

Long memory or shifting means in geophysical time series?

William Rea^{a,b,*}, Marco Reale^a, Jennifer Brown^a, Les Oxley^b

^a Department of Mathematics and Statistics, University of Canterbury, Christchurch, New Zealand

^b Department of Economics and Finance, University of Canterbury, Christchurch, New Zealand

Received 8 October 2009; received in revised form 11 May 2010; accepted 9 June 2010

Available online 23 June 2010

Abstract

In the literature many papers state that long-memory time series models such as Fractional Gaussian Noises (FGN) or Fractionally Integrated series (FI(d)) are empirically indistinguishable from models with a non-stationary mean, but which are mean reverting. We present an analysis of the statistical cost of model mis-specification when simulated long memory series are analysed by Atheoretical Regression Trees (ART), a structural break location method. We also analysed three real data sets, one of which is regarded as a standard example of the long memory type. We find that FGN and FI(d) processes do not account for many features of the real data. In particular, we find that the data sets are not H-self-similar. We believe the data sets are better characterized by non-stationary mean models.

© 2010 IMACS. Published by Elsevier B.V. All rights reserved.

Keywords: Long-range dependence; Strong dependence; Global dependence; Hurst phenomena

1. Introduction

Many geophysical time series exhibit the property of statistical long memory. Statistical long memory, also known as long-range dependence, strong dependence, global dependence, or the Hurst phenomenon, can be defined in more than one, not necessarily equivalent, way. The two most commonly used definitions of long memory are: (1) that the serial correlations decay hyperbolically and so are not summable, that is

$$\rho(k) = c_\rho k^\alpha; \quad -1 < \alpha < 0 \quad (1)$$

$$\sum_{k=0}^{\infty} \rho(k) = \infty \quad (2)$$

where c_ρ is a constant and (2) that the spectrum of the series obeys a power law which has a pole at zero frequency, that is

$$f(\lambda) = c_f |\lambda|^{-\gamma}; \quad 0 < \gamma < 1 \quad (3)$$

* Corresponding author at: Department of Economics and Finance, University of Canterbury, Private Bag 4800, Christchurch 8140, New Zealand. Tel.: +64 3 364 3474; fax: +64 3 364 2635.

E-mail addresses: bill.rea@canterbury.ac.nz (W. Rea), marco.reale@canterbury.ac.nz (M. Reale), jennifer.brown@canterbury.ac.nz (J. Brown), les.oxley@canterbury.ac.nz (L. Oxley).

where c_f is a constant. Of course, in finite series the sum of the correlations and the power at the lowest frequencies are both finite and it is a subjective judgement whether the conditions of Eqs. (1)–(3) hold.

What subsequently became known as long-memory series were discussed by Pearson [30,39] but they were brought to prominence by Hurst [21] who studied a large number of hydrological records and reported long-memory to be pervasive in them. Subsequently, Mandelbrot [27] reported long-memory in many geophysical records such as river flows, precipitation records, mud varve sequences, tree-ring indices, earthquake frequencies and the solar sunspot cycle. Long-memory has also been reported in such diverse series as high precision weight measurements [16], stock market volatilities [17], telecommunications network traffic [25] and sea surface temperatures [28].

Series which are either characterised by long memory or have structural breaks in the mean (which is a simple type of non-stationary mean model) are easily confused [12,17,23,37]. Distinguishing between series with true long memory and series with structural breaks in the mean is difficult because their finite sample properties are similar and so standard methodologies fail [36]. Structural break detection techniques often report breaks in long memory series where none exist. Similarly, long memory estimators often report long memory when only structural breaks are present.

The statistical properties of procedures for detecting and quantifying long memory series when, in fact, the series has structural breaks has been well studied. The statistical properties of procedures for detecting and locating structural change when, in fact, there is only long memory is less well understood and literature on this problem is somewhat sparse [36].

This paper addresses the problem of analysing a series with long memory as if it possessed structural breaks by examining the statistical properties of the “regimes” (defined in Section 2 below) reported by Atheoretical Regression Trees (ART) [9,33], a structural break technique, when ART is applied to pure long memory time series.

The remainder of this paper is organized as follows. Section 2 presents a brief overview of the common models applied to long memory time series and the problem of distinguishing between them in real data. Section 3 describes the method of analysis we applied to our real and simulated data sets. Section 4 describes our example data sets, Section 5 presents the results of our investigation into their H-self-similarity properties, Section 6 presents the results for other properties of the data sets, Section 7 contains the discussion and Section 8 the conclusions.

2. Models

A number of models have been proposed to account for the extraordinary persistence of the correlations across time found in long memory series. Here we consider fractionally integrated series and series with structural breaks in the mean.

2.1. Fractional Gaussian noises and fractionally integrated series

The most widely used models for long memory series are the Fractional Gaussian Noises (FGNs), a continuous time process, and their discrete time counterparts, the fractionally integrated series of order d . FGNs were introduced into applied statistics by Mandelbrot and van Ness [26] in an attempt to model the exceptionally slow decay in the autocorrelations observed in long memory series. FGNs are the stationary increments of an Gaussian H self-similar stochastic process which we now define.

Definition 1. A real-valued stochastic process $\{Z(t)\}_{t \in \mathcal{R}}$ is self-similar with index $H > 0$ if, for any $a > 0$,

$$\{Z(at)\}_{t \in \mathcal{R}} =_d \{a^H Z(t)\}_{t \in \mathcal{R}}$$

where \mathcal{R} denotes the real numbers and $=_d$ denotes equality of the finite dimensional distributions. H is also known as the Hurst parameter.

Definition 2. A real-valued process $Z = \{Z(t)\}_{t \in \mathcal{R}}$ has stationary increments if, for all $h \in \mathcal{R}$

$$\{Z(t+h) - Z(h)\}_{t \in \mathcal{R}} =_d \{Z(t) - Z(0)\}_{t \in \mathcal{R}}.$$

It follows from Definition 1 that H is constant for all subseries of an H self-similar process.

The stationary increments of an H-self-similar series with $0.5 < H < 1$ exhibit long memory. It can be shown (see [4, Section 2.3]) that the autocorrelations of the increments of an H self-similar process are of the form

$$\rho(k) = \frac{1}{2}[(k + 1)^{2H} - 2k^{2H} + (k - 1)^{2H}]$$

where $k \geq 0$ and $\rho(k) = \rho(-k)$ for $k < 0$. Asymptotically

$$\lim_{k \rightarrow \infty} \rho(k) \rightarrow H(2H - 1)k^{2H-2}.$$

Thus for $1/2 < H < 1$ the autocorrelations decay so slowly that they are not summable, i.e.

$$\lim_{N \rightarrow \infty} \sum_{k=-N}^{k=N} \rho(k) \rightarrow \infty. \tag{4}$$

A consequence of the autocorrelations summing to infinity is that the spectral density diverges towards infinity at low frequencies. If we recall (see [10, p. 113]) that the spectral density is given by

$$f(\omega) = \frac{1}{\pi} [1 + 2 \sum_{k=1}^{\infty} \rho(k) \cos \omega k]$$

then

$$\lim_{\omega \rightarrow 0} f(\omega) \rightarrow \infty. \tag{5}$$

Independently Granger, Joyeux and Hosking [18,20] obtained the discrete time counter-parts to FGNs, the Fractionally Integrated processes of order d (FI(d)). These are also a generalization of the “integration” part of the Box–Jenkins ARIMA(p, d, q) (AutoRegressive Integrated Moving Average) models to non-integer values of the integration parameter, d and are often referred to as ARFIMA models (AutoRegressive Fractionally Integrated Moving Average). They have the added flexibility of being able to model any additional short term autocorrelation present in the series not accounted for by the long memory component through the use of AR and/or MA components.

Denoting the backshift operator by B , for a non-integer value of d the operator $(1 - B)^d$ can be expanded as a Maclaurin series into an infinite order AR representation

$$(1 - B)^d X_t = \sum_{k=0}^{\infty} \frac{\Gamma(k - d)}{\Gamma(k + 1)\Gamma(-d)} X_{t-k} = \varepsilon_t \tag{6}$$

where $\Gamma(\cdot)$ is the gamma function, and ε_t is a innovation term drawn from an $\mathcal{N}(0, \sigma^2)$ distribution. The operator in Eq. (6) can also be inverted and X_t written in an infinite order MA representation. It can be shown (see [4, Section 2.5]) that the autocorrelations for a fractionally integrated process are, asymptotically,

$$\rho(k) \approx c|k|^{2d-1}. \tag{7}$$

where c is a constant which depends on d , and on any AR(p), and MA(q) components used in the model. For $0 < d < 1/2$ the autocorrelations decay so slowly that they are not summable and exhibit the same divergence of the spectral density towards infinity (Eq. (5)) as an FGN. It is this slow rate of decay of the autocorrelations and divergence of the spectral density that makes them attractive as models for long memory time series.

The two parameters H and d are related by the simple formula $d = H - 1/2$.

Both FGNs and FI(d)s have been extensively applied as models for long memory time series and their theoretical properties studied. See the volumes by [4,14,29] and the collections of [13,34] and the references therein.

2.2. Constrained non-stationary models

Klemes [23] argued that long memory in hydrological time series was a statistical artifact caused by analyzing non-stationarity time series with statistical tools which assume stationarity. Often series which display the long memory property are constrained for physical reasons to lie in a bounded range, but beyond that we have no reason to believe that

they are stationary. For example, in the series we study in this paper (stalagmite rings, tree rings and varve thicknesses, see Section 4), cannot have an unbounded increasing or decreasing trend.

Models of this type which have been proposed typically have stochastic shifts in the mean, but overall are mean reverting about some long term average.

We define the break model as follows:

$$y_t = \sum_{i=1}^p I_{t_{i-1} < t \leq t_i} \mu_i + \varepsilon_t \quad (8)$$

where y_t is the observation at time t , $I_{t \in S}$ is an indicator variable which is 1 only if $t \in S$ and 0 otherwise, t is the time, t_i , $i = 1, \dots, p$, are the breakpoints and μ_i is the mean of the regime i , and ε_t is a noise term. In this case, a regime is defined as the period between breakpoints.

It is important to note that Eq. (8) is just a way to represent a sequence of different models (i.e. models subjected to structural breaks). This model only deals with breaks in mean. It can be generalized for any kind of break. In series with a structural break the noise process, ε_t , may also undergo a change. We are considering this class of model because it has been used by others when studying long memory processes. However, given a true break each regime must be modeled separately.

3. Method

To motivate the new procedures let us consider the differences in the breaks reported in series generated by a true long memory process and those reported when there are genuine structural breaks in the data. The claim that a time series is generated by a true long memory process is a claim that the data generating process (DGP) is uniform throughout the entire series. Thus if a structural break location method is mistakenly applied to the series any reported breaks are completely spurious. These spurious breaks will yield a number of spurious regimes of differing lengths but, in fact, these spurious regimes will only be subsamples of a single population. It follows that the subsamples will have the same statistical properties as the full series, that is, the data in the subsamples has been drawn from a single population. Any estimated differences will be the result of randomness and long range serial correlation and do not reflect actual differences between the samples.

On the other hand, the claim that a time series has one or more structural breaks is a claim that the series has one or more discontinuities in the DGP. In this case a structural break location method will report a number of breaks which will divide the series into a number regimes which are different subpopulations rather than simply subsamples. Consequently the statistical properties of these subpopulations in the individual regimes will need to be estimated. In this case, any estimated differences will be the result of actual differences between the samples and not simply due to randomness and serial correlation.

Based on these considerations it seems reasonable to suspect that it may be possible to distinguish between a true long memory process and a series with structural breaks by examining the statistical properties of the reported regimes. Because it is possible to simulate FGN processes we can obtain bivariate null distributions for these series for comparison with real data for which the DGP is unknown. For example, for an FGN series we can obtain through simulation the bivariate null distribution of H estimate with regime length by breaking a large number of simulated series into spurious regimes with a structural break location procedure and estimating H for each regime using an established estimator of H . Such a method can usefully be thought of as a parametric bootstrap.

We could then apply the same procedure to a real series and examine whether the bivariate data points for the real series fit the null distribution obtained by simulation. If there was a good fit then we would not reject the null of an FGN series. If the fit was poor then we would accept the alternative that the series has one or more structural breaks.

We use Atheoretical Regression Trees (ART) because of their computational speed. ART was recently introduced by Cappelli et al. [9]. An Atheoretical Regression Tree is a non-parametric approach used to find breaks in the level of a stochastic process based on least squares regression trees which are widely used in statistical analysis, see [8] for a detailed description.

In this paper we present a procedure for determining whether the data is consistent with a fractional Gaussian noise. As indicated above the null hypothesis is that the series is an FGN. The alternative is that the series has one or more structural breaks.

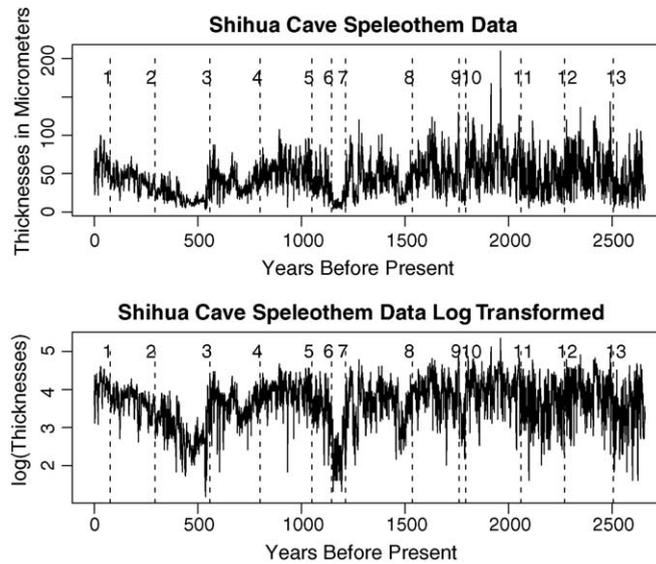


Fig. 1. Plot of Shihua Cave speleothem data with breaks reported by ART.

A procedure for estimating the bivariate distribution of H with regime length:

1. Estimate H for the full series.
2. Through simulation obtain the bivariate probability distribution of H with regime length.
 - (a) Simulate a large number of FGN series with the same H and series length as the series under test.
 - (b) Use ART to break the series into “regimes”.
 - (c) Estimate H (or other statistical quantity) within these regimes.
 - (d) Calculate the empirical median, 75%, 95%, and 99% confidence intervals for the bivariate H and regime length distribution.
 - (e) Plot the bivariate distribution.
3. Apply ART to the full series to obtain the candidate break points.
4. Estimate H for each regime reported by ART.
5. For each regime overplot the H estimate and regime length on the previously determined empirical bivariate distribution.
6. Assess whether H is constant. Details of how the assessment was made are given in below.

In step 2a of the procedure we found that typically it was necessary to simulate 10,000–25,000 FGN series to obtain a good estimate of the bivariate distribution.

Unless otherwise stated all H estimates were obtained using the Whittle estimator as implemented in the R software [31] package `fSeries` of [41]. The choice of estimator was guided by [32]. The Beran test was evaluated using functions implemented in `longmemo` of [7]. FGN series were simulated with the function `fgnSim` in `fSeries` of [41].

The graphical procedure can also be used to assess whether the mean, standard deviation, skewness, or kurtosis are constant.

4. The data sets

We selected three example data sets for analysis. The Campito data set is regarded as a standard example of long memory series. The data sets are as follows.

1. Shihua Cave Speleothems are a set of thicknesses (in microns) of the annual layers of a stalagmite from Shihua Cave, Beijing, China. It covers a 2650-year period from 665BC to 1985AD [40]. Fig. 1 presents the raw data and a log transform together with the ART breaks. The dates are in years before present (BP).

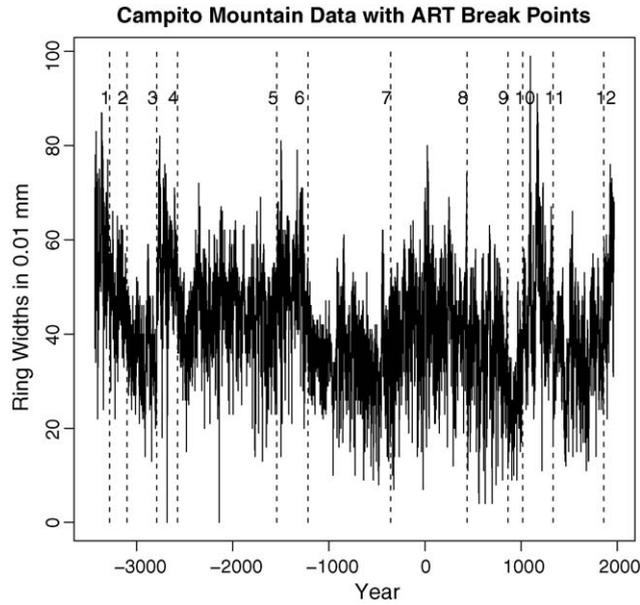


Fig. 2. Plot of Campito Mountain data with breaks reported by ART.

We obtained $H = 0.830$ for the FGN and $d = 0.319$ for an $FI(d)$. Beran’s goodness-of-fit test [3] showed the FGN model provided an adequate fit to the data ($p = 0.15$).

- The Campito Mountain bristlecone pine tree ring sequence [24] is a series of 5405 annual ring widths (in units of 0.01 mm) from bristlecone pines (species *pinus longeava*) on Campito Mountain, California. They cover the period 3435BC to 1969AD. Fig. 2 presents the data with the breaks reported by ART marked.

We obtained $H = 0.876$ for the FGN and $d = 0.447$ for an $FI(d)$. Beran’s goodness-of-fit test [3] showed the FGN model provided an adequate fit to the data ($p = 0.54$).

- The Elk Lake varve sequence [1] is a sequence of 10,224 annual varve (sediment layers) thicknesses (in millimeters) from Elk Lake in Itasca Park, Minnesota. The dates are in years before present. It has been used in studies of paleoclimatology [11]. It was chosen because of its length to test the claim of [27] that varve sequences can be adequately modeled by FGNs. Fig. 3 presents the raw data and a log transform together with the breaks reported by ART.

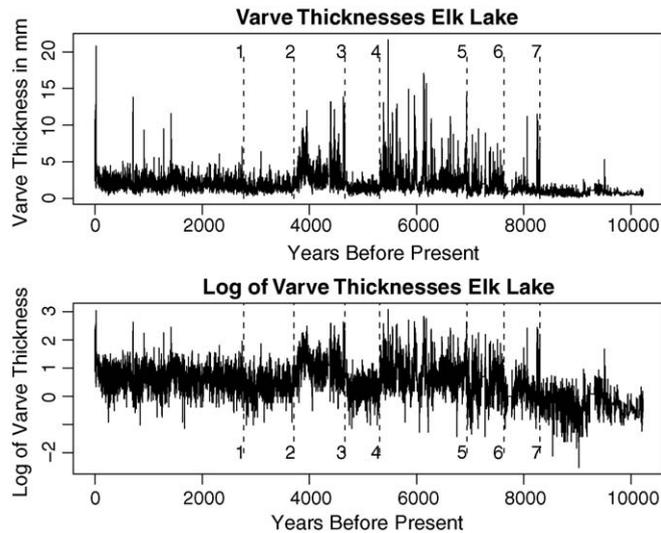


Fig. 3. Plot of Elk Lake data with breaks reported by ART.

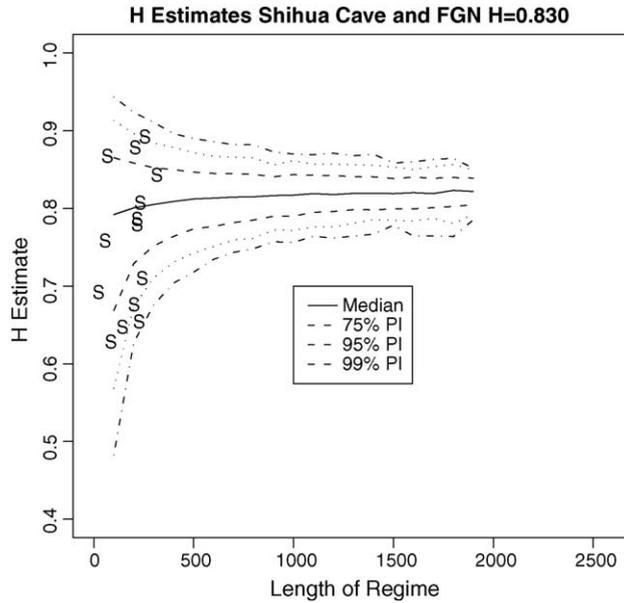


Fig. 4. Empirical distribution and percentage intervals of H estimates against regime length for FGNs with $H=0.83$. “S” indicates a Shihua Cave data point.

We obtained $H=0.860$ for the FGN and $d=0.999$ for an $FI(d)$. Beran’s goodness-of-fit test [3] showed the FGN model provided an adequate fit to the data ($p=0.09$).

5. H -self-similarity results

5.1. Shihua Cave

The results for the Shihua Cave are presented in Fig. 4. The “S”-symbols are the H estimates for each of the Shihua Cave regimes. The percent interval (PI) lines were obtained empirically from 10,000 simulated series by aggregating the data in intervals of 100 data points (e.g. 1–100, 101–200, and so on). This gives a visual presentation of the bivariate distribution of regime lengths (the horizontal axis) and the H estimates (the vertical axis) of the regimes.

ART reported 13 breaks giving 14 regimes. Of the 14 regimes three estimates of H for that regime lie outside the 95% PI.

For a null hypothesis of H constant for the series we can conservatively estimate the p -value by assuming the H -estimates are binomially distributed. The probability of obtaining at least three points outside the 95% PI is thus

$$1 - \sum_{i=0}^2 \binom{14}{i} 0.95^{14-i} 0.05^i.$$

We obtain $p < 0.004$.

We subjected each regime to the [3] goodness-of-fit test. The second to last regime (2271–2506BP) is statistically significantly different from an FGN with the value of H estimated from either the whole series or the regime using the Whittle estimator. We obtain $p=0.007$ for H estimated from the regime and $p=0.0001$ for H estimated from the series.

It seems unreasonable to claim that single value of the H parameter characterizes the long memory properties of this series. This suggests the series is not H -self-similar.

5.2. Campito Mountain

The results are presented in Fig. 5. The “C”-symbols are the Campito data. ART reported 12 breaks giving 13 regimes.

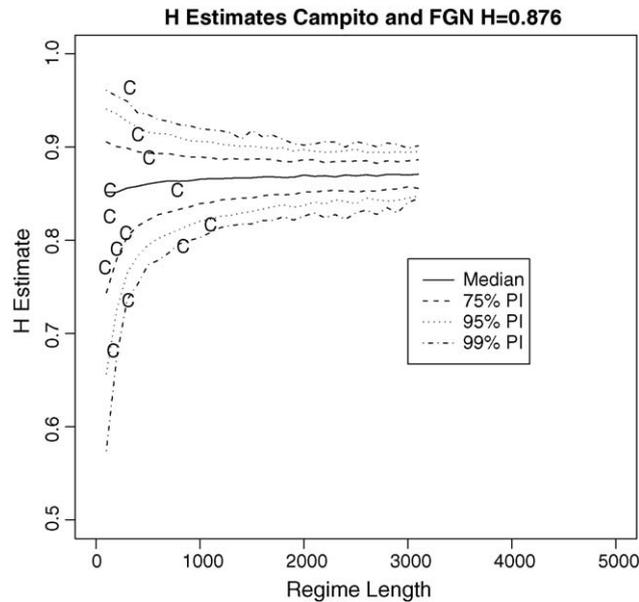


Fig. 5. Empirical distribution and percentage intervals of H estimates against regime length for FGNs with $H=0.876$. “C” indicates a Campito data point.

Four of the 13 data points lie below the lower 95% PI, two of which lie below the lower 99% PI. A further one lies above the upper 99% PI. Using similar reasoning to the Shihua Cave data the probability of finding five of the 13 data points outside the 95% PI yields $p < 0.0002$.

Panel (a) of Table 1 presents the results of Beran [3] goodness-of-fit test for the Campito regimes where the null hypothesis of long memory is not accepted at at least the $p=0.05$ level. The longest regime (2573 to 1544 BC) is statistically significantly different from an FGN with H estimated either from the series or the regime. Two other regimes are statistically significantly different from an FGN on one of these measures. As seen from the table, the period 1018 to 1362 AD can fit an FGN provided the value of H is estimated from the regime rather than from the series.

These results suggested we break the Campito series into three regimes before and after the anomalous 2573 to 1544 BC period. The results are presented in panel (b) of Table 1. At the five percent level the null of long memory is not accepted for the first two periods which collectively represent approximately 35% of the data.

Table 1

Panel (a): individual Campito regimes where the null hypothesis of the Beran [3] test is not accepted at at least the $p=0.05$ level. Panel (b): p -values reported by the Beran [3] goodness-of-fit test for three periods in the Campito data divided by the anomalous 2573–1544 BC period.

(a) Campito mountain individual regimes

Period	$H=H(t)$ p -Value	$H=0.876$ p -Value
2791–2574 BC	0.030	0.128
2573–1544 BC	0.018	0.003
1018–1362 AD	0.329	0.013

(b) Campito mountain aggregated regimes

Period	H	Beran’s p -value
3435–2574 BC	0.868	0.047
2573–1544 BC	0.819	0.018
1543 BC to 1969 AD	0.883	0.641

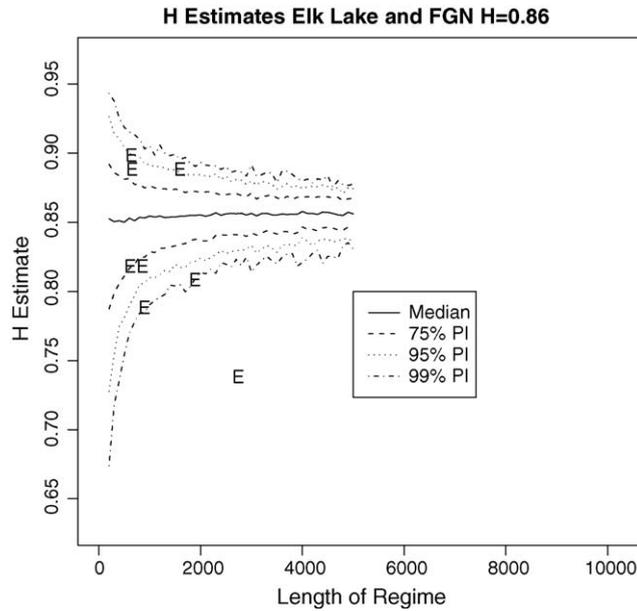


Fig. 6. Empirical distribution and percentage intervals of H estimates against regime length for FGNs with $H=0.86$. “E” indicates a Elk Lake data point.

It seems unreasonable to claim that a single value of the H parameter characterizes the long memory properties of this series. Further, over one-third of the data does not have a spectral density of the form (3). These results suggest the series is not H -self-similar.

5.3. Elk Lake

The results for the Elk Lake data is presented in Fig. 6. ART reported seven breaks giving eight regimes. The “E”-symbols are the Elk Lake data points. While five of the eight values lie outside the 95% PI, the outlier value of H is so far from the empirical distribution that it is impossible to assess its p -value even conservatively. Again, it is unreasonable to claim H is constant for this data set.

All Elk Lake regimes fitted an FGN without the need to allow H to vary according to Beran [3] test.

6. Other results

In this section we continue to use simulated FGNs for comparison purposes.

For reasons of space we only present some of the results for the Campito Mountain and Elk Lake data sets. The remainder are available on request from the authors. Recall that our chosen structural break location method, ART, breaks the series into regimes based on local changes in the series mean.

6.1. Campito Mountain

Of the quantities estimated only the regime means and the magnitudes of the slopes within the regimes could be adequately modeled by FGN or FI(d) processes.

Fig. 7 presents the standard deviations of the regimes. It is expected that subseries of the simulated FGNs will have smaller standard deviations than the whole series. Comparing the 13 Campito data points with the simulated data, one lies above the 99% PI while five lie below the lower 95% PI. This distribution strongly suggests the regimes reported by ART are distinct sub-populations rather than just samples of varying lengths from a population with a uniform data generating process throughout the series.

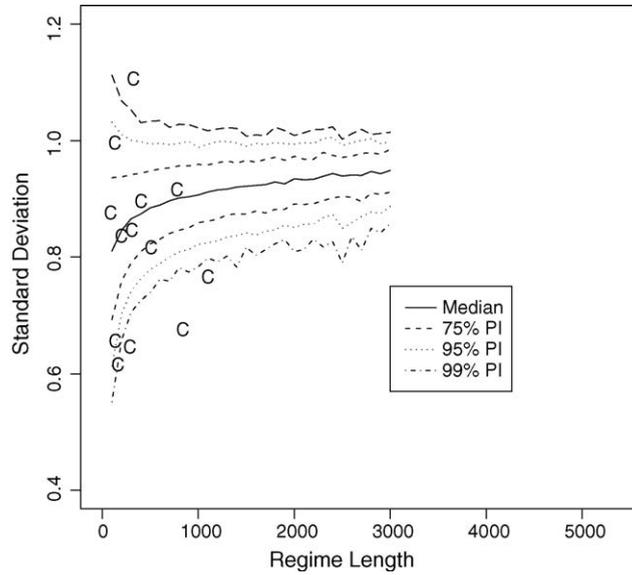


Fig. 7. Empirical distribution and percentage intervals of standard deviation estimates for FGNs with $H=0.88$. “C” indicates a Campito Mountain data point.

Tree ring data are used as proxies for past climatic conditions. The four most recent periods correspond to known climatic shifts giving us confidence the other breaks reflect real changes in climatic conditions in the Sierra Nevada Mountains of California. These are presented in Table 2. The two most recent periods correspond to human-induced changes in the concentration of greenhouse gases in the atmosphere [35].

6.2. Elk Lake

Of the results for the Elk Lake data only the number of breaks in the series could be adequately modeled by an FGN or an $FI(d)$ process.

Fig. 8 presents the results for the mean. Fig. 9 presents the results for the standard deviations. Comparing the Elk Lake regimes with the simulated data we see that six of the eight regime standard deviations lie below the empirical 99% PI of the FGNs. Comparing these two sets of results we can see that the series variability arises because of the distribution of the regime means, not because of the within regime variability of the data. This strongly suggests the regimes reported by ART are different sub-populations.

The varve thicknesses are a result of chemical, physical and biological processes within the lake and atmospheric processes in the surrounding region. Table 3 presents the break dates and causes of the change where these are known.

7. Discussion

We have selected three geophysical data sets of the type which past literature has made claims or suggests that FGNs or $FI(d)$ series provide an adequate model. In practice the additional flexibility available in being able to add AR and MA components to $FI(d)$ models is often not required. For example, for the Campito data [2] reported that

Table 2
The climatic conditions corresponding to the regimes reported by ART for the Campito Mountain tree ring data.

Period (AD)	Climatic conditions
1863–1969	Industrialization and global warming [22]
1333–1862	The Little Ice Age [22]
1018–1332	The Medieval Climate Optimum [22]
862–1017	Extreme drought in the Sierra Nevadas [38]

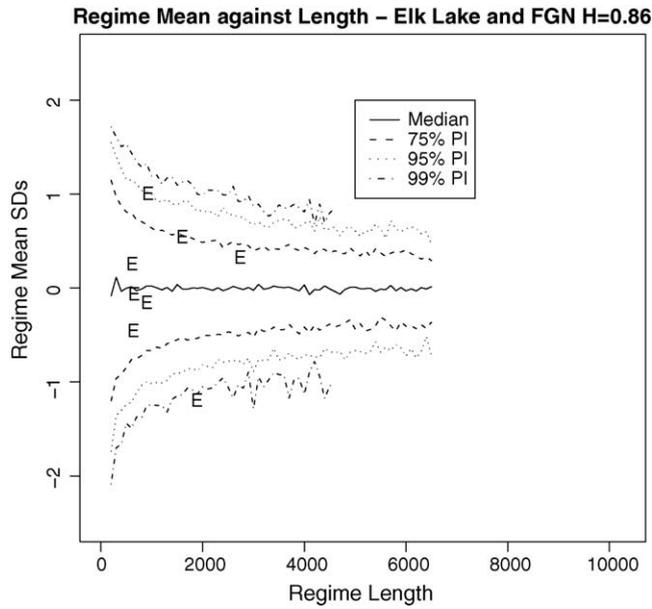


Fig. 8. Empirical distribution and percentage intervals of estimates of the mean for FGNs with $H=0.86$. “E” indicates a Elk Lake data point.

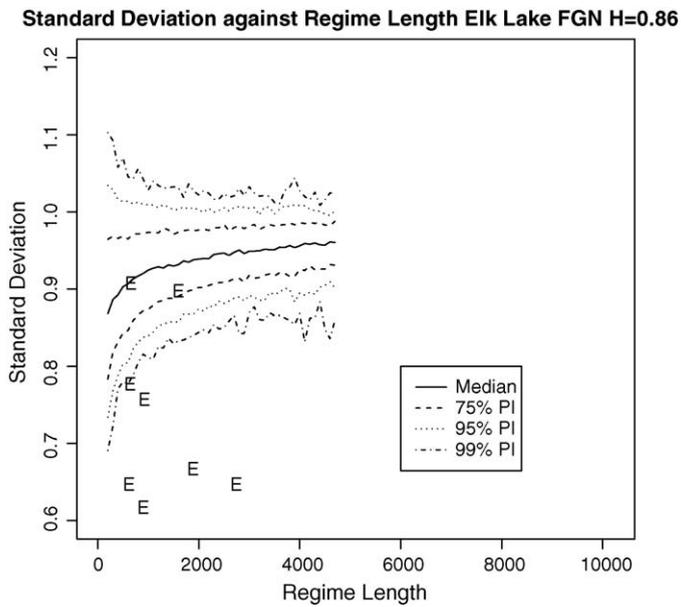


Fig. 9. Empirical distribution and percentage intervals of standard deviation estimates for FGNs with $H=0.86$. “E” indicates a Elk Lake data point.

Table 3

Breaks in the Elk Lake data which correspond to known climate changes or physical changes in the lake.

Change date (BP)	Change
4662	Drought [15]
5311	Chemical change in the lake water [15]
8302	Collapse of Laurentide ice sheet [11]

Table 4

p -Values for the three example data set under the null-hypothesis that H is constant using our graphical results.

Data set	No. data points	p -Value
Shihua Cave	2560	<0.004
Campito Mountain	5405	<0.0002
Elk Lake	10,224	≈ 0

the best fitting ARFIMA model was an ARFIMA(0, 0.44, 0). The lack of additional AR and MA components to the basic fractional integration essentially means the researcher could use either an FGN or and FI(d) model according to their preference. We have examined the hypothesis that changes in the value of H are associated with changes in the local mean of the series. Mean shifts are the simplest type of non-stationary mean model. We have demonstrated the data generating process of these series is not H -self-similar to high levels of statistical confidence without having to exhaustively search all possible binary divisions of the data such as is required in the [5,6] test. We have also been able to show that the fit gets worse with increasing series length. This is summarized in Table 4.

Other authors [5,16,19] seem comfortable with the idea of a time-varying H (or d) parameter despite such a quantity violating the definition of an H -self-similar process. As a parameter only has meaning in the context of a model, if H is allowed to be a function of time we are led to the so-called multifractal models. However, a time-varying H nullifies one of the great strengths of the FGNs or FI(d) models, namely that a single parameter can describe the long-range dependence properties of the whole series. In FI(d) models, a time varying d parameters requires a model of the form

$$(1, B)^{d(t)} X_t = \varepsilon_t$$

where instead of a constant d we have d as a function of time. It is not at all clear what mathematical meaning can be assigned to the operator $(1 - B)^{d(t)}$.

Given the difficulty of proposing a physically reasonable process by which, say, a sediment layer in a lake continues to influence future sediment layer thicknesses over time scales of several millenia, once a time-varying H is admitted, and the model is of necessity non-stationary in H , it seems reasonable to consider whether the apparent long-memory is caused by non-stationary mean. For simplicity we consider the series to have undergone structural breaks in the mean which divide the series into a number of “regimes” where the regimes are short memory processes. The evidence presented in Tables 2 and 3 clearly suggests at least some of the breaks reported by ART are indeed structural breaks.

8. Conclusions

If the goal of the analysis is to find a single statistically parsimonious and mathematically elegant model for all of the data then either an FGN or an FI(d) are good choices. They are quick to fit, require few parameters, and are straight forward to use in forecasting. For pragmatic time series analysis FGNs and FI(d) processes can provide useful modeling tools provided the series is not too long.

While the three data sets examined in detail are clearly continuous processes (stalagmite ring growth, tree growth, sediment accumulation) some of the structural breaks reported by ART do correspond to known climatic or physical changes. Thus a purely continuous model, such as an FGN, can obscure useful information about past states of the data generating process.

In the data sets examined the Hurst phenomenon can be attributed to non-linear time series composed of infrequently changing regimes with different means being incorrectly analysed by linear methods.

The problem of structural break methods finding breaks in FGN or FI(d) series has been regarded as an obstacle in distinguishing between true long-memory and processes with a non-stationary mean in real data. We have shown that this can be exploited by a computationally fast structural break method, such as ART, to obtain distributions of a range of statistical quantities for simulated FGN or FI(d) series which can be compared to the data points obtained for the real data in a method which can be thought of as a parametric bootstrap. Non-equality in distribution leads to a rejection of the FGN or FI(d) as a candidate model for the data.

While the FGN or FI(d) models fail to adequately describe the data in all three cases, we believe it is unwise to make any general claim about the applicability of FGNs and FI(d)s to other long memory time series.

References

- [1] R.Y. Anderson, Elk Lake, Minnesota Database. IGBP PAGES/World Data Center-A for Paleoclimatology Data Contribution Series No. 91-001 (1991).
- [2] R.T. Baillie, S.-K. Chung, Modeling and forecasting from trend-stationary long memory models with applications to climatology, *International Journal of Forecasting* 18 (2002) 215–226.
- [3] J. Beran, A goodness-of-fit test for time series with long range dependence, *Journal of the Royal Statistical Society B* 54 (3) (1992) 749–760.
- [4] J. Beran, *Statistics for Long Memory Processes*, Chapman & Hall/CRC Press, 1994.
- [5] J. Beran, N. Terrin, Testing for a change in the long-memory parameter, *Biometrika* 83 (3) (1996) 627–638.
- [6] J. Beran, N. Terrin, Testing for a change in the long-memory parameter, *Biometrika* 86 (1) (1999) 233.
- [7] J. Beran, B. Whitcher, M. Maechler, longmemo: Statistics for Long-Memory Processes (Jan Beran) – Data and Functions, R package version 0.9-3 (2006).
- [8] L. Breiman, J. Friedman, R. Olshen, C. Stone, *Classification and Regression Trees*, Chapman & Hall/CRC, 1993.
- [9] C. Cappelli, R.N. Penny, W.S. Rea, M. Reale, Detecting multiple mean breaks at unknown points with Theoretical Regression Trees, *Mathematics and Computers in Simulation* 78 (2–3) (2008) 351–356.
- [10] C. Chatfield, *The Analysis of Time Series*, 6th edition, Chapman & Hall/CRC Press, 2004.
- [11] W.E. Dean, R.M. Forester, J.P. Bradbury, Early Holocene change in atmospheric circulation in the Northern Great Plains: an upstream view of the 8.2ka cold event, *Quaternary Science Reviews* 21 (2002) 1763–1775.
- [12] F.X. Diebold, A. Inoue, Long memory and regime switching, *Journal of Econometrics* 105 (2001) 131–159.
- [13] P. Doukhan, G. Oppenheim, M. Taqqu, *Theory and Applications of Long-Range Dependence*, Birkhäuser, 2003.
- [14] P. Embrechts, M. Maejima, *Selfsimilar Processes*, Princeton University Press, 2002.
- [15] S.L. Forman, R. Oglesby, R.S. Webb, Temporal and spatial patterns of Holocene dune activity on the Great Plains of North America: megadroughts and climate links, *Global and Planetary Change* 29 (2001) 1–29.
- [16] H.P. Graf, E.R. Hampel, J. Tacier, The problem of unsuspected serial correlations, in: J. Franke, W. Härdle, R. Martin (Eds.), *Robust and Nonlinear Time Series Analysis*, vol. 26 of *Lecture Notes in Statistics*, Springer, 1985.
- [17] C.W.J. Granger, N. Hyung, Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns, *Journal of Empirical Finance* 11 (2004) 213–228.
- [18] C.W.J. Granger, R. Joyeux, An introduction to long-range time series models and fractional differencing, *Journal of Time Series Analysis* 1 (1980) 15–30.
- [19] T. Higuchi, Approach to an irregular time series on the basis of fractal theory, *Physica D* 31 (1988) 277–283.
- [20] J.R.M. Hosking, Fractional differencing, *Biometrika* 68 (1) (1981) 165–176.
- [21] H.E. Hurst, Long-term storage capacity of reservoirs, *Transactions of the American Society of Civil Engineers* 116 (1951) 770–808.
- [22] P.D. Jones, M.E. Mann, Climate over the past millennia, *Reviews of Geophysics* 42 (2004) 1–42, rG2002.
- [23] V. Klemes, The Hurst phenomenon—a puzzle? *Water Resources Research* 10 (4) (1974) 675–688.
- [24] V.C. LaMarche, Paleoclimatic inferences from long tree-ring records, *Science* 183 (1974) 1043–1048.
- [25] W.E. Leland, M.S. Taqqu, W. Willinger D.V. Wilson, Ethernet traffic is self-similar: Stochastic modeling of packet traffic data (1993) preprint, Bellcore, Morristown.
- [26] B.B. Mandelbrot, J.W. van Ness, Fractional brownian motions, fractional noises and applications, *SIAM Review* 10 (4) (1968) 422–437.
- [27] B.B. Mandelbrot, J.R. Wallis, Global dependence in geophysical records, *Water Resources Research* 5 (1969) 321–340.
- [28] J.E. Overland, D.B. Percival, H.O. Mofjeld, Regime shifts and red noise in the North Pacific, *Deep-Sea Research I* 53 (2006) 582–588.
- [29] W. Palma, *Long-Memory Time Series Theory and Methods*, Wiley-Interscience, 2007.
- [30] E.S. Pearson, On the variations in personal equation and the correlation of successive judgments, *Biometrika* 14 (1922) 23–102.
- [31] R Development Core Team, R: A language and environment for statistical computing, R Foundation for Statistical Computing, Vienna, Austria, ISBN 3-900051-07-0 (2005). URL <http://www.R-project.org>.
- [32] W. Rea, L. Oxley, M. Reale, J. Brown, The Empirical Properties of Some Popular Estimators of Long Memory Processes, Working Paper 13/2008, Department of Economics and Finance, University of Canterbury (2008).
- [33] W.S. Rea, M. Reale, C. Cappelli, J. Brown, Identification of changes in mean with regression trees: an application to market research, *Econometric Reviews* (2010), forthcoming.
- [34] P.M. Robinson, *Time Series with Long Memory*, Oxford University Press, 2003.
- [35] W.F. Ruddiman, The anthropogenic greenhouse era began thousands of years ago, *Climatic Change* 61 (2003) 261–293.
- [36] P. Sibbertsen, Long memory versus structural breaks: an overview, *Statistical Papers* 45 (4) (2004) 465–515.
- [37] A.D. Smith, Level shifts and the illusion of long memory in economic time series, *Journal of Business and Economic Statistics* 23 (3) (2005) 321–335.
- [38] S. Stine, Extreme and persistent drought in California and Patagonia during mediaeval times, *Nature* 369 (3) (1994) 546–549.
- [39] Student, Errors of routine analysis, *Biometrika* 19 (2) (1927) 151–164.
- [40] M. Tan, T. S. Liu, J. Hou, X. Qin, H. Zhang, T. Li, 2650-year Beijing Stalagmite Layer Thickness and Temperature Reconstruction IGBP PAGES/World Data Center for Paleoclimatology Data Contribution Series No. 2003-050 (2001).
- [41] D. Wuertz, fSeries: Financial Software Collection, R package version 220. 10063 (2005). URL <http://www.itp.phys.ethz.ch/econophysics/R/2.1>.