: Prentice Hall.

α	Sample size								
	8	16	32	64	128	256	512	1024	∞
0.10	1.109	1.135	1.157	1.182	1.193	1.197	1.206	1.209	1.224
SE	0.003	0.003	0.003	0.003	0.003	0.003	0.003	0.003	
0.05	1.232	1.265	1.293	1.313	1.326	1.329	1.345	1.341	1.358
SE	0.004	0.004	0.004	0.004	0.004	0.004	0.004	0.004	
0.01	1.459	1.508	1.553	1.584	1.596	1.596	1.630	1.617	1.628
SE	0.007	0.008	0.008	0.009	0.008	0.010	0.008	0.007	

Table 1: Monte Carlo critical values for the test statistic $(N/2)^{\frac{1}{2}} D$, using the Haar wavelet filter, for a level α test. These values are based upon 10,000 replicates. The standard error (SE) is provided for each estimate, and was computed via $SE = \{\alpha(1 - \alpha)/(10,000f^2)\}^{\frac{1}{2}}$ where f is the histogram estimate of the density at the $(1 - \alpha)$ th quantile using a bandwidth of 0.01 (Inclán and Tiao 1994). Quantiles of a Brownian bridge process are given at the far right for comparison.

Scale	Haar	D(4)	LA(8)
1	-0.0626	-0.0797	-0.0767
2	-0.0947	-0.1320	-0.1356
4	-0.1133	-0.1511	-0.1501
8	-0.1211	-0.1559	-0.1535

Table 2: Lag one autocorrelations for wavelet coefficients of scales 1, 2, 4, and 8 for an FD process with $d = 0.45$ using the Haar, D(4) and LA(8) wavelet filters.

Scale	D	10% critical level	5% critical level	1% critical level
1	0.1559	0.0945	0.1051	0.1262
2	0.1754	0.1320	0.1469	0.1765
4	0.1000	0.1855	0.2068	0.2474
8	0.2313	0.2572	0.2864	0.3436

Table 3: Results of testing the Nile River water levels for homogeneity of variance ($N = 663$) using the Haar wavelet filter with Monte Carlo critical values. As shown in the table, the test statistic at scale 1 is significant at the 1% level, and the test statistic at scale 2 is significant at the 5% level.

List of Figures

1. Estimated Haar wavelet variances for the Nile River minimum water levels before and after the year 721 AD, along with 95% confidence intervals based upon a chi-square approximation given in Percival (1995).
2. Estimated locations of variance change points for FD processes ($N = 663$, $d = 0.4$) using $\widetilde{W}_{j,t}$, with known variance change at $t = 100$ (dotted line). From bottom to top, the variance ratio between the first 100 and remaining observations is 1.5:1, 2:1, 3:1 and 4:1, respectively.
3. Nile River minimum water levels (upper panel) and normalized cumulative sum of squares using $\widetilde{W}_{j,t}$ based upon D(4) wavelet filter for the Nile River minimum water levels (lower two panels). The vertical dotted line marks the year 715 AD, after which a nilometer on Roda Island was used to record the water levels. (The source of the Nile River is Toussoun (1925). The data can be obtained via the World Wide Web at <http://lib.stat.cmu.edu/S/> under the title ‘beran’. This is the address for StatLib, a statistical archive maintained by Carnegie Mellon University.)




Figure 1:




Figure 2:




Figure 3: