Multiscale Detection and Location of Multiple Variance Changes in the Presence of Long Memory

Brandon Whitcher
EURANDOM
whitcher@eurandom.tue.nl

Peter Guttorp
University of Washington
peter@stat.washington.edu

Donald B. Percival
University of Washington
dbp@apl.washington.edu

July 1, 1999

Abstract

Procedures for detecting change points in sequences of correlated observations (e.g., time series) can help elucidate their complicated structure. Current literature on the detection of multiple change points emphasizes the analysis of sequences of independent random variables. We address the problem of an unknown number of variance changes in the presence of long-range dependence (e.g., long memory processes). Our results are also applicable to time series whose spectrum slowly varies across octave bands. An iterated cumulative sum of squares procedure is introduced in order to look at the multiscale stationarity of a time series; that is, the variance structure of the wavelet coefficients on a scale by scale basis. The discrete wavelet transform enables us to analyze a given time series on a series of physical scales. The result is a partitioning of the wavelet coefficients into locally stationary regions. Simulations are performed to validate the ability of this procedure to detect and locate multiple variance changes. A 'time' series of vertical ocean shear measurements is also analyzed, where a variety of nonstationary features are identified.

Some key words: Cumulative sum of squares; Discrete wavelet transform; Homogeneity of variance; Maximal overlap discrete wavelet transform; Vertical ocean shear.

1 Introduction

In many fields, such as the physical sciences and economics, the hypothesis that a process is composed of many features which occur at different scales is quite natural. Recently, attention has been given to identifying and modelling so-called long memory processes. A common model for such processes is that the observations $Y_0, \ldots, Y_{N-1}$ constitute one portion 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 the long memory parameter $|d| < \frac{1}{2}$, 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 $|f| \leq \frac{1}{2}$. When $0 < d < \frac{1}{2}$, the SDF has an asymptote at zero, 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 practice, one may question if the process are truly stationary, or composed of several stationary segments. We prefer to view this problem by considering a time series with an unknown number of variance change points. A number of methods have appeared in the literature for investigating the presence of multiple change points. Hinkley (1971) studied the use of cumulative sum tests for detecting a change point and then estimating its location, which are adaptable to the presence of multiple change points through a sequential algorithm. Haccou and Meelis (1988) proposed a hierarchical likelihood ratio test for determining the number of change points in a sequence of exponentially distributed random variables, and Meelis *et al.* (1991) generalized the well-known test by Pettitt (1979) into a Kruskal–Wallis test for determining the number of change points in a sequence of inde-

pendent random variables. Inclán and Tiao (1994) investigated the detection and location 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. Chen and Gupta (1997) utilized the Schwartz information criterion in order to both detect and locate multiple changes of variance. Recently, Khalil and Duchêne (1999) compared autoregressive and multiscale approaches to identify multiple events in piecewise stationary time series with application to uterine electromyogram signals.

The discrete wavelet transform (DWT) decomposes a process into different scales, or bands of frequency. The DWT has already proven useful for investigating various 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, Ogden and Parzen (1996) used wavelet coefficients to develop data-dependent thresholds for removing noise from a signal and Whitcher *et al.* (1998) investigated a single change of variance in long memory processes. 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 times. For FD and related long memory processes, the wavelet coefficients for a given scale are approximately uncorrelated; see, Tewfik and Kim (1992), McCoy and Walden (1996) and Whitcher *et al.* (1998). We propose to test for multiple variance change points in FD processes using an iterated cumulative sum of squares statistic applied to the output from the DWT. By testing the output from the DWT we also gain the ability to identify the scale at which the inhomogeneity occurs. Using a variation of the DWT, the 'non-decimated' DWT, we can estimate the times at which the variance change points occur.

Section 2 provides a brief description of the DWT and 'non-decimated' DWT. Section 3 establishes our procedure for testing multiple variance change points. The cumulative sum

of squares test statistic is defined and a bisection algorithm is outlined in order to recursively test a vector of coefficients. Simulation results are provided to determine the empirical size and power of our procedure. Using an auxiliary testing procedure, we also provide a method for locating multiple variance change points. We apply this methodology to a series of vertical ocean shear measurements. Not only does the procedure isolate two visually obvious nonstationary segments, but also a much more subtle region.

2 Discrete Wavelet Transforms

Let $h_1 = \{h_{1,0}, \ldots, h_{1,L-1}, 0 \ldots, 0\}$ 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

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

be the discrete Fourier transform (DFT) of h_1. Let $g_1 = \{g_{1,0}, \ldots, g_{1,L-1}, 0, \ldots, 0\}$ be the zero padded scaling filter coefficients, defined via $g_{1,l} = (-1)^{l+1} h_{1,L-1-l}$ for $l = 0, \ldots, L-1$, 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

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

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 $Y_0, \ldots, Y_{N-1}$ 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

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

where

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

The $W_{j,t}$ coefficients are associated with changes on a scale of length τ_j and are obtained by subsampling every 2^jth 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). 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, 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 Multiple Variance Changes

If $Y_0, \ldots, Y_{N-1}$ 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 multiple variance change points using wavelet coefficients for FD processes.

3.1 The Test Statistic

Let $X_0, \ldots, X_{N-1}$ be a time series that can be regarded as a sequence of independent Gaussian (normal) random variables with zero means and variances $\sigma_0^2, \ldots, \sigma_{N-1}^2$. We would like to test the hypothesis $H_0 : \sigma_0^2 = \cdots = \sigma_{N-1}^2$. A test statistic that can discriminate between this null hypothesis and a variety of alternative hypotheses (such as $H_1 : \sigma_0^2 = \cdots = \sigma_k^2 \neq \sigma_{k+1}^2 = \cdots = \sigma_{N-1}^2$, 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 *et al.* (1975), Hsu (1977) and Inclán and Tiao (1994). To define D, let

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

The desired statistic is given by $D \equiv \max(D^+, D^-)$. Critical levels for D under the null hypothesis can be readily obtained through Monte Carlo simulations for arbitrary sample size. Inclán and Tiao (1994) established the asymptotic distribution of D and also provided a relationship with the usual F-statistic for testing the equality of variances between two independent samples and the likelihood ratio.

3.2 Iterated Cumulative Sums of Squares Algorithm

In practice, a given time series may exhibit more than one change in variance. A natural approach is to test the entire series first, split at a detected change point and repeat until no change points are found. This is known as a 'binary segmentation' procedure studied by Vostrikova (1981), who proved its consistency. Inclán and Tiao (1994) and Chen and Gupta (1997) have both recently used this in order to detect and locate variance change points. The test statistic used by Chen and Gupta (1997) is based on the Schwarz information criterion (SIC). All subsequent simulation studies were attempted to be duplicated using the SIC procedure, but we were unable to replicate the empirical size of the test in both Lisp and C even after correspondence with one of the authors.

We propose to use the iterated cumulative sum of squares (CSS) algorithm to test for multiple variance changes from the output of the DWT – in effect partitioning the series into stationary segments. Thus, each series of wavelet coefficients are put through the iterated CSS algorithm producing a multiscale analysis of the original series. In order to reduce the computational time we use the asymptotic approximation for the critical values of D.

This has been shown to be an adequate approximation for sample sizes of 128 or greater (Whitcher *et al.* 1998). For a time series $Y_0, \ldots, Y_{N-1}$, the iterated CSS algorithm proceeds as follows (Inclán and Tiao 1994):

1. Compute the partial DWT (order J) of $Y_0, \ldots, Y_{N-1}$, as defined in Section 2 using a Daubechies family wavelet filter;

2. Discard all coefficients on each scale that make explicit use of the periodic boundary conditions;

3. Determine the test statistic D, via the equations in Section 3.1, and record the point k_1 at which D is attained. If D exceeds its critical value for a given level of significance α, then proceed to the next step. If D is less than the critical value, the algorithm terminates.

4. Determine the test statistic D for the new time series $Y_0, \ldots, Y_{k_1-1}$. If D exceeds its critical value, then repeat this step until D is less than its critical value.

5. Determine the test statistic D for the new time series $Y_{k_1}, \ldots, Y_{N-1}$. Repeat this step until D is less than its critical value.

6. Go through the potential change points as outlined in the following paragraph.

Inclán and Tiao (1994) included an additional step when detecting multiple variance change points. After the bisection algorithm had terminated, each potential change point was tested again using only those observations between its two adjacent change points. For example, a vector of length 128 containing potential change points at 26, 69, and 108, would again test 26 using only observations $1, \ldots, 69$, test 69 using observations $26, \ldots, 108$ and test 108 using observations $69, \ldots, 128$. This was to compensate for an apparent overestimation

of the number of variance change points. Simulations were run both with and without this additional procedure. The rejection rates for the first two scales were found to change up to 4% for low variance ratios and up to 1% for larger variance ratios. All tables using the iterated CSS algorithm include this extra step.

3.3 Empirical Size and Power

Whitcher *et al.* (1998) established the empirical size of the cumulative sum of squares test under the alternative hypothesis of a single variance change. Since the iterated procedure simply utilizes this specific hypothesis test recursively throughout the separate scales of the DWT, the empirical size for the iterated CSS algorithm follows immediately.

The procedure outlined in Section 3.2 was repeated a large number of times for a specific sample size $N = 656$, significance level $\alpha = 0.05$ and long memory parameter $d = 0.40$, with a partial DWT of order $J = 4$. A vector of independent Gaussian random variables was added to the first 100 observations of the FD processes. Instead of adjusting the long memory parameter, the variance of the first 100 observations was adjusted – producing variance ratios between the first 100 and subsequent observations of $\Delta \in \{1.5, 2, 3, 4\}$.

Table 1 displays simulation results for the iterated CSS method when detecting one unknown variance change point for the Haar, D(4) and LA(8) wavelet filters and scales $\tau_j = 1$, 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). We see the iterated CSS procedure does quite well at locating the single variance change point for all variance ratios. With a ratio of $\Delta = 2$ or greater, it errs only towards multiple change points – always indicating at least one change point in the first few scales. For larger variance ratios ($\Delta \geq 3$) the procedure produces rejection rates around 90% or greater in the first two scales, and it errs on the side of three or more change-points with

		Haar			D(4)			LA(8)	
	Level	0	1	≥ 2	0	1	≥ 2	0	1	≥ 2
$\Delta = 1.5$	1	9.5	**85.2**	5.3	7.6	**87.8**	4.6	7.3	**88.8**	3.9
	2	58.9	**39.8**	1.2	58.9	**39.7**	1.4	60.4	**38.5**	1.2
	3	87.5	**12.2**	0.3	88.2	**11.6**	0.2	90.3	**9.5**	0.2
	4	95.1	**4.9**	0.0	95.2	**4.8**	0.0	95.9	**4.1**	0.0
$\Delta = 2$	1	0.1	**93.0**	6.9	0.1	**93.5**	6.4	0.1	**94.2**	5.7
	2	17.7	**79.6**	2.8	17.9	**79.5**	2.6	20.6	**77.1**	2.3
	3	69.2	**30.1**	0.7	71.2	**28.2**	0.6	77.2	**22.4**	0.4
	4	90.8	**9.2**	0.0	91.2	**8.8**	0.0	93.5	**6.5**	0.0
$\Delta = 3$	1	0.0	**92.9**	7.1	0.0	**93.6**	6.4	0.0	**93.8**	6.2
	2	1.3	**95.4**	3.3	1.2	**95.5**	3.3	1.9	**95.0**	3.2
	3	33.8	**64.7**	1.5	38.3	**60.4**	1.2	49.4	**49.6**	1.0
	4	79.4	**20.6**	0.0	80.9	**19.1**	0.0	87.2	**12.3**	0.0
$\Delta = 4$	1	0.0	**92.9**	7.1	0.0	**92.7**	7.3	0.0	**93.2**	6.8
	2	0.1	**96.3**	3.6	0.1	**96.3**	3.6	0.2	**96.5**	3.3
	3	16.6	**81.5**	1.9	19.1	**79.1**	1.8	29.1	**69.4**	1.5
	4	66.0	**34.0**	0.0	69.1	**30.9**	0.0	79.4	**20.6**	0.0

Table 1: Empirical power of iterated CSS algorithm for FD processes ($N = 512, d = 0.4$) with one variance change ($k = 100$). Variance ratios are given by Δ.

greater frequency.

Table 2 displays simulation results for the iterated CSS when detecting two unknown variance change points. Again, all tests were performed at the $\alpha = 0.05$ level. Gaussian random variables, of length 100, were added to the middle of the series creating two variance changes at $k_1 = 250$ and $k_2 = 350$. The iterated CSS method once again performs quite well for small variance ratios $\Delta = 1.5$, with a slight increase in power as the wavelet filter increases in length. For larger variance ratios, the first scale gives a maximum rejection rate of 94% and then hovers around 90% for very large Δ. All errors in the first scale, for higher variance ratios, are towards overestimating the number of variance changes. The second scale, which exhibits almost no power for smaller variance ratios, rapidly approaches the 90–95% range

		Haar				**D(4)**				**LA(8)**			
	Level	0	1	2	≥ 3	0	1	2	≥ 3	0	1	2	≥ 3
$\Delta = 1.5$	1	14.5	9.3	**71.8**	4.3	11.7	10.4	**73.6**	4.2	11.2	11.4	**74.2**	3.2
	2	67.4	20.8	**11.7**	0.1	67.3	22.6	**10.0**	0.1	67.7	23.8	**8.5**	0.1
	3	88.7	10.1	**1.1**	0.0	88.8	10.6	**0.6**	0.0	90.4	9.2	**0.4**	0.0
	4	94.6	5.4	**0.0**	0.0	100.0	0.0	**0.0**	0.0	100.0	0.0	**0.0**	0.0
$\Delta = 2$	1	0.1	0.2	**91.8**	7.9	0.0	0.2	**92.4**	7.3	0.0	0.3	**94.0**	5.8
	2	26.7	17.2	**55.4**	0.7	26.2	22.1	**51.1**	0.6	27.2	26.4	**45.8**	0.5
	3	75.8	18.5	**5.6**	0.0	77.3	19.0	**3.7**	0.0	78.8	18.3	**2.9**	0.0
	4	91.5	8.5	**0.0**	0.0	100.0	0.0	**0.0**	0.0	100.0	0.0	**0.0**	0.0
$\Delta = 3$	1	0.0	0.0	**90.4**	9.6	0.0	0.0	**91.4**	8.6	0.0	0.0	**92.5**	7.5
	2	1.6	2.2	**93.8**	2.3	1.6	4.0	**92.5**	1.9	2.2	5.5	**90.4**	1.9
	3	48.0	24.0	**27.8**	0.2	49.7	29.8	**20.3**	0.2	55.2	29.4	**15.3**	0.1
	4	83.5	16.5	**0.0**	0.0	100.0	0.0	**0.0**	0.0	100.0	0.0	**0.0**	0.0
$\Delta = 4$	1	0.0	0.0	**89.5**	10.5	0.0	0.0	**90.5**	9.5	0.0	0.0	**91.4**	8.6
	2	0.1	0.2	**96.5**	3.2	0.1	0.7	**96.6**	2.6	0.1	0.9	**96.3**	2.8
	3	26.3	18.8	**54.4**	0.6	29.6	28.0	**42.1**	0.4	34.9	28.9	**35.7**	0.5
	4	74.0	26.0	**0.0**	0.0	100.0	0.0	**0.0**	0.0	100.0	0.0	**0.0**	0.0

Table 2: Empirical power of the iterated CSS algorithm for FD processes ($N = 512, d = 0.4$) with two variance changes ($k_1 = 250, k_2 = 350$). Variance ratios are given by Δ.

for $\Delta \geq 3$ and errs primarily towards overestimating the number of variance changes also. The 100% rejection rates for the D(4) and LA(8) wavelet filters in the fourth scale occurs because of a reduction, due to boundary affects, in the number of wavelet coefficients below a minimum established threshold.

3.4 Locating Multiple Variance Changes

As discussed in Whitcher *et al.* (1998), the DWT suffers from a loss of resolution by down-sampling at each level of the transform. An alternative test statistic $\widetilde{D}$, using the MODWT, is utilized to estimate the location of multiple variance changes. The procedure involves much more bookkeeping than in the single variance change scenario, but is easily manageable. For each iteration of the algorithm, estimated locations of the variance change points for both the DWT and MODWT are recorded. The DWT estimates are used to test for homogeneity of variance and the MODWT estimates are used to determine the location of the variance change. Obviously, the estimated location of the variance change is discarded if the test statistic is found not to be significant.

Figures 1–4 displays the estimated location of variance change for various FD processes with one change of variance. We see more and more of the area of the histogram centered at $k = 100$ as the variance ratio increases. There also appears to be a small percentage of rejections to the right of the main peak across all levels and wavelet filters. This is to be expected, since we are not forcing the testing procedure to stop at only one change of variance. The small percentage of second or third variance changes (cf. Table 1) in the same scale appear as an increase in the right tails of these histograms. With this feature in mind, the procedure still performs quite well when the variance ratio is relatively large ($\Delta \geq 2$), especially in the first two scales. The third and fourth scales are quite spread out and not recommended for estimating simple variance change points in practice.

Figures 5–8 display the estimated location of variance change for various FD processes with two variance changes. For small variance ratios ($\Delta = 1.5$) the iterated CSS procedure does a good job of locating the variance change points in the first scale, with mixed results for the second scale. As before, we do not expect much information to come from looking at higher scales. Although, as the magnitude of the variance ratio increases the higher scales ($j = 3, 4$) do exhibit structure similar to the first two. Regardless, we shall strictly use the first two scales for inference in the future. With respect to the first two scales, as the variance ratio increases to, say, $\Delta = 3$ or 4, then the bimodality is readily apparent. As was the case for a single variance change, the longer wavelet filters give a slightly more spread out distribution for the locations of the variance changes. To be more precise, the estimated locations appear to be skewed to the right at $k_1 = 250$ and $k_2 = 350$ for the D(4) and LA(8) wavelet filters, especially in the second scale.

4 Vertical Ocean Shear Measurements

Percival and Guttorp (1994) analyzed a set of vertical ocean shear measurements using the wavelet variance. The data were collected by dropping a probe into the ocean which records the water velocity every 0.1 meter as it descends. Hence, the 'time' index is really depth (in meters). The shear measurements (in s^{-1}) are obtained by taking a first difference of the velocity readings over 10 meter intervals and applying a low-pass filter to the difference readings. Figure 9 shows all 6875 observations available for analysis. We see two sections of greater variability, one around 450m and the other around 1000m, with a fairly stationary section in between. Percival and Guttorp (1994) commented on this fact and only looked at 4096 observations ranging from 489.5m to 899.0m in their paper.

The influence of the ends of the time series (i.e., the observations outside the vertical

dotted lines in Figure 9) is most evident when comparing its wavelet variance to the wavelet variance between the middle 4096 observations; see Figure 10. The bursts of increased variability observed in the first 5 scales make a significant contribution to the wavelet variance. For those scales, the confidence intervals do not overlap between the full and truncated time series, whereas the confidence intervals do overlap for all subsequent scales. This feature hints at a possible heterogeneity of variance in the first 5 scales.

Figure 11 shows the MODWT wavelet coefficients for the first five scales of the vertical ocean shear measurements. The vertical dotted lines are the estimated locations of variance change points using the DWT to test and the MODWT to locate with asymptotic critical values ($\alpha = 0.05$). The procedure does a good job of isolating the two regions of increased variability at 450m and 1000m in each scale, except for the second scale. There, the first burst has been 'picked apart' by the procedure with 10 distinct stationary regions. This does not seem appropriate and it is unclear why this only occurred on the second scale when the third scale appears to be similar in changing variability with time. Besides the two obvious regions of increased variability, there appears to be a third burst around 800m. It is present, to differing degrees, in the first four scales whereas most other bursts disappear after the first and second scale. This is a much more subtle type of nonstationarity, compared to the obvious bursts at 450m and 1000m, and not particularly visible in the original time series with the naked eye.

5 Conclusions

We have presented an iterated CSS algorithm for detecting and locating multiple variance changes in time series with long-range dependence. The first scale of wavelet coefficients is quite powerful for detecting single or multiple variance change points when the variance

ratio is a factor of 2 or greater. The second scale is also equally powerful, but when the variance ratio is a factor of 3 or greater. This procedure also performs well at locating single or multiple variance change points using the auxiliary test statistic computed via the MODWT.

Currently, the procedure for detecting and locating multiple variance changes via the DWT is in its infancy. More work is needed in order to refine the procedure and further investigate its properties. Areas for future research include obtaining exact critical values for the DWT test statistic and adapting the MODWT test statistic to not only locate but also test for a variance change point. Given the ability of the DWT to remove heavy amounts of autocorrelation in time series, this method has wide application in many fields. The point being that this test can handle high amounts of autocorrelation, as found in long memory processes, through the fact that limited assumptions are made with respect to the underlying spectrum of the observed time series.

References

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

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 B 37*, 149–163.

Chen, J. and A. K. Gupta (1997). Testing and locating variance changepoints with application to stock prices. *Journal of the American Statistical Association 92*(438), 739–747.

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.

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.

Haccou, P. and E. Meelis (1988). Testing for the number of change points in a sequence of exponential random variables. *Journal of Statistical Computation and Simulation 30*, 285–298.

Hinkley, D. V. (1971). Inference about the change point from cumulative sum tests. *Biometrika 58*, 509–525.

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.

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.

Khalil, M. and J. Duchêne (1999). Detection and classification of multiple events in piecewise stationary signals: Comparison between autoregressive and multiscale approaches. *Signal Processing 75*(3), 239–251.

Mallat, S. (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.

Meelis, E., M. Bressers, and P. Haccou (1991). Non-parametric testing for the number of

change points in a sequence of independent random variables. *Journal of Statistical Computation and Simulation 39*, 129–137.

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. 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.

Pettitt, N. A. (1979). A non-parametric approach to the change point problem. *Applied Statistics 28*, 126–135.

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.

Vostrikova, L. J. (1981). Detecting 'disorder' in multidimensional random processes. *Soviet Mathematics Doklady 24*, 55–59.

Wang, Y. (1995). Jump and sharp cusp detection by wavelets. *Biometrika 82*(2), 385–397.

Whitcher, B., S. D. Byers, P. Guttorp, and D. B. Percival (1998). Testing for homogeneity of variance in time series: Long memory, wavelets and the Nile River. Under revision for the *Journal of the American Statistical Association*. Available for download at `http://www.eurandom.tue.nl/whitcher/papers/`.

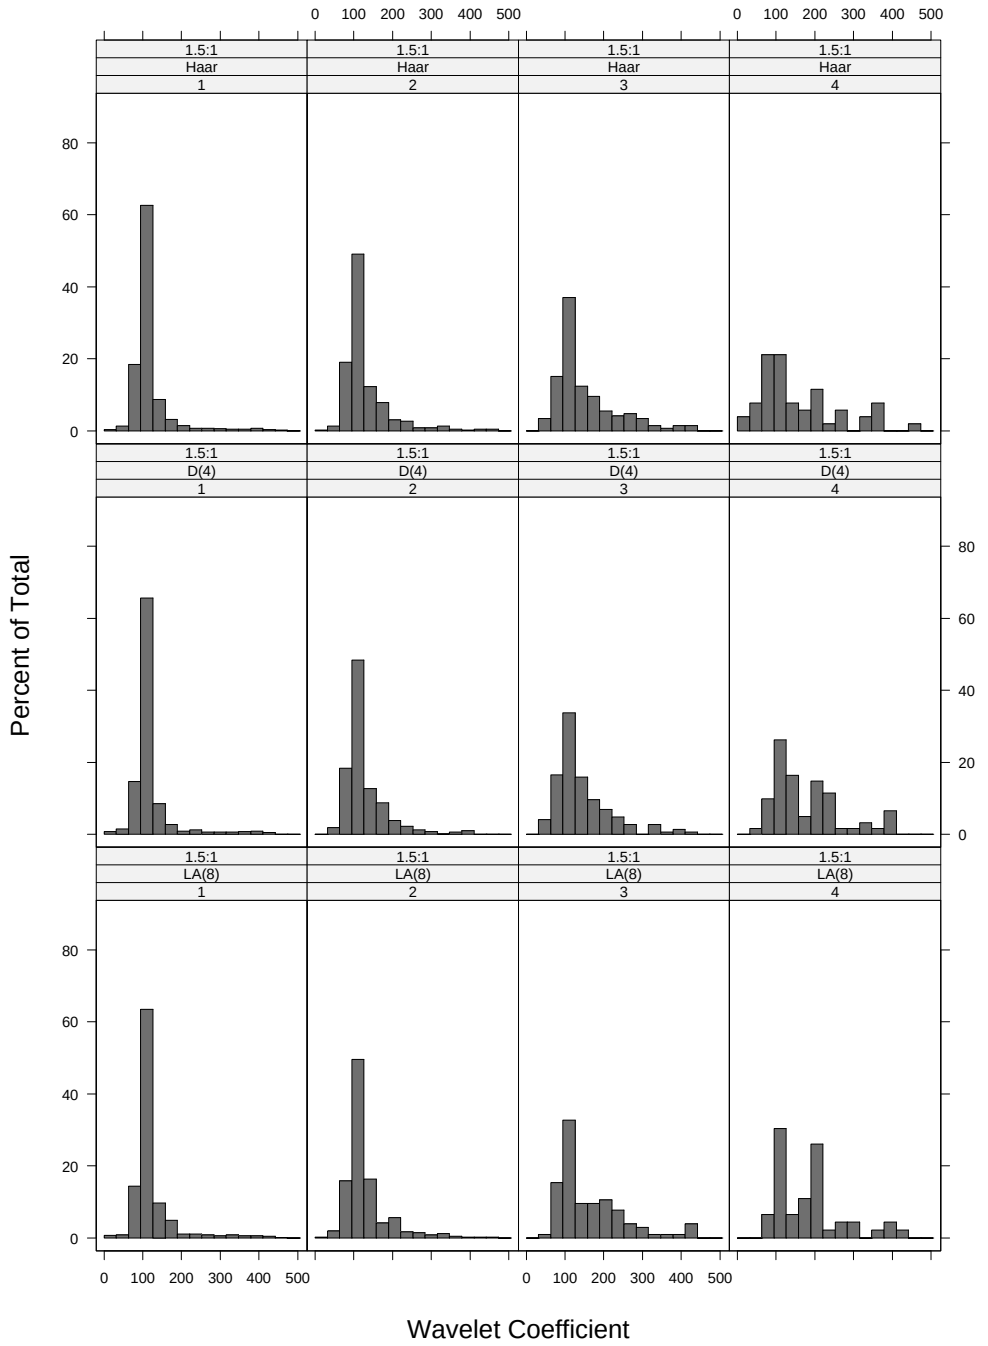

Figure 1: Estimated locations of a single variance change ($k = 100$) for FD processes ($N = 656, d = 0.4$) using the iterated CSS procedure and MODWT. The rows denote the various wavelet filters, the columns denote the levels of the MODWT and the variance ratio is $\Delta = 1.5$.

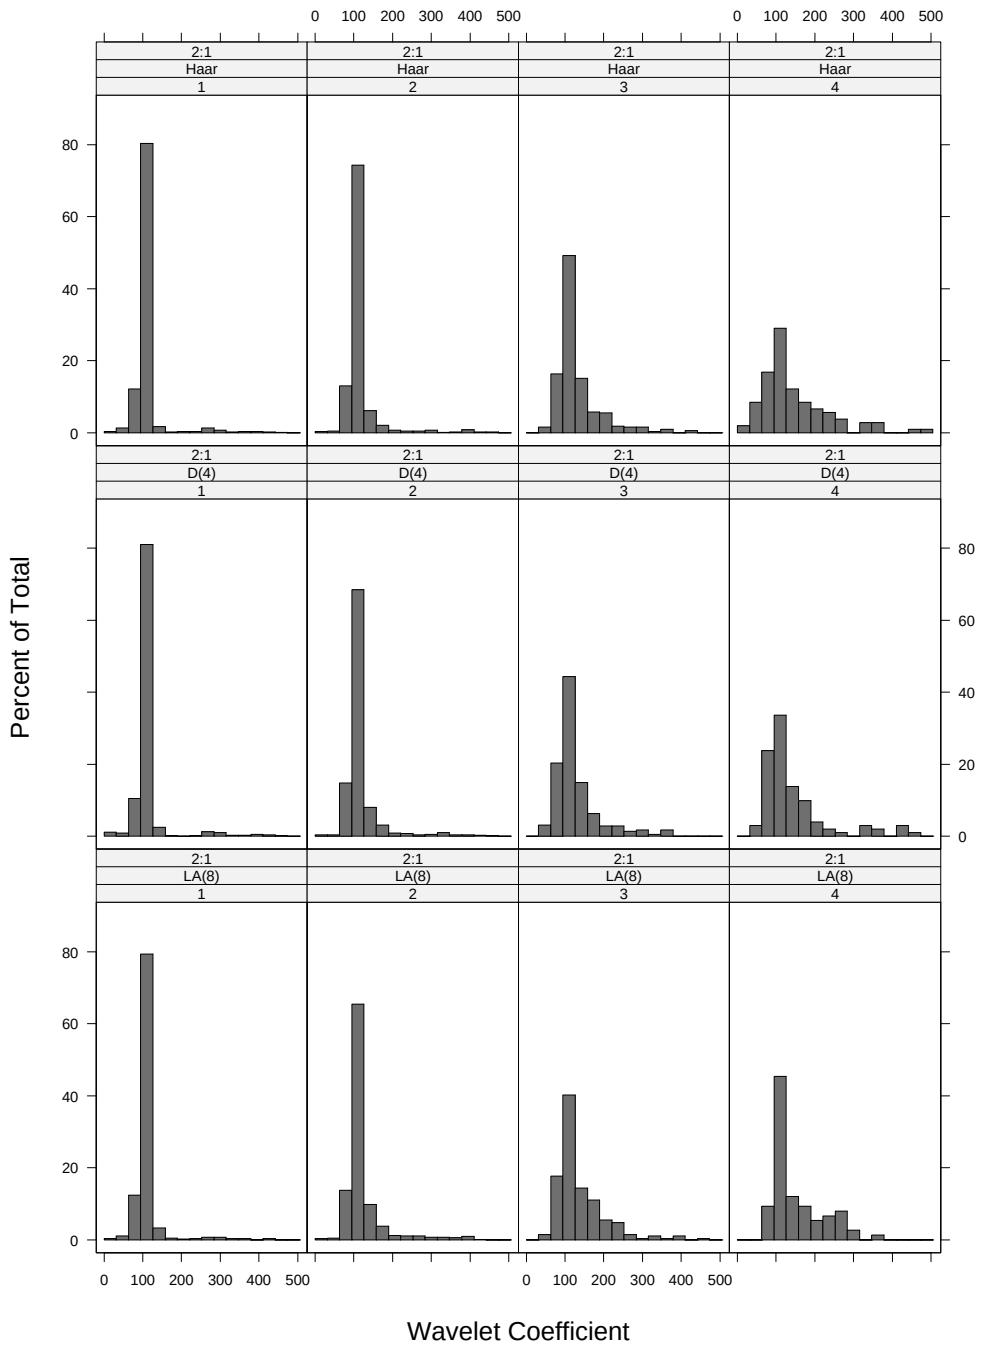

Figure 2: Same as Figure 1 with variance ratio $\Delta = 2$.

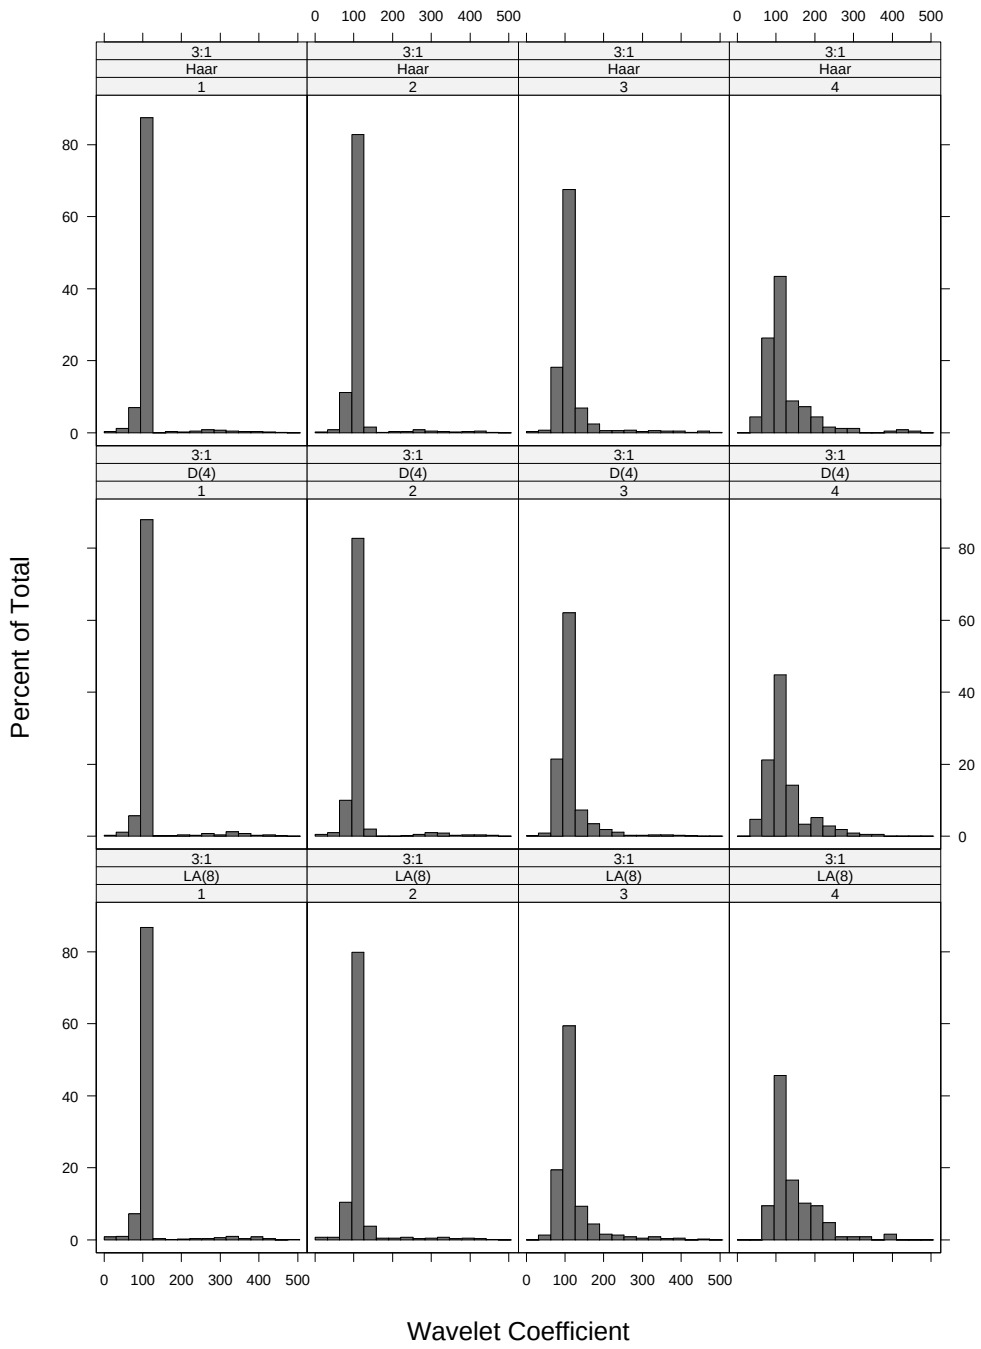

Figure 3: Same as Figure 1 with variance ratio $\Delta = 3$.

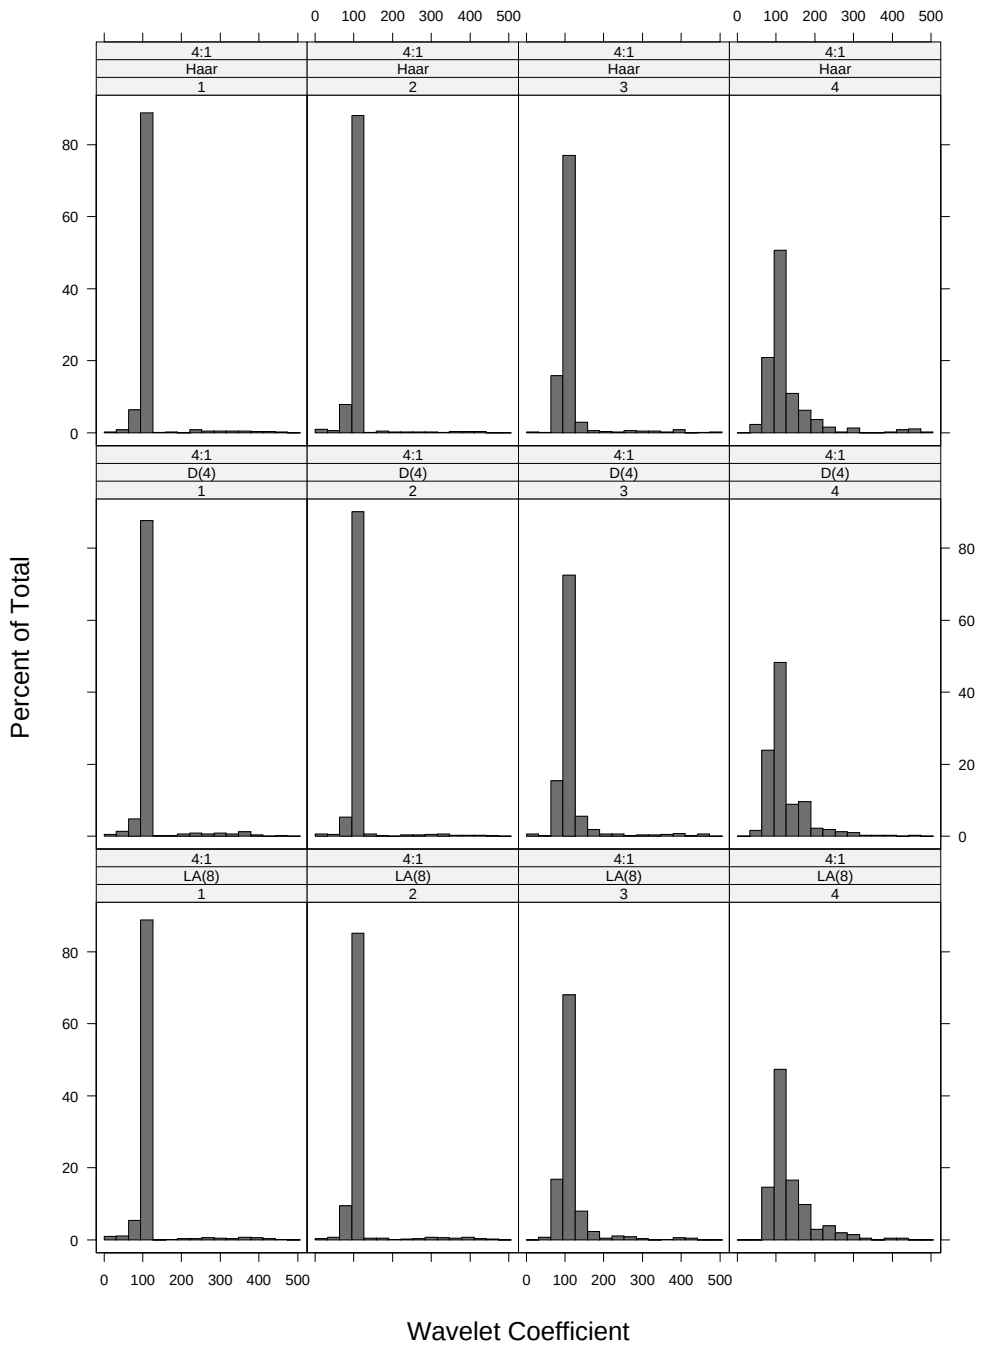

Figure 4: Same as Figure 1 with variance ratio $\Delta = 4$.

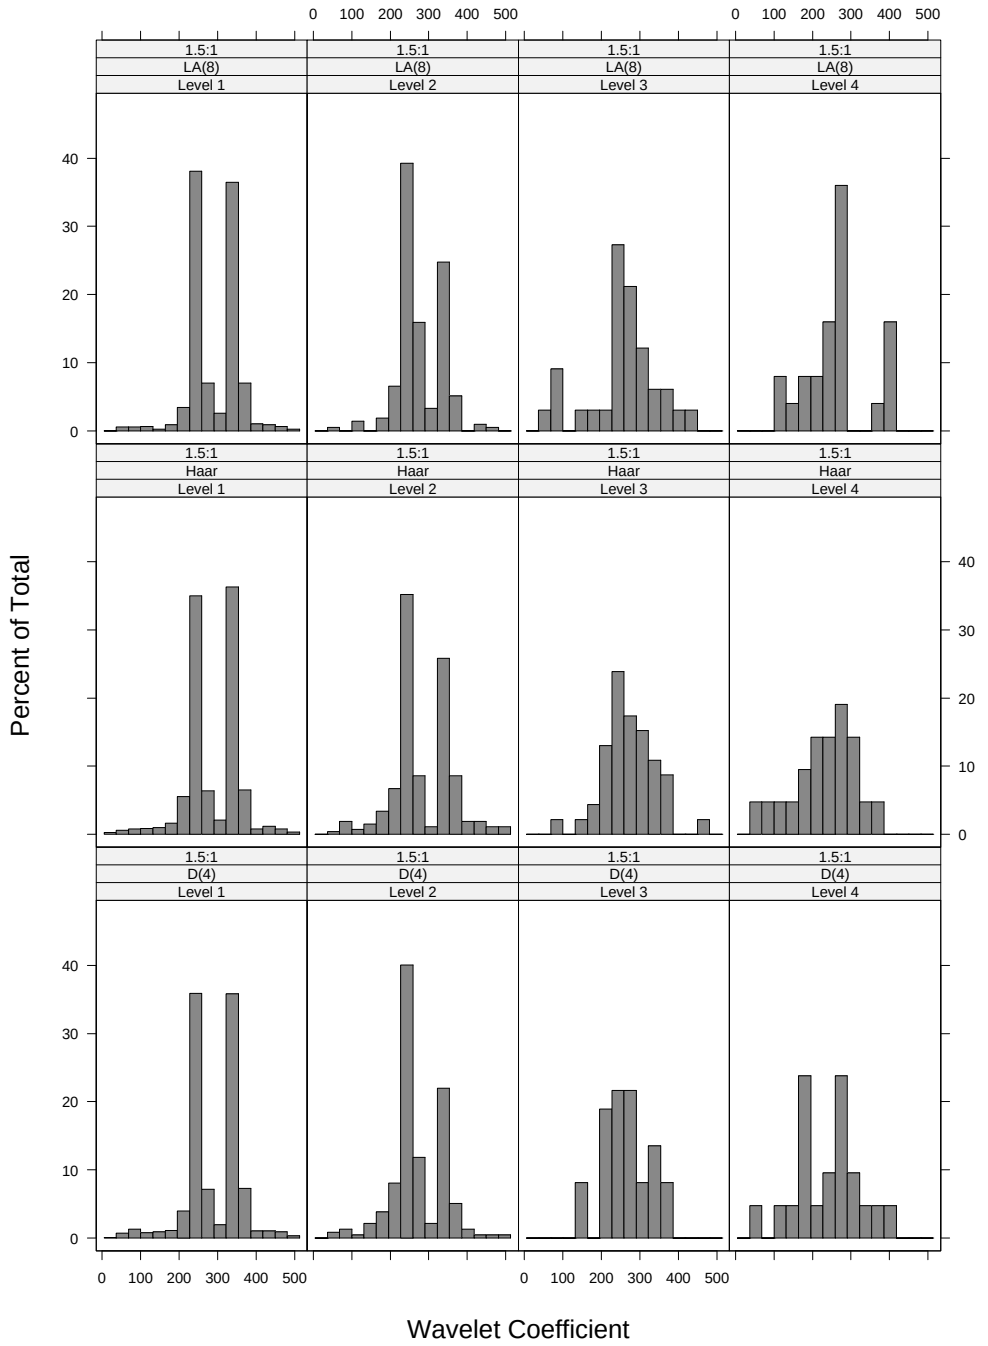

Figure 5: Estimated locations of two variance changes ($k_1 = 251, k_2 = 350$) for FD processes ($N = 656, d = 0.4$) using the iterated CSS procedure and MODWT. The rows denote the various wavelet filters, the columns denote the levels of the MODWT and the variance ratio is $\Delta = 1.5$.

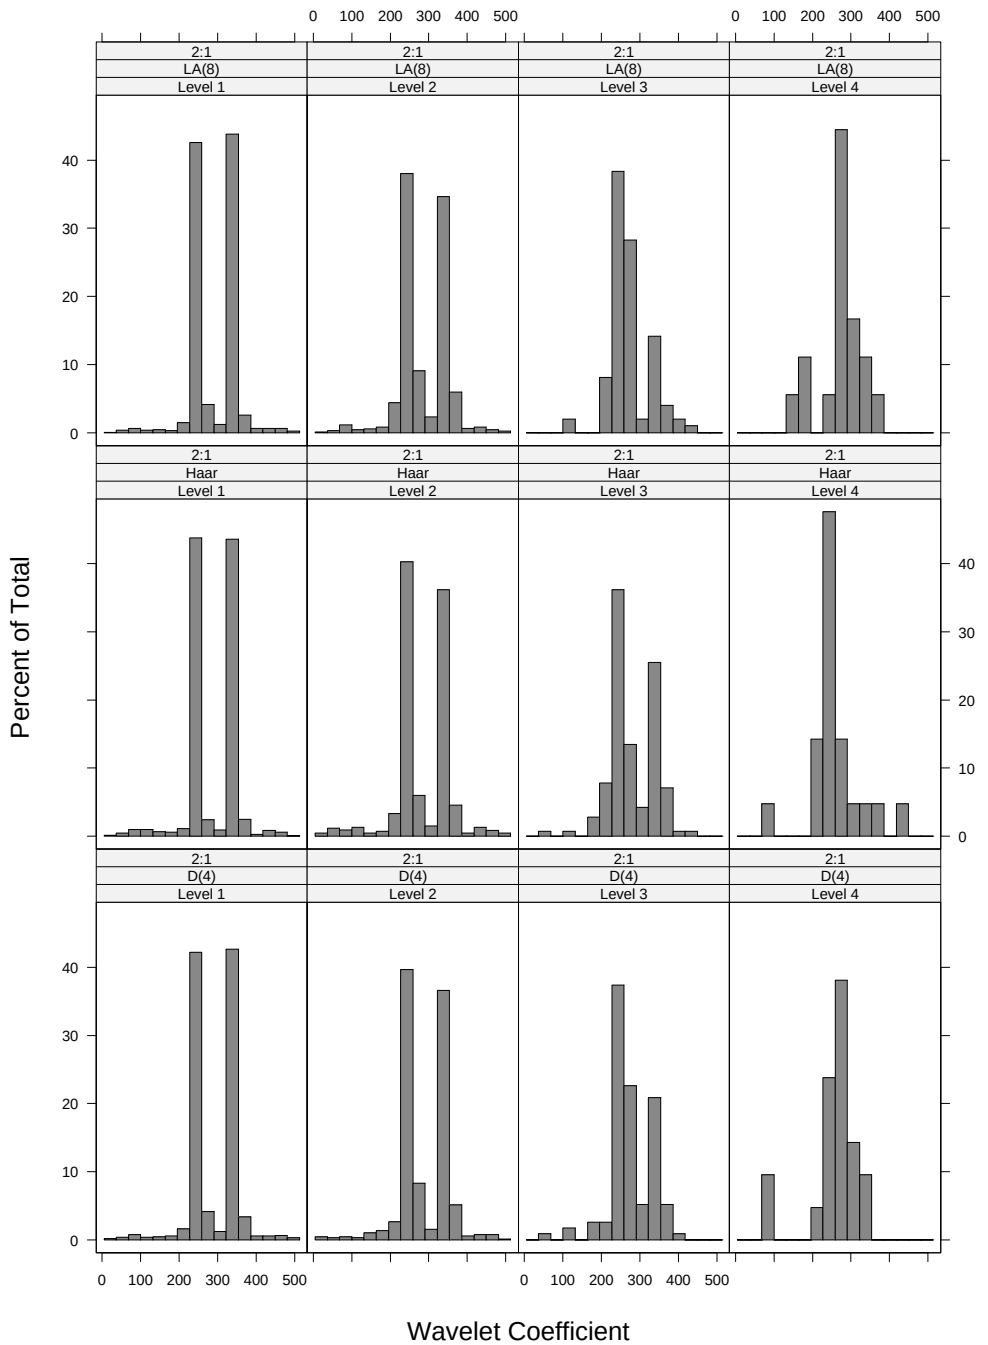

Figure 6: Same as Figure 5 with variance ratio $\Delta = 2$.

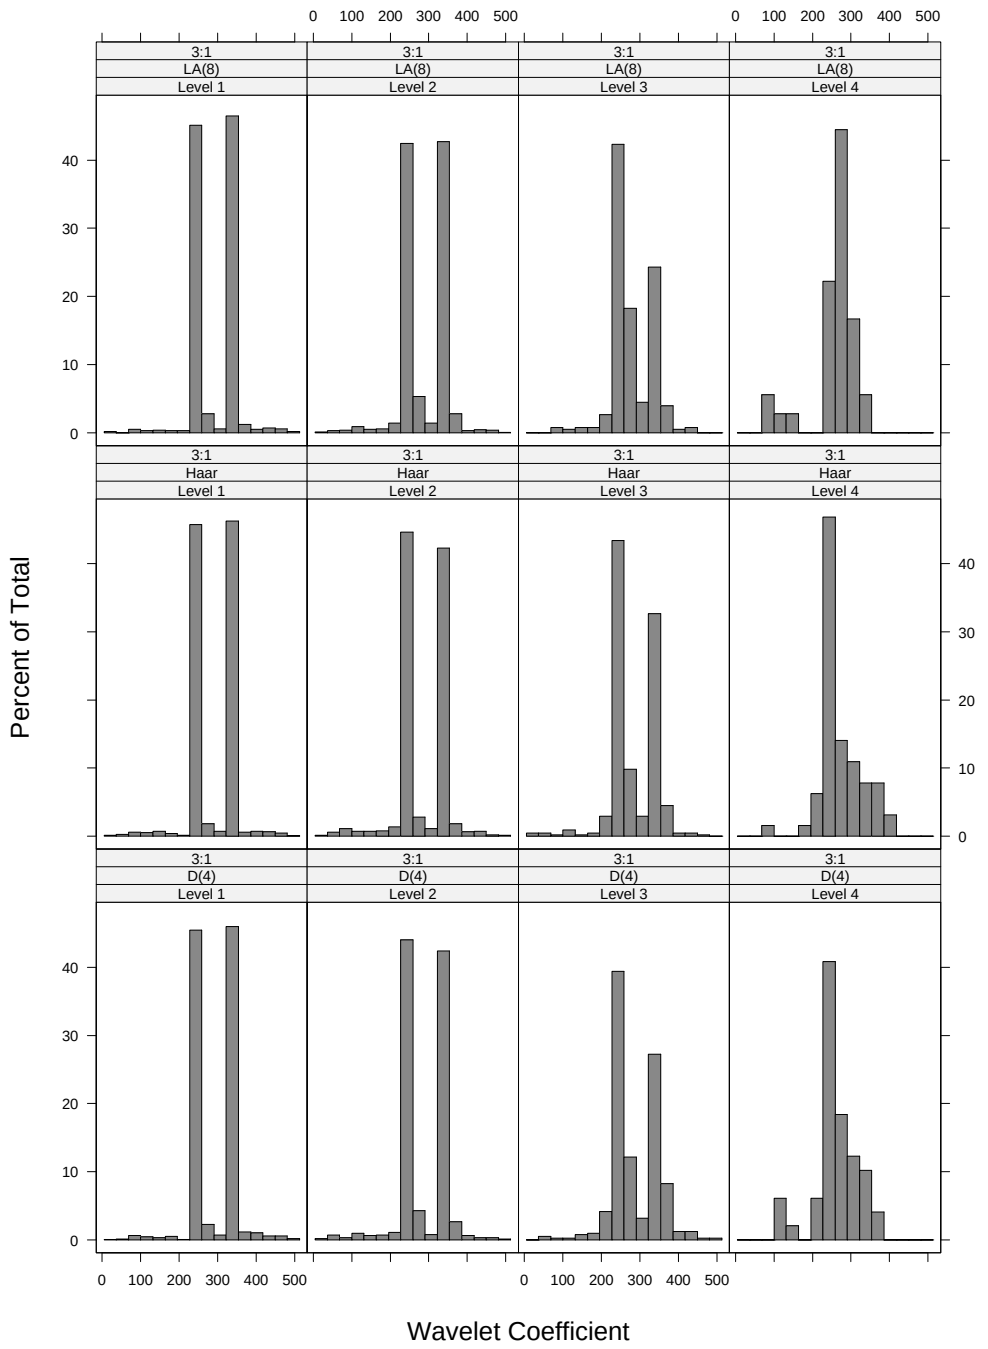

Figure 7: Same as Figure 5 with variance ratio $\Delta = 3$.

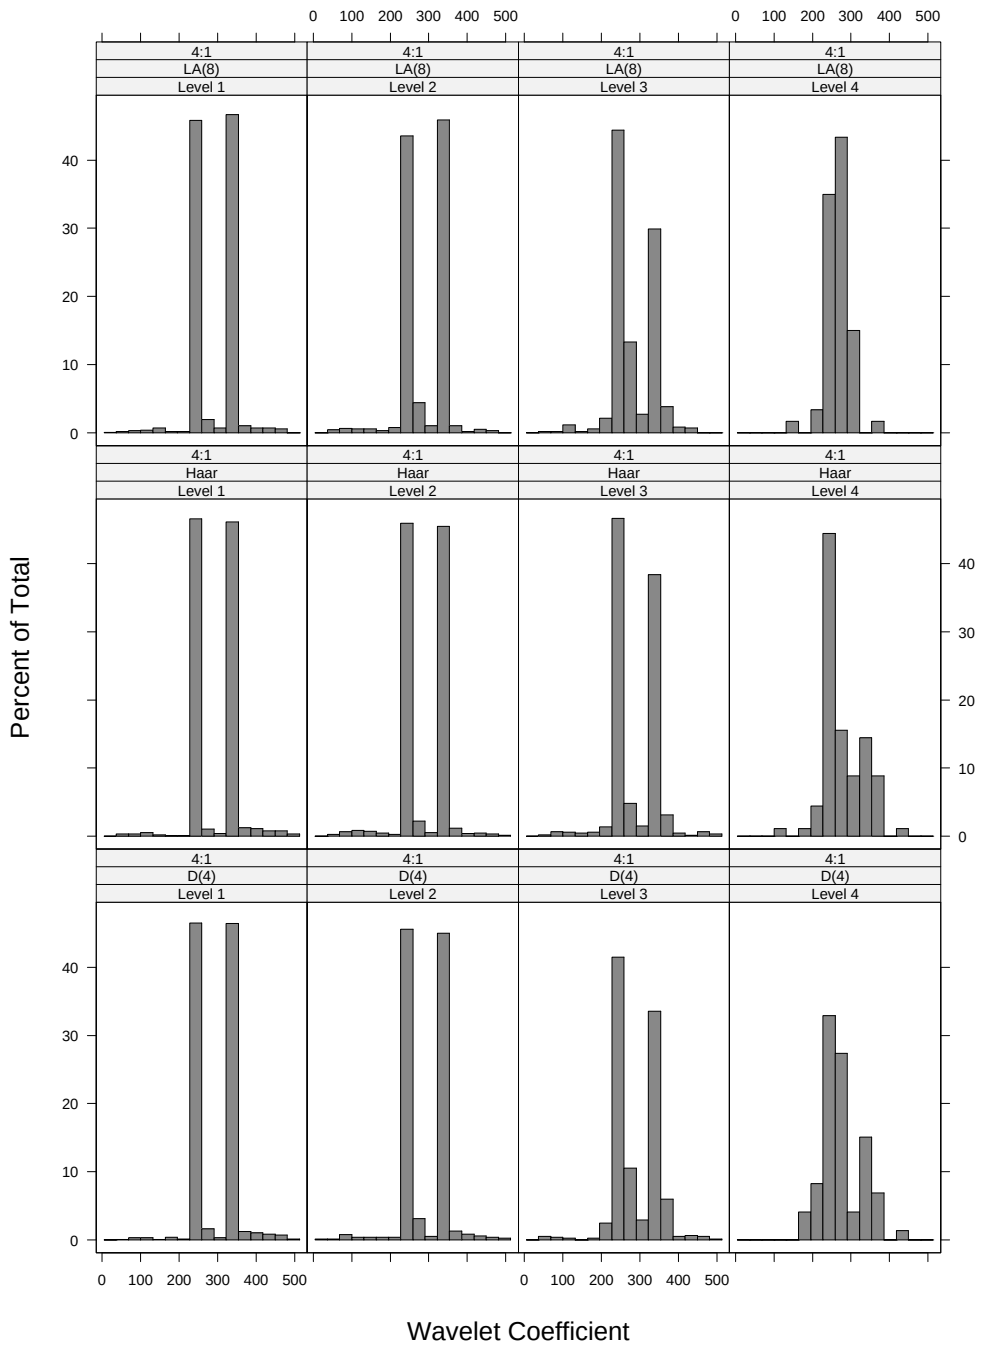

Figure 8: Same as Figure 5 with variance ratio $\Delta = 4$.

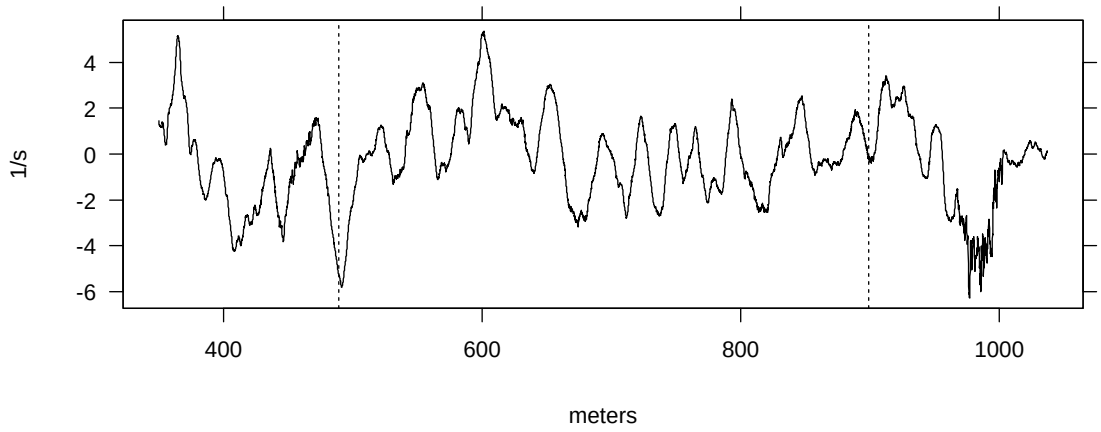

Figure 9: Plot of vertical shear measurements (inverse seconds) versus depth (meters). The two vertical lines are at 489.5m and 899.0m, and denote the roughly stationary series used by Percival and Guttorp (1994). This series can be obtained via the World Wide Web at `http://lib.stat.cmu.edu/datasets/` under the title '`lmpavw`'.

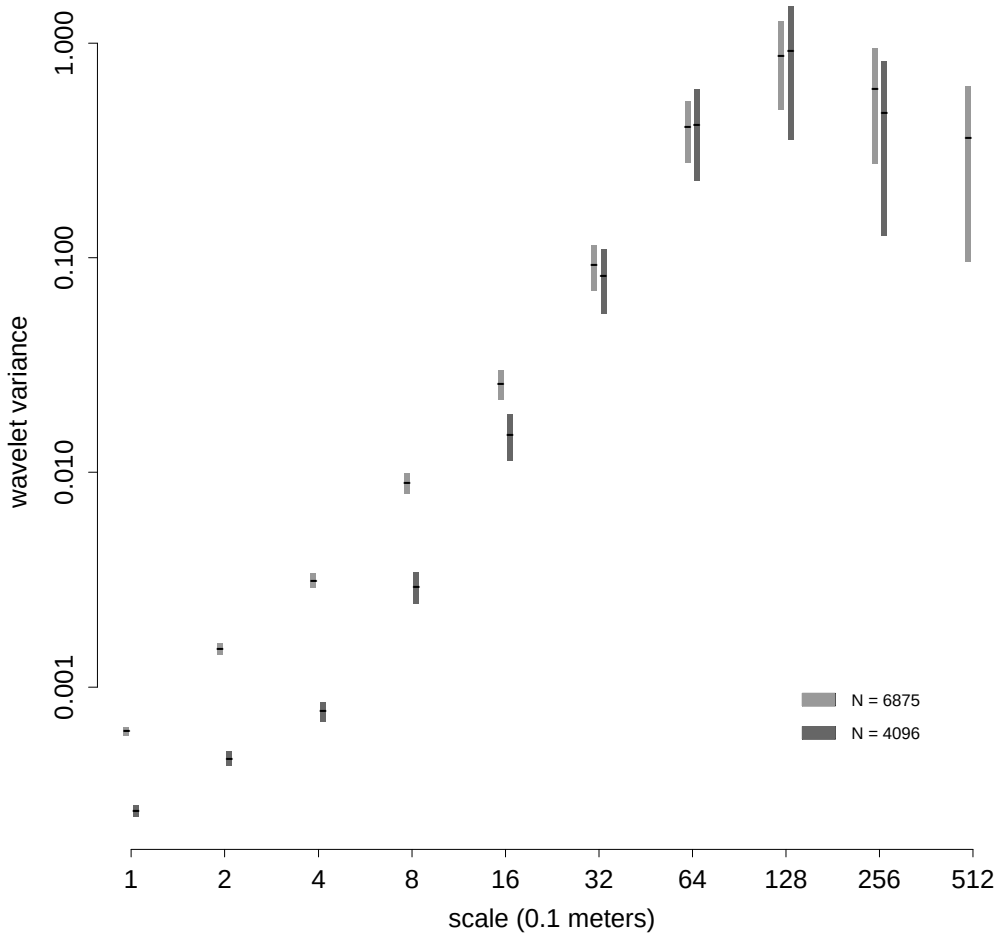

Figure 10: Estimated wavelet variance of the vertical ocean shear measurements using the D(4) wavelet filter and MODWT. The light grey confidence intervals correspond to all 6875 observations, while the dark grey confidence intervals correspond to the middle 4096 observations as analyzed in Percival and Guttorp (1994).

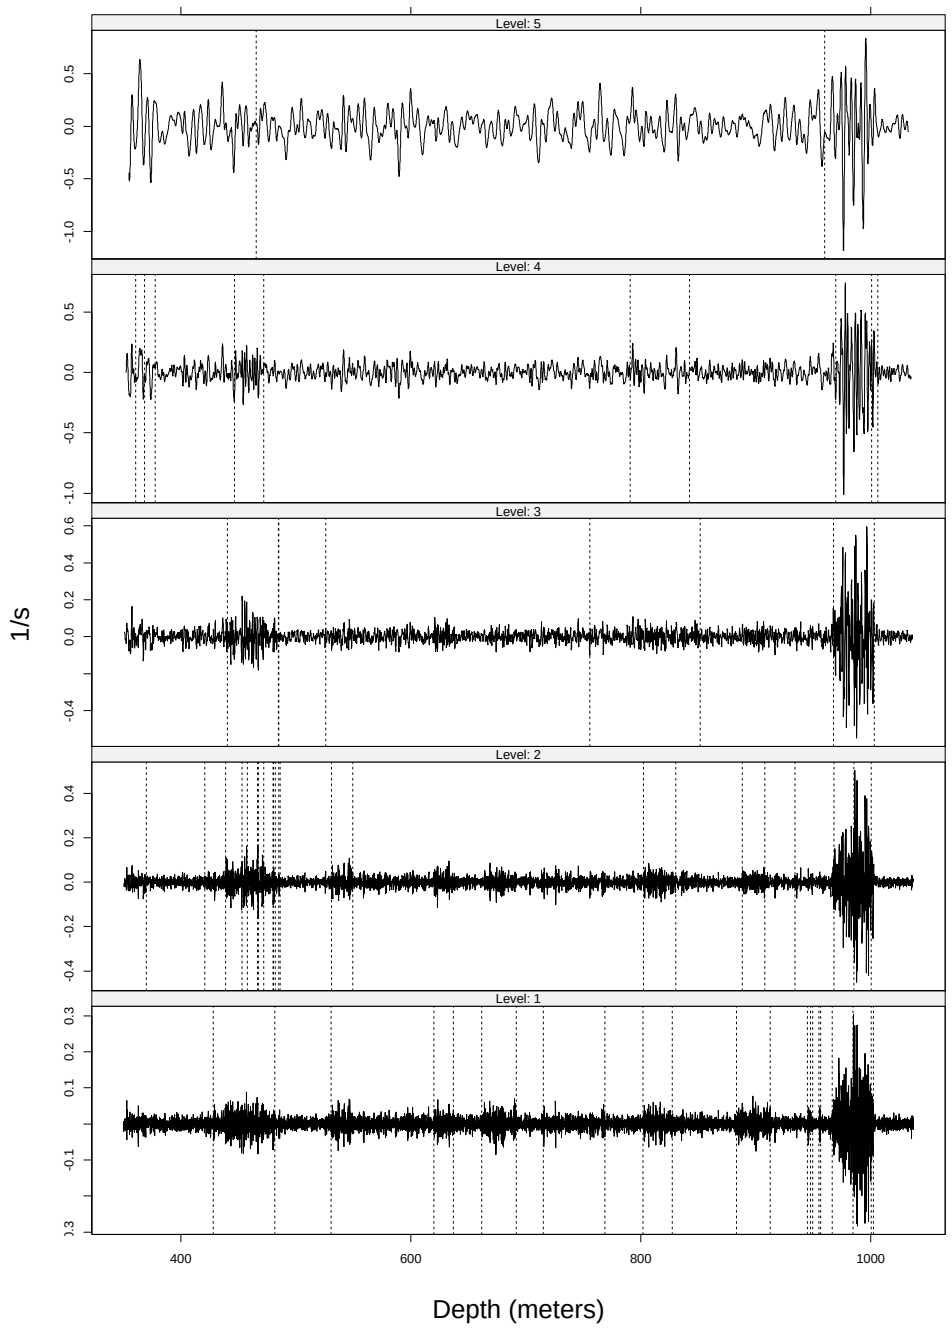

Figure 11: Estimated locations of variance change for the vertical ocean shear measurements using the D(4) wavelet filter displayed on the MODWT wavelet coefficients. Only the first five scales were found to have significant changes of variance. Asymptotic critical values were used for the hypothesis testing at the $\alpha = 0.05$ level of significance.