



This is a repository copy of *Gradual multifractal reconstruction of time-series: Formulation of the method and an application to the coupling between stock market indices and their Hölder exponents*.

White Rose Research Online URL for this paper:
<http://eprints.whiterose.ac.uk/125982/>

Version: Accepted Version

Article:

Keylock, C.J. orcid.org/0000-0002-9517-9896 (2018) Gradual multifractal reconstruction of time-series: Formulation of the method and an application to the coupling between stock market indices and their Hölder exponents. *Physica D: Nonlinear Phenomena*, 368. pp. 1-9. ISSN 0167-2789

<https://doi.org/10.1016/j.physd.2017.11.011>

Article available under the terms of the CC-BY-NC-ND licence
(<https://creativecommons.org/licenses/by-nc-nd/4.0/>).

Reuse

This article is distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivs (CC BY-NC-ND) licence. This licence only allows you to download this work and share it with others as long as you credit the authors, but you can't change the article in any way or use it commercially. More information and the full terms of the licence here: <https://creativecommons.org/licenses/>

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Gradual multifractal reconstruction of time-series: Formulation of the method and an application to the coupling between stock market indices and their Hölder exponents

Christopher J. Keylock^{a,b},

^aDepartment of Civil and Structural Engineering, University of Sheffield, Mappin Building, Mappin Street, Sheffield, S1 3JD, U.K.

^bVisiting Researcher, Department of Aeronautics, Imperial College London, London, SW7 2AZ, U.K.

Abstract

A technique termed gradual multifractal reconstruction (GMR) is formulated. A continuum is defined from a signal that preserves the pointwise Hölder exponent (multifractal) structure of a signal but randomises the locations of the original data values with respect to this ($\varphi = 0$), to the original signal itself ($\varphi = 1$). We demonstrate that this continuum may be populated with synthetic time series by undertaking selective randomisation of wavelet phases using a dual-tree complex wavelet transform. That is, the $\varphi = 0$ end of the continuum is realised using the recently proposed iterated, amplitude adjusted wavelet transform algorithm [Keylock, C. J. 2017. Phys. Rev. E 95, 032123] that fully randomises the wavelet phases. This is extended to the GMR formulation by selective phase randomisation depending on whether or not the wavelet coefficient amplitudes exceeds a threshold criterion. An econophysics application of the technique is presented. The relation between the normalised log-returns and their Hölder exponents for the daily returns of eight financial indices are compared. One particularly noticeable result is the change for the two american indices (NASDAQ 100 and S&P 500) from a non-significant to a strongly significant (as determined using GMR) cross-correlation between the returns and their Hölder exponents from before the 2008 crash to afterwards. This is also reflected in the skewness of the phase difference distributions, which exhibit a geographical structure, with asian markets not exhibiting significant skewness in contrast to those from elsewhere globally.

Keywords: Wavelets, multifractal, Hölder exponent, surrogate data, econophysics

1. Introduction

A wide range of data have been analysed in terms of their multifractal properties. These include turbulence [1, 2], economic time series [3, 4, 5], rainfall data and landscape surfaces [6, 7], images [8], and medical time series [9, 10]. Characterizing the multifractality in terms of the variation of pointwise Hölder exponents, $\alpha_p(\mathbf{x}, t)$ provides

Email address: c.keylock@sheffield.ac.uk (Christopher J. Keylock)
Preprint submitted to Physica D

local information (in time or space) on the scaling behavior of a signal and, thus, how the energy of the phenomenon under consideration is transferred to small scales and dissipated. A globally uniform, constant Hölder exponent, arises in a Brownian motion where $\langle \hat{\alpha}_p(\mathbf{x}, t) \rangle = 0.5$, or fractional Brownian motion where $0 < \hat{\alpha}_p(\mathbf{x}, t) < 1$ [11], where the hat indicates a value that is constant everywhere (either theoretically or to the precision of the estimation procedure). Significant departure from such a scenario leads to intermittent behaviour where large changes can suddenly become manifest because $\alpha_p \ll \langle \alpha_p \rangle$, where the angled braces indicate the average of a quantity that varies more than estimation precision. Stochastic processes with a range of Hölder exponents can be realised by multifractional Brownian motions [12], stochastic cascading models for wavelet coefficients [13], or from a Fokker-Planck model for the signal's increments [14].

Being able to predict these periods of uncertainty is of obvious importance in haemodynamics, or econophysics, and there has been a significant corpus of work on the multifractal analysis of financial time series [15]. Broadly speaking these fall into five main types of study:

1. Multifractal structure of markets [16, 17];
2. Market and informational efficiency [18, 19];
3. Volatility forecasting [20, 21];
4. Portfolio allocation [22]; and,
5. High frequency trading analyses [23], leading to volatility hedging models [5].

The technique developed in this paper is perhaps of greatest relevance to work in the first, third and fifth of these.

More generally, characterizing intermittency can provide clues as to the physical processes governing system behavior, permitting models of partially specified or complex systems to be constructed [24, 25, 26]. Such nonlinear phenomena are not always readily amenable to analysis by conventional statistical hypothesis testing, yet given a dataset of observations, particularly if they are expensive to obtain or are intrinsically unique (for example a rainfall record at a point or the history of trading prices for a company), it is not always possible to obtain replicates. Therefore, the method of surrogate data was developed by Theiler and co-workers [27] within nonlinear physics to provide a means to obtain bootstrapped confidence intervals on observations that would permit the testing for nonlinear effects such as the existence of multiple values for the Hölder exponents in one dataset. From the perspective of the refined surrogate data algorithm of Schreiber and Schmitz [28], surrogates are constructed to match the histogram of the data values and to converge asymptotically on the Fourier amplitude spectrum of the original data, with the Fourier phases randomized. Hence, any nonlinear structure in the data such as intermittency, will only be preserved by chance in the surrogates. Adopting a significance level, a , and a two-tailed statistical test, if $(2/a) - 1$ surrogates are generated and the value for the data on the chosen metric (perhaps some measure of the dispersion of the Hölder exponents such as their range or standard deviation) is outside that for the surrogates, then a significant difference is deemed to exist.

Such hypothesis tests yield a binary result so there either is, or is not, a significant difference between the data and linearised surrogates. Given two datasets where the

null hypothesis has been rejected, is one of them more complex than the other based on the metric used to determine nonlinearity [29]? This question can be answered using the gradual wavelet reconstruction (GWR) approach [30], where rather than solely employing a Fourier domain-based randomisation, a set amount of phase information is locked in place before randomisation using a wavelet transform [31]. Varying the proportion of wavelet energy fixed in the surrogate data, ρ , from 0 (Schreiber and Schmitz-type surrogate) to 1 (the original data without any change) and generating surrogates at each ρ means that the hypothesis can be tested against a continuum defined by ρ . With this GWR method it is then possible to

- Test the relative complexity of different data series using the same metric [32, 33];
- Bootstrap data with the right degree of nonlinearity preserved in the surrogates [34];
- Define boundary conditions for numerical models and examine the effects on the simulation as a function of ρ [35]; and
- Test the relative sensitivity of different metrics of nonlinearity on a particular dataset [30].

Given that the $\rho = 0$ surrogates in GWR remove multifractal characteristics, this formulation is very useful for studying the properties of multifractal signals as a function of ρ . For example, an analysis of the velocity increments in turbulence highlighted the importance of two (from four) parameters in a Fokker-Planck model for these increments for the nonlinear structure of turbulence [36]. Gradual wavelet reconstruction showed both that these terms, which are an order of magnitude smaller than the other two terms, are significantly different to zero in turbulent flows and that they are crucial for controlling the multifractal behaviour.

On the other hand, there are important research questions that one might seek to address where the hypothesis concerns additional properties of a signal, conditioned on a given multifractal structure. In which case, the Fourier-domain randomisation at $\rho = 0$ in GWR does not provide an optimal starting point for analysis, and a new form of surrogate generation algorithm is needed. An example of such analyses that is of practical relevance in a number of fields, including econophysics, is the extent to which the multifractality is coupled to the values of the signal itself. If this is the case, prediction of intermittent events is potentially facilitated. Note too that such processes are no longer just multifractal processes even though they have a measurable multifractality. Instead, processes where there is some functional relation between the Hölder exponents, α_p of the signal and the signal itself are termed self-regulating processes [37, 38, 39].

Given the starting point of an algorithm that preserves the multifractal structure of the data but generates no significant coupling between the $\alpha_p(t)$ and the values of the time series, y , except by chance, a procedure similar to GWR can be formulated. We term this Gradual Multifractal Reconstruction or GMR and it builds up any such a coupling in a systematic way. Hence, the relative strength of self-regulation may be defined in terms of a threshold value for the control parameter for GMR, φ , that demarcates a change from a significant difference between data and surrogates, to no significant difference. Clearly, such a framework is not restricted to the analysis of the degree of coupling seen in self-regulating processes, but is sufficiently generic to permit the analysis of any

higher order property of a time series conditioned on the multifractal structure. Hence, another potential application would be to the study of oscillating singularities in time series [40, 41, 42, 43], with GMR able to detect differences between the number or strength of such oscillations. However, the focus of this paper is the development of the GMR framework and its utility for analysing self-regulation in financial markets. More specifically, having formulated GMR we study the degree of self-regulation in the log-returns from eight international financial indices and consider the geography of the interconnected financial system on the basis of our results.

2. Gradual multifractal reconstruction (GMR)

A decade ago, an algorithm was proposed for generating surrogate data that respects the multifractal characteristics of a data set [44]. More recently, we have proposed an algorithm that goes further and seeks to not merely replicate the distribution function of the Hölder exponents, but their local values. This new algorithm also fulfils the constraint in the Schreiber and Schmitz algorithm that the histogram of the original data values is preserved, while also more accurately mimicking the multifractal characteristics [45]. It is this algorithm that we use as the basis for GMR, and we briefly recount how it works here. First, however, we define the pointwise Hölder function, $\alpha_p(t)$, of a time-series, which is the quantity that multifractal analyses based on structure functions [46, 47], wavelet transform modulus maxima [2, 48] or multifractal detrended fluctuation analysis [49] seek to characterize. We then consider how multifractality may be determined from wavelet coefficients because the principle behind our base algorithm for GMR is the constrained phase randomisation of complex-valued wavelet coefficients.

2.1. Pointwise Hölder exponents and multifractality

Given a signal, $y(t)$, a position in time, t_0 , and assuming that $\alpha_p \in \{0, \dots, 1\}$ (singularity structure governed by first derivatives), we seek the β_p , for which

$$|y(t) - y(t_0)| \leq c_p |t - t_0|^{\beta_p} \quad (1)$$

The value for $\alpha_p(t_0)$ is then the supremum of this set of legitimate β_p values. The pointwise Hölder function of Hölder series is then the time series of these values, $\alpha_p(t)$.

The algorithm developed by Schreiber and Schmitz for testing the nonlinear structure of time series [28] was termed the iterated, amplitude-adjusted Fourier transform or IAAFT. In order to preserve the original values in the time series, our base algorithm also includes an iterated amplitude adjustment step, leading to an approach we termed the iterated, amplitude adjusted wavelet transform (IAAWT) [45]. There are very strong similarities between the structure of the two algorithms but there is a fundamental difference: while the IAAFT approach is based on preservation of Fourier amplitudes and randomization of Fourier phases, the IAAWT is based on preservation of the amplitudes of wavelet coefficients and randomization of their phases. As a well-known stochastic method for generating multifractal signals is based on the hierarchical generation of wavelet coefficients on a dyadic tree that are then inverted using a suitable real-valued wavelet transform [13], and as the real, first derivative-of-Gaussian wavelet transform underpins the wavelet transform modulus maxima approach to characterising multifractality [2], we can adopt a complex-valued wavelet transform and retain the

wavelet amplitudes while randomising the phases to preserve the Hölder function. This is analogous to the retention of Fourier amplitudes and randomisation of phases in the IAAFT. Key concepts in wavelet-based multifractal analysis are reviewed by Venugopal et al. [7] while mathematical details are given by Jaffard [50]. The key idea follows from the Frisch-Parisi conjecture [46] that

$$D(\alpha_p) = \min_m (\alpha_p m - \xi_m + 1). \quad (2)$$

Thus, the singularity spectrum that characterises the set of Hölder exponents in the signal, $D(\alpha_p)$, is given by an inverse Legendre transform involving moment order m , and ξ_m the structure function for the time series for moment, m . Given two values for the time series, y_t and $y_{t+\tau}$, the structure function is given by

$$\langle |y_t - y_{t+\tau}|^m \rangle \propto |\tau|^{\xi_m}. \quad (3)$$

With a wavelet transform providing a multiresolution analysis as a function of the scale, j , (a surrogate for τ) and the coefficients of a first derivative-of-Gaussian wavelet at scale j characterising the increments, $w_{t,j} \sim |y_t - y_{t+\tau}|$ then following Jaffard (1997), if we are close to a singularity of order α_p , we will find that in a window, $|\tau|$, that the local behaviour scales as

$$|y_{t+\tau} - y_t|^m \approx |\tau|^{\alpha_p m}. \quad (4)$$

With a dimension to these singularities of $D(\alpha_p)$ there are approximately $|\tau|^{-D(\alpha_p)}$ boxes with size $|\tau|$ that cover the signal. Hence, the contribution of this singularity to the integral used to evaluate the structure function $\langle |y_t - y_{t+\tau}|^m \rangle$ is approximately $|\tau|^{\alpha_p m + 1 - D(\alpha_p)}$. The key result to link the $w_{t,j}$ to the increments explicitly is the partition function [51], which for the wavelet transform modulus maxima method [2], is given by

$$Z(q, j) = \sum_{\ell} |w(j, \xi_{\ell})|^q, \quad (5)$$

where $q \in \mathfrak{R}$ is a selected power that measures the scaling behaviour of $Z(q, j)$, ξ is a maximum of the wavelet transform modulus maxima, and ℓ is an index for these maxima. Scaling exponents are calculated by

$$\phi(q) = \liminf_{j \rightarrow 0} \frac{\log Z(q, j)}{\log j}, \quad (6)$$

and these scaling exponents can be related to the support of the multifractal distribution via a Legendre transform [48, 50]:

$$\phi(q) = \min[q(\alpha_p + 1/2) - D(\alpha_p)]. \quad (7)$$

2.2. The base algorithm for GMR - the IAAWT algorithm

The key idea underpinning the IAAWT is therefore the phase-randomization of complex-valued wavelet coefficients. The analytic signal of some real signal, $y(t)$ is given by $y_a(t) = y(t) + i y_{(H)}(t)$, where $y_{(H)}(t)$ is the Hilbert transform of $y(t)$, which is a convolution operator with a filter given by $h(t) = 1/(\pi t)$:

$$y_{(H)}(t) = \int_{-\infty}^{\infty} h(\beta) y(t - \theta) d\beta. \quad (8)$$

Because the Fourier transform of $h(t)$ lies completely in the imaginary plane, it follows that a Hilbert transform approach can be used to perform a complex-valued wavelet transform. A pair of dyadic wavelet trees may be designed to form a Hilbert transform pair [52] and we employ the dual-tree complex discrete wavelet transform [53, 54]. Given two filters $g(t)$ and $h(t)$ and their Fourier transforms $G(\omega)$ and $H(\omega)$ then it may be shown that if $G(\omega) = H(\omega)e^{-i\omega/2}$ for $|\omega| < \pi$ then their associated wavelets form a Hilbert pair [55], which can be achieved for orthogonal wavelets by offsetting the scaling filters by one half sample. In order to accomplish this operation effectively, Kingsbury proposed the Q-shift dual tree where, below the coarsest scale, all filters are even length, but no longer linear in phase [53]. By designing the filters to have a delay of $\frac{1}{4}$ sample and by using the time reverse of one set of filters in the other tree, the required $\frac{1}{2}$ sample delay can be achieved. In this paper we use symmetric, biorthogonal filters with support widths of 13 and 19 values for the first level of the algorithm and Q-shift filters with a support of 14 values for all other levels on the dual tree. The Q-shift dual tree approach retains properties that make undecimated transforms advantageous for use in surrogate generation, such as shift invariance [56], but at a computational cost that is merely double that for a standard discrete wavelet transform (DWT).

The IAAWT algorithm may now be stated. Given a signal, $y(t)$ of length $N = 2^J$:

1. Undertake the dual-tree complex DWT and obtain amplitudes, $A_{k,j}$ and phases, $\omega_{k,j}$ over all J scales for the $k = 1, \dots, 2^{J-j}$ complex valued $w_{k,j}^{(0)}$ at each j ;
2. Randomly sort the original time series to give a time-series, $z^{(0)}$;
3. Take its dual-tree complex DWT to derive randomised wavelet phases, $\omega_{k,j}^{(0)}$ for each scale and position;
4. Produce new $w_{k,j}^{(1)}$ by combining the original amplitudes with the randomised phases:

$$w_{k,j}^{(1)} = A_{k,j} \exp(i\omega_{k,j}^{(0)}) \quad (9)$$

5. Iterate the following steps until a convergence criterion is met, where at each step, s :
 - (a) Take the inverse wavelet transform to give a new time series, $z^{(s)}(t)$ and then apply the amplitude adjustment step used in the IAAFT algorithm, where a mapping is established between $y(t)$ and $z^{(s)}(t)$ by rank-order matching to permit the values of $z^{(s)}$ to be replaced by the value in $y(t)$ with the same rank;
 - (b) Take the dual-tree complex DWT and obtain the new phases, $\omega_{k,j}^{(s)}$. Combine these with the original amplitudes, $A_{k,j}$ to give the $w_{j,k}^{(s+1)}$ using the s 'th iterated variant of eq. (9).

2.3. Relevant principles from gradual wavelet reconstruction

Given the IAAWT base model [45], instead of the IAAFT [28], the ethos of GMR is the same as that for GWR. Hence, the key principles underpinning GWR are stated:

- a. A control parameter (ρ for GWR) defines a continuum between 0 (Fourier phase randomised data) and 1 (the original data);

- b. Surrogate data are derived for various choices of ρ using a maximal overlap discrete wavelet transform [57];
- c. We express ρ as a proportion of the total energy content, E , of the wavelet coefficients of signal:

$$E_\rho = \sum_{k=1}^N \sum_{j=1}^J |w_{k,j}|^2. \quad (10)$$

It is important to note that with such an undecimated transform, each $w_{k,j}$ has the same weight in the analysis and there are a total of $J \times N$ coefficients. With all $|w_{k,j}|^2$ placed in descending rank order (with their position indexed by r), we seek the smallest value for r such that

$$\frac{\sum_{r=1}^{J \times N} |w_r|^2}{E_\rho} \geq \rho; \quad (11)$$

- d. These coefficients are fixed in place in the template of wavelet coefficients and the others are randomised for each scale, j in turn in an analogous way to the IAAFT algorithm but applied in the wavelet domain [56, 30];
- e. The inverse wavelet transform of this set of coefficients then gives a time-series with a constrained degree of phase randomisation;
- f. With a set of suitably constrained surrogate time series at different ρ , given a metric for nonlinearity or multifractality, the value for the data is compared to that for the surrogates. The values for ρ that yield no significant difference and those where there is a significant difference are delimited.

An example of gradual wavelet reconstruction for the increment skewness of the height of bed-forms advected beneath a probe in a hydraulic flume [32] is given in Fig. 1. The saw-tooth nature of these bed-forms is an intrinsic part of their morphology and is destroyed in the IAAFT surrogates at $\rho = 0$. For $0.6 \leq \rho < 0.9$ there is a sudden increase in the increment skewness as the scales of features responsible for the saw-tooth pattern begin to get ‘locked in’ to the surrogates. Subsequent analysis showed that these features were ripple-like perturbations sitting on top of the primary bed form [32]. For $\rho \geq 0.9$ there are minimal differences between data and surrogates. Undertaking similar analyses for different water discharges (as a surrogate variable for hydraulic forcing) showed that this threshold value for ρ increased with discharge.

2.4. The GMR algorithm

From 2.2 and 2.3 the key aspects of the GMR algorithm are apparent: We will define a control parameter (φ instead of ρ) and with the IAAWT algorithm used to populate $\varphi = 0$ surrogates, we will modify this algorithm to increasingly fix in place the wavelet phase information until, at $\varphi = 1$, no randomisation is possible and we recover the original data exactly. Hence, the final part to the algorithm is the formulation of the means to fix in place the phase information. Here we draw upon some of the concepts previously




Figure 1: : Gradual wavelet reconstruction of the normalized increment skewness for elevations of a water-worked bed created in a hydraulic flume [32] with a water discharge, Q , of 1600 l s^{-1} . The time increment of separation is 5s, or one sampling interval. As is the case throughout this study, boxplots show the median value for the surrogates (heavy line), the first and third quartiles (edges of the box), with whiskers extending up to the limit for the surrogates or 1.5 times the interquartile deviation. Values outside this limit are shown as crosses. In order to replicate the observed degree of increment skewness, given by the horizontal dotted line, a threshold of $\rho_{thresh} = 0.90$ is required. These results are extracted from Fig. 7 of Keylock, C. J., A. Singh, and E. Foufoula-Georgiou (2014), The complexity of gravel bed river topography examined with gradual wavelet reconstruction, *J. Geophys. Res. Earth Surf.*, 119, 682–700, doi:10.1002/2013JF002999 (copyright American Geophysical Union) and are reproduced with the permission of the AGU.

adopted in GWR, but need to modify the energy measure to account for the decimated nature of the dual tree complex transform:

$$E_\varphi = \sum_{k=1}^K \sum_{j=1}^J \frac{|w_{k,j}|^2}{2^j} \quad (12)$$

That is, with $j = 1, \dots, J$ scales, there are $K = 2^{J-j}$ coefficients at each scale, meaning that more energy will be associated with each $w_{k,j}$ on average at the larger j , necessitating the introduction of the denominator. By a similar ranking procedure to (11) we may then seek the smallest r such that $\frac{\sum_{r=1}^{K \times N} |w_r|^2}{E_\varphi} \geq \varphi$. This selected set of $w_{k,j}$ are fixed in place on the wavelet coefficient template, while the others are phase randomised using eq. (9). Hence, while all $w_{k,j}^m$ in the IAAWT algorithm have the same amplitudes as the original $w_{k,j}$, the degree of phase randomisation reduces as φ increases and the number of coefficients exceeding the energy criterion increases.

3. Application to studying financial indices as self-regulating processes

We study the normalized log-returns of the closing prices from eight international financial indices and attempt to place confidence on any observed coupling between the

Table 1: The closing price indices studies in this paper using GMR.

Index	Country	$\langle \vec{y}(t) \rangle$ $\times 10^{-3}$	$\sigma[\vec{y}(t)]$ $\times 10^{-3}$
AORD	Australia	0.06	4.3
BSE Sensex	India	0.19	6.5
GDAX	Germany	0.05	6.7
Hang Seng	China	0.02	6.5
JKSE Composite	Malaysia	0.22	6.1
NASDAQ 100	USA	0.02	7.6
Nikkei 225	Japan	0.01	6.8
Standard and Poor's 500	USA	0.04	5.4

returns, $y(t)$ and their pointwise Hölder exponents, $\alpha_p(t)$ using GMR. We consider two measures of any such relation: the cross-correlation between $y(t)$ and $\alpha_p(t)$; and, the Hilbert-transform-based phase difference. With $N = 4096$ days in each case, the cross-correlation function will be of length $2N - 1$ and the phase difference time-series will be of length N . We therefore need a summary metric or set of metrics that characterize these measures. As detailed below, we use the maximum and minimum values of the cross-correlation function for the former case, and the mean and skewness of the phase difference distribution in the latter case. In this section we detail the calculation procedures applied to the data and GMR surrogates. The results of the analysis are considered in section 4.

3.1. The indices and calculation of normalized log-returns

The eight financial indices examined in this study are listed in Table 1 together with statistical moments before normalization. In each case, we took the closing prices from 5th September 2000 to 15th December 2016, to give 4097 values. We obtain the normalized log returns, $y'(t)$, [58] according to

$$\begin{aligned} \vec{y}(t) &= \ln[y(t + \Delta)/y(t)] \\ y'(t) &= \frac{\vec{y}(t) - \langle \vec{y}(t) \rangle}{\sigma[\vec{y}(t)]}, \end{aligned} \quad (13)$$

where the angle braces indicate a time average, $\sigma[\dots]$ is the standard deviation of the variable and we set Δ to one day. These data are shown in Fig.2.

Previous work on the analysis of market efficiency with multiple indices over a similar time period has demonstrated a change in the Hölder exponent with the 2008 change in market conditions for both corporate bond indices and sovereign bond indices [59]. Hence, our study explores a complementary hypothesis regarding a change in coupling between Hölder exponent and values for the log-returns.

3.2. Hölder exponent estimation, and measures and metrics of the normalized log-returns

The method we use to derive $\alpha_p(t)$ is a time-domain method [60] that has seen several applications to nonlinear time-series [61, 62] and follows from the form of eq. (1): we




Figure 2: The normalized log returns, $y^l(t)$ for the eight indices considered in this study.

undertake a log-log regression of the signal oscillations, $O_{T\pm\tau}$, within some distance τ of T against τ , where $O_{T\pm\tau}$ is given by:

$$O_{T\pm\tau} = \max(y_{l \in (T-\tau, \dots, T+\tau)}) - \min(y_{l \in (T-\tau, \dots, T+\tau)}) \quad (14)$$

and τ is distributed logarithmically (from 2^1 to 2^{10} in this study).

Figure 3 illustrates the normalized log-returns, $y^l(t)$ and the accompanying Hölder exponents, $\alpha_p(t)$ for the NASDAQ 100 index for $\varphi \in \{0.0, 0.3, 0.6\}$ (surrogate data given as grey lines). With the original data (black lines), by definition at $\varphi = 1.0$ four approximately equal sub-divisions of φ for the surrogates are shown over the permitted range. The selected region of the data highlights the 2008 crash. Note that the $\alpha_p(t)$ are replicated extremely well for all φ , which is the essential property of the IAAWT algorithm. As with the IAAFT algorithm, this is not quite perfect because of the imposition of the original data values following the inverse transform [28]. As φ increases, the $y^l(t)$ become more similar to the original values. One can see that the positive spike at $t = 2047$ days is negative in the surrogates (dotted line) for $\varphi = 0$. As φ increases, the sign (at $\varphi = 0.3$) and then shape (at $\varphi = 0.6$) of this feature is replicated in the surrogates.

In order to characterize the relation between the $y^l(t)$ and the $\alpha_p(t)$, we use two different measures: the cross-correlation function and the phase differences. The former is given for a lag time, τ by

$$R(\tau) = \frac{\sum_{\tau=-(N-1)}^{N-1} (y^l(t+\tau) - \langle y^l \rangle)(\alpha_p(t) - \langle \alpha_p \rangle)}{\sigma(y^l)\sigma(\alpha_p)}. \quad (15)$$




Figure 3: : Gradual multifractal reconstruction algorithm applied to the NASDAQ 100 index for three choices of φ , showing the central region of the 4096 values, and the crash of 2008. The left hand column shows the normalized log-returns, $y'(t)$ and the right-hand column their $\alpha_p(t)$ (black lines) while the grey lines are the surrogate series with the median convergence error. The surrogate series are offset vertically by -15 in the left-hand column and -0.1 in the right-hand column. The vertical dotted line at $t = 2047$ is discussed in the text.

The metrics we adopt in this study to summarize the behaviour of this measure are the maximum and minimum of $R(\tau)$ over all $2N - 1$ values for τ , which we denote by R_{max} and R_{min} , respectively.

Our phase difference approach uses the Hilbert transform (8) to obtain complex valued variants of the y' and corresponding α_p . Thus, the phases for our two signals are

$$\begin{aligned}\psi_y(t) &= \tan^{-1} \frac{y'_{(H)}(t)}{y'(t)} \\ \psi_\alpha(t) &= \tan^{-1} \frac{\alpha_{p(H)}(t)}{\alpha_p(t)},\end{aligned}\tag{16}$$

with their difference given by $\theta(t) = \psi_y - \psi_\alpha$. The metrics we adopt for this are based on the first (mean) and third (skewness) moments of the distribution function, $p(\theta)$.

4. Results

4.1. Correlation structure in the data summarised using proper orthogonal decomposition (POD)

We examine the correlation structure in the data using proper orthogonal decomposition/ principal components analysis [63, 64] as a means to summarize relations be-




Figure 4: : Proper Orthogonal Decomposition (POD) of the normalized log returns (top row) and the Hölder series (bottom row) of the eight indices analysed in this paper. The POD is based on the correlation matrix and the first three components are shown, explaining 49.6% of the total variance for the returns and 86.6% for the Hölder series.

tween indices. All indices are already suitably normalized for a correlation rather than covariance-based analysis as shown in (13). We looked at the inter-relations between the returns and between the Hölder functions in the top and bottom rows of Fig. 4, respectively. The correlations were stronger in the latter case as reflected in the higher proportion of the total variance accounted for in the first three components. The first component in the top row of Fig. 4 clearly separates the two American indices from the others, while component 2 contrasts the Nikkei and BSE with the Hang Seng (and JKSE to a lesser extent). The third component in the top row, and to a lesser extent, the second, highlights relations between the German GDAX and Australian AORD indices. The first two components of the Hölder POD isolate this latter index as particularly distinct, with the largest magnitude values on component 1 again being the closely correlated American indices. However, the GDAX and Hang Seng indices are not dissimilar to the American indices on this component. Component 2 contrasts the Australian and Japanese indices, while component 3 contrasts the small Asian markets (BSE and JKSE) with the European (GDAX) and, to a lesser extent, American markets. Hence, over both forms of classification, and the first three components, all the different indices, except the two from the USA, can be separated on some combination of axes.

4.2. Cross-correlations between normalized log-returns and Hölder series

Figure 5 shows the cross-correlation function (in black) for four of the eight indices, with results shown for three time series windows, corresponding to the periods before, during and after the stock market crash of 2008. The equivalent functions for each of the thirty nine surrogate datasets is shown in grey, where $\varphi = 0$ (i.e. the generation is based on the IAAWT algorithm [45]). In none of these cases, or for the eight indices, was the minimum of $R(\tau)$, R_{min} , less than the values for the surrogates, indicating there is no

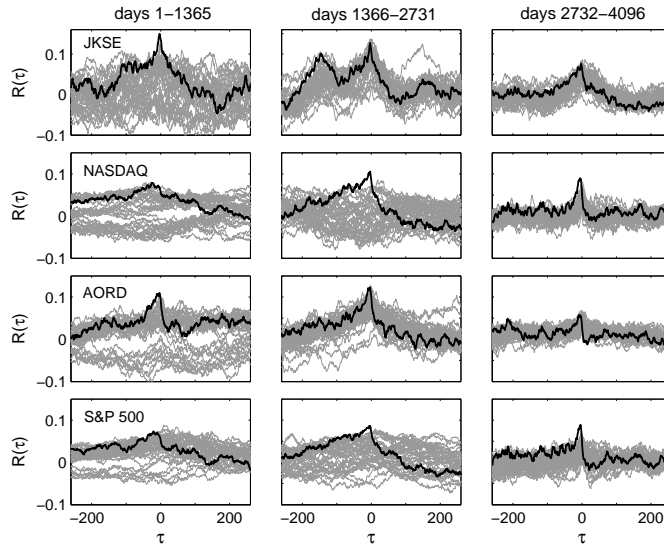


Figure 5: The cross-correlation functions, $R(\tau)$ (in black) for four of the eight indices and three different time windows. The cross-correlation functions for each of the 39 surrogates generated at $\varphi = 0$ are shown as grey lines in each panel. The range of values shown for the lag, τ correspond to \pm a year of trading.

significant negative correlation between the y^l and $\alpha_p(t)$. This is interesting because the maximum and minimum observed cross-correlations are given in Fig. 6, and an effect of the crash appears to be a stronger negative correlation for days 1366-2731 than is seen either before or afterwards. Clearly, given the values in the time series, it is possible to replicate these correlations at random during this period given the multifractal structure, which has a smaller Hölder exponent on average than the earlier and later data as is discerned from Fig. 2. In contrast, six significant differences for the maximum value, R_{max} can be seen in Fig. 5 and these cases, along with the three instances that arose in the other four datasets, are highlighted by solid symbols in Fig.6.

There are three primary aspects of the results for R_{max} :

1. There is no simple relation between the magnitude of R_{max} and whether or not the observed result is significant;
2. The two American indices exhibit a similar behaviour, with no significant R_{max} before the crash, and then a significant R_{max} emerging during the crash and persisting afterwards; and,
3. The three smallest markets (AORD, BSE Sensex, JKSE composite) had the highest R_{max} before the crash (two of which were significant), while since the crash, the two U.S. indices have the largest (and significant) R_{max} .

Of course, all data have different autocorrelation scales for both the log returns and their Hölder exponents, so the first conclusion is not unexpected. However, note that, for a




Figure 6: Values for the maximum and minimum of the cross-correlation function as a function of each index and three non-overlapping windows with a time base of 1365 or 1366 days. Solid symbols (or for the case of the JKSE index, outer circles) indicate that a significant difference was determined at $\varphi = 0$. The first value for φ at which R_{max} for the surrogates exceeds that for the data (φ_{thresh}) is stated next to these points.

given window on the data, there is a relation between R_{max} and φ_{thresh} , the value of φ from the GMR at which a significant difference is no longer observed. The second conclusion is useful because we would expect that the NASDAQ 100 and S&P 500 indices exhibit similar behaviour because they are trying to characterise the same market, and the POD analyses demonstrated this empirically. Hence, this result indicates a certain robustness to our method. Note too that these values for R_{max} since the crash are, along with the JKSE value before the crash, the most significant according to GMR. The NASDAQ value for R_{max} exceeds that for the maximum of the surrogates even at $\varphi = 0.99$ in the most recent period.

A distinct feature of the results for the smaller asian markets in particular, which is seen in Fig. 5 for the JKSE index, is the existence of a secondary peak to the cross-correlation function at negative τ in the middle data window, which is not seen in the Australian, European or American markets. All the middle window $R(\tau)$ functions are given in Fig. 7, with the asian markets in the right-hand column. This feature is very prominent in the BSE and JKSE indices at $\tau \sim -150$ days, with a less pronounced occurrence in the larger Hang Seng and Nikkei indices at $\tau \sim -150$.

4.3. Phase differences between normalized log-returns and Hölder series

Histograms for the phase differences between the y^l and the $\alpha_p(t)$ are given in Fig. 8 for each of the three data windows. For the Hang Seng, NASDAQ, S&P 500 and, to a lesser extent, the GDAX and Nikkei, before the crash there is a very broad peak to the θ values extending from $0 < \theta < \pi$ (solid black lines). During the crash, the degree of skewness drops giving an almost symmetrical distribution with a mode at $\theta \sim 0$ for the smaller markets (AORD, BSE, JKSE) and slightly positive for the other cases. More




Figure 7: The cross-correlation functions for the middle window (days 1366-2731) for all eight indices. The asian markets are in the right-hand column.

recently, the phase differences are primarily negative for the asian and american markets, with this tendency less clear in the AORD and GDAX markets.

These observations suggest that the mean and skewness of θ are useful test statistics to study the significance of these patterns using GMR. Significant differences for the mean of θ at $\varphi = 0$ were found for the middle window of the S&P 500 index and the middle and latter window for the Nikkei 225 index. These differences for the Nikkei index were resolved by $\varphi = 0.1$ as is shown by the full GMR of these cases in Fig. 9, while that for the S&P 500 persisted until $\varphi = 0.9$.

For the phase skewness, significant differences were observed for the GDAX index for the first and second time periods ($\varphi_{thresh} = 0.7$, $\varphi_{thresh} = 0.95$, respectively), the most recent period for the JKSE Composite ($\varphi_{thresh} = 0.3$), the middle period for the Nikkei 225 index ($\varphi_{thresh} = 0.1$), and then the most recent two periods for the two american indices. The full GMR for these two cases is shown in Fig. 10 and the absence of a significant difference for the first time period in both cases is clear. Results for the period of the crash appear qualitatively similar at first glance, with the boxplots converging on the skewness of the data by $\varphi = 0.8$ for the NASDAQ 100 index and $\varphi = 0.95$ for the S&P 500 index. However, for the latter case, there is an outlier surrogate dataset at $\varphi = 0.3$ that attains the skewness in the data. For the most recent time period, differences persist at $\varphi = 0.99$ and what the surrogate series show for both of these indices is that, conditioned on the observed multifractal structure, the anticipated skewness of the phase differences is positive for not just random data values ($\varphi = 0.0$), but also with a large proportion of the wavelet energy fixed in place. In contrast, the observations have a negative skewness.

Finally, we consider the GMR of the phase difference skewness calculated over the full 4096 days rather than in the three windows used previously. We are able to demonstrate a marked geography to the results as is shown by contrasting the four sets of results in Fig. 11 with those in Fig. 12. For the JKSE Composite and BSE Sensex indices we

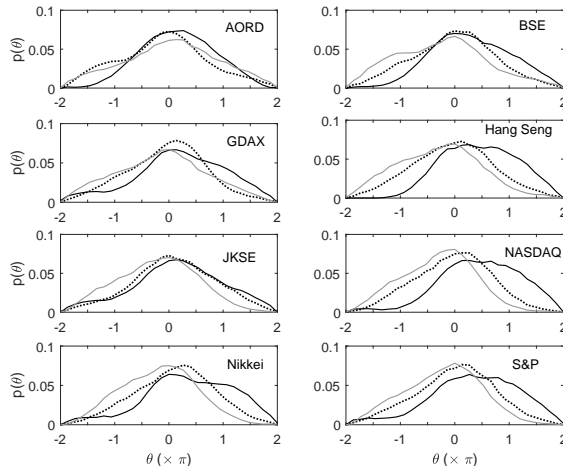


Figure 8: Histograms for the phase differences, θ for the eight indices. The black solid, black dotted and thin grey lines correspond to analysis over days 1-1365, 1366-2731 and 2732-4096, respectively.

see in Fig. 11 that there is no significant difference between the θ skewness values at $\varphi = 0$ for data and surrogates. For the Nikkei 225 and Hang Seng indices, a significant difference at $\varphi = 0$ does not persist and is absorbed by $\varphi_{thresh} = 0.1$. Hence, the skewness of the phase differences is either not significant or small for the asian markets. In very clear contrast, the bottom three rows of Fig. 12 show that there is a significant difference that persists beyond $\varphi = 0.99$ for the european and american markets. The australian case would appear to be intermediate in nature: at $\varphi = 0$ it is possible to generate outlier surrogates with the appropriate negative skew to the phase differences. However, imposition of some wavelet phase structure results in significant differences that persist until $\varphi = 0.95$. How these differences reflect differences in trading behaviour as well as the nature and strength of inter-market coupling would be a fascinating topic for future research.

5. Conclusion

This paper has formulated a framework for testing hypotheses with regards to the coupling between a time-series and its accompanying series of Hölder exponents, phenomena that have been termed self-regulating processes [37, 39]. The formulation is termed gradual multifractal reconstruction (GMR) and is not dissimilar to the gradual wavelet reconstruction (GWR) surrogate data approach to studying nonlinearity in time series [30]. In both cases a continuum is populated with surrogate data that converge on the relevant properties of the original data as a control parameter increases in value. However, while GWR uses the iterated, amplitude adjusted, Fourier transform model [28] to populate the phase-randomised end of the continuum, GMR uses a recent algorithm that sets the Hölder series as constant but randomises the wavelet phases [45] as the base model for the $\varphi = 0$ end of the continuum. The GMR algorithm then systematically




Figure 9: Gradual multifractal reconstruction of the mean phase difference, $\langle \theta \rangle$ for the Nikkei 225 index determined for three time periods: days 1-1365, 1366-2731 and 2732-4096, respectively.

reduces the degree of phase randomisation as $\varphi \rightarrow 1$, until no randomisation arises in the surrogates at the limit of $\varphi = 1$. Thus, for a particular time-series we may ask not only is there a statistically significant degree of coupling between the time-series and its Hölder function, given the latter and the original distribution of values for the former, but we can determine if two time-series have different strengths of coupling by checking to see if the value for φ at which a significant difference no longer arises (φ_{thresh}) itself differs.

Our example application is to 4096 normalized log returns, $y^l(t)$ for eight financial indices from around the globe centred on the 2008 crash, and their Hölder series, $\alpha_p(t)$. GMR helps us to show that while it would appear that the minimum of the cross-correlation between $y^l(t)$ and $\alpha_p(t)$ increases in strength around the period of the crash, this is not a significant effect. What is particularly significant according to GMR is the transition for the two US indices from a lack of a significant maximum cross-correlation from before the crash to a very significant maximum cross-correlation at the present time. When the Hilbert transform-based phase difference histograms are examined this is seen as a change from a mean positive phase difference to one where the mode of the distribution is close to zero but with strong negative skew. Indeed, while the surrogates tend to give a positive skewness for even high values of φ , the observed skewness for all the markets, when all 4096 days are considered, is negative. However, GMR reveals a marked geography to this result, with this negative skewness statistically significant for the GDAX, NASDAQ 100 and S&P 500 indices at $\varphi = 0.99$, but not for the asian indices beyond $\varphi = 0.1$. The Australian AORD index exhibits characteristics of both regions as might be anticipated. There would appear to be a weak market size effect as the Nikkei 225 and Hang Seng indices are the two Asian markets where $\varphi_{thresh} = 0.1$ with $\varphi_{thresh} = 0$ for the smaller Indian and Malaysian markets. However, the regional effect clearly dominates this relatively small difference.

Consequently, we can suggest that GMR provides a potentially useful tool in many




Figure 10: Gradual multifractal reconstruction of the skewness of the phase differences for the NASDAQ 100 and S&P 500 indices and for three time periods: days 1-1365, 1366-2731 and 2732-4096, respectively.

areas of nonlinear physics where the significance of any potential coupling between a time series and its Hölder regularity may be important both for process insights and for predicting extreme events. Turbulence physics and biomedical time series are two additional fields where this is likely to be the case [45].

6. Acknowledgements

This research was supported by NERC grant NE/F00415X/1, EPSRC Grant EP/K007688/1, and Royal Academy of Engineering/Leverhulme Senior Research Fellowship LTSRF1516-12-89.

References

- [1] C. Meneveau, K. Sreenivasan, *Phys. Rev. Lett.* 59 (1987) 1424–1427.
- [2] J. F. Muzy, E. Bacry, A. Arnéodo, *Phys. Rev. Lett.* 67 (1991) 3515–3518.
- [3] K. Matia, Y. Ashkenazy, H. E. Stanley, *EPL* 61 (2003) 422–428.
- [4] S. Drozd, J. Kwapien, P. Oswiecimka, R. Rak, *EPL* 88 (2009) 60003.
- [5] Y. Wei, Y. Wang, D. Huang, *Physica A* 390 (2011) 4260–4272.
- [6] J. Gagnon, S. Lovejoy, D. Schertzer, *Europhys. Lett.* 62 (2003) 801–807.
- [7] V. Venugopal, S. G. Roux, E. Foufoula-Georgiou, A. Arneodo, *Water Resour. Res.* 42 (2006).
- [8] N. Sarkar, B. B. Chadhuri, *Signal Process.* 42 (1995) 181–190.
- [9] P. C. Ivanov, L. A. N. Amaral, A. L. Goldberger, S. Havlin, M. G. Rosenblum, Z. R. Struzik, H. E. Stanley, *Nature* 399 (1999) 461–465.
- [10] L. A. N. Amaral, P. C. Ivanov, N. Aoyagi, I. Hideka, S. Tomono, A. L. Goldberger, H.E. Stanley, Y. Yamamoto, *Phys. Rev. Lett.* 86 (2001) 6026–6029.




Figure 11: GMR of the phase difference skewness for the asian markets for all 4096 trading days studied.

- [11] B. Mandelbrot, J. W. van Ness, *SIAM Review* 10 (1968) 422–437.
- [12] R. Peltier, J. Lévy Véhel, *Multifractional Brownian motion: definition and preliminary results*, Technical Report 2645, INRIA, 1995.
- [13] R. Benzi, L. Biferale, A. Crisanti, G. Paladin, M. Vergassola, A. Vulpiani, *Physica D* 65 (1993) 352–358.
- [14] A. P. Nawroth, J. Peinke, *Phys. Lett. A* 360 (2006) 234.
- [15] T. D. Matteo, *Quantitative Finance* 7 (2007) 21–36.
- [16] A. Serletis, I. Andreadis, *Energy Economics* 26 (2004) 389–399.
- [17] J. Barunik, T. Aste, T. Di Matteo, R. Liu, *Physica A* 391 (2012) 4234–4251.
- [18] L. Zunino, B. Tabakd, A. Figliola, D. Pérez, M. Garavaglia, O. Rosso, *Physica A* 387 (2008) 6558–6566.
- [19] A. F. Bariviera, M. B. Guercio, L. B. Martinez, *The Econ. Soc. Rev.* 45 (2014) 349–369.
- [20] L. Calvet, A. Fisher, *J. Econometrics* 105 (2001) 27–58.
- [21] T. Lux, *J. Bus. Econ. Stat.* 26 (2008) 194–210.
- [22] J. F. Muzy, D. Sornette, J. Delour, A. Arneodo, *Quantitative Finance* 1 (2001) 131–148.
- [23] Y. Wei, P. Wang, *Physica A* 387 (2008) 1585–1592.
- [24] U. Frisch, P. L. Sulem, M. Nelkin, *J. Fluid Mech.* 87 (1978) 719–736.
- [25] D. Schertzer, S. Lovejoy, *J. Geophys. Res.* 92 (1987) 9693–9714.
- [26] Z. She, E. Leveque, *Phys. Rev. Lett.* 72 (1994) 336–339.
- [27] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, J. D. Farmer, *Physica D* 58 (1992) 77–94.
- [28] T. Schreiber, A. Schmitz, *Phys. Rev. Lett.* 77 (1996) 635–638.
- [29] T. Schreiber, A. Schmitz, *Phys. Rev. E* 55 (1997) 5443.
- [30] C. J. Keylock, *Nonlinear Proc. Geophys.* 17 (2010) 615–632.
- [31] C. J. Keylock, *Physica D* 225 (2007) 219–228.
- [32] C. J. Keylock, A. Singh, E. Foufoula-Georgiou, *J. Geophys. Res.* 119 (2014) 682–700.
- [33] C. J. Keylock, S. N. Lane, K. S. Richards, *J. Geophys. Res.* 119 (2014) 264–286.
- [34] C. J. Keylock, *Water Resour. Res.* 48 (2012).
- [35] C. J. Keylock, T. E. Tokyay, G. Constantinescu, *J. Turbul.* 12 (2011) N45.
- [36] C. J. Keylock, R. Stresing, J. Peinke, *Phys. Fluids* 27 (2015) 025104.
- [37] O. Barrière, A. Echelard, J. Lévy Véhel, *Electron. J. Prob.* 17 (2012) 1–30.
- [38] J. Lévy Véhel, *Nonlin. Proc. Geophys.* 20 (2013) 643–655.
- [39] A. Echelard, J. Lévy Véhel, A. Philippe, *Scand. J. Stat.* 42 (2015) 485–503.
- [40] A. Arnéodo, E. Bacry, S. Jaffard, J. F. Muzy, *J. Fourier Anal. Appl.* 4 (1998) 159–174.
- [41] F. Nicolleau, J. C. Vassilicos, *Phil. Trans. R. Soc. Lond., Ser. A* 357 (1999) 2439–2457.




Figure 12: GMR of the phase difference skewness for the non-asian markets for all 4096 trading days studied.

- [42] J. M. Aubry, S. Jaffard, *Comm. Math. Phys.* 227 (2002) 483–514.
- [43] S. Seuret, *Math. Nachr.* 279 (2006) 1195–1211.
- [44] M. Paluš, *Phys. Rev. Lett.* 101 (2008) 134101.
- [45] C. J. Keylock, *Phys. Rev. E* 95 (2017) 032123.
- [46] U. Frisch, G. Parisi, in: M. Ghil, R. Benzi, G. Parisi (Eds.), *Turbulence and Predictability in Geophysical Fluid Dynamics and Climate Dynamics*, North-Holland, Amsterdam, 1985, pp. 84–88.
- [47] R. Benzi, S. Ciliberto, R. Tripicciono, C. Baudet, S. Succi, *Phys. Rev E* 48 (1993) 29.
- [48] E. Bacry, J. Muzy, A. Arnéodo, *J. Stat. Phys.* 70 (1993) 635–674.
- [49] J. W. Kantelhardt, S. A. Zschiegner, E. Koscielny-Bunde, S. Havlin, A. Bunde, H. E. Stanley, *Physica A* 316 (2002) 87–114.
- [50] S. Jaffard, *SIAM J. Numerical Anal.* 28 (1997) 944–998.
- [51] M. Holschneider, *L’analyse d’objets fractals et leur transformation en ondelettes*, Ph.D. thesis, Univ. of Aix-Marseille II, Marseille, France, 1989.
- [52] I. Selesnick, *IEEE Trans. Sig. Proc.* 50 (2002) 1144–1152.
- [53] N. Kingsbury, *Appl. Comput. Harmon. Anal.* 10 (2001) 234–253.
- [54] I. Selesnick, R. Baraniuk, N. Kingsbury, *IEEE Signal Proc. Mag.* 22 (2005) 123.
- [55] I. Selesnick, *IEEE Sig. Proc. Lett.* 8 (2001) 170–173.
- [56] C. J. Keylock, *Phys. Rev. E* 73 (2006) 036707.
- [57] D. B. Percival, A. T. Walden, *Wavelet Methods for Times Series Analysis*, Cambridge University Press, Cambridge, U.K., 2000.
- [58] S. Kumar, N. Deo, *Physica A* 388 (2009).
- [59] A. F. Bariviera, M. B. Guercio, L. B. Martinez, *Econ. Lett.* 116 (2012) 426–428.
- [60] K. Kolwankar, J. Lévy Véhel, *J. Fourier Anal. Appl.* 8 (2002) 319–334.
- [61] C. J. Keylock, *Geophys. Res. Lett.* 35 (2008) L11804.
- [62] C. J. Keylock, K. S. Chang, G. S. Constantinescu, *J. Fluid Mech.* 805 (2016) 656–685.
- [63] G. Berkooz, P. Holmes, J. Lumley, *Annu. Rev. Fluid Mech.* 25 (1993) 539–75.
- [64] I. T. Jolliffe, *Principal Component Analysis*, Springer, Berlin, 2002.