DIRECTIONALITY IN TIME SERIES AND ITS APPLICATIONS

Mohd Mahayaudin bin Mansor

June 2017

Thesis submitted for the degree of
Doctor of Philosophy
in
Statistics
at The University of Adelaide
Faculty of Engineering, Computer and Mathematical Sciences
School of Mathematical Sciences

Contents

Contents iii

List of Tables

List of Figures

Signed Statement

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name, in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name, for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint-award of this degree.

I give consent to this copy of my thesis when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I acknowledge that copyright of published works contained within this thesis resides with the copyright holder(s) of those works.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Signed by: Mohd Mahayaudin bin Mansor
Date: 12 April 2018

Published (or submitted) works within the thesis

Mansor, M. M., Glonek, M. E, Green, D. A. and Metcalfe, A. V. (2016), *Time Series Analysis and Forecasting: Contributions to Statistics*, Springer International Publishing Switzerland, chapter "Threshold Autoregressive Models for Directional Time Series", pp. 13–25.

Mansor, M. M., Mohd. Isa, F. L, Green, D. A. and Metcalfe, A. V. (2016), "Modelling Directionality for Paleoclimatic Time Series", *ANZIAM J.* **57** *(EMAC 2015)*, C66–C81.

Mansor, M. M., Green, D. A. and Metcalfe, A. V. (2016), "Directionality and Volatility in Electroencephalogram Time Series", *American Institute of Physics (AIP) Conf. Proc.* **1739**, 020080:1–8.

Mansor, M. M., Glonek, M. E, Green, D. A. and Metcalfe, A. V. (2015), "Modelling Directionality in Stationary Geophysical Time Series", *in Proceedings of the ITISE 2015*, pp. 755–766.

Mansor, M. M., Green, D. A. and Metcalfe, A. V. (2015), "Modelling and Simulation of Directional Financial Time Series, *in Proceedings of the MOD-SIM 2015*, pp. 1022–1028.

Mansor, M. M., Green, D. A. and Metcalfe, A. V., "Directionality in Time Series of Bank Log-return Share Prices", *Journal of Business & Economic Statistics*, submitted.

Mansor, M. M., Green, D. A. and Metcalfe, A. V., "Directionality and Volatility in High Frequency Time Series", *High Frequency*, submitted.

Mansor, M. M., Green, D. A. and Metcalfe, A. V., "Detecting Directionality in Time Series", *The American Statistician*, submitted.

Acknowledgements

I consider myself fortunate to be surrounded by people who are brilliant, successful and passionate about what they set out to do during the Study period. They are humble, helpful and very kind to me. They became part of my motivation and inspiration to achieve my dreams. I am grateful and thankful for that.

I would like to express my sincere appreciation to my colleagues during the Study period, in particular to Associate Professor Andrew V. Metcalfe, my principal supervisor, for his outstanding supervision and unwavering commitment to his duties. Not forgotten is my co-supervisor, Dr. David A. Green, for his encouragement and guidance throughout my candidature period.

I also gratefully acknowledge the Majlis Amanah Rakyat (MARA) for sponsoring my research studies at the School of Mathematical Sciences, The University of Adelaide.

To my beloved parents, Manshor bin Mat and Sharipah binti Hj. Harun, thank you for your prayers, unconditional love and continuous support over the years.

To my wife, Norirwani binti Abdul Halim, mother of our precious daughter, Medina Zahra Arosa binti Mohd Mahayaudin, thank you for all your everything.

To all my beloved siblings and friends in Malaysia and Australia, thank you very much for the support, advice, hugs and smiles.

To Judith Burgemeister, thank you for providing helpful editorial support at such short notice.

May longevity and good health always be with us. May we all be blessed and continue to be kind to one another.

Thank you.

Dedication

I dedicate this thesis to *Abah* and *Mak*.

Abstract

A suite of seven statistics to detect directionality in time series is presented. Applications from various disciplines including business, environmental science, finance and medicine are considered. Models that allow for directionality are proposed, and methods of fitting these models are investigated. Time series models that incorporate directionality provide more precise prediction limits and more realistic simulations than the models that do not. Potential practical applications include: providing evidence to support physical interpretations; directionality trading rules for investment portfolio; prediction of unstable financial periods; and possible early warning of epileptic seizures.

Chapter 1

Contextual Statement

1.1 Overview

The analysis of time series typically involves fitting a model. A time series model $\{X_t\}$ for $t = 0, \pm 1, \pm 2, \cdots$ is reversible if the joint distribution of

$$(X_t, X_{t+1}, \cdots, X_{t+r})$$

is identical to the joint distribution of

$$(X_{t+r}, X_{t+r-1}, \cdots, X_t)$$

for all $r = 1, 2, \cdots$ and $t = 0, \pm 1, \pm 2, \cdots$. It follows that a reversible time series model is necessarily stationary. A time series model is directional if it is not reversible (or time irreversible).

A crucial result (Weiss, 1975) is that an autoregressive moving average ARMA(p, q) defined by

$$X_t = \alpha_0 + \alpha_1 X_{t-1} + \cdots + \alpha_p X_{t-p} + \beta_0 \varepsilon_t + \beta_1 \varepsilon_{t-1} + \cdots + \beta_q \varepsilon_{t-q}$$

where $p, q > 0$ and $\{\varepsilon_t\}$ is a sequence of independent random variables with mean 0 (errors) is reversible if, and only if, the errors have an identical Gaussian distribution. The proof depends on the fact that a multivariate Gaussian distribution is completely defined by its mean and variance-covariance matrix. When $p = 0$ the result holds provided that $\beta_i \neq \pm\beta_{q-i}$ for all i. An example of these exceptional cases is:

$$X_t = \varepsilon_t + 2\varepsilon_{t-1} + \varepsilon_{t-2}$$

which is reversible for any distribution of errors, because the distribution of $\varepsilon_t + 2\varepsilon_{t-1} + \varepsilon_{t-2}$ is indistinguishable from the distribution of $\varepsilon_{t-2} + 2\varepsilon_{t-1} + \varepsilon_t$. Non-linear time series models are generally, but not necessarily, directional (Lawrance, 1991). For example,

$$X_t = \varepsilon_t + 2\varepsilon_{t-1}^3 + \varepsilon_{t-2}$$

is non-linear and reversible. It follows that a directional model requires non-linear terms or errors that are not independently and identically distributed (*iid*) with a Gaussian distribution.

This study focuses on stationary time series modelled by random variables. A stationary random process is directional if it is possible to distinguish between a realisation in its time order and the same realisation in

reverse time order (time-to-go). For some time series, it is possible to discern directionality by comparing a plot against time with a plot against time-to-go. An example is the annual sunspot numbers shown in Figure 1. The sunspots series is plotted twice. In the upper frame the number of sunspots is plotted against year from 1700 until 2014 (time order), and in the lower frame is reversed and the number of sunspots is plotted against year from 2014 until 1700 (reverse time order or time-to-go). A striking feature of these plots is that they are distinguishable; in the upper plot there is a tendency for rapid increases to a peak to be followed by slower recessions whereas in the lower plot the tendency is for slower increases and rapid recessions. The sunspot time series is said to show directionality.
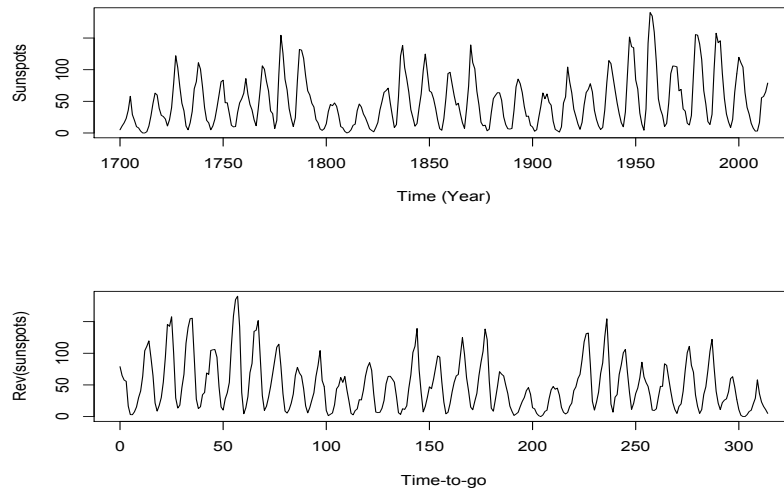
Figure 1: The sunspot series (1700-2014) has sharp increases before the peaks followed by slow recessions to the troughs in time order (above), and it has slow increases before the peaks followed by sharp decreases to the troughs in reverse time order (below).

There is no evidence of any trend or deterministic periodic pattern in the sunspots series, but many time series do have changes in the mean which generally indicates a direction in time. For such time series a trend and seasonal pattern can be identified and removed to give a time series of residuals that can reasonably be considered a realisation of a first order stationary model. In the following chapters we assume all time series are realisations of first order stationary time series models.

However, directionality is not always clearly visible a time series plot (Figure 2), for example, when the points are closely plotted or when the directionality is intermittent (Mansor et al., 2018). In general, a statistical measure is needed to provide empirical evidence of directionality in time series. Moreover, directionality exists in forms that are not necessarily identical and a single measure may not be able to detect all the various possible forms of directionality in time series.
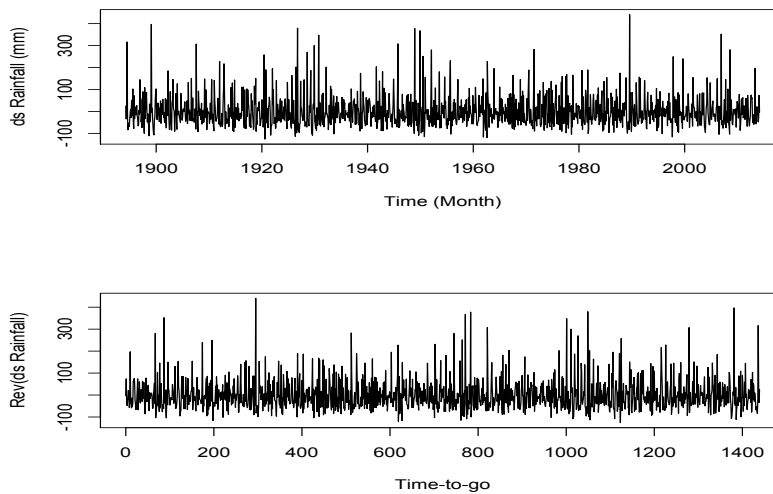
Figure 2: There is no clear visible directionality in the deseasonalized rainfall (Brisbane, Kinross station) series between April 1894 and March 2015. However, statistical analysis gives evidence of directionality.

Directional features include, but are not limited to, sharp increases followed by slow recessions in time order (Figure 1) or slow increases followed by sharp decreases in time order (Lawrance, 1991). In such cases the series is asymmetric with respect to time, and also with respect to its mean. In contrast, a time series may exhibit both sharp increases and sharp decreases, followed by more gradual returns to the mean value. Such time series are asymmetric with respect to time but statistically symmetric with respect to the mean. Statistically, symmetric here refers to a statistic that detects and distinguishes by sign, steep rises and slow recessions from slow rises and steep recessions will tend to average zero. Therefore, different statistics are appropriate for detecting directionality in series that are asymmetric or symmetric with respect to the mean.

This study aims to develop a general procedure, involving a suite of statistics, to detect a range of possible directional features in univariate time series and help classify its nature. The statistical significance for each directionality measure is determined by using Monte-Carlo procedures, corresponding to a null hypothesis of reversibility (no directionality). The time series used in this research are collected from a range of processes and disciplines covering business, environmental science, finance and medicine.

Detecting directionality is important for several reasons. One is that it can provide insight into the physical processes underlying the time series. For example, directionality can be interpreted as evidence of complex feedback processes after a shock or occasional extreme events. Soubeyrand et al. (2014) suggest that intense rainfall influences future rainfall because biological particles released by rain aid the coalescence mechanism of rain formation, and demonstrate directionality in rainfall time series. A second reason is that directionality can indicate the condition of a physical system, for example, directionality in an electrocardiogram (ECG) is a good sign in humans (van Prehn et al., 2009) and directionality in electroencephalogram (EEG) time series has potential for early warning of an epileptic seizure (Mansor et al., 2016b). Another reason is that directionality indicates when non-linear time series models, or at least non-Gaussian errors, are appropriate leading to more accurate forecasts and more realistic scenarios (Lawrance, 1991).

Several types of non-linear time series models that can be associated with directionality. An early study of directionality in non-linear models is discussed by Tong & Lim (1980) in their paper on threshold autoregressive (TAR) models. Other authors such as Granger & Anderson (1978) and Rao (1981) discuss bi-linear models, and Priestly (1980) discusses state-dependent models. More recent studies by Wild et al. (2014) demonstrate that the detection of directionality and non-linearity has significant implications in the selection of models for financial time series.

If non-linear models such as threshold autoregressive (TAR) models provide a better fit than autoregressive (AR) linear models, then the forecast of these non-linear models will be more accurate. In general, TAR models provide a piecewise linear approximation to a wide range of non-linear processes. TAR models are appropriate for time series with a directionality component because these models use threshold space to capture several non-linear characteristics which are commonly observed in practise as asymmetry in declining and rising patterns (directionality) of a process (Tsay,

2005). These ideas are generalised to multivariate time series models that can show directionality in simulations through either non-Gaussian errors or non-linearity, or both.

Accordingly, this study investigates the modelling aspects of directionality using TAR models and AR models with non-Gaussian errors for univariate time series. This investigation includes TAR models with Gaussian, non-Gaussian, or resampled residuals from the fitted models distributions, for the errors. Several symmetric and asymmetric non-Gaussian distributions are used in the simulations to investigate their effects on directionality. There is a scope for more realistic modelling of directionality by using a penalized least squares (PLS) approach. Thus, this study aims to leverage the prospect of modelling directionality in time series, by introducing directionality as a fitting criterion when fitting TAR models to a directional time series. The effects of the PLS approach on predictions are considered, and the differences between models in terms of short term forecasts and future scenarios are investigated.

This study also considers a five-dimensional threshold vector autoregressive (TVAR) model for a portfolio of five high frequency time series, and proposes a strategy for using directionality as a regime-switching criterion between two-regime TVAR models in the model fitting. The performance of this innovation is investigated by comparing the forecasts from the TVAR models with a vector autoregressive (VAR) model. It is also worthwhile to consider modelling the non-Gaussian errors as a Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) processes for time series. A weighted-average measure of the forecasting errors to the estimated conditional variance is used to evaluate forecasting performances by the VAR-GARCH and TVAR-GARCH models.

In addition to considering the univariate and multivariate forecasting applications, this study investigates the behaviour of directionality in the high frequency time series of daily log-returns in relation to financial crisis and non-crisis periods. A moving directionality index is defined to observe and monitor directionality in response to the shocks or extreme irregularities inputs to the system. The results are compared with volatility and moving volatility, respectively. The result, an investment simulation for portfolios of shares, is conducted to explore the use of directionality as a criterion for buying and selling the shares, in comparison to volatility trading rules. Other applications include detecting directionality in time series in patients

diagnosed with epilepsy and subjects without epilepsy. Implications and potential benefits of detecting, modelling and monitoring directionality in time series are summarized in the chapter's concluding remarks.

1.2 Thesis Structure

The results of this study are presented by a series of published or submitted manuscripts in a thesis comprising five chapters. This chapter, Chapter 1, provides an overview of the study and a link between the manuscripts together with the aims for each publication. Chapter 2 consists of a paper on detecting directionality in time series. Chapter 3 comprises three papers focused on the modelling aspects of directionality using non-linear threshold autoregressive models fitted by least squares and penalized least squares approach. Practical perspectives of detecting directionality in time series become apparent. Chapter 4 extends the work from the previous chapters through four papers expanding on mphasizing on the practical implications of directionality. A summary of findings and future direction for the study are given in Chapter 5.

The findings are presented in four themes:

Theme I: Detecting directionality in time series;

Theme II: Applications from the detection of directionality;

Theme III: Time series models that explicitly emulate directionality; and

Theme IV: Applications from modelling directionality.

Appendix A gives information about testing the suite of directionality statistics on the deterministic and probabilistic structures. Appendix B provides a list of time series used in the study, and Appendix C is a listing of the R codes for detecting directionality in univariate time series and a typical penalized least squares fit.

This thesis consists of eight manuscripts distributed over Chapters 2, 3 and 4. The link between manuscripts is illustrated in Figure 3 using a tree diagram.
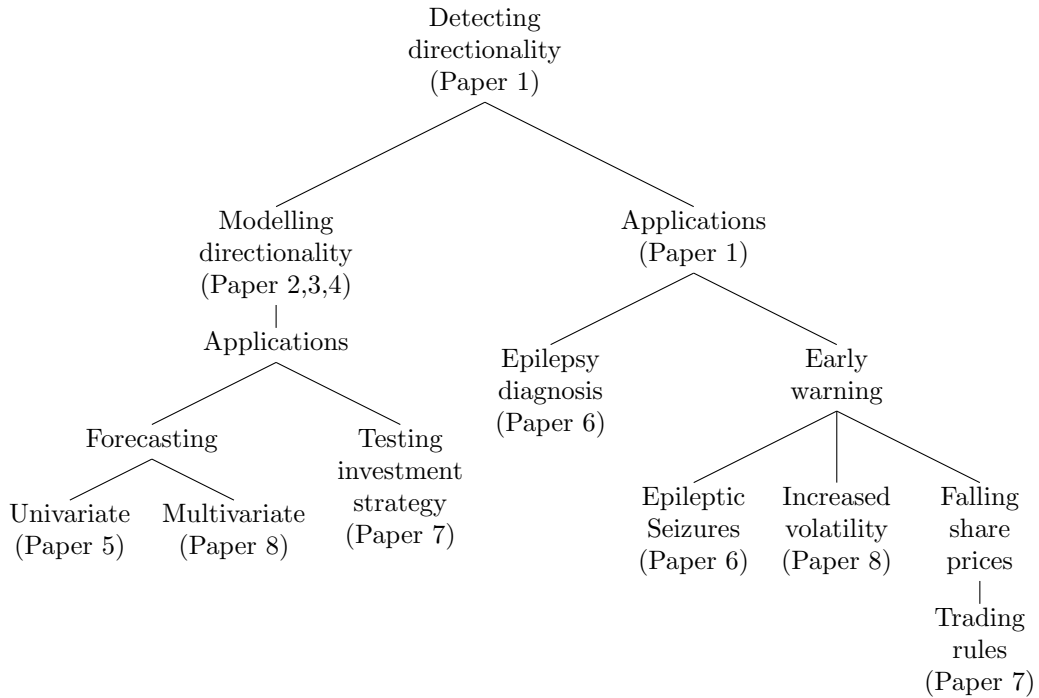
Figure 3: A link between the eight manuscripts.

The definition of directionality in time series is given in a paper entitled "Detecting Directionality in Time Series" (Paper 1, Section 2.1), and applied consistently throughout this thesis.

Chapter 2 consists of Paper 1, in which a general procedure involving using a suite of statistics to detect the various forms of directionality in univariate time series, and to describe a range of directional qualitative features is provided. The time series used in Paper 1 are: the yearly sunspot numbers; the daily Australian dollar against United States of America dollar (AUDUSD) exchange rate; the monthly Southern Oscillation Index; the daily log-returns of JP Morgan Chase & Co. and Wells Fargo & Co.; the electroencephalogram (EEG) records from healthy volunteers and epilepsy patients; the monthly rainfall series of Brisbane; the monthly visitor arrivals to Australia; and the North Greenland Ice Core Project (NGRIP) ice-core series.

The proposed general procedure for detecting directionality in deseasonalized and detrended time series has been set to allow a complete separation of directionality from trend and seasonal patterns. The suite of statistics used in the proposed procedure consists of three statistics from the literature, three alternatives and one modified test. A directionality measure with p-value, estimated by using the Monte-Carlo procedures, of less than 0.05 indicates that the time series is statistically directional at 5% significance level. The suite of directionality statistics has been tested on deterministic and probabilistic processes with symmetry and asymmetry in time properties to verify the veracity of directionality values, before testing them on the observed time series. Detailed information regarding the analysis and results for the deterministic and probabilistic structures are reported separately from Paper 1, in Appendix A.

In Paper 1, the corresponding author also aims to explore some practical importances of directionality in time series using the sunspots, the EEG records and the log-returns of Wells Fargo. Work based on these preliminary practical findings for the EEG and the Wells Fargo time series continues in Chapter 4.

There are three papers in Chapter 3: Modelling Directionality, namely "Modelling Directionality in Stationary Geophysical Time Series" (Paper 2); "Modelling and Simulation of Directional Financial Time Series" (Paper 3); and "Modelling Directionality for Paleoclimatic Time Series (Paper 4).

Work reported in these papers aims to model directionality found in the time series; the directionality value is given by the best directionality statistic from the suite of statistics. The best indicator refers to the most sensitive statistic for detecting directionality in term of directionality value relative to its standard error under the null hypothesis of reversibility. Before modelling the directionality for the sunspots (from Paper 1), Paper 2 investigates how directionality can be produced using Beta and t-distributions (symmetric non-Gaussian errors) as well as exponential distribution (asymmetric non-Gaussian errors) on an autoregressive (AR) model, and a threshold autoregressive (TAR) model with Gaussian errors. Directionality in the sunspots is then modelled by the TAR model with Gaussian errors. Findings from Paper 2 raise the question of whether a measure of matching directionality in the optimization criteria when fitting the TAR model can improve the fit in terms of directionality, and provide better forecasts for the sunspots. This work is continued in Paper 5 (Chapter 4). Apart from modelling the directionality, Paper 2 provides an insight to support physical interpretations of

directionality in the sunspots.

In Paper 3, the corresponding author investigates directionality in the quarterly Great British pound against United States of America dollar (GB-PUSD) exchange rates, the monthly Australian two-year bond yield rates, and the yearly U.S. unemployment rates. In all cases, the directionality in the realisation of an AR model, as a reasonable first approximation to many economic time series, with the resampled residuals from the fitted model for the errors, indicates the need for TAR model. Here, the aim is to improve this aspect of modelling directionality by using a penalized least squares (PLS) approach where the strategy is to use directionality as a fitting criterion in the optimization function when fitting the TAR model with resampled residuals errors. The performance differences between the models are observed by comparing the simulated directionality relative to the observed directionality and the least squares error criterion. Paper 3 also compares the distribution of extremes from the models with Gumbel and back-to-back Weibull errors (asymmetric non-Gaussian) for the U.S. unemployment rates.

The work reported in Paper 4 aims to model the directional ice-core time series of NGRIP (from Paper 1) and Vostok, and to demonstrate that TAR models with PLS strategy are reproducible models for modelling directionality in time series. This paper investigates TAR models fitted using the PLS with two different resampled residuals, from AR models and from the TAR models itself, for the errors. A brief statement linking the directionality in the ice-core time series with the ancient extreme events is suggested, together with a climate change simulation using an AR model and the TAR model with back-to-back Weibull errors for the NGRIP series.

Chapter 4 consists of: "Threshold Autoregressive Models for Directional Time Series" (Paper 5); "Directionality and Volatility in Electroencephalogram Time Series" (Paper 6); "Directionality in Time Series of Bank Share Prices" (Paper 7); and ends with "Directionality and Volatility in High Frequency Time Series" (Paper 8).

The work reported in Paper 5 extends the work in Paper 2, and proposes a modification to the PLS strategy used in Papers 3 and 4. The aims are to improve the fit in terms of directionality for the sunspot time series, and investigate the effects of the modification on predictions. Also considered in Paper 5 are simulations of 15-year extreme values for the sunspots and comparisons made between the results from AR model with resampled residuals errors, and TAR models with Gaussian, resampled residuals, Gumbel, and

back-to-back Weibull distributions for the errors.

The next paper, Paper 6, features the initial work illustrated in Paper 1 on EEG time series. Besides directionality, further comparisons of EEG records from healthy subjects and subjects with epilepsy are made in terms of volatility, the proportion of variability accounted for by time series models, and the skewness and the kurtosis of the residuals. Directionality is modelled by a TAR model fitted using the PLS strategy introduced in Paper 3, and volatility is modelled by a GARCH model. An additional aim here is to suggest potential implications of detecting directionality in the EEG time series.

Paper 7 and Paper 8 report further work on the practical findings for the Wells Fargo log-returns in Paper 1. Both papers use high-frequency time series of bank share prices of: JP Morgan, Wells Fargo, Bank of America, Citibank Inc. and Goldman Sachs Group. These institutions are the five largest banks in the U.S. by market capitalization for 2017. The aims are to demonstrate that moving directionality results in increases during financial crises, and it has potential as an early warning indicator of falling share prices.

Moreover, work reported in Paper 7 focuses on the speculative investment simulations using directionality-based trading rules for a portfolio of shares from the five U.S. banks, and a single asset investment for a bank in China, U.K., Australia, Canada and Japan. Portfolio performance comparison between trading rules based on directionality and volatility for 3-month, 6-month and 9-month moving series are also investigated in this paper. Thus, comparison between directionality trading rules when changing the threshold values of directionality are investigated. Subsequently, a criterion for choosing an investment strategy between directionality trading rules is determined by using conditional value at risk (CVaR). A time series model giving the best fit in terms of directionality will be used to calculate CVaR for JP Morgan log-returns. This includes TAR-GARCH models fitted by the PLS approaches.

Conversely, Paper 8 focuses more on the multivariate forecasting for the five largest U.S. banks using VAR-GARCH and TVAR-GARCH models featuring directionality and volatility as a regime-switching criterion. In-sample and out-of-sample forecasts accuracy is defined as the mean relative error squared and the mean weighted of the forecasting errors to the estimated conditional variances, for the multivariate models without and with GARCH

processes, respectively. This paper also aims to investigate directionality in high frequency financial time series using: 1-, 3-, 6-minutes data and 2-, 5-, 10-day moving average and resampled series of the five daily log-returns.

Key findings from each paper are summarised thematically in Chapter 5, and outlines of possible paths for future work in this area are provided. The $\boldsymbol{R}$ codes for the suite of directionality statistics and the PLS fit are given in Appendix C. Appendix B is the time series data.

Chapter 2

Detecting Directionality

2.1   Detecting Directionality in Time Series

<div style="border:1px solid black; padding:1em;">

Statement of Authorship

Detecting Directionality in Time Series

Submitted for publication, currently under review. This version includes corrections following referees comments.

Mahayaudin M. Mansor (Candidate)
Constructed statistical models, conducted analyses and interpreted results, wrote manuscript and acted as corresponding author.
Overall percentage: 80%

This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.

Signed: .... Date: 14 June 2017

By signing the Statement of Authorship, each co-author certifies that: (i) the candidate's stated contribution to the publication is accurate; (ii) permission is granted for the candidate in include the publication in the thesis; and (iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

David A. Green
Supervised development of work and provided critical evaluation

Si Date: 14/6/17

Andrew V. Metcalfe
Supervised development of work and provided critical evaluation

Signed: . Date: 19/6/17

</div>

Detecting Directionality in Time Series

Mahayaudin M. Mansor, David A. Green
and Andrew V. Metcalfe*
School of Mathematical Sciences,
University of Adelaide, Adelaide 5005 SA, Australia.

Abstract

Directionality can be seen in many stationary time series from various disciplines, but it is overlooked when fitting linear models with Gaussian errors. Moreover, we cannot rely on distinguishing directionality by comparing a plot of a time series in time order with a plot in reverse time order. In general, a statistical measure is required to detect and quantify directionality. There are several quite different qualitative forms of directionality, and we distinguish: rapid rises followed by slow recessions; rapid increases and rapid decreases from the mean followed by slow recovery towards the mean; and directionality above or below some threshold. The first objective is to develop a suite of statistical measures that will detect directionality and help classify its nature. The second objective is to illustrate the potential of directionality for early warning of changes in a process. We demonstrate the proposed procedure for the detection of directionality with examples form business, environmental science, finance and medicine. Time series data is collected from many processes, both natural and anthropogenic, by a wide range of organizations, and directionality can easily be monitored as part of routine analysis. We suggest that doing so may provide new insights to the processes.

Keywords: time reversible, time irreversibile; directional non-linear time series; Monte-Carlo; stationary time series.

1   Introduction

The time series of annual sunspot numbers is plotted twice in Figure 1. In the upper frame the sunspot number is plotted against year from 1700 until 2014 (time), and in the lower frame the time order is reversed and the sunspot number is plotted against year from 2014 until 1700 (time-to-go). A striking feature of these plots is that they are distinguishable, in the upper plot there is a tendency for rapid increases to a peak to be followed by slower recessions whereas in the lower plot the tendency is for slower increases and rapid recessions. The sunspot time series is said to show directionality. The analysis of time series typically involves fitting a model, and a time series model is directional if it has probabilistic properties that depend on the direction of time (e.g. Lawrance 1991). A time series model is reversible if it has no probabilistic properties that depend on the direction of time. Linear models with constant variance independent Gaussian errors are reversible, but if the errors are not Gaussian or the model is non-linear, then it is said to be directional.
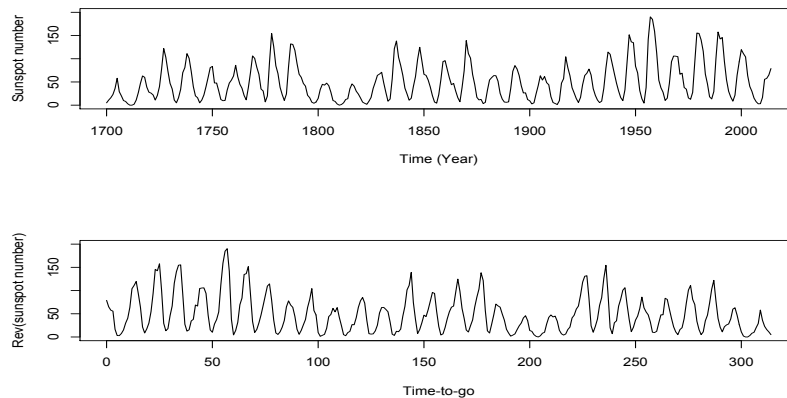
Figure 1: Graphical inspection of directionality shows the sunspot series (1700-2014) has sharp increases before the peaks followed by slow recessions to the troughs in time order (above); and it has slow increases before the peaks followed by sharp decreases to the troughs in reverse time order (below).

Detecting directionality is important for several reasons. One is that it can provide insight into the physical processes underlying the time series. For example, directionality can be interpreted as evidence of complex feedback processes after a shock or occasional

extreme events. Soubeyrand et al. (2014) suggest that intense rainfall influences future rainfall because biological particles released by rain aid the coalescence mechanism of rain formation, and demonstrated directionality in rainfall time series. A second reason is that directionality can indicate the condition of a physical system, for example an asymmetric electrocardiogram (ECG) is a good sign in humans (van Prehn et al. 2009) and directionality in electroencephalogram (EEG) time series has potential for early warning of an epileptic seizure (Mansor et al. 2016b). Another reason is that directionality indicates when non-linear time series models, or at least non-Gaussian errors, are appropriate leading to more accurate forecasts and more realistic scenarios (Lawrance 1991). For example, Wild et al. (2014) demonstrated that the detection of non-linearity and directionality has significant implications in the selection of models for financial time series.

There is a considerable literature on non-linear models for time series but the practical aspects of directionality are not always emphasized. An early study of directionality in non-linear models is discussed by Tong & Lim (1980) in their paper on threshold autoregression models. However, other authors such as Granger & Andersen (1978) and Rao (1981), when discussing bi-linear models, and Priestley (1980), when discussing state-dependent models, do not specifically emphasise the applicability of non-linear time series model for modelling directional series. Nevertheless, these non-linear models including exponential autoregressive models and random coefficient autoregressive models possess directionality as a property of their higher order dependence structure (Lawrance 1991).

In practice, directionality is visible in many time series from various disciplines. For example, the logarithms of Canadian lynx trappings (1821-1934) considered by Lawrance (1991); the fragmented time series of rainfall (Soubeyrand et al. 2014); paleoclimatic time series of the North Greenland Ice Core Project (NGRIP) and the Vostok ice core records (Mansor et al. 2016d); financial time series of the U.S. unemployment rate (1969-2014) and the GBP-USD exchange rates (March 1990-March 2015) (Mansor et al. 2015b); and the selected EEG signals from subjects who are prone to epileptic seizures (Mansor et al. 2016b).

In other cases directionality is difficult to observe from a single time series and a statistical test is required to detect and quantify this feature. But, there are various forms

of probabilistic asymmetry in time, and a single measure of directionality may not have power to detect all forms of directionality. Therefore, it is advisable to use a variety of statistical measures to investigate directionality for a specific time series. Many tests to detect directionality in time series have been proposed but they can be described under three categories. Methods that focus on the original series, methods that focus on the distribution of differences, and indirect methods that provide evidence of non-linearity or non-Gaussian errors or both. Lawrance (1991) suggested the differences between lag one correlations between the variable and the variable squared and the lag one correlations between the variable squared and the variable, and Beare & Seo (2014) have generalized this approach by fitting copulas. Soubeyrand et al. (2014) investigate the behaviour of the time series around peaks above some threshold. Cox (1981) noted that the skewness of differences could be used as a measure of skewness and Wild et al. (2014) use trispectral analysis to investigate asymmetry of the distribution of differences.

In this paper we propose a general procedure, involving a suite of tests, to detect directionality in univariate time series from several disciplines that exhibit a range of the many possible forms of directionality. We measure directionality in these series using seven different statistics, and estimate their standard errors and compare the performance of the associated tests of reversibility using Monte-Carlo procedures. The aim is to determine which tests are more powerful at detecting specific forms of directionality in time series.

This article is organized as follows. In Section 2 we describe three tests from the literature and three alternative methods for detecting directionality in time series, together with a modification for time series that are stochastically symmetric about the mean. We assess the statistical significance with Monte-Carlo procedures. Examples of applying the suite of tests to a variety of time series are given in Section 3. In Section 4 we discuss some implications and potential benefits of monitoring directionality, in the context of three case studies. The concluding remarks are given in Section 5.

4

2 Methods for Detecting Directionality

2.1 Defining directionality

We begin by excluding directional features that are a consequence of a deterministic or stochastic trend or of seasonal variation by restricting the discussion to first order stationary time series models.

A stationary time series model $\{X_t\}$ for $t = 1, 2, ..., n$ is first order stationary if $E[X_t] = \mu$ where μ is the constant mean of the process. It is said to be reversible if the joint distribution of $X_t, X_{t+1}, ..., X_{t+r}$ is equal to the joint distribution of $X_{t+r}, X_{t+r-1}, ..., X_t$ for all $r = 1, 2, ...$ (Lawrance 1991). A stationary time series model is directional if these joint distributions differ. If a stationary time series model is reversible, then it will not be possible to distinguish realizations plotted against time order from those realizations plotted in reverse time order, that is against time-to-go. If a stationary time series model is directional, then there will, in principle, be qualitative differences between realizations plotted in time order and the same realizations plotted against time-to-go. However the difference may not be discernible from a single time series, so further statistical tests are required to detect and quantify the directionality in these cases. The crucial theoretical result concerning directionality, due to Weiss (1975), is that stationary linear time series models are reversible if and only if the errors have an identical Gaussian distribution. Non-linear time series models are generally, but not necessarily, directional (Lawrance 1991). It follows that a directional model requires non-linear terms or errors that are not independently and identically distributed (*iid*) with a Gaussian distribution.

There are many possible directional features in time series, and these include, for example, sharp increases followed by slow recessions or slow increases followed by sharp decreases. In such cases the series is asymmetric with respect to time and also with respect to its mean. In contrast a time series may exhibit both sharp increases and sharp decreases, followed by more gradual returns to the mean value. Such time series are asymmetric with respect to time but statistically symmetric with respect to the mean. By statistically symmetric we mean that any statistic that detects, and distinguishes by sign, steep rises and slow recessions from slow rises and steep recessions will tend to average zero. Different statistics

5

are appropriate for detecting directionality in series that are asymmetric or symmetric with respect to the mean.

In practice, many time series appear to have trend or seasonal variation, defined as any deterministic periodic pattern. For such time series, a model for the trend or seasonal variation should be fitted so that the residuals can reasonably be considered as realization of a first order stationary time series model. There are several strategies for detecting and estimating deterministic trends and seasonal effects (e.g. Chatfield 2004). Some time series, particularly financial time series such as share prices, are reasonably modelled as a random walk (Dickey-Fuller unit root tests) and show stochastic trends. If a random walk is an appropriate model, then log returns ($\ln x_t - \ln x_{t-1}$ where x_t is the share price at time t) are independently distributed and cannot be directional. It follows that evidence of directionality in log-return series is evidence against a hypothesis of a random walk.

2.2 Methods Based on the Original Series

2.2.1 Visual Inspection

A comparison of plots of the time series in its original order and in reverse time order may show a qualitative difference.

2.2.2 Analysis of non-Gaussian Residuals

An indirect method of testing a null hypothesis of reversibility is to fit a linear time series model to the data and to test either a null hypothesis of linearity, or a null hypothesis of Gaussian errors or both. In particular, evidence of volatility, typically modelled as a generalized autoregressive conditional heteroskedasticity (GARCH) process, is evidence against Gaussian white noise because the marginal distribution of errors will have higher kurtosis than a Gaussian distribution. However, there may, for example, be evidence against a hypothesis of Gaussian errors in cases when the effect on directionality is negligible. In practice, it is more relevant to test for directional phenomena, and quantify them, directly.

2.2.3 Difference in Linear Quadratic Lagged Correlations

It follows from the definition that a reversible time series model has

$$Corr(X_t, X_{t+1}^2) = Corr(X_t^2, X_{t+1}).\tag{1}$$

A measure of directionality can be based on the difference in the sample estimates of these correlations (Lawrance 1991). We use the non-dimensional measure

$$DLQC = \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})(x_{t+1} - \bar{x})^2}{[\sum_{t=1}^{n}(x_t - \bar{x})^2]^{3/2}} - \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})^2(x_{t+1} - \bar{x})}{[\sum_{t=1}^{n}(x_t - \bar{x})^2]^{3/2}}.\tag{2}$$

The rationale behind this statistic is as follows. Consider, for example, a time series for which sharp increases tend to be followed by slow recessions. Suppose the sharp increase occurs between x_t and x_{t+1}, then $(x_t - \bar{x})$ could be negative or positive but $(x_{t+1} - \bar{x})$ is very likely to be positive. It follows that $(x_t - \bar{x})(x_{t+1} - \bar{x})^2$ can be negative or positive whereas $(x_t - \bar{x})^2(x_{t+1} - \bar{x})$ is very likely to be positive and hence DLQC will tend to be negative. Both terms in DLQC are correlations, so bounds for the DLQC are [-2,2], but typical values in directional time series are smaller by two orders of magnitude.

2.2.4 Markov Chain Detailed Balances

A stationary finite state Markov chain is reversible if the detailed balance equation is

$$\pi_i P_{ij} = \pi_j P_{ji}\tag{3}$$

for every $i,j = 1, 2, ..., \upsilon$ where υ is the number of states, π_i is the stationary distribution and P_{ij} is the transition probability for, state i to state j (e.g. Kroese et al. 2011). In physical terms the flux from i to j equals the flux from j to i. An example of a chain that is not reversible is one in which it is possible to go from S_1 to S_2 in one step, but not possible to go from state S_2 to state S_1 in one step – in any realization we will notice S_1 followed by S_2 but no S_2 followed by S_1.

Any time series can be converted into a Markov chain by classifying the continuous variable into a finite set of states, by quantile for example, although the state space becomes

large if there are many lags in the model. For example, an AR(2) model with marginal distribution discretized at quantiles would require a 4^2 state Markov chain to describe transitions from (X_{t-1}, X_t) to (X_t, X_{t+1}). However, there are only $4^2 \times 4 = 64$ non-zero transition probabilities. If the Markov chain is directional, then so is the time series.

We implicitly make a lag 1 approximation and divide a time series into a four state Markov chain by up crossings and down crossings of the time series between its 25% (first quartile, Q_1), 50% (second quartile, Q_2) and 75% (third quartile, Q_3), then we label each point in the data set as: 1 if $x_t \leqslant Q_1$ (state 1); 2 if $Q_1 < x_t \leqslant Q_2$ (state 2); 3 if $Q_2 < x_t \leqslant Q_3$ (state 3); and 4 if $x_t > Q_3$ (state 4). We calculate P_{ij} for each entry and determine steady state probability, π for each transition in the Markov chain then calculate detailed balance, $\pi_i P_{ij} = \pi_j P_{ji}$. Directionality exists if $\pi_i P_{ij} - \pi_j P_{ji}$ is not equal to zero for at least one pair (i, j). In practice we have a realization from the Markov chain. The following test statistic

$$MCDB = \sum_{(i,j)} abs(\hat{\pi}_i \hat{P}_{ij} - \hat{\pi}_j \hat{P}_{ji}) \tag{4}$$

will be used and its significance will be determined through Monte-Carlo procedure corresponding to a null hypothesis of a reversible Markov chain. Sharifdoost et al. (2009) proposed a somewhat different statistic based on detailed balance for testing reversibility of finite state Markov chains.

2.2.5   Peaks Over Threshold Test

In some cases, directionality can also be a consequence of feedback in a system following extreme events. For example, a link between extreme values and directionality has been discussed by Soubeyrand et al. (2014) in the context of intense rainfall, and Mansor et al. (2015a) in the context of vigorous magnetic fields of the Sun.

A method for identifying peaks in a time series is to consider excursions above a threshold. A peak over threshold (POT) is defined as the greatest observation above the threshold, and we check for the asymmetry about the peaks. For example, there may be a tendency for a rapid rise to the peak to be followed by a slow recession or for a slow rise to the peak to be followed by a rapid decline.

We start by choosing a threshold (T) as an upper percentile of the marginal distribution of the time series. Here we have used 80% for this parameter. The test is based on the differences (d) between the mean of h observations before a peak and the mean of h observations following a peak, where the parameter h is typically between 2 and 5 time intervals. The asymmetry of peaks can be calculated by taking the difference between the averaging values of h observations before and after the peaks.

We define peaks iteratively in the following way, to obtain differences that are approximately independently distributed. The greatest observation in the time series, more than h steps from the ends, is found and defined as a peak. The difference between the means before and after h observations is calculated, the $2h+1$ consecutive points are removed from the sequences of observations, and the process is continued until all remaining points, other than the first or last h points, are below the threshold.

If the process is reversible, then the mean of the differences is 0. We can test this hypothesis using

$$t_{paired\ statistic} = \frac{\bar{d}}{(s_d/\sqrt{k})} \tag{5}$$

where $\bar{d}$ is the mean of differences between the mean of h observations before a peak and the mean of h observations after the peak, s_d is the standard deviation of the differences and k is the number of paired samples, for a paired comparison. There are many possible variations on POT, but any is valid provided the significance is determined by a Monte-Carlo procedure. The rationale for the POT test is the same as that given by Soubeyrand et al. (2014). The difference is in the definition of peaks, we consider the maximum value in independent excursions over a threshold whereas the FeedbackTS test of Soubeyrand et al. (2014) consider all points over a threshold and overlapping intervals.

2.3 Methods Based on the First Differences

More intuitive measures of directionality can be based on the distribution of lag one differences. For example, if there are sharp increases and slow recessions, then there will be fewer large positive differences and more small negative differences. Also, the distribution of differences will be positively skewed. Let the observed time series of length n be X_t and

define the lag one first order differences (differences) as

$$Y_t = X_t - X_{t-1} \quad \text{for} \quad t = 2, 3, ...n. \tag{6}$$

For example, consider an AR(1) process with parameter α close to 1 and errors $\{\varepsilon\}$. Then $X_{t+1} = \alpha X_t + \varepsilon_t$ subsequently $Y_t \approx \varepsilon_t$ and skewed white noise will lead to similarly skewed differences. However, if there are sharp increases followed by slow recessions to the mean and sharp decreases followed by slow return towards the mean, then the distribution of differences will be symmetric. To detect directionality in such cases, a modification is to consider a derived time series, obtained as the absolute values of deviations of the original series from the mean, before taking differences.

As all the tests in this section are based on the skewness of the distribution of first differences, a histogram or boxplot should be drawn. In particular, extreme outliers can have an excessive influence on the product moment skewness.

2.3.1   Percentage of Positive Differences

If there are sharp increases followed by slow recessions, then the proportion of positive differences will be less than 0.5. In contrast, if there are slow increases followed by sharp decreases, then the proportion of positive differences will exceed 0.5. The percentage of positive of differences is

$$P_d^+ = \frac{number\ of\ positive\ y_t}{number\ of\ positive\ y_t + \ number\ negative\ of\ y_t} \times 100. \tag{7}$$

This formula excludes possible zero differences. It is robust against extreme outliers.

2.3.2   Product Moment Skewness of Differences

A potentially more sensitive test for directionality is to consider the skewness of the distribution of differences (Lawrance 1991) given by

$$\hat{\gamma}_d = \frac{\sum_{t=1}^{n}(y_t - \bar{y})^3/(n-1)}{[\sum_{t=1}^{n}(y_t - \bar{y})^2/(n-1)]^{3/2}}. \tag{8}$$

10

If $\hat{\gamma}_d$ is statistically significantly different from zero, then we have evidence of skewness of differences and hence of directionality. This statistic is highly sensitive to extreme outliers.

2.3.3 L-skewness of Differences

An alternative measure of the skewness of the distribution of differences Y_t is L-skewness. Sankarasubramanian & Srinivasan (1999) found that while conventional moments perform better for distribution with relatively small skewness, particularly for smaller samples, L-moments are preferable for distributions with relatively high skewness, for all sample sizes. Moreover, L-skewness is less sensitive to outliers than the product moment skewness.

The L-moments are given by linear combinations of the order statistics (Hosking 1990). The first L-moment λ_1 is the mean of the distribution, the second is a measure of dispersion defined as $\lambda_2 = (E[Y_{2:2}] - E[Y_{1:2}])/2$ and the third is a measure of asymmetry defined as $\lambda_3 = (E[Y_{3:3}] - 2E[Y_{2:3}] + E[Y_{1:3}])/3$ where Y is the i-th order statistic in a sample of size n and E is the expected value of distribution Y. The L-moments can be estimated by $\hat{\lambda}_1 = b_0$, $\hat{\lambda}_2 = 2b_1 - b_0$, and $\hat{\lambda}_3 = 6b_2 - 6b_1 + b_0$ where $\hat{b}_0 = n^{-1} \sum_{j=1}^{n} y_j$, $\hat{b}_1 = n^{-1} \sum_{j=2}^{n} y_j(j-1)/(n-1)$, and $\hat{b}_2 = n^{-1} \sum_{j=3}^{n} y_j[(j-1)(j-2)]/[(n-1)(n-2)]$.

The L-skewness (LSK) can be estimated by

$$LSK = \frac{\hat{\lambda}_3}{\hat{\lambda}_2}. \tag{9}$$

A time series is said to be directional if the LSK is statistically significantly different from zero skewness.

The percentage of positive differences and the two measures of skewness of differences should lead to similar conclusions. In contrast, the next statistic is designed to identify both sharp increases and sharp decreases with slow return to mid-values.

2.3.4 Product Moment Skewness of Differences of Absolute Values about the Mean

If a time series has both sharp increases and sharp decreases with relatively slow return to the mean, then the distribution of differences is likely to be nearly symmetric. However,

11

we can change the sharp decreases into sharp increases by calculating a new series $x_t^* = \left| x_t - x_{(n+1)/2:n} \right|$ where the $x_{(n+1)/2:n}$ is the mean of the original series. Define the differences $y_t^* = x_t^* - x_{t-1}^*$. We adapt the definition of Section 2.3.2 to be

$$\hat{\gamma}_{dab} = \frac{\sum_{t=1}^n (y_t^* - \bar{y^*})^3/(n-1)}{[\sum_{t=1}^n (y_t^* - \bar{y^*})^2/(n-1)]^{3/2}}. \tag{10}$$

For a series with sharp increases and relatively slow decreases we expect $\hat{\gamma}_d$ to be substantially positive and $\hat{\gamma}_{dab}$ to be near zero. For a series with sharp increases and sharp decreases followed by slow return to the mean we expect $\hat{\gamma}_d$ to be near zero and $\hat{\gamma}_{dab}$ to be substantially positive. The only limitation of this statistic $\hat{\gamma}_{dab}$ is that it is less sensitive than $\hat{\gamma}_d$ because the noise to signal ratio will be doubled by taking absolute values of mean corrected observations.

2.4   Testing Significance of Directionality

Statistical tests for detecting directionality are applied to time series that are either assumed to be realizations of a stationary process or have had seasonal effects and a trend identified and removed. The approximate statistical significance of the calculated value of a test statistic for a particular time series is determined by the following simulation (Monte-Carlo) procedure for long time series:

1. Fit an autoregressive model of order p, $AR(p)$, to the time series of length n using the Akaike information criterion (AIC) to select p. Calculate the variance of the residuals;

2. Simulate a large number, N typically at least 10^3, time series of length n from the fitted $AR(p)$ model, using a Gaussian distribution, with mean 0 and variance equal to the variance of the residuals, for the errors. For each simulated time series calculate the value of the test statistic for directionality; and

3. Determine the proportion of the simulated values of the test statistic that are more extreme than the value calculated for the original time series. Return this proportion as the P-value.

Figure 2: A flowchart representing a general procedure for detecting directionality in univariate time series.
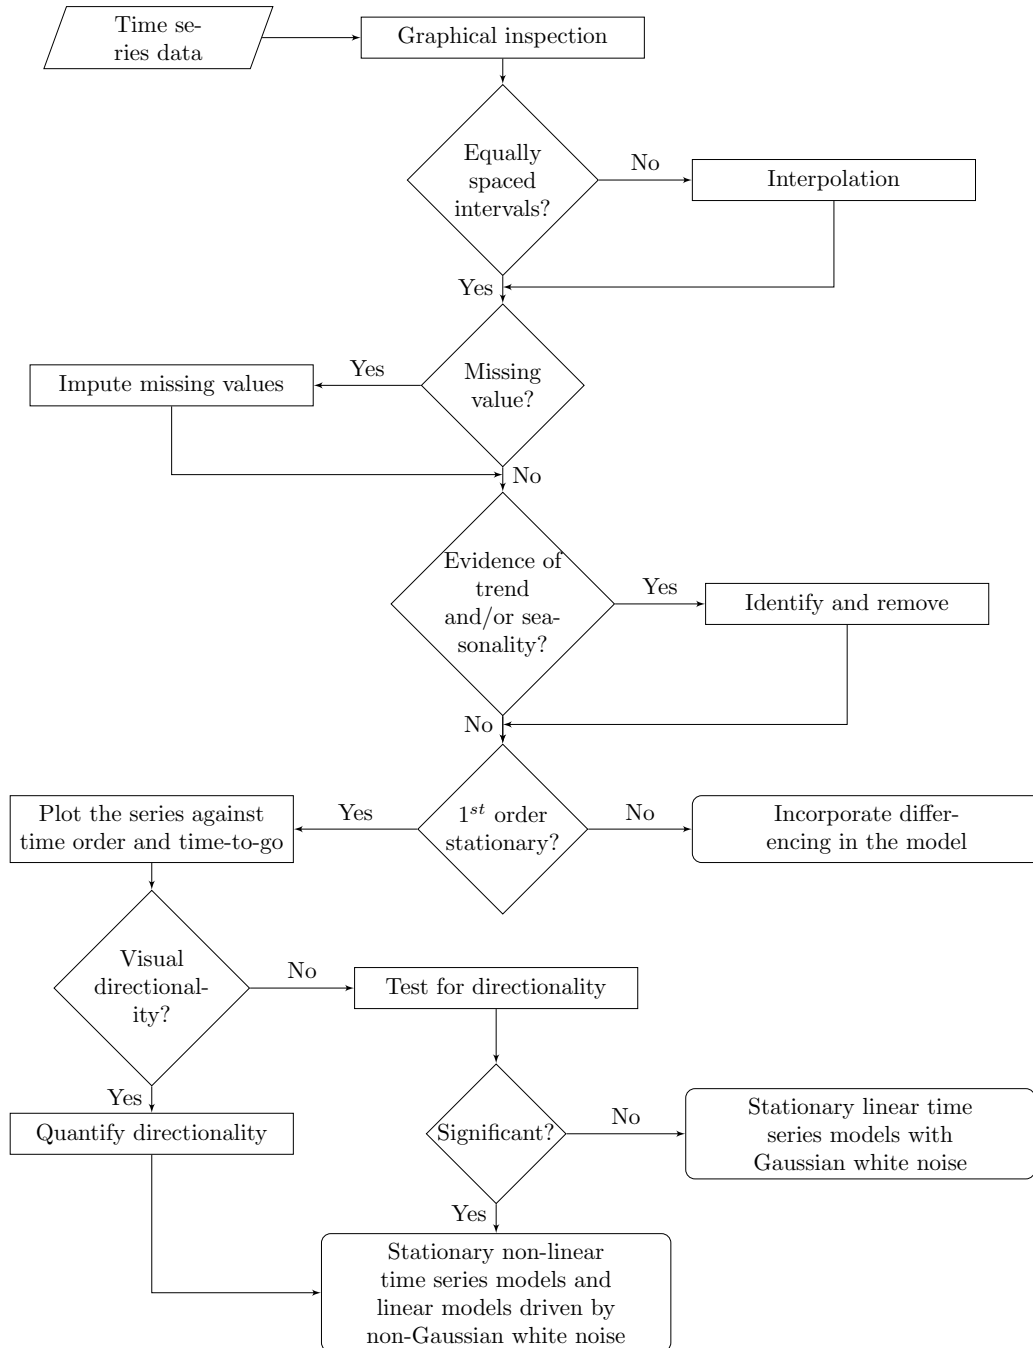
In the case of short time series the Monte-Carlo procedure process should be modified by sampling from posterior multivariate distribution of coefficients for the $AR(p)$ model, subject to stability constraints, between simulation of time series. This will, to some extent, increase the variability of the sampling distribution of the statistic for detecting directionality.

A flowchart summarizing a general process for detecting directionality in univariate time series is shown in Figure 2. The process assumes an evenly spaced intervals time series that contains no missing values. If a time series is not equally spaced, then interpolation, with for example cubic splines, could be used.

3   Examples

We investigate directionality in nine time series from different fields of application, based on the proposed procedure illustrated in the flowchart of Figure 2. In all cases the sampling interval is constant, and except for the rainfall time series there are no missing values. The observed time series plots are given in Figure 3.

3.1   Description of Series

There is no physical reason to expect, or statistical evidence of, trends or seasonal variation in the six series.

- *Sunspots.* Sunspots are the dark spots that appear on the photosphere of the Sun, arising from the different surface temperatures induced by intense magnetic fields. Sunspots are regions of relatively low temperature. The time series is yearly average sunspot number from 1700 to 2014. Data obtained from the Sunspot Index and Long-term Solar Observations (SILSO), World Data Center, Belgium.

- *Exchange rate.* Daily exchange rate of the U.S. dollar against the Australian dollar from 1 January 1990 to 20 April 2015, obtained from Bloomberg L.P.

- *SOI.* The Southern Oscillation Index (SOI) is the standardized difference in mean atmospheric pressure between Tahiti and Darwin (Tahiti minus Darwin). It indicates

the development and intensity of El Niño, when substantially negative, or La Niña, when substantially positive. The series is monthly from from January 1876 to June 2015. Data obtained from the Australian Bureau of Meteorology.

- *Log-returns.* Daily log-return of stock price for JPMorgan Chase & Co. from 3 January 1990 to 17 April 2015. The log-return at time t is the natural logarithm of the ratio between the last-traded price at time t and the last-traded price at time $t-1$, multiplied by 100. This transformation will not affect the time reversibility of the process generating the original series of last traded prices (Wild et al. 2014). The stock prices were obtained from Bloomberg L.P.

- *EEG (normal).* EEG (normal) is a series of length 4096 of electrical activity of the brain from a subject without the condition of epilepsy. The EEG (normal) was recorded during a conscious relaxed state with the eyes closed using a standard multi-location placement of sensors (Andrzejak et al. 2001) sampling rate of 173.61 Hz. The sample courtesy of the Klinik für Epileptologie, Universität Bonn, Germany.

- *EEG (epilepsy).* EEG (epilepsy) is the natural logarithm of the sum of the original EEG signal and 400. The original EEG signal of length 4096 was recorded from epileptogenic focal, a specific recording region site of the patient's brain (Andrzejak et al. 2001). The other details are the same as EEG (normal). The original signal had a minimum of -356, the transformation gave an approximately symmetric marginal distribution.

The following time series do have seasonal components.

- *ds Rainfall.* Monthly rainfall (mm) series from April 1894 to March 2015 was recorded at the Raymond Terrace (Kinross) station, Brisbane. Data obtained from the Australian Bureau of Meteorology. It has 22 missing values which were replaced by the means of adjoin non-missing values. Rainfall series typically contains seasonal variation that is relatively constant over time, so we remove this effect by fitting a centred moving average with additive seasonals, implemented in the ***R*** software (R_Core_Team 2013), to obtain the deseasonalized rainfall series (ds Rainfall).
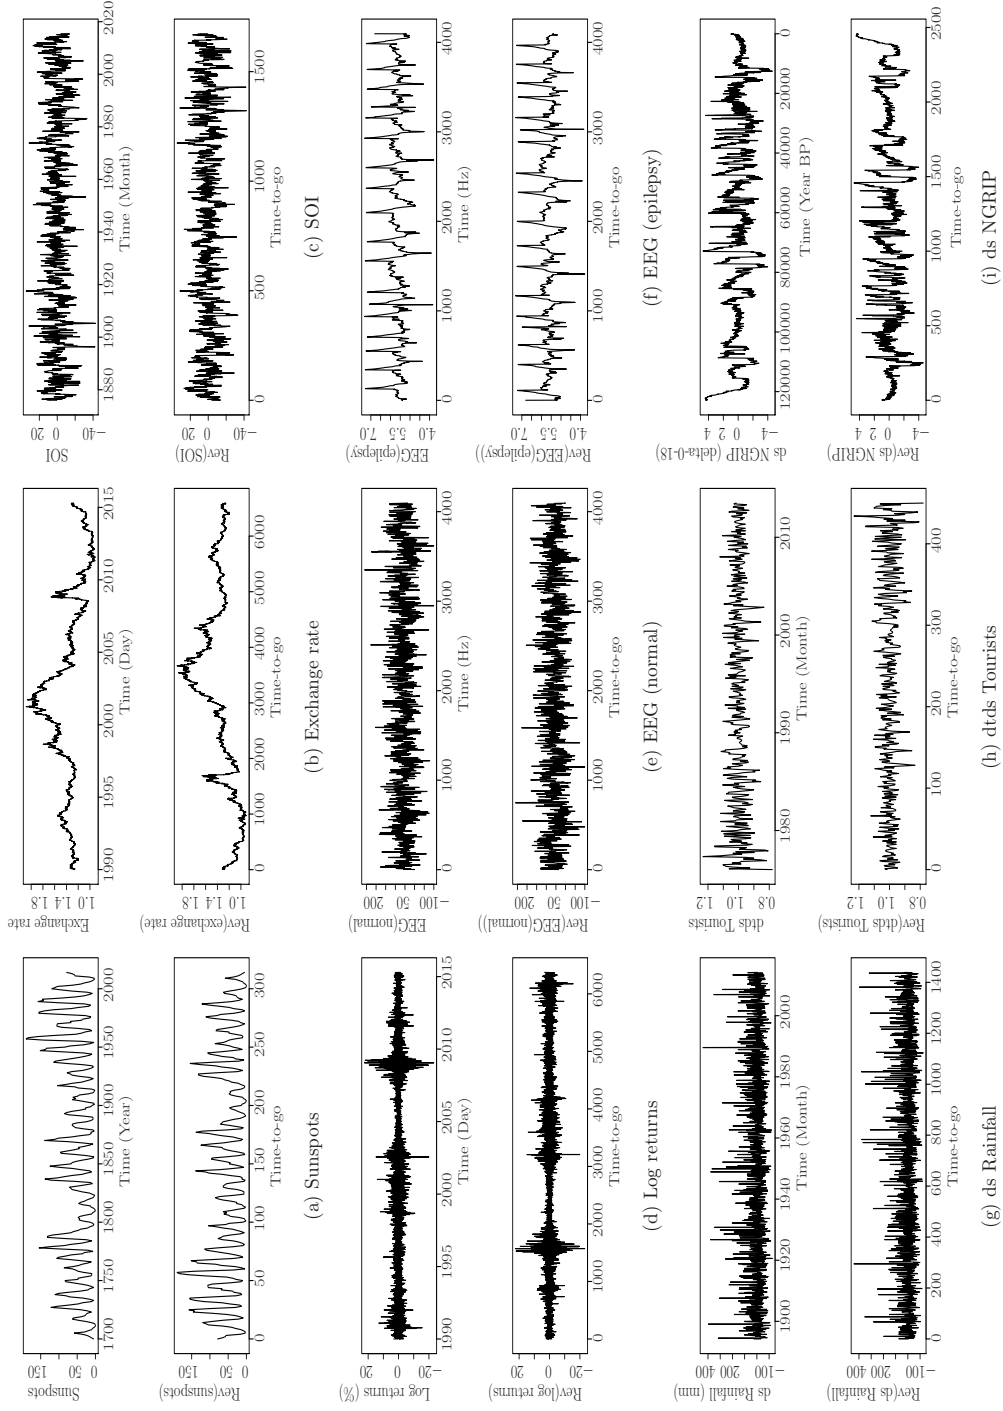
Figure 3: Stationary time series plotted in time order (upper sub-figure) and in time-to-go (lower).

- *dtds Tourists.* Monthly series of short-term (less than one year) visitor arrivals to Australia from January 1976 to July 2014, courtesy of the Australian Bureau of Statistics. It has trend and seasonal variations that increases over time. We estimate and remove these effects by fitting a centred moving average with multiplicative seasonals, implemented in the **R** software (R_Core_Team 2013), to acquire the detrended and deseasonalized tourists series (dtds Tourists).

- *ds NGRIP.* North Greenland Ice Core Project (NGRIP) series begins 122,900 years before present (BP) obtained from NGRIP_members et al. (2004), which consists the ratio ($\delta^{18}O$) of oxygen isotope-18 (^{18}O) to oxygen isotope-16 (^{16}O) at 50 year intervals. The series contains seasonal variation that corresponds to the Milankovitch cycles, long term variations in the Earth's orbit that have been affecting the Earth's climate change for aeons. The effect of Milankovitch cycles in the series is removed by fitting a multiple regression to obtain the deseasonalized time series (ds NGRIP) (Mansor et al. 2016d).

All time series are plotted in time order (upper sub-figure) and in time-to-go (lower) in Figure 3. All series appear to be realizations of stationary time series models, as expected from their physical contexts. If AR(p by AIC) models are fitted, then the Anderson-Darling statistic provides evidence against an assumption of Gaussian errors in all cases (P-value < 0.01) except dtds Tourists (P-value $= 0.63$). So we have indirect evidence against an assumption of reversibility for all the series except for the dtds Tourists.

Autocorrelation plots of the residuals after fitting the AR(p by AIC) models show that there is no statistically significant correlations at any lags except for EEG (epilepsy). In the case of EEG epilepsy there is a lag 1 autocorrelation of 0.36, suggesting that ARMA(10,1) would be better than AR(10). This suggests that the Monte-Carlo procedure using AR(p) with p selected by AIC is suitable in most cases. We compare the estimated standard deviation of the residuals $\hat{\sigma}_{err}$ with the marginal standard deviation s of the time series (Table 1 and 2, first column). In most cases the AR(p) model accounts for a substantial component of the variability in the original time series.

The EEG (epilepsy) is well modelled as ARMA(10,1), and if this model is used for the simulation, then the ratios of the observed statistic to its standard error changes slightly

17

(Table 1(f), last row). But, the practical significance of the change is negligible.

3.2   Results and Discussion

In some cases there are physical reasons to expect directionality in time series, in particular for a time series with a complex feedback process following a shock or occasional extreme events (Soubeyrand et al. 2014). The detection of directionality can provide support for physical theories about the generating process (Mansor et al. 2015a).

Results from formal statistical tests of directionality for the time series plotted in Figure 3 are given in Table 1 and 2, together with a summary of the results in Table 3. There are no extreme outliers that would distort $\hat{\gamma}_d$ in any of the box plots in Figure 4, and the conclusions drawn using it or L-skewness are consistent.

- *Sunspots.* The sunspot time series is clearly directional from the plots. The rapid increases of sunspots correspond to the increase in solar activity due to the extreme electromagnetic polarities of the sun – the difference between high and low temperatures becomes more apparent, followed by slow recedes (e.g. Mansor et al. 2015a). All the statistics detect directionality, including the alternatives MCDB, POT, LSK and $\hat{\gamma}_{dab}$. The $\hat{\gamma}_d$ is found to be the most sensitive test, inasmuch as it has the largest ratio of absolute value of the difference between the observed statistic and its mean to its standard error [6.64]. This has been denoted by * in Table 3.

- *Exchange rate.* Graphical inspection for the exchange rate shows a few small sharp increases followed by relatively slow recessions (small peaks) around 1998-2003, and one marked increase in between 2008 and 2009 – after the global financial crisis of 2007-08, followed by slow recession (big peak) in the upper frame, and otherwise in the lower frame. All the statistics based on the first differences detect directionality, while none of those based on the original time series does, including the POT even with the threshold T set at 90-th, 70-th and 60-th percentiles. The most sensitive test for the exchange rate is the LSK.

(a) Sunspots

(b) Exchange rate

(c) SOI

(d) Log-returns

(e) EEG (normal)

(f) EEG (epilepsy)

(g) ds Rainfall

(h) dsdt Tourists

(i) ds NGRIP
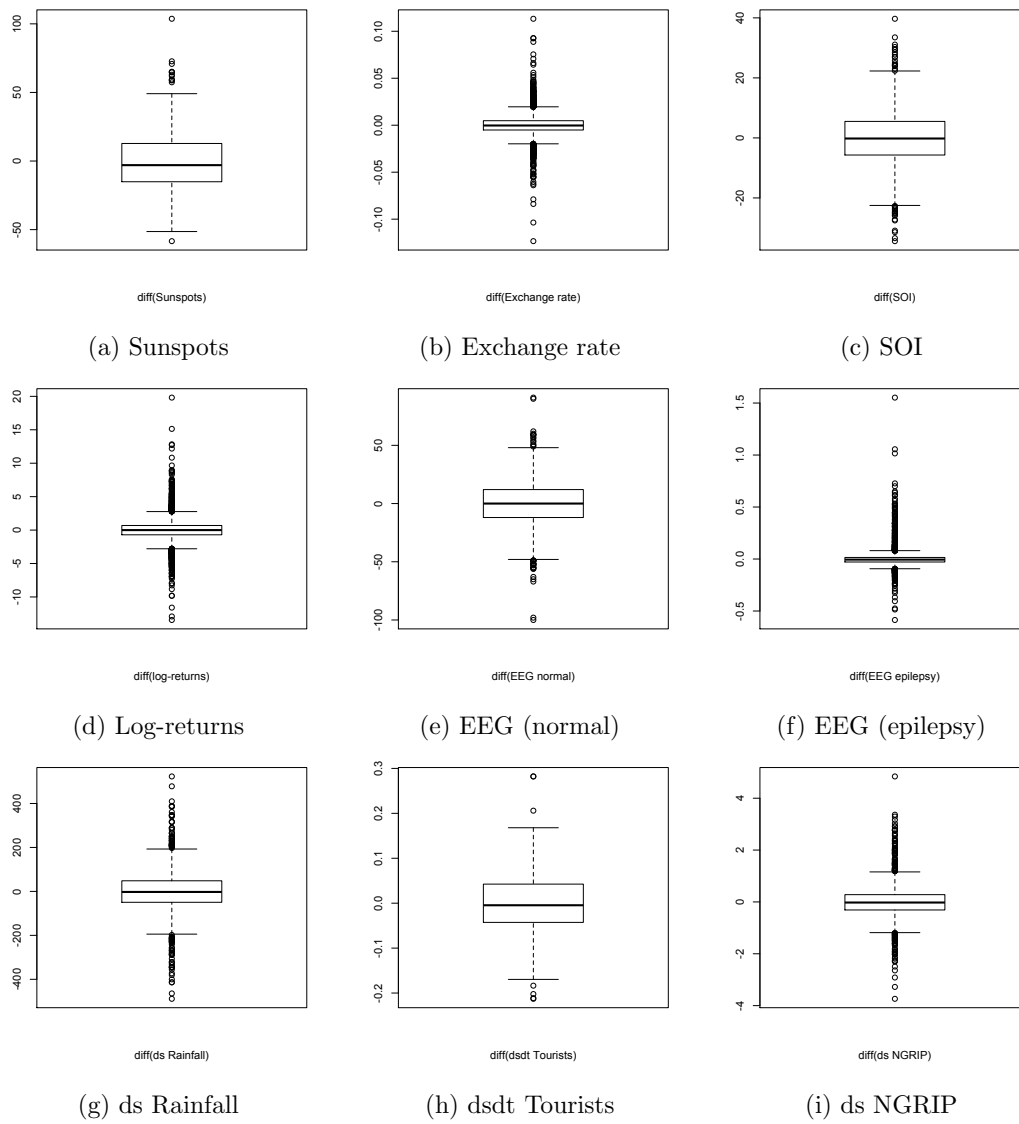
Figure 4: Boxplots of the first differences of the time series.

Table 1: Summary table of test statistics of directionality in the observed time series.

Method	Based on the original series					Based on the first differences			
Time series	DLQC	MCDB	POT(2)	POT(3)	POT(5)	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
(a) Sunspots	-0.06	0.14	4.71	5.36	4.83	42.5%	0.86	0.15	0.66
n=315, s=40.24		{0.04}	{	T=81.0	}	{0.32%}			
AR(9), $\hat{\sigma}_{err}$=15.54	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
	[5.57]	[5.43]	[4.43]	[4.97]	[4.38]	[4.25]	[6.64]	[6.15]	[5.84]
(b) Exchange rate	0.00	0.00	1.44	1.03	0.71	48.4%	0.13	0.03	-0.15
n=6594, s=0.25		{0.00}	{	T=1.53	}	{1.03%}			
AR(34), $\hat{\sigma}_{err}$=0.01	(0.91)	(0.97)	(0.12)	(0.26)	(0.43)	(0.00)	(0.00)	(0.00)	(0.00)
	[0.01]	[0.00]	[1.57]	[1.08]	[0.77]	[4.14]	[4.18]	[5.50]	[4.56]
(c) SOI	-0.02	0.01	0.61	0.38	-0.003	49.3%	0.08	0.01	0.15
n=1674, s=10.48		{0.02}	{	T=9.10	}	{0.42%}			
AR(14), $\hat{\sigma}_{err}$=7.83	(0.19)	(0.93)	(0.53)	(0.70)	(1.00)	(0.32)	(0.16)	(0.34)	(0.01)
	[1.30]	[1.36]	[0.64]	[0.39]	[0.00]	[1.02]	[1.38]	[0.95]	[2.53]
(d) Log-returns	-0.58	0.02	1.23	1.96	3.43	49.7%	0.58	0.01	0.04
n=6372, s=2.494		{0.02}	{	T=1.41	}	{0.93%}			
AR(36), $\hat{\sigma}_{err}$=2.479	(0.00)	(0.17)	(0.24)	(0.05)	(0.00)	(0.39)	(0.00)	(0.03)	(0.17)
	[22.5]	[0.98]	[1.20]	[1.97]	[3.44]	[0.89]	[22.5]	[2.27]	[1.42]
(e) EEG (normal)	0.001	0.001	1.96	1.73	2.14	49.4%	-0.03	0.003	0.07
n=4096, s=46.99		{0.005}	{	T=95.0	}	{2.47%}			
AR(36), $\hat{\sigma}_{err}$=9.38	(0.49)	(0.93)	(0.05)	(0.09)	(0.03)	(0.23)	(0.54)	(0.73)	(0.00)
	[0.66]	[1.25]	[1.94]	[1.70]	[2.17]	[1.18]	[0.61]	[0.34]	[1.85]
(f) EEG (epilepsy)	-0.01	0.004	4.02	4.86	5.87	41.9%	5.40	0.20	2.32
n=4096, s=0.48		{0.00}	{	T=5.98	}	{4.86%}			
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
AR(10), $\hat{\sigma}_{err}$=0.06	[15.9]	[45.8]	[4.01]	[4.89]	[6.00]	[16.3]	[113.5]	[25.5]	[50.8]
ARMA(10,1), $\hat{\sigma}_{err}$=0.05	[15.5]	[45.7]	[4.16]	[5.13]	[6.05]	[15.4]	[108.1]	[24.3]	[49.7]

Notes for Table 1:
1. In the series description: n is the length of the time series; s is the marginal standard deviation; p, in AR(p), is the order of the AR model fitted using AIC; and $\hat{\sigma}_{err}$ is the estimated standard deviation of residuals in the AR(p) model.
2. (number) refers to P-value, [number] refers to the ratio of the test statistic to its standard deviation and {number} in MCDB column is the mean of the test statistic from Monte-Carlo procedure. {number} in P_d^+ refers to proportion of zero in the distribution of differences in original series. T in POT columns refers to threshold value at 80-th percentile in the original series. POT(2), POT(3) and POT(5) are the POT with averaging of 2, 3 and 5 observations before and after the peak.

- *SOI.* It is difficult to discern directionality from the plots but $\hat{\gamma}_{dab}$ detects directionality in SOI, and is the only measure to do so. This is consistent with occasional sharp increases and decreases that correspond to the beginning of El Niño and La Niña events.

- *Log-returns.* There is a little qualitative difference between plots in time order and in time-to-go. The series does show high volatility and the clear increase in variability around 2008 and 2009 corresponds to the global financial crisis. DLQC, POT(5) and $\hat{\gamma}_d$ all detect directionality. The easiest to interpret is the $\hat{\gamma}_d$ which is positive, corresponding to rapid increases followed by relatively slow recessions. But, the statistic $\hat{\gamma}_{dab}$ is near 0, so there is little evidence that the occasional sharp increases are mirrored by occasional sharp decreases.

- *EEG (normal).* Directionality often indicates non-linearity, which in itself may result from complex feedback mechanisms that could be a cause of epilepsy (Mansor et al. 2016b). In the case of EEG (normal), an EEG sample from a subject without the condition of epilepsy, there is less reason to expect directionality, particularly as the signal is an average from a multi-location placement of sensors. However, POT(5), and to a lesser extent POT(2) and POT(3), do detect directionality at 5% level. The series is long, so the test is quite powerful and the directionality is a relatively small effect.

- *EEG (epilepsy).* In time order there is a tendency for sharp increases followed by gradual recessions, whereas in time-to-go, slow increases are followed by plummets. As expected all measures detect directionality in the series. The directional spikes suggest feedbacks during seizures. Directionality detection in this case may have

21

potential as a diagnostic or indication of onset for epilepsy (Mansor et al. 2016b). Directionality in the series is best detected by the $\hat{\gamma}_d$.

- *ds Rainfall.* Asymmetry in time between plots in time order and in time-to-go is difficult to distinguish, but can be detected by MCDB, P_d^+, and $\hat{\gamma}_{dab}$. The most sensitive test for ds Rainfall is $\hat{\gamma}_{dab}$. ds Rainfall is another example of directional time series. Brisbane has the highest rainfall in summer typically between January and March and occasionally flooded. Intense rainfall might influence future rainfall because, for example, biological particles released by rain aid the coalescence mechanism of rain formation and such effects can lead to intermittent (fragmented time) directionality (Soubeyrand et al. 2014).

- *dtds Tourists.* There are few well defined rapid increases in the first quarter and few well defined plummets in the remaining quarters in the upper frame, and otherwise in the below frame. However, none of the statistics detects directionality in this series. There is no evidence to reject a null hypothesis of reversibility at the 5% level of significance, as indicated by the indirect method of non-Gaussian residuals using the Anderson-Darling statistic (P-value = 0.63).

- *ds NGRIP.* The series has a tendency for sharp increases to the peaks to be followed by gradual decreases to the troughs when plotted against time, and gradual increases to the peaks to be followed by sharp decreases to the troughs when plotted against time-to-go. All statistics, except for MCDB, detect directionality. The $\hat{\gamma}_d$ outperforms the other measures for ds NGRIP.

Table 2: Summary table of test statistics of directionality in the detrended and/or deseasonalized time series.

Method	Based on the original series					Based on the first differences			
Time series	DLQC	MCDB	POT(2)	POT(3)	POT(5)	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
(g) ds Rainfall	-0.01	0.07	0.43	-0.47	0.20	48.4%	0.01	0.02	-0.20
n=1440, s=69.97		{0.03}	{	T=39.9	}	{0.00%}			
AR(29), $\hat{\sigma}_{err}$=60.09	(0.84)	(0.01)	(0.67)	(0.64)	(0.83)		(0.83)	(0.09)	(0.00)
	[0.19]	[2.91]	[0.40]	[0.46]	[0.19]		[0.20]	[1.71]	[3.67]
(h) dtds Tourists	-0.09	0.07	1.17	0.13	-0.07	47.8%	0.19	0.03	0.04
n=451, s=0.055		{0.06}	{	T=1.04	}	{0.00%}			
AR(14), $\hat{\sigma}_{err}$=0.047	(0.21)	(0.29)	(0.27)	(0.91)	(0.96)		(0.07)	(0.20)	(0.71)
	[1.24]	[0.49]	[1.15]	[0.12]	[0.06]		[1.76]	[1.33]	[0.36]
(i) ds NGRIP	-0.02	0.01	2.46	3.38	3.75	47.7%	0.61	0.05	0.54
n=2459, s=1.51		{0.01}	{	T=1.23	}	{0.00%}			
AR(5), $\hat{\sigma}_{err}$=0.62	(0.00)	(0.13)	(0.01)	(0.00)	(0.00)		(0.00)	(0.00)	(0.00)
	[11.4]	[1.08]	[2.60]	[3.44]	[3.86]		[12.1]	[5.15]	[9.98]

Notes for Table 2 are the same as the notes for Table 1.

Overall, the POT method seems particularly useful as it detects asymmetry restricted to around peak values and can also, as it is restricted to peaks, detects directionality when a time series is symmetric about the mean. However $\hat{\gamma}_{dab}$ seems more sensitive in this respect for the SOI time series.

Table 3: Summary of results.

Method	Based on the original series					Based on the first differences			
Time series	DLQC	MCDB	POT(2)	POT(3)	POT(5)	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
Sunspots	✓	✓	✓	✓	✓	✓	✓*	✓	✓
Exchange rate	✗	✗	✗	✗	✗	✓	✓	✓*	✓
SOI	✗	✗	✗	✗	✗	✗	✗	✗	✓*
Log-returns	✓*	✗	✗	✗	✓	✗	✓*	✓	✗
EEG (normal)	✗	✗	✗	✗	✓*	✗	✗	✗	✗
EEG (epilepsy)	✓	✓	✓	✓	✓	✓	✓*	✓	✓
ds Rainfall	✗	✓	✗	✗	✗	✓	✗	✗	✓*
dtds Tourists	✗	✗	✗	✗	✗	✗	✗	✗	✗
ds NGRIP	✓	✗	✓	✓	✓	✓	✓*	✓	✓

Note:✓ refers to statistic that is statistically directional at 5% significance level from Monte-Carlo procedure, otherwise ✗. ✓* is the statistic with the smallest P-value for the time series.

4   Applications

In this section, we discuss some practical implications of directionality in time series, and suggest that it has potential as a diagnostic and predictive measure. The time series chosen are the yearly sunspots from 1700 to 2014, randomly selected electroencephalogram (EEG) from subjects with and without a diagnosis of epilepsy, and daily log-returns of Wells Fargo & Co from May 4, 1999 to February 17, 2017.

4.1   The sunspots model

The sunspots time series is a clearly directional time series (Lawrance 1991). The NASA Solar Physics website states that the number of sunspots visible on the sun waxes and wanes with an approximate 11-year cycle (NASA 2017), and there has been much discussion of underlying cycles (e.g. Martens et al. 2011, Nandy et al. 2011). We show an ACF, spectrum and Fourier line spectrum for the sunspot time series in Figure 5. We modelled the mean annual number of sunspots, X_t, as a sequence of over-dispersed Poisson random variable

with a mean at t given by

$$\text{E}[X_t] = \beta_0 + \beta_1 X_{t-1} + \beta_2 X_{t-2} + \sum_{j=1}^{H} \beta_{1j} cos(2\pi\omega_j t) + \beta_{2j} sin(2\pi\omega_j t) \tag{11}$$

where $t = 1, 2, ..., n$ and H is the number of dominant frequencies. The frequencies are ω_j for $j = 1, 2, \cdots, H$ and their amplitudes are $\sqrt{\beta_{1j}^2 + \beta_{2j}^2}$. Nine frequencies were identified as the highest spikes in the Fourier line spectrum, and in order of magnitude these are: 30, 31, 32, 4, 29, 27, 7, 3 and 33.

(a)   (b)
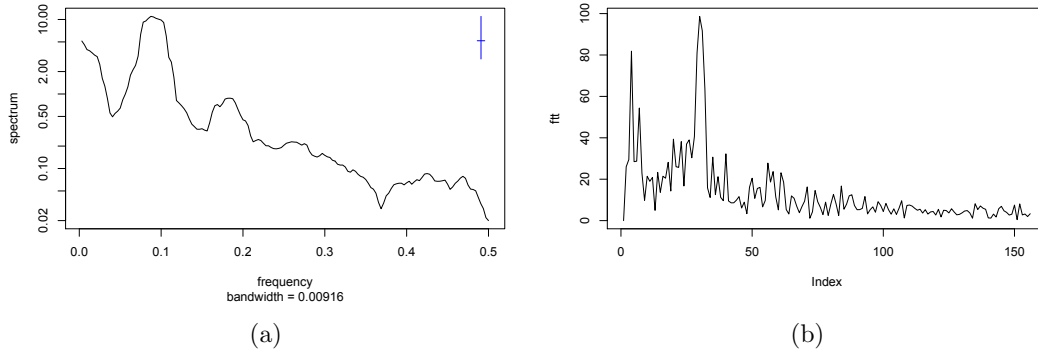
Figure 5: Spectrum (a) and the FTT (b) of the sunspots.

The over dispersion factor was 2.3. The reason for taking two lags in Y_t is that Equation 11 then represents the difference equation of the differential equation form for a linear system with a single mode of vibration with sinusoidal forcing. A Poisson model was chosen to allow for the non-negative variable (modeling the logarithm of the number gave too many unrealistically high values after exponentiation in simulations). The model provided a reasonable empirical fit to the sunspot time series, but realizations from the model shows no sign of variability. This suggests that there are significant non-linear effects that are lost in the simple model.

4.2   EEG records

The EEG time series from normal subjects are well modelled by stationary linear models driven by Gaussian white noise (e.g. Steyn-Ross et al. 1999), and are so reversible. However,

EEG time series from patients who have had brain injuries are typically quite different. Epileptic seizures are caused by temporary disruption of normal brain activity. People who are diagnosed as prone to epileptic seizures would benefit greatly from an early warning of a seizure because they could move to a relatively safe location before its onset.

(a) Healthy A_40   (b) Healthy B_98

(c) Epilepsy C_37   (d) Epilepsy D_76   (e) Epilepsy E_56
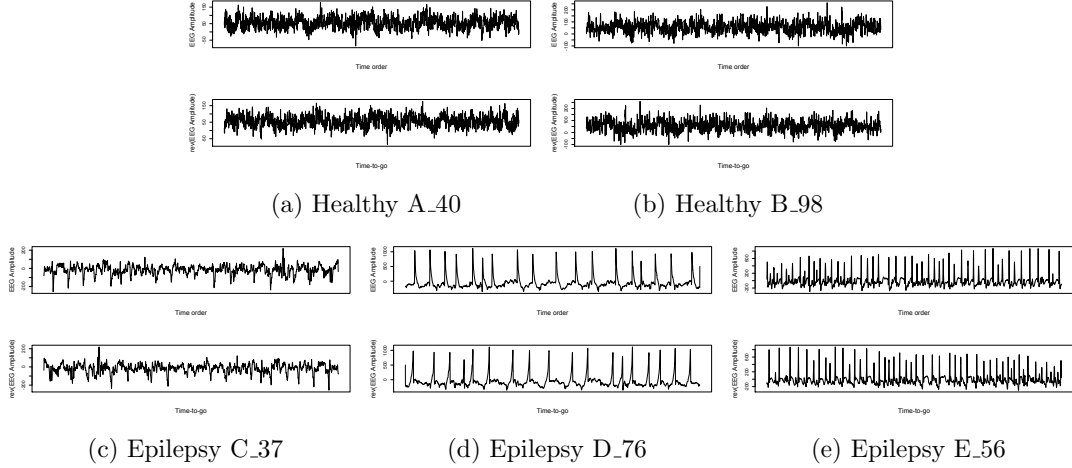
Figure 6: Selected EEG time series from healhy volunteers and subjects diagnosed with epilepsy in time order (upper frames) and time-to-go (lower frames).

In Figure 6 we randomly selected EEG time series of length 4097 observations from two normal subjects (A_40 and B_98) and three subjects who are prone to epilepsy (C_37, D_76, and E_56). The time series are not directly comparable. The time series A_40 and B_98 are based on signals from sensors placed around the skull, and A_40 is with eyes open whereas B_98 is with eyes closed. In contrast C_37, D_76, and E_56 are from different specific regions of the brain: namely hippocampal formation, epileptogenic region, and seizure region respectively.

Table 4: Directionality in the selected EEG time series.

Series	A_40	B_98	C_37	D_76	E_56
Directionality	0.060	-0.029	-0.654	6.425	-0.167
P-value	0.10	0.55	0.00	0.00	0.00

The time series from D_76 and E_56 are dramatically different from those from the

normal subjects. But, the time series from C_37 is not qualitatively different from those from the normal subjects. However, it is directional whereas those from A_40 and B_98 are not (Table 4).

We speculate that monitoring of EEG in subjects who have been diagnosed as prone to epileptic seizures might provide early warning of a seizure. It could be investigated as part of ongoing research projects.

4.3 Well Fargo log-returns

Here we investigate the daily log-returns of Wells Fargo & Co. The series of length 4479 and is considered as evenly spaced, with no missing values. The log-returns can reasonably be considered as a realisations of a stationary time series model, and Table 5 summarizes the suite of directionality indicators. The results provide convincing statistical evidence of directionality although the phenomenon is not clearly visible in the time series plot.

Table 5: Summary table of the suite of directionality indicators in the Wells Fargo daily log-returns.

Method based on	original series			first differences			
Time series	$DLQC$	$MCDB$	POT	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
Wells Fargo	−0.40	0.02	2.06	49.2%	0.37	0.01	0.32
n=4290, s=2.463	(0.00)	(0.23)	(0.04)	(0.04)	(0.00)	(0.16)	(0.00)
AR(35), $\hat{\sigma}_{err}$=2.392	[12.2]	[0.75]	[2.09]	[1.99]	[12.1]	[1.42]	[10.2]

Notes: n is the length of the log-return time series from 1-Feb-2000 to 17-Feb-2017; s is the marginal standard deviation; $\hat{\sigma}_{err}$ is the estimated standard deviation of residuals in the AR(p by AIC) model; (number) refers to P-value; and [number] refers to the ratio of the directionality indicator to its standard deviation from Monte-Carlo simulation. POT refers to the POT(5).

The U.S. banking industry has experienced a series of severe financial and economic crises which have affected the performance of banks (e.g. Dell'Ariccia et al. 2008, Allen & Christa 2013). Therefore, we are interested in comparing directionality during crises with directionality during periods.

Allen & Christa (2013) identified two banking crises: credit crunch (1990:Q1 to 1992:Q4), and subprime lending crisis (2007:Q3 to 2009:Q4); three market crises: stock market crash (1987:Q4), Russian debt crisis, and Long-Term Capital Management bailout (1998:Q3 to 1998:Q4); together with dot.com bubble and September 11 attack on the World Trade

Table 6: Stable and unstable periods between May-1999 to Feb-2017.

Period	Status
May to Dec-1999	Stable
Jan-2000 to Dec-2002	Unstable
Jan-2003 to Dec-2006	Stable
Jan-2007 to Dec-2012	Unstable
Jan-2013 to Feb-2017	Stable

Center (2000:Q2 to 2002:Q3), in the U.S. as crises periods. Here, we combine the dot.com bubble and the September 11 attack with the 2001-2002 recession (Wikipedia 2016*b*) as major events that happened in the U.S between 2000 and 2002. We also combine the subprime mortgage crisis with the 2007-2008 GFC, with effects continuing into 2009, the collapse of Lehman Brothers in 2008 and the Great Recession of 2008-2012 (Wikipedia 2016*a*) as a second period of major shocks in the U.S. during 2007 to 2012. Our definition of stable and unstable periods is summarized in Table 6.

Table 7: Summary statistics of the Wells Fargo log-returns by stable and unstable sub-series.

Period	Mean	Range	Skewness	Kurtosis	Volatility	Directionality
Stable	-0.041	14.0	0.15	3.4	2.32	0.11
Unstable	0.028	21.5	0.22	6.1	1.98	0.46
Stable	0.054	7.0	0.13	4.2	0.88	0.04
Unstable	0.007	55.6	0.73	16.0	3.74	0.26
Stable	0.063	12.5	0.15	5.8	1.18	0.05

We compare mean, range, skewness, kurtosis, volatility and directionality in the stable and unstable sub-series (Table 7). The volatility is here defined as the standard deviation of the log-returns, and the directionality is measured by skewness of the first differences $\hat{\gamma}_d$. It seems reasonable to expect higher mean values during the stable periods, but the first stable period is quite anomalous. The skewness is somewhat higher during the unstable periods. Directionality is positive and markedly higher during all of the unstable periods than in any of the stable periods. The range is higher in the unstable periods but the volatility is higher in the first stable period than in the following unstable period. These preliminary

results lead us to question whether monitoring directionality, as well as volatility, would give early warning of unstable periods.

To investigate this we use a 9-month moving directionality series (MD), 21 days by 9 months is 189 trading days, which is defined in terms of skewness of the first differences $\hat{\gamma}_d$ in the log-returns of Wells Fargo. The formula used to calculate the MD is given in Equation (12),

$$\hat{\gamma}_{d,t} = \frac{\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta x}_t)^3/m}{[\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta x}_t)^2/m]^{3/2}}, \tag{12}$$

where $\Delta x_t = x_t - x_{t-1}$, $t = 189, ..., 4479$, $m = 189$, $\bar{\Delta x}_t = \sum_{i=0}^{188}\Delta x_{t-i}/m$ and x_t is the full-length time series of log-returns. Similarly for the volatility, given by the marginal standard deviation of log-returns $\{x_t\}$, the 9-month moving volatility series MV is defined by Equation (13).

$$\hat{\sigma}_i = \sqrt{\frac{\sum_{i=0}^{188}(x_{t-i} - \bar{x}_t)^2}{m}}. \tag{13}$$

where $\bar{x}_t = \sum_{i=0}^{188} x_{t-i}/m$.

The plots of the MD, $\{\hat{\gamma}_{d,t}\}$, and the MV ,$\{\hat{\sigma}_i\}$,for the Wells Fargo together with its share prices and log-returns are illustrated in Figure 7. We use vertical lines in the plots to show the dates when the crises are reported in news. For example, 15-Jan-2001 for the 2001 dotcom crash, 11-Sep-2001 for the 2001 Sept 11 attack on the World Trade Centre, the subprime mortgage crisis was first widely reported on 1-Jul-2007 and the collapse of Lehman Brothers was reported on 15-Sep-2008. No exact dates are found to indicate the recessions, but the U.S. unemployment rate rose to 4.2% on 1-Feb-2001, the National Association of Securities Dealers Automated Quotations (NASDAQ) reached a 6-year low on 24-Sept-2002 and U.S. stock market fell dramatically on 1-Aug-2011 for the 2011 recession. These vertical lines are particularly helpful in this case as the share prices move in short to long-term downward trends after the shocks in Figure 7(a). In particular, major declines after the 1-Jul-2007 and the 15-Sep-2008 during the crisis period. Similarly, more variability in the log-returns during the 2007-2012 unstable period in Figure 7(b).

The MD appears to be less smooth than the MV, and the MD soars on two occasions during crises periods in Figure 7(c). In particular, the fourth vertical line of 24-Sept-2002 and the second last vertical line of 14-Sept-2008. In contrast, MV appears to be less

influenced by crises. We have seen similar effects in the moving directionality at 3-month, 6-month and 12-month.
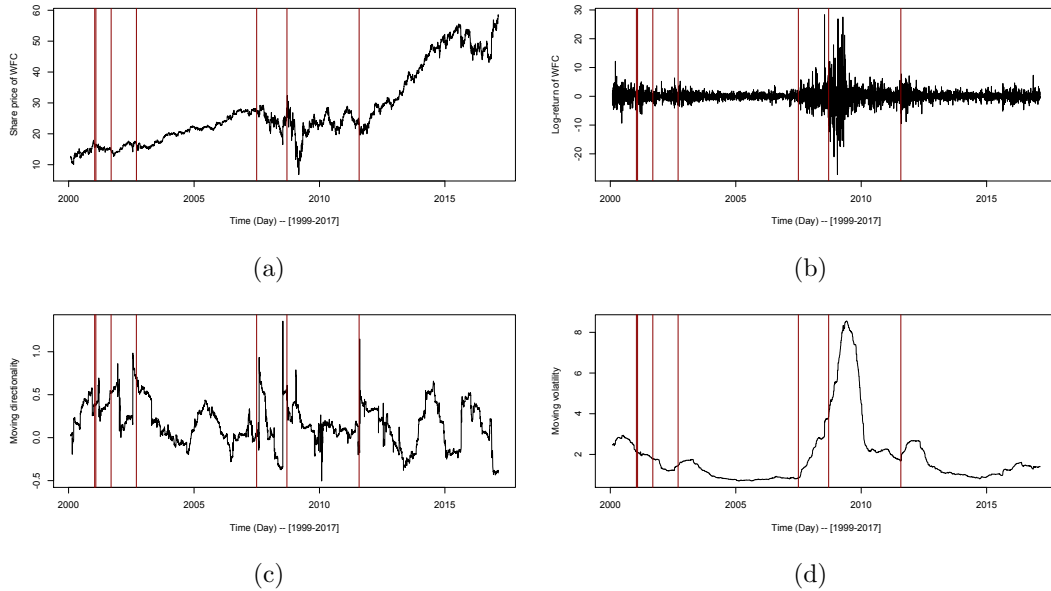
Figure 7: Daily share price (a), daily log-return (b), 9-month moving directionality (c), and 9-month moving volatility (d) of Wells Fargo.
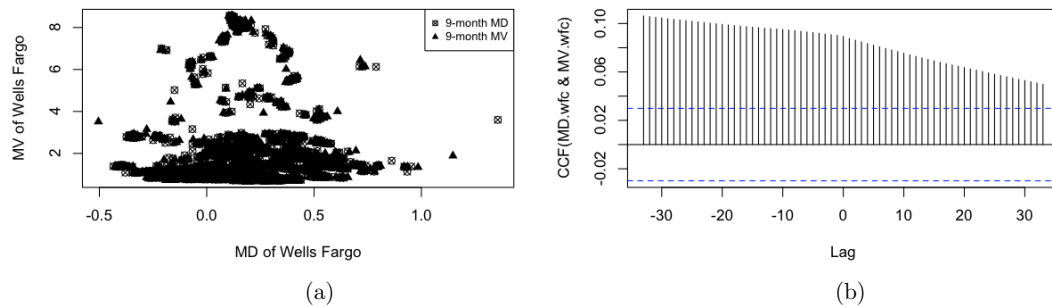
Figure 8: Scatter plot (left) and cross-correlation function (CCF) (right) of the 9-month moving directionality and the 9-month moving volatility of Wells Fargo.

We illustrate the relationship between directionality and volatility for the Wells Fargo daily log-return series using scatter plot of the $\{\hat{\gamma}_{d,t}\}$ versus the $\{\hat{\sigma}_i\}$ in Figure 8(a). There is no obvious linear pattern from the plots, and the correlation coefficient at a lag of 0 is

0.09. This low correlation coefficient indicates that directionality could provide additional information to volatility. Moreover, in the case of Wells Fargo, the CCF in Figure 8(b) shows directionality leads the occurrence of volatility over this period.

These features of directionality suggest that it may have potential as an early warning signal of financial crises and consequent falling share prices. We have carried out preliminary work on an investment strategy of buying when the MD is high and selling when it is low for a portfolio of U.S. bank shares. Over the period of May 3, 1999 to February 17, 2017 the return was 2.15% per annum, which was higher than achieved with MV. This work is continuing.

5 Concluding Remarks

The first stage in a time series analysis is usually to identify any trend and seasonal effects and to remove these before fitting stationary time series models. It is straightforward to check for directionality after removing trend and seasonal effects, and before fitting any stationary models. In particular, the plot and reverse order plot, and the statistics DLQC, P_d^+, $\hat{\gamma}_d$ and $\hat{\gamma}_{dab}$ are straightforward and quick to implement. If directionality is detected, then this indicates that the errors should be modelled by some non-Gaussian distribution or that the model should be non-linear or both.

There are many ways in which a time series can exhibit directionality, and different formal quantitative tests are essential to detect different forms of directionality in time series. The suite of directionality tests that we propose distinguishes four categories of qualitative directional characteristics.

- *A time series with rapid rises followed by slow recessions, or slow increases followed by fast recessions.*
  Four directional time series fall into this category: the sunspots; the exchange rate; the EEG (epilepsy); and the ds NGRIP. This form of directionality is best detected by the methods based on the first differences. The estimator $\hat{\gamma}_d$appears to be the best for detecting directionality but it is heavily influenced by any extreme outliers and it is advisable to use it in conjunction with a box plot of the differences. The

L-skewness has an advantage that it is less influenced by outliers.

- *A time series with rapid rises above the mean and rapid recessions below the mean (or rapid returns to the mean).*
  The statistic $\hat{\gamma}_{dab}$ was introduced to detect this category of directionality. However the DLQC and POT statistics should detect ot if it is marked. The SOI series and the ds Rainfall appear to be in this category.

- *A time series with directionality or asymmetric patterns above a threshold.*
  Example of time series for this category are: (i) the time series in category 1, except the exchange rate; and (ii) the time series in category 2 because POT ignores information below the threshold. This form of directionality is best detectable by the POT but sensitive to the chosen averaging parameter.

- *Intermittent directionality.*
  In some time series, share returns in particular, the directionality seems to be intermittent. This can be monitored by calculating a moving directionality index. This could be based on any of the measures but we have focused on a moving form of the product moment skewness of differences that appears to work well with shorter series.

Testing the normality of residuals after fitting ARMA models is an indirect method of testing a null hypothesis of reversibility. There was evidence against this null hypothesis for all time series except dtds Rainfall. All the series we considered had length greater than 200 and for all but one there was convincing statistical evidence of directionality. This leads to the more subjective decision of whether the directionality is sufficiently noticeable to be allowed for in the modelling, and raises the question of the potential benefits from modelling or monitoring directionality.

We use seven formal quantitative tests to detect directionality time series, based on the general procedure for detecting directionality illustrated in Figure 2. We conclude that certain tests are more sensitive at detecting specific forms of directionality. For the time series we considered, $\hat{\gamma}_d$, $\hat{\gamma}_{dab}$ and POT appeared to be more sensitive than the other statistics.

However, calculation of the suite of seven statistics may provide useful information and we recommend this as a general procedure.

In Section 4, we considered three case studies that made use of directionality. In the first we demonstrated how directionality might be used as a criterion for choosing a suitable time series model. The second was highly speculative, but the possibility that directionality might provide, or contribute towards, early warning of an epileptic seizure is sufficiently important to warrant further investigation. In the third we demonstrate that the monitoring of directionality has potential as an early warning of unsettled financial markets and that it is a useful complement to the monitoring of volatility. Time series data is collected from many processes, both natural and anthropogenic, by a wide range of organizations, and directionality can easily be monitored as part of routine analysis. We suggest that doing so may provide new insights to the processes.

SUPPLEMENTARY MATERIAL

R Code for the Suite of Directionality Tests: R Code for plotting and calculating the seven statistics of directionality may be found in the online version of this article. (pdf file)

References

Allen, N. B. & Christa, H. S. B. (2013), 'How does capital affect bank performance during financial crises?', *Journal of Financial Economics* **109**(1), 146–176.

Andrzejak, R. G., Lehnertz, K., Mormann, F., Rieke, C., David, P. & Elger, C. E. (2001), 'Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state', *Physical Review E* **64**(061907), 1–8.

Beare, B. K. & Seo, J. (2014), 'Time irreversible copula-based markov models', *Econometric Theory* **30**(5), 923–960.

Chatfield, C. (2004), *The Analysis of Time Series: An Introduction, Sixth Edition*, Chapman & Hall/CRC.

Cox, D. R. (1981), 'Statistical analysis of time series: Some recent developments', *Scandinavian Journal of Statistics* **8**, 93–115.

Dell'Ariccia, G., Detragiache, E. & Rajan, R. (2008), 'The real effect of banking crises', *Journal of Financial Intermediation* **17**(1), 89 – 112. Financial Contracting and Financial System Architecture.

Granger, C. W. J. & Andersen, A. (1978), 'On the invertibility of time series models', *Stochastic Processes and their Applications* **8**(1), 87–92.

Hosking, J. R. M. (1990), 'L-moments: analysis and estimation of distributions using linear combinations of order statistics', *Journal of the Royal Statistical Society. Series B (Methodological)* **52**(1), 105–124.

Kroese, D. P., Taimre, T. & Botev, Z. I. (2011), *Handbook of Monte Carlo Methods*, John Wiley & Sons.

Lawrance, A. J. (1991), 'Directionality and reversibility in time series', *International Statistical Review/Revue Internationale de Statistique* **59**(1), 67–79.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2015a), Modelling directionality in stationary geophysical time series, *in* 'Proceedings of the International work-conference on Time Series (ITISE 2015)', pp. 755–766.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2015b), Modelling and simulation of directional financial time series, *in* 'Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)', pp. 1022–1028.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2016b), 'Directionality and volatility in electroencephalogram time series', *American Institute of Physics (AIP) Conf. Proc.* **1739**, 020080:1–8.

Mansor, M. M., Mohd. Isa, F. L., Green, D. A. & Metcalfe, A. V. (2016d), 'Modelling directionality for paleoclimatic time series', *ANZIAM J.* **57** (EMAC2015), C66–C81.

Martens, P. C., Nandy, D. & Munoz-Jaramillo, A. (2011), Meridional surface flows and the recent extended solar minimum, *in* 'Bulletin of the American Astronomical Society', Vol. 1, p. 1705.

Nandy, D., Muñoz-Jaramillo, A. & Martens, P. C. (2011), 'The unusual minimum of sunspot cycle 23 caused by meridional plasma flow variations', *Nature* **471**(7336), 80–82.

NASA (2017), 'The sunspot cycle', https://solarscience.msfc.nasa.gov (last accessed 17 March 2017).

NGRIP_members, Andersen, K. K., Azuma, N., Barnola, J.-M., Bigler, M., Biscaye, P., Caillon, N., Chappellaz, J., Clausen, H. B., Dahl-Jensen, D., Fischer, H., Flückiger, J., Fritzsche, D., Fujii, Y., Goto-Azuma, K., Grønvold, K., Gundestrup, N. S., Hansson, M., Huber, C., Hvidberg, C. S., Johnsen, S. J., Jonsell, U., Jouzel, J., Kipfstuhl, S., Landais, A., Leuenberger, M., Lorrain, R., Masson-Delmotte, V., Miller, H., Motoyama, H., Narita, H., Popp, T., Rasmussen, S. O., Raynaud, D., Röthlisberger, R., Ruth, U., Samyn, D., Schwander, J., Shoji, H., Siggard-Andersen, M.-L., Steffensen, J. P., Stocker, T., Sveinbjörnsdottir, A. E., Svensson, A., Takata, M., Tison, J.-L., Thorsteinsson, T., Watanabe, O., Wilhelms, F. & White, J. (2004), 'High-resolution record of northern hemisphere climate extending into the last interglacial period', *Nature* **431**, 147–151.

Priestley, M. B. (1980), 'State-dependent models: a general approach to non-linear time series analysis', *Journal of Time Series Analysis* **1**(1), 47–71.

Rao, T. S. (1981), 'On the theory of bilinear time series models', *Journal of the Royal Statistical Society. Series B (Methodological)* **43**(2), 244–255.

R_Core_Team (2013), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

Sankarasubramanian, A. & Srinivasan, K. (1999), 'Investigation and comparison of sam-

pling properties of {L}-moments and conventional moments', *Journal of Hydrology* **218**(1), 13–34.

Sharifdoost, M., Mahmoodi, S. & Pasha, E. (2009), 'A statistical test for time reversibility of stationary finite state markov chains', *Applied Mathematical Sciences* **52**, 2563–2574.

Soubeyrand, S., Morris, C. E. & Bigg, E. K. (2014), 'Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data', *Environmental Modelling & Software* **61**, 78–86.

Steyn-Ross, M. L., Steyn-Ross, D. A., Sleigh, J. W. & Liley, D. T. J. (1999), 'Theoretical electroencephalogram stationary spectrum for a white-noise-driven cortex: Evidence for a general anesthetic-induced phase transition', *Phys. Rev. E* **60**, 7299–7311.

Tong, H. & Lim, K. S. (1980), 'Threshold autoregression, limit cycles and cyclical data', *Journal of the Royal Statistical Society. Series B (Methodological)* **42**(03), 245–292.

van Prehn, J., Vincken, K. L., Sprinkhuizen, S. M., Viergever, M. A., van Keulen, J. W., van Herwaarden, J. A., Moll, F. L. & Bartels, L. W. (2009), 'Aortic pulsatile distention in young healthy volunteers is asymmetric: Analysis with {ECG}-gated {MRI}', *European Journal of Vascular and Endovascular Surgery* **37**(2), 168 – 174.

Weiss, G. (1975), 'Time-reversibility of linear stochastic processes', *Journal of Applied Probability* **12**(4), 831–836.

Wikipedia (2016a), 'Financial crisis of 2007–2008', https://en.wikipedia.org (last accessed 11 November 2016).

Wikipedia (2016b), 'List of economic crises', https://en.wikipedia.org (last accessed 5 October 2016).

Wild, P., Foster, J. & Hinich, M. J. (2014), 'Testing for non-linear and time irreversible probabilistic structure in high frequency financial time series data', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **177**(3), 643–659.

Chapter 3

Modelling Directionality

3.1   Modelling Directionality in Stationary Geophysical Time Series

<div style="border:1px solid">

Statement of Authorship

Modelling Directionality in Stationary Geophysical Time Series

Published in *Proceedings of the International Work-Conference on Time Series Analysis (ITISE 2015)*, 2015.

Mahayaudin M. Mansor (Candidate)

Constructed statistical models, conducted analyses and interpreted results, wrote manuscript and acted as corresponding author.

Overall percentage: 80%

This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.

Signed: .... Date: 14 June 2017

By signing the Statement of Authorship, each co-author certifies that: (i) the candidate's stated contribution to the publication is accurate; (ii) permission is granted for the candidate in include the publication in the thesis; and (iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Max E. Glonek

Assisted in performing an analysis to observe the effect of excess kurtosis of symmetrical error distributions for a reversible process on directionality. Reviewed the draft manuscript.

Signed: .... Date: 14/6/17

David A. Green

Supervised and ... of work and provided critical evaluation.

Signed   Date: 14/6/17

Andrew V. Metcalfe

Supervised development of work and provided critical evaluation.

Signed   .... Date: 14/6/17

</div>

Modelling Directionality in Stationary Geophysical Time Series

Mahayaudin M. Mansor, Max E. Glonek,
David A. Green, and Andrew V. Metcalfe

School of Mathematical Sciences
University of Adelaide, Adelaide 5005 SA, Australia
`{mohdmahayaudin.mansor,max.glonek,david.green,`
`andrew.metcalfe}@adelaide.edu.au`

Abstract. Many time series show directionality inasmuch as plots against time and against time-to-go are qualitatively different, and there is a range of statistical tests to quantify this effect. There are two strategies for allowing for directionality in time series models. Linear models are reversible if and only if the noise terms are Gaussian, so one strategy is to use linear models with non-Gaussian noise. The alternative is to use non-linear models. We investigate how non-Gaussian noise affects directionality in a first order autoregressive process AR(1) and compare this with a threshold autoregressive model with two thresholds. The findings are used to suggest possible improvements to an AR(9) model, identified by an AIC criterion, for the average yearly sunspot numbers from 1700 to 1900. The improvement is defined in terms of one-step-ahead forecast errors from 1901 to 2014. A physical interpretation of directional time series for the sunspots is suggested.

Keywords: Directional time series·reversibility·sunspot numbers·non-linear time series·non-Gaussian errors.

1   Introduction

Directionality, defined as asymmetry in time [3], enables us to tell a difference between a sequence of observations plotted in time order (time series) and the sequence plotted in reverse time order (time-to-go). A clear example of directionality can be seen in the average yearly sunspot numbers 1700-2014 (Fig. 1). A time series model is reversible if, and only if, its serial properties are symmetric, with respect to time and time-to-go. A linear time series model with Gaussian errors (LGE model) is reversible, but a linear time series model with non-Gaussian errors is directional (LNGE model) [3]. Non-linear time series models are also directional, whether or not the errors are Gaussian (NLGE and NLNGE models respectively).

Directionality is important for forecasting because it indicates that models other than LGE should be considered [2,3]. If the error distribution is better modelled as non-Gaussian, this will lead to more precise limits of prediction

for forecasts. If non-linear models provide a better fit than linear models, the forecast from non-linear models will be more accurate.
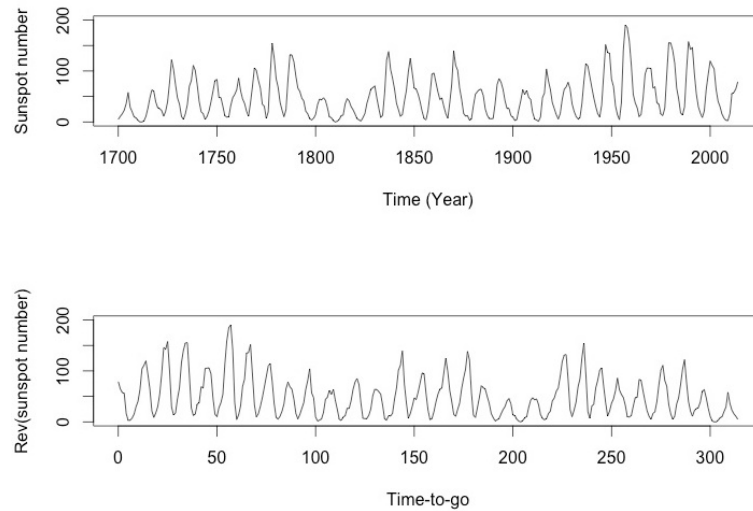
Fig. 1. Graphical inspection of directionality shows the sunspot observations rise more quickly than they fall in time order (above) and rise more slowly than they fall in reverse time order (below)

The paper is arranged as follows. Section 2 describes well-established procedures for detecting directionality in time series, together with a modification for directional time series that are symmetrical about the time axis. In Section 3 we consider some LNGE and NLGE models for time series and investigate the performances of various measures of directionality by simulation. In Section 4 we provide evidence of directionality in the sunspot series, model it with a threshold autoregressive (TAR) model of order 2 [2], and demonstrate improved predictions. Section 5 is a discussion.

2   Detecting Directionality

In general any trend or seasonality should be removed before investigating directionality in the stationary series. There are many possible directional features in time series, and these include, for example, sharp increases followed by a slow recession or slow increases followed by a plummet. In such cases the series is asymmetric with respect to time and also with respect to its mean or its median. In contrast a time series may exhibit both sharp increases and sharp decreases, followed by more gradual returns to the median value. Such time series are asymmetric with respect to time but symmetric with respect to the median. Different statistics are appropriate for detecting directionality in series that are asymmetric or symmetric with respect to the median.

In this paper, we employ relatively simple and well-established tests [3] to detect directionality in time series: difference in linear quadratic lagged correlations; proportion of positive differences; skewness of differences; and tests based on comparisons of time from threshold to peak with time from peak to threshold. More recent tests are based on properties of Markov chains [1]; spectral estimation of kurtosis of differences in financial time series [6]; and the FeedbackTS package in R to detect time directionality occurring in specific fragments of time series [5].

2.1   Difference in Linear Quadratic Lagged Correlations

Directionality has the complementary concept of reversibility. Demonstrating evidence that a series is not reversible, is another way of expressing that the series is directional. Following [3], a time series modelled by random variables $\{X_t\}$ for $t = 0, \pm 1, \pm 2, ...$ is reversible if the joint distribution of $X_t, X_{t+1}, ..., X_{t+r}$ and $X_{t+r}, X_{t+r-1}, ..., X_t$ is the same for all $r = 1, 2, ...$ In particular, a time series that is reversible has

$$Corr(X_t, X_{t+1}^2) = Corr(X_t^2, X_{t+1}). \tag{1}$$

A measure of directionality can be based on the difference in the sample estimates of these correlations. The non-dimensional measure used here is

$$DLQC = \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})(x_{t+1} - \bar{x})^2}{[\sum_{t=1}^{n}(x_t - \bar{x})^2]^{3/2}} - \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})^2(x_{t+1} - \bar{x})}{[\sum_{t=1}^{n}(x_t - \bar{x})^2]^{3/2}}. \tag{2}$$

We explain the rationale behind this statistic if, for example, there are sharp increases followed by slow recessions. Suppose the sharp increase occurs between x_t and x_{t+1}, then $(x_t - \bar{x})$ could be negative or positive but $(x_{t+1} - \bar{x})$ is very likely to be positive. It follows that $(x_t - \bar{x})(x_{t+1} - \bar{x})^2$ is negative or positive whereas $(x_t - \bar{x})^2(x_{t+1} - \bar{x})$ is positive and hence DLQC will tend to be negative.

Both terms in DLQC are correlations, so bounds for the DLQC are [-2,2], but typical values in directional time series are smaller by two orders of magnitude. For the sunspot series, which exhibits clear directionality DLQC is -0.06.

2.2   Methods Based on First Differences

More intuitive measures of directionality can be based on the distribution of lag one differences, Y_t. If, for example, there are sharp increases and slow recession there will be fewer large positive differences and more small negative differences. The distribution will be positively skewed. Let the observed time series be X_t and define the lag one first order differences[1] as

$$Y_t = X_t - X_{t-1} \quad \text{for} \quad t = 2, 3, ...n. \tag{3}$$

---

[1] In the following, "differences" refers to these lag one differences.

2.2.1 Percentage of Positive Differences

The percentage of positive of differences is

$$P^+ = \frac{number\ of\ positive\ Y_t}{number\ of\ positive\ Y_t\ +\ number\ negative\ of\ Y_t} \times 100 \qquad (4)$$

This formula excludes possible zero differences. If the time series is symmetric about the median and there are both sharp increases and sharp decreases then there will tend to be both large positive and large negative differences, which tend to negate. Therefore, we adjust the measure to

$$P_{abm} = \frac{(P^+_{above}) + (P^-_{below})}{(P^+_{above}) + (P^-_{above}) + (P^+_{below}) + (P^-_{below})} \times 100 \qquad (5)$$

where a difference is classified as above or below according to whether X_t is above or below the median. Also, P^+_{above}, and P^-_{below}, are the proportions of differences above, or below the median that are positive, or negative.

If a series is reversible the expected values of the percentages in Equations (4) and (5) are 50%. If differences are treated as independent, in a time series of length 1000, P_{abm} would need to be differ from 50% by at least 3.2% to be statistically significant at the 0.05 level.

2.2.2 Product Moment Skewness of Differences

A potentially more sensitive test for directionality is to consider the skewness of the distribution of differences ($\hat{\gamma}$) [3].

$$\hat{\gamma} = \frac{\sum_{t=1}^n (y_t - \bar{y})^3}{[\sum_{t=1}^n (y_t - \bar{y})^2]^{3/2}} \qquad (6)$$

If a time series has both sharp increases and sharp decreases and is symmetric about the median we adapt the definition to be

$$\hat{\gamma}_{abm} = |\hat{\gamma}_{above}| + |\hat{\gamma}_{below}|. \qquad (7)$$

Significant non-zero skewness of either $\hat{\gamma}$ or $\hat{\gamma}_{abm}$ are evidence of directionality.

2.3   Threshold-Peak or Threshold-Trough Test

Consider a threshold (H) set, in this investigation, at the upper quintile of the marginal distribution of the time series x_t. Suppose that $x_{j-1} < H$, $x_j > H$, and that x_t remains above H until $x_{j+k+1} < H$. Denote the time when x_t is greatest (peak value) for $j \leq t \leq (j+k)$ as $(j+p)$. Define the difference between time from threshold to peak and time from the peak to the threshold as

$$DHPPH_j = (k - p) - p. \qquad (8)$$

A similar definition can be constructed for a threshold-trough test, using least values (troughs) of series of observations below the lower quintile (L). Denote

the difference between time from threshold to trough and time from trough to threshold as $DLTTL_j$. Calculate $DHPPH$ and $DLTTL$ as the average of $DHPPH_j$ and $DLTTL_j$ respectively for all exceedances of H and excursions below L. The expected value of $DHPPH$ and $DLTTL$ is 0 for a reversible series.

2.4 Evidence of Directionality

In general, a directional time series will not show directionality on all of these statistics. Any one test being statistically significant at the α level, where $\alpha < 0.05$ say, is evidence of directionality. A conservative allowance for multiple testing would be to use a Bonferroni inequality and claim overall significance at less than an $m\alpha$ level where m is the number of tests.

3 Models for Directionality

3.1 LNGE Model with Symmetric Distribution of Errors

Directionality can be modeled using LNGE, NLGE or NLNGE. In this subsection we simulate directional series using LNGE by taking an autoregressive process of order 1, AR(1) with Beta and t-distributions for the errors. Each series is an AR(1) process

$$X_t = \alpha X_{t-1} + \epsilon_t , \tag{9}$$

with $\alpha = 0.9$ to give substantial serial correlation. We create 14 directional time series of length 1 million. Seven processes have errors with a Beta distribution parameterised by (α, β) and seven processes with t-distribution errors (Table 1).[2] A series using AR(1) with Gaussian errors is a symmetric control.

3.1.1 Measures of Directionality of LNGE Model

We calculated P^+_{above} and P^-_{below} for each of the 15 simulated AR(1) processes. Table 1, and Fig. 2, show the changes in proportions of differences as the excess kurtosis of the error distribution changes. The proportions are 0.50 when the excess kurtosis is 0, as the sampling error in time series of length 1 million is negligible. Significant deviation from 0.500 are evidence of directionality.

3.1.2 Power of Tests of Directionality of LNGE Model

To compare the power of the P^+_{above}, P^-_{below}, $DHPPH$, $DLTTL$ and DLQC statistics we split each time series of length 1,000,000 values into 1000 time series of 1000 observations. For each time series of 1000 observations we calculate whether or not a directionality statistic (U) is statistically significant at a nominal 0.05 level as follows. Calculate the value of U for the twenty consecutive sub-series, u_i for $i = 1, 2, ..., 20$, each of length 50 and compute

$$\bar{u} = \frac{\sum_{i=1}^{20} u_i}{20} \quad \text{and} \quad sd(u_i) = \sqrt{\frac{\sum_{i=1}^{20} (u_i - \bar{u})^2}{19}}. \tag{10}$$

---

[2] All series start at 0, and are simulated with a 30-observation burn-in time.

Table 1. Proportion of positive differences above the median (P_{above}^+), negative differences below the median (P_{below}^-), and excess kurtosis of the error distributions.

Error dist.	P_{above}^+	P_{below}^-	Excess kurtosis
β(1e-6,1e-6)	0.6109	0.6112	-2.0000
β(0.25,0.25)	0.5665	0.5670	-1.7143
β(0.5,0.5)	0.5471	0.5461	-1.5000
β(1,1)	0.5292	0.5285	-1.2000
β(3.5,3.5)	0.5107	0.5105	-0.6000
β(6,6)	0.5059	0.5056	-0.4000
β(28.5,28.5)	0.5011	0.5015	-0.1000
N(0,1)	0.4993	0.4998	0.0000
t, 30df	0.4976	0.4967	0.2308
t, 15df	0.4947	0.4946	0.5455
t, 8df	0.4892	0.4884	1.5000
t, 6df	0.4833	0.4835	3.0000
t, 5df	0.4791	0.4793	6.0000
t, 4.25df	0.4743	0.4730	24.000
t, 4.06df	0.4724	0.4719	100.000
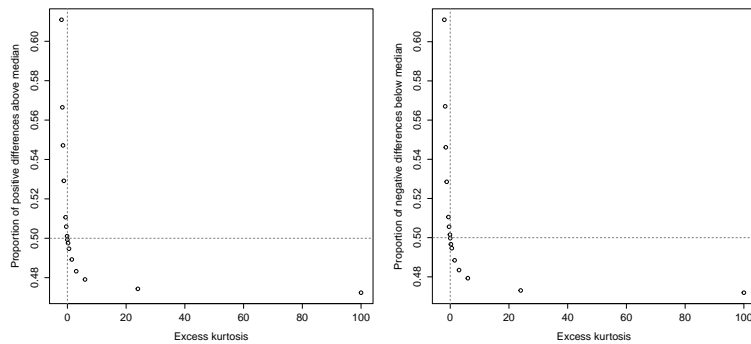
Fig. 2. Proportion of positive differences above the median vs. excess kurtosis (left) and proportion of negative differences below the median vs. excess kurtosis (right)

Under the null hypothesis (H_0) that the time series is reversible $E[U] = u_0$, where, for example, if U is the proportion of positive differences $E[U] = 0.5$. Calculate

$$t = \frac{\bar{u} - u_0}{sd(u_i)/\sqrt{20}}. \tag{11}$$

Reject H_0, against the alternative hypothesis that $E[U] \neq u_0$, if $|t| > t_{19,0.025}$ where $t_{19,0.025}$ is the upper 0.025 quantile of the t-distribution with 19 degrees of freedom (df).

The most striking feature of Table 2 is the proportions when the errors are N(0,1), which demonstrate that the actual significance level is lower than 0.05, especially for the DLQC test. This is because the correlation structure in the

time series has a consequence that the u_i are not independent. The P^+_{above} and P^-_{below}, and $DHPPH$ and $DLTTL$, give almost identical results, as they should for a time series with a symmetric error distribution.

Table 2. Proportion of series, of 1000 observations, that show significant directionality at nominal 5% level with two-sided alternative hypothesis.

Error dist.	P^+_{above}	P^-_{below}	$DHPPH$	$DLTTL$	DLQC
β(1e-6,1e-6)	0.969	0.984	0.738	0.732	0.002
β(0.25,0.25)	0.725	0.752	0.440	0.484	0.004
β(0.5,0.5)	0.486	0.504	0.300	0.297	0.001
β(1,1)	0.193	0.191	0.166	0.147	0.005
β(3.5,3.5)	0.048	0.052	0.050	0.053	0.007
β(6,6)	0.044	0.039	0.042	0.048	0.007
β(28.5,28.5)	0.043	0.041	0.043	0.045	0.003
N(0,1)	0.036	0.031	0.039	0.041	0.011
t, 30df	0.031	0.040	0.053	0.040	0.008
t, 15df	0.043	0.052	0.040	0.058	0.016
t, 8df	0.044	0.064	0.079	0.060	0.018
t, 6df	0.088	0.087	0.098	0.081	0.017
t, 5df	0.091	0.115	0.105	0.132	0.024
t, 4.25df	0.138	0.149	0.158	0.178	0.024
t, 4.06df	0.144	0.158	0.160	0.166	0.039

All tests show an approximate four fold increase in significant results when the t-distribution of errors has 4.06 df, although the P^+_{above} and P^-_{below} are the most powerful tests for this case. When the excess kurtosis is negative the DLQC test is ineffective and the P^+_{above} and P^-_{below} are slightly more powerful than the $DHPPH$ and $DLTTL$ tests. The excess kurtosis of the symmetric distributions of the errors used in the AR(1) model needs to be quite extreme to see substantial directionality in the linear time series.

3.2 LNGE Model for Assymmetric Distribution of Errors

We create a directional series using AR(1) with $\alpha = 0.9$ and exponential errors with mean 1.

Table 3. Test statistics of directionality for simulated series given by AR(1) with exponential errors.

Mean	sd	DLQC	P^+	P_{abm}	$\hat{\gamma}$	$\hat{\gamma}_{abm}$
10.03	2.4614	-0.0313	38.17%	49.60%	1.8148	2.6927

Results in Table 3 shows clear directionality in the series with DLQC, P^+, $\hat{\gamma}$ and $\hat{\gamma}_{abm}$ statistically significantly different from 0, 50%, 0 and 0 respectively well beyond the 0.05 level [4]. Exponential errors in a linear model are an effective means of modelling one type of directionality.

3.3   NLGE Model

We simulate a series of length 2000 using a piecewise linear function for TAR(1) with Gaussian errors and 2 thresholds.

$$X_t = \begin{cases} \alpha_U X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T_U \\ \alpha_M X_{t-1} + \epsilon_t & \text{if } T_U < X_{t-1} < T_L \\ \alpha_L X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T_L \end{cases} \tag{12}$$

where $\epsilon_t \sim N(0,1)$. The upper threshold (T_U) is set at 3 and the lower threshold (T_L) is set at -3. The upper parameter (α_U), middle parameter (α_M), and lower parameter (α_L) are set at 0.5, 0.8 and -0.5 respectively.

Table 4. Test statistics of directionality for simulated series given by TAR(1) with Gaussian errors.

Mean	sd	DLQC	P^+	P_{abm}	$\hat{\gamma}$	$\hat{\gamma}_{abm}$
0.2004	1.5322	-0.2781	47.22%	52.65%	1.2417	0.7670

The realisation from TAR(1) with Gaussian errors model shows clear directionality with a substantial skewness ($\hat{\gamma}$) and a remarkably large absolute value of DLQC. The $\hat{\gamma}_{abm}$, P^+ and P_{abm} measures seem less effective at detecting directionality in this case.

4   The Sunspot Series

We provide formal evidence of directionality in the average yearly sunspots (1700-2014) [8], and fit time series models to the first 200 points (1700-1900). We use the remaining 115 points (1901-2014) to compare the one-step-ahead forecast errors.

4.1   Directionality in Sunspots

The DLQC, P^+ and $\hat{\gamma}$ all indicate directionality in the series and all have P-values of 0.00 (two sided P-values calculated from a parametric bootstrap procedure)[4].

Table 5. Summary table of test statistics of directionality for the sunspot series (1700-2014).

Series	Length	Mean	sd	DLQC	P^+	P_{abm}	$\hat{\gamma}$	$\hat{\gamma}_{abm}$
Sunspot numbers	315	46.68	40.24	-0.0598	42.49%	51.61%	0.8555	1.1047

We first consider AR(p) models of orders 1,2 and 9. AR(9) corresponds to the lowest Akaike Information Criterion (AIC) for AR(p) models. The results are summarised in Table 6. The AR(2) model is a substantial improvement on

the AR(1) model in terms of standard deviation of the errors. An AR(9) and ARIMA(2,0,1) give some further improvement on AR(2). The TAR models are of the form

$$(X_t - \mu) = \begin{cases} \alpha_{1U}(X_{t-1} - \mu) + \alpha_{2U}(X_{t-2} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) > T_U \\ \alpha_{1L}(X_{t-1} - \mu) + \alpha_{2L}(X_{t-2} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) < T_L \end{cases} \quad (13)$$

We consider three thresholds (T) set at the 70%, 80% and 90% percentiles respectively. The two parameter TAR(2) models use the AR(2) coefficients for α_{1L}, α_{2L} and then α_{1U}, α_{2U} which are estimated by non-linear least squares. The 4 parameter models estimate all of α_{1L}, α_{2L}, α_{1U} and α_{2U} by non-linear least squares from the mean adjusted (to 0) time series. Of all the models considered, TAR(2)_90% with 4 estimated parameters is the best with an estimated standard deviation of the errors ($\hat{\sigma}_{error}$) of 13.94.

Table 6. Time series models for the sunspot series (1700-1900) compared by σ_{est}.

AR model:	AR(1)	AR(2)	AR(9)
$\hat{\sigma}_{error}$	20.34	15.36	14.84
ARIMA model:	ARIMA(0,0,1)	ARIMA(1,0,1)	ARIMA(2,0,1)
$\hat{\sigma}_{error}$	21.80	16.60	14.72
TAR model, 2 parameters:	TAR(2)_70%	TAR(2)_80%	TAR(2)_90%
$\hat{\sigma}_{error}$	14.77	14.75	14.30
TAR model, 4 parameters:	TAR(2)_70%	TAR(2)_80%	TAR(2)_90%
$\hat{\sigma}_{error}$	14.67	14.61	13.94

We compare the one-step-ahead forecasting performance of AR(2), which is substanstially simpler than AR(9) and more practical for forecasting, with the 4 parameter TAR(2)_90% model for the years 1901-2014.

4.2 Details of Fitting AR(2) and TAR(2)_90% Models

Table 7. Sample mean, 2 coefficients and $\hat{\sigma}_{error}$ of AR(2).

$\hat{\mu}$	$\hat{\alpha}_1$	$\hat{\alpha}_2$	$\hat{\sigma}_{error}$
44.11	1.3459	-0.6575	15.36

Table 8. Sample mean, 4 estimated parameters and $\hat{\sigma}_{error}$ of TAR(2).

$\hat{\mu}$	T	α_{1L}	α_{2L}	α_{1U}	α_{2U}	$\hat{\sigma}_{error}$
44.11	51.79	1.5643	-0.8117	0.9978	-0.3158	13.94

The upper and the lower regimes of the TAR(2)_90% 4 parameter model in Table 8 is a stable AR(2) process which satisfied the requirements of stationary-triangular region $\alpha_2 > -1$, $\alpha_1 + \alpha_2 < 1$ and $\alpha_1 - \alpha_2 > -1$ [2]. Based on the Q-Q plots (Fig. 3), an assumption of normal errors seems a reasonable approximation.
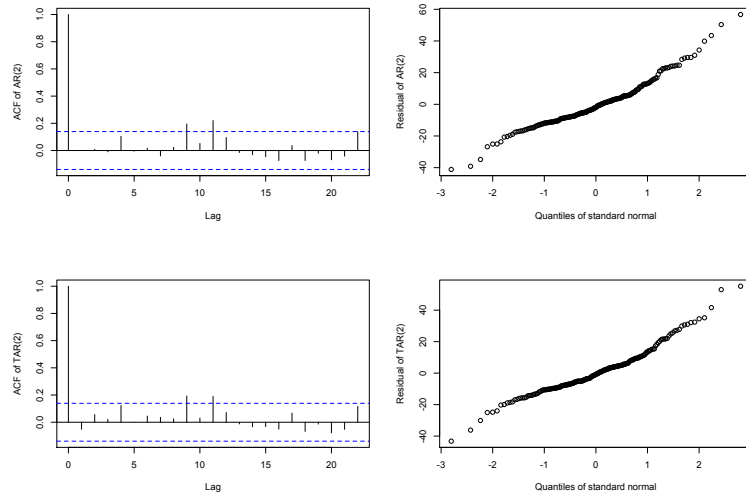
Fig. 3. Residuals of AR(2) and TAR(2)_90% with 4 estimated parameters

4.3   Simulation to Validate TAR Model

A simulation of 2000 points from the TAR(2)_90% 4 parameter model with Gaussian errors gave the statistics shown in Table 9. The TAR(2)_90% 4 parameter model shows some directionality from P_{abm} but unlike the sunspot series not on the other measures.

Table 9. Test statistics of directionality in the simulated sunpots and the sunspot numbers (1700-1900).

Series	Length	Mean	sd	DLQC	P^+	P_{abm}	$\hat{\gamma}$	$\hat{\gamma}_{abm}$
Simulated	2000	47.11	32.35	0.0069	50.67%	47.75%	-0.0892	0.2969
Sunspot numbers	200	44.11	34.76	-0.0609	43.94%	48.48%	0.8344	1.0646

5   Results and Discussion

5.1   Comparison of Predictions with TAR(2) and AR(2)

We compare the one-step-ahead forecasting performance of AR(2) with 4 parameter TAR(2)_90% model for the years 1901-2014. We define the forecasting performance by measuring the relative errors given by formula (14) and (15).

$$Relative\ error\ (E_{rel}) = \frac{(actual - predicted)}{actual} \qquad (14)$$

$$Absolute\ relative\ error\ (|E_{rel}|) = \frac{|(actual - predicted)|}{actual} \qquad (15)$$

Table 10. Forecasting measures of predicted sunspots to the sunspot series from 1901 to 2014.

Series	Mean(E_{rel})	sd(E_{rel})	Mean($|E_{rel}|$)	sd($|E_{rel}|$)
Predicted sunspots by AR(2)	-0.3831	1.8260	0.6730	1.7394
Predicted sunspots by TAR(2)	-0.2314	1.7044	0.6045	1.6095
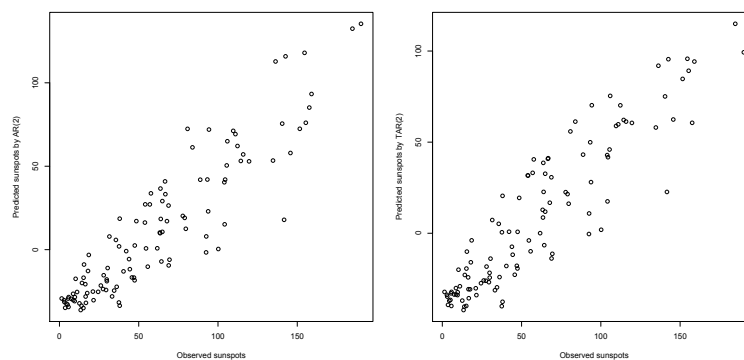
Fig. 4. Comparison of forecast values given by AR(2) and TAR(2) at 90 % threshold

The TAR(2) model offers some improvement over the AR(2) in terms of one-step-ahead predictions. However, there is a scope for more realistic modelling of directionality. This might be achieved by adding more thresholds, and including a measure of matching directionality in the optimisation criteria.

5.2 Physical Interpretation

In some cases there are physical reasons to expect directionality in time series. For example, stream flow will increase rapidly following a storm because of the immediate overland run-off, but decrease slowly because of the rain that has infiltrated the catchment and augmented the ground water flow to the river. However, directionality can also be a consequence of feedback in the system following extreme events [5]. For example, intense rainfall might influence future rainfall because biological particles released by rain aid the coalescence mechanism of rain formation. Such effects can lead to intermittent (fragmented time) directionality. Once directionality has been detected in a time series, a detailed analysis of its characteristics can either provide support for physical theories about the generating process.

Sunspots are dark regions of intense magnetic field that are associated with solar flares and coronal mass ejections. They are relatively cool because their intense magnetic fields inhibit the rise of heat from the solar interior, but are relatively bright at higher frequencies of radiation. Sunspots can appear alone, or in close connection to other sunspots making an active region. The energy

that sunspots lack in heat is compensated for by energy of the magnetic field. The magnetic fields rise above the surface and remain strong, while the rest of the sun has weak overlying magnetic fields. The strong magnetic fields form into loops that confine solar plasma and heat it to extreme temperatures in excess of 1 million K. It is highly plausible that such extreme electromagnetic events will promote feedback mechanisms [7].

6 Conclusion

There are many ways in which a time series can exhibit directionality, and different measures are needed to identify different characteristics. Similarly, different time series models can exhibit different directional characteristics. The most versatile modelling strategy is TAR models, which provide a piecewise linear approximation to a wide range of NLGE processes.
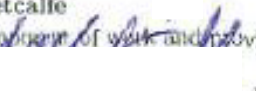
The sunspot series shows clear directionality, and a non-linear TAR(2) model gives a slight improvement on, out of sample, one-step-ahead predictions made with an AR(2) model. However, fitting the TAR(2) model by least squares did not, in this case at least, result in close reproduction of the directional characteristics of the sun spot series. In further work, we will explicitly include a penalty for deviations from the sun spot directionality measures when fitting the TAR model and investigate the effect on predictions.

References

1. Beare, B. K., Seo, J.: Time Irreversible Copula-Based Markov Models. Econometric Theory, 1–38 (2012)
2. Chatfield, C.: The Analysis of Time Series: An Introduction, Sixth Edition. Chapman and Hall/CRC Publication, 6 Edition, pp. 218–219, 223–224, 44–45 (2004)
3. Lawrance, A.: Directionality and Reversibility in Time Series. International Statistical Review/Revue Internationale de Statistique, 67–79 (1991)
4. Mahayaudin, M. M., Green, D. A., Metcalfe, A. V.: Detecting Directionality in Time Series. *in preparation*
5. Soubeyrand, S., Morris, C. E., Bigg, E. K.: Analysis of Fragmented Time Directionality in Time Series to Elucidate Feedbacks in Climate Data. Environmental Modelling & Software, 61:78–86 (2014)
6. Wild, P., Foster, J., Hinich, M.: Testing for Non-Linear and Time Iirreversible Probabilistic Structure in High Frequency Financial Time Series Data. Journal of the Royal Statistical Society: Series A (Statistics in Society), 177(3):643–659 (2014)
7. National Oceanic and Atmospheric Administration, `http://esrl.noaa.gov`
8. Solar Influences Data Analysis Center, Sunspot Index and Long-term Solar Observations, `http://www.sidc.be/silso`

3.2 Modelling and Simulation of Directional Financial Time Series

Statement of Authorship

Modelling and Simulation of Directional Financial Time Series

Published in *Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)*, 2015.

Mahayaudin M. Mansor (Candidate)

Constructed statistical models, conducted analyses and interpreted results, wrote manuscript and acted as corresponding author.

Overall percentage: 80%

This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.

Signed: . . . . Date: 14 June 2017

By signing the Statement of Authorship, each co-author certifies that: (i) the candidate's stated contribution to the publication is accurate; (ii) permission is granted for the candidate in include the publication in the thesis; and (iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

David A. Green

Supervised development of work and provided critical evaluation.

Signed: Date: 14/6/17

Andrew V. Metcalfe

Supervised development of work and provided critical evaluation.

Signed: . Date: 14/6/17

Modelling and Simulation of Directional Financial Time Series

Mahayaudin M. Mansor, David A. Green and Andrew V. Metcalfe

School of Mathematical Sciences, University of Adelaide, Adelaide 5005 SA, Australia
Email: mohdmahayaudin.mansor@adelaide.edu.au

Abstract: A stationary time series model is directional if it has properties that are not symmetric with respect to time, for example if, after the expected value is adjusted to 0, the expected value of the product of the variable squared and the following variable differs from the expected value of the product of the variable and the following variable squared. Marked directionality in time series is apparent from a comparison of time series plots in chronological order and in reverse order (time-to-go), but formal quantitative tests are needed to determine the strength of evidence against an assumption of reversibility (no directionality). Linear time series models with Gaussian noise (LGN models) are reversible whereas linear models with other noise distributions and non-linear models are directional. Evidence of directionality in time series can be found in many disciplines including financial time series. Examples of directional financial time series are presented, together with values of directionality statistics. A method of calculating the p-values of the directional statistics, corresponding to a null hypothesis of reversibility, using simulation is presented. The time series are then modelled by threshold autoregressive models, with noise distributions, which can be non-Gaussian, fitted to match the empirical distribution of the residuals. The fitting criterion is a weighted sum of the error sum of squares and a measure of discrepancy between values of the directionality statistics in the time series and long simulations from the potential model. Strategies for determining suitable weights for the criterion will be investigated. The directional time series models will be assessed in terms of the distribution of extreme values of the time series, which is a feature that was not used in the fitting process. The assessment will be made by comparing the distribution of extremes from: the observed time series; a long simulation from the fittted directional model; and a long simulation from a LGN model fitted to the time series on the basis of the Akaike information criterion (AIC).

Keywords: *Directional time series, reversibility, financial time series, randomization test, threshold autoregressive models*

M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.

1 INTRODUCTION

Many financial time series show non-stationarity inasmuch as there is a trend, or there are seasonal effects. For example, sales of summer clothing are seasonal and may be increasing in line with an increasing population. In contrast there is no reason to suppose that there is any seasonality or long term trend in exchange rates. In any case, if a trend or seasonality are identified in a time series they can be estimated and removed to give a deseasonalized and detrended series that can be considered a realisation of a stationary random process.

In this investigation we assume that the time series is a realisation of a stationary random process. A stationary random process is reversible if it is not possible to distinguish between a realisation in its time order and the same realisation in reverse time order (time-to-go). If a stationary random process is not reversible it is directional. Linear causal models with independent Gaussian errors are reversible. Linear causal models with non-Gaussian errors are directional, as are non-linear models. For some time series it is possible to discern directionality by comparing a plot against time with a plot against time-to-go. For example, the monthly exchange rate of the British pound sterling against the U.S. dollar (GBPUSD XRT), shown in the upper frame of Figure 1, does show three marked plummets and relatively slow recoveries. The reverse time plot (time-to-go) is clearly different in these respects, in fact we see rapid increases, followed by slow decreases.

2 FINANCIAL DATA

Figure 1. Quarterly trading records of GBPUSD XRT from March 1990 to March 2015 [Data courtesy of Bloomberg]
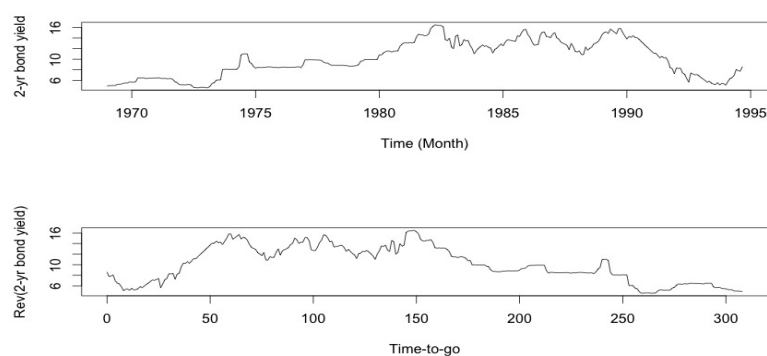
Figure 2. Interest rates (per annum) for Australian Government Bond Yield 2-yr securities, monthly from Jan 1969 to Sep 1994 [Data courtesy of Reserve Bank of Australia]

M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.
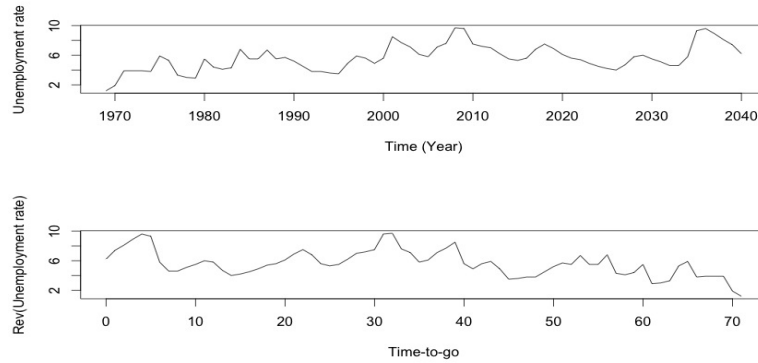
Figure 3. Yearly U.S. unemployment rate from 1969 to 2014 [Data courtesy of the U.S. Bureau of Labor Statistics]

Graphical inspection for the financial series show some degree of directionality in its time order (above) and in its reverse time order (below). The exchange rates (Figure 1) show a few well defined plummets. In contrast there is little qualitative difference between the bond interest rate (Figure 2) in time order and time-to-go order. In the case of the unemployment (Figure 3) increases generally appear to be steeper than the decreases.

3   TESTING FOR DIRECTIONALITY

There are many statistics that have been designed to detect directionality in time series (e.g. Lawrance, 1991, Beare and Seo, 2012 and Wild et al., 2014). The methods can be categorised into methods based on original series and methods based on first differences. A general, and intuitive, indicator is the skewness of the first order lag one differences. A time series modelled by random variables is given by

$$\{X_t\} \quad \text{for} \quad t = 1, 2, ..., n \tag{1}$$

and the first order lag one differences are

$$Y_t = X_t - X_{t-1} \quad \text{for} \quad t = 2, 3, ..., n. \tag{2}$$

The mean of the differences is $(x_n - x_1)/n$ which is approximately 0 for any long time series. If there is a tendency for rapid increases to be followed by more gradual recessions there will be rather more negative differences which will on average be less than the average of the positive differences. However there will be fewer positive differences. It follows that the distribution of differences will be positively skewed: there will be some outlying positive differences associated with the rapid increases; and the median will be negative and to the left of the mean which is near 0. In contrast, if there is a tendency for rapid decreases to be followed by relatively slow recoveries the distribution of differences will have negative skewness. We use the product moment measure of skewness, and refer to the skewness of the differences, γ_Y as the directionality.

The skewness in Y_t can be estimated by

$$\hat{\gamma}_Y = \frac{\sum_{t=1}^n (y_t - \bar{y})^3}{[\sum_{t=1}^n (y_t - \bar{y})^2]^{3/2}}. \tag{3}$$

There is evidence of directionality if the estimate $\hat{\gamma}_Y$ of γ_Y is statistically significantly different from 0.

M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.

Table 1. Statistics of time series

Statistic	GBPUSD XRT	2-yr Bond Yeild	Unemployment Rate
Time interval	quarterly	monthly	yearly
Length, n	101	309	72
Median, M	1.62	9.90	5.50
Mean, $\bar{x}$	1.65	10.06	5.63
Standard dev., s	0.1525	3.3542	1.7596
Lag 1 ACF	0.8490	0.9840	0.7650
Directionality, $\hat{\gamma}$	-0.7566	0.3039	1.0246
Two-sided P-value	(0.004)	(0.028)	(0.00)
Ratio	[3.18]	[2.18]	[3.89]

The directionality of the exchange rates (Figure 1), which does show clear rapid decreases, is -0.75. The directionality of the bond yields (Figure 2) and the unemployment rate (Figure 3) is 0.30 and 1.02 respectively. Inspection of time series plot and time-to-go plot of the unemployment rates (Figure 3) does show steeper ascents than descents in the former, and steeper descents than escents in the latter.

All three directionalities shown to be statistically different from 0 (two-sided P-values in Table 1) using a simulation procedure (Mansor et al., 2015). In this simulation procedure, we fit an AR(p) model to the time series of length n using the AIC criterion to select p, which is given in Table 2. The null hypothesis is that the time series is reversible. We simulate 1000 time series of length n from the fitted AR(p) model, using a Gaussian distribution with mean 0 and variance equal to the variance of the residuals for the errors. For each of the simulated time series, we calculate $\hat{\gamma}_Y$. We determine the proportion of the simulations for which the test statistic is absolutely more extreme than the calculated $\hat{\gamma}_Y$ from the original time series. We return this proportion from the simulation as the two-sided P-value. The ratio in Table 1 is the test statistic in absolute terms to its standard deviation from the simulation.

4   MODELLING DIRECTIONALITY

For all three time series the residuals from the AR(1) model have a much smaller standard deviation (σ_{est}), given in Table 2 (row 2, Order 1), than that of the original time series (s), given in Table 2 (row 1, Order 0 i.e. the original time series). In contrast, the further reductions in σ_{est} when AR(p) models, as selected by the AIC, are fitted, given in Table 2 (row 3, AIC) is slight.

Table 2. Estimated standard deviation of errors for AR models

Order	GBPUSD XRT	2-yr Bond Yeild	Unemployment Rate
0	$s = 0.1525$	$s = 3.3542$	$s = 1.7596$
1	AR(1), $\sigma_{est} = 0.0811$	AR(1), $\sigma_{est} = 0.6029$	AR(1), $\sigma_{est} = 1.1409$
AIC	AR(3), $\sigma_{est} = 0.0798$	AR(10), $\sigma_{est} = 0.5934$	AR(2), $\sigma_{est} = 1.1214$

These results suggest that the simple AR(1) model is a reasonable first approximation for all three time series. However, in all cases, the directionality of the realisation of the fitted AR(1) model, using random permutation of the residuals as errors, does not match the directionality calculated from the observed time series (Table 3). The objective is to improve this aspect of the modeling by introducing a non-linear TAR(1) model. The TAR(1) model is given by

$$X_t = \begin{cases} \alpha_U X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T_U \\ \alpha_L X_{t-1} + \epsilon_t & \text{if } X_{t-1} < T_L \end{cases} \tag{4}$$

We fit the TAR(1) model by finding the values of α_L and α_U that minimize the discrepancy,

$$\omega = \sum_{t=2}^{n} r_t^2 + \phi(\hat{\gamma}_Y - \tilde{\gamma}_Y)^2, \tag{5}$$

M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.

where $\{r_t\}$ are the residuals, $\tilde{\gamma}_Y$ is an estimate of the directionality of the TAR(1) model, and ϕ is the weight given to agreement between the $\tilde{\gamma}_Y$ and the target directionality value $\hat{\gamma}_Y$. The estimate of directionality of the TAR(1) model is calculated from a long simulation (length N = 10^5) of the TAR(1) model using random sampling, with replacement from the residuals. We use $\hat{\alpha}$ from the AR(1) model as the inital value for α_L and α_U when fitting the TAR(1) model. The constraints for stability are that $-1 < \alpha_L, \alpha_U < 1$. The optimization is carried out with the *optim()* function in *R* using the Nelder-Mead option (e.g. Nash, 2014). We consider an TAR(1) with 80% threshold for each series.

Table 3. Directionality measure modelled by AR(1) and TAR(1)_80% using resampled residuals

Financial Series	Target $\hat{\gamma}$	AR(1)	TAR(1)_80%	ϕ
Quarterly exchange rates of GBPUSD	-0.7566	$\tilde{\gamma} = -0.4031$ $\alpha = 0.8487$ $\sigma_{est} = 0.0811$	$\tilde{\gamma} = -0.7273$ $\alpha_L = 0.9981, \alpha_U = 0.9361$ $\sigma_{est} = 0.0805$	10^1
Monthly interest rates of 2-yr Government bond yield	0.3039	$\tilde{\gamma} = 0.4164$ $\alpha = 0.9838$ $\sigma_{est} = 0.6029$	$\tilde{\gamma} = 0.3040$ $\alpha_L = 0.7587, \alpha_U = 0.7310$ $\sigma_{est} = 0.9467$	10^8
Yearly United States unemployment rates	1.0246	$\tilde{\gamma} = 0.9932$ $\alpha = 0.7652$ $\sigma_{est} = 1.1409$	$\tilde{\gamma} = 1.0231$ $\alpha_L = 0.6170, \alpha_U = 0.8349$ $\sigma_{est} = 0.9959$	10^7

- In the case of the exchange rates, the TAR(1) gives a close agreement with the observed directionality, and a slightly lower estimate of the standard deviation of errors.
- In the case of the bond rates, which has a value of α close to 1 for the AR(1) model, the closer agreement of the TAR(1) model with the observed directionality is at the expense of an increase in estimated standard deviation of the errors.
- In the case of the U.S. unemployment rates, with the lowest value of α in the AR(1) model, both AR(1) and TAR(1) give a close fit to the observed directionality. The TAR(1) has a slightly lower estimated standard deviation of the errors.
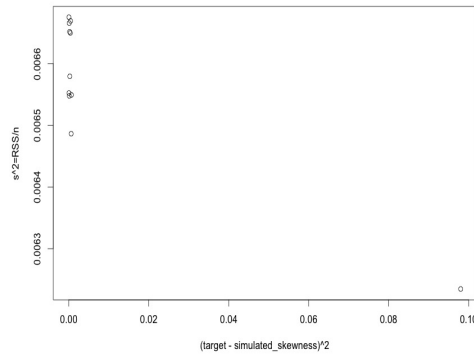
Figure 4. Trade-off between minimising the sum of squared residuals and minimising the discrepancy in directionality by increasing ϕ for the GBPUSD exchange rates time series

We illustrate in Figure 4 the balance achieved between the sum of squared residuals and the discrepancy between the target directionality and the simulated directionality of the quarterly GBPUSD exchange rates for N=10^5.
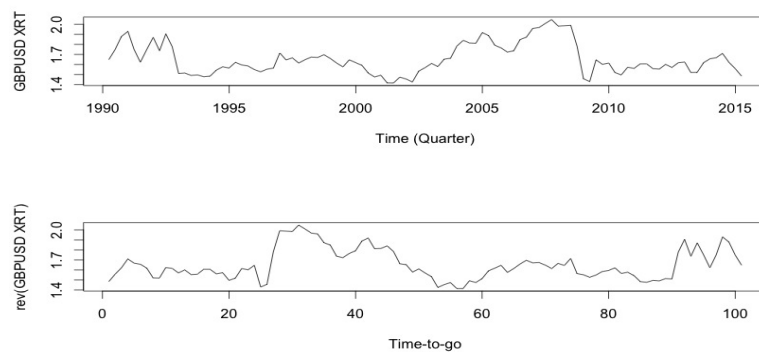
M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.

5   SIMULATION OF DIRECTIONAL TIME SERIES

We simulated 1000 time series of length 72 years, the length of the original time series for the U.S. unemployment rates using an AR(1) and a TAR (1; 80% threshold, 2 parameter) models. The upper and lower coefficients for TAR (1) are the optimised parameters of α_L=0.6197 and α_U=0.8359 (Table 3). For, both models two different error distributions were fitted to the residuals: an Extreme Value Type 1 (Gumbel) distribution of minima (GD); and back-to-back Weibull distribution, that is one Weibull distribution was fitted to the positive residuals and another to the absolute values of the negative residuals, (WD). The maxima, over 72 years, from the 1000 simulations, with GD and with WD, were stored. Their boxplots are shown in Figure 5, and the quantile-quantile (Q-Q) plots are shown in Figure 6.
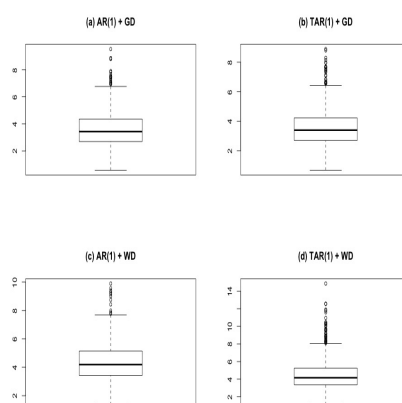
Figure 5. Boxplot of the simulated maxima in the U.S. unemployment rate for each model
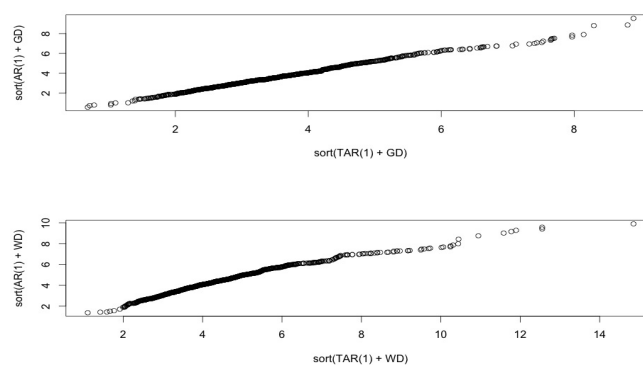
Figure 6. Q-Q plot of the simulated maxima in the U.S. unemployment rate, between TAR(1) versus AR(1)

There is no discernible difference in the boxplots for the distributions of maxima with the GD distribution (Figure 5 upper). In both cases the distribution of the maxima is positively skewed as is typical for extreme values. However, there is a noticeable difference between TAR(1) and AR(1) with WD errors, such that the maxima from the TAR(1) are more skewed and more extreme. The Q-Q plots (Figure 6 lower) show that the TAR(1) with WD has generally higher values than AR(1) with WD. Moreover, the Q-Q plots (Figure 6 upper) also show that TAR(1) has slightly higher values than AR(1) with GD.

M.M. Mansor, D.A. Green and A.V. Metcalfe, Modelling and simulation of directional financial time series.

6 CONCLUSIONS

We focus on directionality which we have defined in terms of skewness of the first differences of the time series. The financial time series of the quarterly GBPUSD exchange rates, the monthly Australian bond yield and the yearly U.S. unemployment rates are shown to be directional and the estimated values of the lag 1 autocorrelation are 0.8490, 0.9840 and 0.7650 respectively. All three series are well modelled as AR(1) if a least squares error criterion is adopted.

Many economic time series are well modelled as an AR(1). If high frequency sampling has the consequence that the parameter of the AR(1) model (α) is close to 1, then the estimate of directionality is very close to the skewnnes of the errors. We choose to model the time series of U.S. unemployment in more detail as it had the lowest value of α (0.76). In contrast, if an AR(1) model has α close to 0, then the skewness of errors has little effect on directionality.

We have demonstrated that in the intermediate range, in particular for GBPUSD exchange rates and U.S. unemployment rates TAR(1) model, fitted with discrepancy strategy, can improve the fit to observed directionality and decrease the estimate standard deviation of the errors (σ_{est}) (Table 3). We have also shown that the TAR(1) model combined with a back-to-back Weibull model (WD) for the residuals gives substantially higher extreme values than the AR(1) model combined with WD for the residuals, in the case of the U.S. unemployment series. A link between extreme values and directionality has been discussed by Soubeyrand et al. (2014) in the context of rainfall, and Mansor et al. (2015) in the context of sunspots.

REFERENCES

Beare, B. K. and J. Seo (2012). Time irreversible copula-based markov models. *Econometric Theory*, 1–38.

Lawrance, A. (1991). Directionality and reversibility in time series. *International Statistical Review/Revue Internationale de Statistique*, 67–79.

Mansor, M. M., M. E. Glonek, D. A. Green, and A. V. Metcalfe (2015). Modelling directionality in stationary geophysical time series. *Proceedings of the international work-conference on Time Series (ITISE 2015)*, 755–766.

Mansor, M. M., D. A. Green, and A. V. Metcalfe (2015). Detecting directionality in time series. *in preparation*.

Nash, J. (2014). On best practice optimization methods in r. *Journal of Statistical Software 60*(2), 1–14.

Soubeyrand, S., C. E. Morris, and E. K. Bigg (2014). Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data. *Environmental Modelling & Software 61*, 78–86.

Wild, P., J. Foster, and M. Hinich (2014). Testing for non-linear and time irreversible probabilistic structure in high frequency financial time series data. *Journal of the Royal Statistical Society: Series A (Statistics in Society) 177*(3), 643–659.

3.3 Modelling Directionality for Paleoclimatic Time Series

Statement of Authorship

Modelling Directionality for Paleoclimatic Time Series

Published in *ANZIAM J.* 57 (EMAC 2015), 2016.

Mahayaudin M. Mansor (Candidate)
Constructed statistical models for: deseasonalized Vostok and NGRIP records. Conducted analyses and interpreted results, wrote manuscript and acted as corresponding author.
Overall percentage: 60%

This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.

Signed: ⬚ Date: 14 June 2017

By signing the Statement of Authorship, each co-author certifies that: (i) the candidate's stated contribution to the publication is accurate; (ii) permission is granted for the candidate in include the publication in the thesis; and (iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

Farah L. Mohd. Isa
Obtained Vostok ice-core time series data, and interpolated time series to allow for uneven time intervals and mining values. Deseasonalized the time series. Reviewed the draft manuscript.

Signed: ⬚ Date: 12/06/2017

David A. Green
Supervised development of work and provided critical evaluation.

Signed: ⬚ Date: 14/6/17

Andrew V. Metcalfe
Supervised development of work and provided critical evaluation.

Signed: ⬚ Date: 14/6/17

Modelling directionality for paleoclimatic time series

M. M. Mansor[1]   F. L. Mohd. Isa[2]   D. A. Green[3]
A. V. Metcalfe[4]

(Received 2 January 2016; revised 26 May 2016)

Abstract

The ice core time series from Vostok Station in Antarctica and
the North Greenland Ice Core Project have seasonal variation corre-
sponding to the Milankovitch cycles. After removing these cycles, and
interpolating to equal time intervals, stationary time series models are
fitted. The series show clear directionality and this feature is mod-
elled by either non-Gaussian errors or non-linear time series models.
Threshold autoregressive models are fitted by penalized least squares
and compared with non-threshold autoregressive models. Since both
ice core time series are reasonably modelled as first order autoregres-
sive series with parameters close to one, directionality will arise from
non-symmetric error distributions. However, two regime threshold

autoregressive models, of order one and two for Greenland and Vostok, respectively, give an improved match to the observed directionality and a reduced sum of squared residuals. Realizations from the threshold autoregressive models are noticeably different from the non-threshold models. Since the non-threshold models are a restricted case of the threshold models, and the threshold models are a better fit to the observed time series, threshold models should provide more realistic realizations.

Contents

1 Introduction

Ice cores contain information about the history of Earth's climate. This information is derived from the ancient impurities trapped in the ice for thousands of years, including air bubbles, volcanic ash and soot. The proportion of dissolved oxygen isotope-18 to dissolved oxygen isotope-16 in ice,

and the proportion of deuterium (heavy water) in ice, are closely related to temperature. The local temperature is deduced from the proportion of oxygen-18 or from the relative amount of deuterium in the water molecules of the ice compared with seawater.

In this article we investigate paleoclimatic time series from the North Greenland Ice Core Project (NGRIP) which begins $122\,900$ years before present (BP) [7], and the Vostok ice core record which begins $422\,766$ years BP [8]. The NGRIP data includes $\delta^{18}O$ which is the ratio of oxygen isotope-18 (^{18}O) to oxygen isotope-16 (^{16}O) at 50 year intervals. The Vostok time series data is the deuterium (an isotope of hydrogen) content δD as a percentage of Standard Mean Ocean Sea Water (SMOW) at approximately 50 year intervals.

Both time series contain seasonal variations that corresponds to the Milankovitch cycles, long term variations in the Earth's orbit that have been affecting the Earth's climate change for aeons. We remove the effect of Milankovitch cycles in both series by fitting multiple regression to obtain the deseasonalized time series. The deseasonalized NGRIP (ds NGRIP series) is equally spaced at 50 year intervals. However, the deseasonalized Vostok (ds Vostok series) is unequally spaced due to missing values, particularly in the early record. As further analyses are based on evenly spaced time series, we applied linear interpolation at 50 year increments to obtain the deseasonalized and interpolated Vostok (dsi Vostok series).

2   Detecting directionality

A stationary time series model $\{X_t\}$ for time $t = 1, 2, \ldots, n$ is reversible if the joint distribution of $X_t, X_{t+1}, \ldots, X_{t+r}$ is equal to the joint distribution of $X_{t+r}, X_{t+r-1}, \ldots, X_t$ for all $r = 1, 2, \ldots$ [2]. A stationary time series model is directional if these joint distributions differ. If a stationary time series model is reversible, then it will not be possible to distinguish realisations plotted against time order from those realisations plotted in reverse time order, that

is, against time-to-go[1]. If a stationary time series model is directional, then, in principle, there will be qualitative differences between realisations plotted in time order and the same realisations plotted against time-to-go. However, the difference may not be discernible from a single time series, so a formal statistical test is required to detect and quantify the directionality.

Stationary linear time series models with Gaussian white noise are reversible, whereas non-linear time series models and linear models driven by non-Gaussian white noise are directional [3]. Directionality is clearly implicit in a trend and in asymmetric seasonal patterns, so directionality is only of interest in itself when considered in its own right for stationary time series models. If a time series appears to have a trend or seasonal effects, then these should be identified and removed before considering directionality.

Directionality is visible in many stationary time series from various disciplines including environmental science. We compare a plot against time with a plot against time-to-go for the ds NGRIP series (Figure 1) and for the dsi Vostok series (Figure 2). The ds NGRIP appears to be a realisation of a stationary time series model, and there is evidence against a null hypothesis of a unit root (Dickey–Fuller test, probability $P < 0.01$) to support this claim. The ds NGRIP series has a tendency for sharp increases before the peaks followed by gradual decreases to the troughs when plotted against time, and gradual increases before the peaks followed by sharp decreases to the troughs when plotted against time-to-go.

The dsi Vostok series also appears to be a realisation of a stationary time series model, and the Dickey–Fuller test again provides evidence against a hypothesis of a unit root ($P < 0.01$). There is a slight tendency for rapid increases to be followed by slower decreases, but this is more apparent in the first half of the series which is considerably smoother because many values are interpolated. To avoid the possibility that our findings are heavily influenced by interpolation, we focus on the second half of the dsi Vostok series (dsi 2H Vostok series) for further analyses.

---

[1] If time t runs from 1 to n, then the time-to-go is $n - t$, and runs from $n - 1$ to 0.

Figure 1: The ds NGRIP series (top) against time with scale in years BP; and (bottom) against time-to-go in 50 year units.
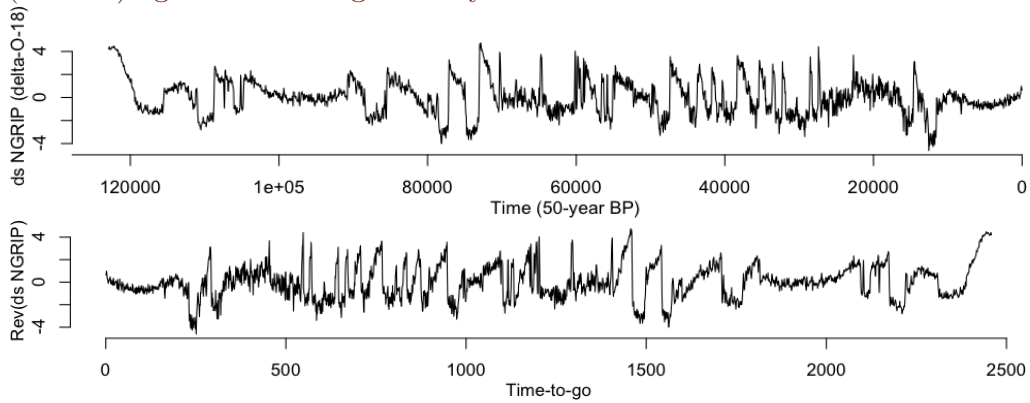
Figure 2: The dsi Vostok series (top) against time with scale in years BP; and (bottom) against time-to-go in 50 year units.
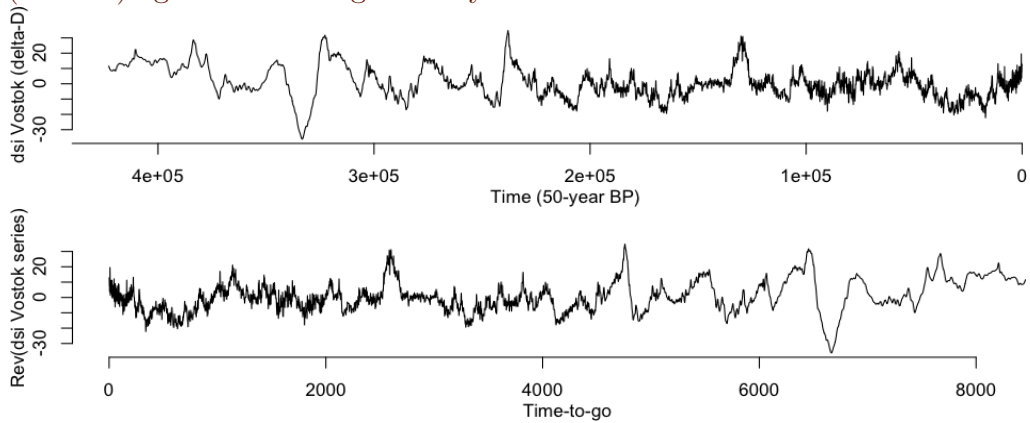
2 *Detecting directionality*   C71

Table 1: Statistics for ds NGRIP and dsi 2H Vostok series.

Statistic	NGRIP(δ^{18}O)	Vostok(δD)
Time interval (years)	50	50
No. obs., n	2459	4229
Mean, $\bar{x}$	0.00	-2.00
Standard dev., s	1.51	7.96
Directionality, $\hat{\gamma}_{\text{dif}}$	0.61	0.18

The next step is to quantify the apparent directionality in the data and determine whether or not it can plausibly be attributed to chance. A general indicator of directionality is the product moment skewness of first differences [3] estimated by

$$\hat{\gamma}_{\text{dif}} = \frac{\sum_{t=1}^{n}(y_t - \bar{y})^3/(n-1)}{\left[\sum_{t=1}^{n}(y_t - \bar{y})^2/(n-1)\right]^{3/2}} \, , \tag{1}$$

where $\{y_t\} = \{x_t\} - \{x_{t-1}\}$ for $t = 2, 3, \ldots, n$, $\bar{y}$ is the mean of $\{y_t\}$, and $\{x_t\}$ is the observed time series. For example, a distribution of first differences y_t of the NGRIP series (that has rapid increases followed by more gradual recessions) tends to have more small negative differences than positive differences. Although the positive differences are fewer, they include outlying positive differences corresponding to the rapid increases. It follows that y_t is positively skewed with a longer tail on the right side of the distribution (negative median and mean value to the right of the median). Non-zero skewness indicates directionality or asymmetry in time of a time series. We refer to $\hat{\gamma}_{\text{dif}}$ as directionality. Directionalities in the ds NGRIP and dsi 2H Vostok series are shown in Table 1, together with the lengths, means and standard deviations of these series. The positive directionalities correspond to the relatively rapid increases and slow recessions seen in Figures 1 and 2.

Next, we determine the significance level of the indicator of directionality using a randomization test based on autoregressive models AR(p), where the p-order is determined by the minimum Akaike information criterion (AIC), with

Gaussian white noise for the errors. The probability of a directionality with absolute magnitude greater than 0.61 in Gaussian white noise in realisations of length 2459 (the number of 50 year intervals in the ds NGRIP data) is less than 0.001, as determined from AR(5) where five is selected by AIC. So there is strong evidence of directionality in the ds NGRIP series ($P < 0.001$). Similarly, there is strong evidence of directionality in the dsi 2H Vostok series ($P < 0.001$, based on AR(36)).

Directionality in time series has several implications. Directionality provides evidence of complex feedbacks after shocks or occasional extreme events [9, 4, e.g.]. In the context of paleoclimatic series, extreme events include earth quakes, volcanoes, and meteorite strikes. Moreover, directionality may indicate that non-linear time series models are appropriate and these should provide more accurate forecasts and more realistic ensembles of scenarios [3, e.g.].

3   Modelling directionality

Reproduction of directionality in time series is done by introducing non-Gaussian errors or by using a non-linear model [2, 3] with either Gaussian or non-Gaussian error distributions. Non-Gaussian errors include: asymmetric probability distributions, such as Gumbel, three-parameter Weibull, beta and back-to-back Weibull; and symmetric probability distributions such as Student's–t. However, the Student's–t distribution needs to have high kurtosis (low degree of freedom) to inculcate noticeable directionality in the linear model [4].

We fit the observed ds NGRIP and dsi 2H Vostok series to AR models,

$$(X_t - \mu) = \sum_{i=1}^{p} \alpha_i (X_{t-i} - \mu) + \epsilon_t \,, \tag{2}$$

where the mean μ and the coefficients $\alpha_1, \ldots, \alpha_p$ are the parameters to be estimated, and ϵ_t is a sequence of independent zero mean random errors. To

Table 2: Estimated standard deviation of residuals $\hat{\sigma}_{\mathrm{err}}$ for AR models.

Order	NGRIP (δ^{18}O)		Vostok (δD)	
0		$s = 1.51$		$s = 7.96$
1	AR(1)	$\hat{\sigma}_{\mathrm{err}} = 0.6203$	AR(1)	$\hat{\sigma}_{\mathrm{err}} = 1.6134$
2	AR(2)	$\hat{\sigma}_{\mathrm{err}} = 0.6200$	AR(2)	$\hat{\sigma}_{\mathrm{err}} = 1.4907$
AIC	AR(5)	$\hat{\sigma}_{\mathrm{err}} = 0.6173$	AR(36)	$\hat{\sigma}_{\mathrm{err}} = 1.3398$

fit the models we use ordinary least squares, and we assess the goodness of fit by comparing the estimated standard deviation of the residuals $\hat{\sigma}_{\mathrm{err}}$ with the marginal standard deviation s of the observed time series.

Table 2 presents the fitting results and suggests that the simple AR(1) model is a reasonable first approximation for the ds NGRIP series, and similarly the AR(2) model for the dsi 2H Vostok series. In both cases, the standard deviation of the residuals $\hat{\sigma}_{\mathrm{err}}$ are considerably lower than the marginal standard deviations of the ice core time series s, and the further decreases in the residuals for AIC model are negligible.

3.1 Modelling directionality for NGRIP series

We model directionality for the ds NGRIP series using the AR(1) model and threshold autoregressive model (TAR) of order one, TAR(1), fitted by (non-linear) least squares, with Gaussian errors (GE) and resampled residuals (RE), randomly with replacement,. The TAR(1) model is

$$(X_t - \mu) = \begin{cases} \alpha_U(X_{t-1} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) > T, \\ \alpha_L(X_{t-1} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) < T, \end{cases} \tag{3}$$

where the mean μ and the coefficients α_U and α_L are parameters to be estimated, ϵ_t is a sequence of independent zero mean random errors, and T is the threshold which is calculated here as the upper 0.80 quantile of the marginal distribution of the ds NGRIP time series.

Table 3: Comparison of directionality $\hat{\gamma}_{\text{sim}}$, mean $\bar{x}_{\text{sim}}$ and standard deviation $\bar{x}_{\text{sim}}$ from realisations (length 10^5) of $\text{AR}(1)$ and $\text{TAR}(1)$ fitted to the ds NGRIP series with target values $\hat{\gamma} = 0.61$, $\bar{x} = 0.00$ and $s = 1.51$.

	$\text{AR}(1) + \text{GE}$	$\text{TAR}(1) + \text{GE}$	$\text{AR}(1) + \text{RE}$	$\text{TAR}(1) + \text{RE}$
$\hat{\gamma}_{\text{sim}}$	-0.004	-0.002	0.630	0.632
$\bar{x}_{\text{sim}}$	-0.016	0.055	-0.005	-0.008
s_{sim}	1.496	1.479	1.508	1.507

The directionality arises from the distribution of the errors rather than the non-linearity of the $\text{TAR}(1)$ model (Table 3), which has almost identical parameters above and below the 0.80 quantile threshold (Table 6, second row).

3.2   Modelling directionality for Vostok series

In the case of the dsi 2H Vostok series, we model directionality using $\text{AR}(2)$ and $\text{TAR}(2)$ models. The $\text{TAR}(2)$ model is

$$(X_t - \mu) = \begin{cases} \alpha_{1U}(X_{t-1} - \mu) + \alpha_{2U}(X_{t-2} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) > T, \\ \alpha_{1L}(X_{t-1} - \mu) + \alpha_{2L}(X_{t-2} - \mu) + \epsilon_t & \text{if } (X_{t-1} - \mu) < T, \end{cases}$$
$$(4)$$

where the mean μ and the coefficients α_{1U}, α_{2U}, α_{1L} and α_{2L} are parameters to be estimated, ϵ_t is a sequence of independent zero mean random errors, and T is the threshold which is calculated as the upper 0.80 quantile of the marginal distribution of the dsi 2H Vostok series.

Results in Table 4 show both the errors and non-linearity of the $\text{TAR}(2)$ induce directionality, but the $\text{TAR}(2)$ with RE has more directionality then the original series.

Table 4: Comparison of directionality $\hat{\gamma}_{sim}$, mean $\bar{x}_{sim}$ and standard deviation s_{sim} of simulation of length 10^5 from AR(2) and TAR(2) fitted to the dsi 2H Vostok series with target values $\hat{\gamma} = 0.18$, $\bar{x} = -2.00$ and $s = 7.96$.

	AR(2) + GE	TAR(2) + GE	AR(2) + RE	TAR(2) + RE
$\hat{\gamma}_{sim}$	−0.008	0.021	0.123	0.318
$\bar{x}_{sim}$	−2.117	−2.725	−1.754	−1.630
s_{sim}	7.875	8.141	7.944	8.295

3.3   TAR model with penalized least squares

In order to improve the agreement between the simulated directionality and the directionality observed in the original time series (target directionality), we estimate parameters of the TAR(1) model given in equation (3) for the ds NGRIP series and parameters of the TAR(2) model given in equation (4) for the dsi 2H Vostok series using penalized least squares (PLS). The objective is

$$\omega = \sum_{t=p+1}^{n} r_t^2 + \phi(\hat{\gamma}_{observed} - \hat{\gamma}_{simulated})^2, \tag{5}$$

where $\{r_t\}$ are residuals from the fitted models and ϕ is the weight given to minimise discrepancy between the target directionality and the simulated directionality. Constraints on the stability for TAR(1)[PLS] are $-1 < \alpha_L$ and $\alpha_U < 1$; and for TAR(2)[PLS] are $\alpha_2 > -1$, $\alpha_1 + \alpha_2 < 1$ and $\alpha_1 - \alpha_2 > -1$ [1]. Errors for the TAR[PLS] models are resampled residuals from the AR model (RE1) and resampled residuals from the TAR model (RE2).

The TAR(1) is fitted using PLS with errors randomly sampled with replacement from the residuals of a previously fitted AR(1) model. For the ds NGRIP series, this TAR(1)[PLS] gives a very close agreement to the target directionality, but at the expense of a marginal standard deviation that is less than that observed (Table 5).

Table 6 shows that the AR(1) model (with resampled residuals that preserves the marginal standard deviation and is close to the target directionality) has

Table 5: Results from realisations (length 10^5) of TAR[PLS] models compared to target values.

ds NGRIP		TAR(1)[PLS, RE1]	TAR(1)[PLS, RE2]
$\hat{\gamma} = 0.61$	$\hat{\gamma}_{sim}$	0.609	0.579
$\bar{x} = 0.00$	$\bar{x}_{sim}$	-0.013	-0.008
$s = 1.51$	s_{sim}	1.317	1.339
	ϕ	10^6	10^6
dsi 2H Vostok		TAR(2)[PLS, RE1]	TAR(2)[PLS, RE2]
$\hat{\gamma} = 0.18$	$\hat{\gamma}_{sim}$	0.177	0.115
$\bar{x} = -2.00$	$\bar{x}_{sim}$	-2.113	-2.084
$s = 7.96$	s_{sim}	7.669	7.793
	ϕ	10^7	10^7

Table 6: Fitting detailed: estimated coefficients and $\hat{\sigma}_{err}$ for each model.

Model	Estimated parameters for ds NGRIP	$\hat{\sigma}_{err}$
AR(1)	$\hat{\alpha} = 0.91$	0.620
TAR(1)	$\hat{\alpha}_U = 0.92,\ \hat{\alpha}_L = 0.90$	0.614
TAR(1)[PLS, RE1]	$\hat{\alpha}_U = 0.86,\ \hat{\alpha}_L = 0.89$	0.617
TAR(1)[PLS, RE2]	$\hat{\alpha}_U = 0.88,\ \hat{\alpha}_L = 0.90$	0.615
Model	Estimated parameters for dsi 2H Vostok	$\hat{\sigma}_{err}$
AR(2)	$\hat{\alpha}_1 = 1.4,\ \hat{\alpha}_2 = -0.38$	1.491
TAR(2)	$\hat{\alpha}_{1U} = 1.2,\ \hat{\alpha}_{2U} = -0.18,\ \hat{\alpha}_{1L} = 1.5,\ \hat{\alpha}_{2L} = -0.50$	1.444
TAR(2)[PLS, RE1]	$\hat{\alpha}_{1U} = 1.6,\ \hat{\alpha}_{2U} = -0.71,\ \hat{\alpha}_{1L} = 1.9,\ \hat{\alpha}_{2L} = -0.91$	1.641
TAR(2)[PLS, RE2]	$\hat{\alpha}_{1U} = 1.3,\ \hat{\alpha}_{2U} = -0.35,\ \hat{\alpha}_{1L} = 1.4,\ \hat{\alpha}_{2L} = -0.37$	1.467

errors with standard deviation 0.620 that is only slightly higher than that of the TAR(1)[PLS,RE2] (0.615). The time series is long and a Monte Carlo simulation test indicates that TAR(1)[PLS,RE2] is a statistical improvement in terms of the standard deviation of errors ($P < 0.001$). However, the simulated marginal standard deviation of 1.34 is substantially lower than the marginal standard deviation of 1.51 in the observed time series. These results highlight the limitations of relying on any single measure of goodness of fit when comparing linear and non-linear models. For NGRIP the AR(1) model, or almost equivalently the TAR(1) model, fitted without penalty, seem the best models of those considered.

For the dsi 2H Vostok series the TAR(2) model fitted by PLS offers a potential improvement on the AR(2) model, in some respects at least. The choice of error distribution affects the fit through the penalty term, because the errors determine the simulated directionality. If resampled errors after fitting the AR(2) model are used, then the directionality is matched (to two decimal places) and the marginal standard deviation of 7.67 is reasonably close to the observed 7.96. For the errors, the estimated standard deviation of 1.641 is rather higher than that of the AR(2) model which is 1.491. If resampled errors after fitting the TAR(2) model are used, then the estimated standard deviation of errors is reduced to 1.467 but the directionality is matched slightly better by the AR(2) model. There is no clear best model amongst the three.

4   Climate change simulation

Although the choice of best fitting model may be equivocal, there is a difference in terms of simulating up to ten steps ahead. This is demonstrated in Table 7 for $1\,000$ simulations up to ten steps ahead using the AR(1) and TAR(1)[PLS,RE2] models.

For the simulation: a back-to-back Weibull distribution was fitted to the residuals of the AR(1) model; back-to-back Weibull distributions were fitted

Table 7: Upper 1%, median and lower 1% of n-step ahead predictions using AR(1) and TAR(1)[PLS,RE2] models for the ds NGRIP series.

Step	Upper 1% AR	Upper 1% TAR	Median AR	Median TAR	Lower 1% AR	Lower 1% TAR
1	2.68	3.31	−0.03	−0.12	−2.83	−3.09
2	3.05	3.41	−0.10	−0.14	−2.78	−3.03
3	4.03	4.01	−0.10	−0.11	−3.37	−3.83
4	5.12	4.55	−0.10	−0.07	−4.83	−4.73
5	6.09	5.13	−0.10	−0.03	−5.92	−5.73
6	6.97	5.60	−0.13	0.01	−6.90	−6.73
7	7.47	6.09	−0.12	0.01	−8.03	−7.38
8	8.09	6.73	−0.11	0.03	−9.13	−8.13
9	8.85	7.05	−0.10	0.05	−10.2	−8.79
10	9.60	7.45	−0.11	0.06	−10.9	−9.39

to the TAR(1) model residuals when x_t was below the threshold, and to the TAR(1) residuals when x_t was above the threshold. The TAR(1) model prediction intervals were wider than the AR(1) up to three steps ahead but narrower for further steps.

5   Conclusion

There is clear directionality in the ds NGRIP series and statistically significant directionality in the dsi 2H Vostok series. The ds NGRIP series is approximated as a realisation of an AR(1) time series model with $\alpha = 0.91$, which has first differences which are close to the errors. It follows that the directionality as measured by skewness of the differences is largely determined by the skewness of the errors. The dsi 2H Vostok series is roughly approximated as a realisation of an AR(1) time series model with $\alpha = 0.98$ and it again follows that the directionality mainly follows the skewness of the errors. The

ds NGRIP series at 50 year time steps is quite well modelled as a realisation of an AR(1) process with non-Gaussian errors. The non-Gaussian errors allow for some catastrophic events and the model, at least, seems stable.

Such a simple model is not satisfactory for the Vostok series. The residuals after fitting an AR(2) model still show some degree of autocorrelation. An AR(36) model is needed to obtain residuals that appear uncorrelated. The partial improvements offered by TAR models suggest that there may be more substantial non-linear effects in the Vostok series. A high order AR(p) model with thresholds for coefficients up to some smaller lag might provide a substantial improvement on the models considered here. However, such a complex empirical model may not provide much insight into the underlying physical processes. We also found that TAR models with penalized least squares are reproducible models for modelling directionality in stationary time series [3, 5, 6, e.g.].

References

[1] C. Chatfield. *The Analysis of Time Series: An Introduction.* CRC Press, 2004. https://www.crcpress.com/ The-Analysis-of-Time-Series-An-Introduction-Sixth-Edition/ Chatfield/p/book/9781584883173 C75

[2] A. J. Lawrance. Directionality and reversibility in time series. *Int. Stat. Rev.* 59(1):67–79, 1991. doi:10.2307/1403575 C68, C72

[3] M. M. Mansor, M. E. Glonek, D. A. Green and A. V. Metcalfe. Threshold autoregressive models for directional time series. In I. Rojas and H. Pomares (Eds.), *Time Series Analysis and Forecasting Selected Contributions from the ITISE Conference (ITISE 2015)*. pp. 13–25, 2016. doi:10.1007/978-3-319-28725-6 C69, C71, C72, C79

[4] M. M. Mansor, M. E. Glonek, D. A. Green and A. V. Metcalfe. Modelling directionality in stationary geophysical time series. *International work-conference on Time Series (ITISE 2015)*. http://www.researchgate.net/publication/281835075 C72

[5] M. M. Mansor, D. A. Green and A. V. Metcalfe. Modelling and simulation of directional financial time series. *Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)*, pp. 1022–1028, 2015. http://www.mssanz.org.au/modsim2015/E4/mansor.pdf C79

[6] M. M. Mansor, D. A. Green and A. V. Metcalfe. Directionality and volatility in electroencephalogram time series. *Proceedings of the 2nd International Conference on Mathematical Sciences and Statistics (ICMSS 2016), AIP Conf. Proc.* 1739:020080, 2016. doi:10.1063/1.4952560 C79

[7] North Greenland Ice Core Project members. High-resolution record of Northern Hemisphere climate extending into the last interglacial period. *Nature*, 431:147–151, 2004. doi:10.1038/nature02805 C68, C79

[8] J. R. Petit, J. Jouzel, D. Raynaud, N. I. Barkov, J.-M. Barnola, I. Basile, M. Bender, J. Chappellaz, M. Davis, G. Delaygue, M. Delmotte, V. M. Kotlyakov, M. Legrand, V. Y. Lipenkov, C. Lorius, L. Pepin, C. Ritz, E. Saltzman and M. Stievenard. Climate and atmospheric history of the past 420,000 years from the Vostok ice core, Antarctica. *Nature*, 399:429–436, 1999. doi:10.1038/20859 C68, C79

[9] S. Soubeyrand, C. E. Morris and E. K. Bigg. Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data.

Environ. Modell. Softw. 61:78–86, 2014. doi:10.1016/j.envsoft.2014.07.003
C72

Author addresses

1. **M. M. Mansor**, School of Mathematical Sciences, University of
   Adelaide, South Australia 5005, Australia.
   `mailto:mohdmahayaudin.mansor@adelaide.edu.au`

2. **F. L. Mohd. Isa**, School of Mathematical Sciences, University of
   Adelaide, South Australia 5005, Australia.

3. **D. A. Green**, School of Mathematical Sciences, University of
   Adelaide, South Australia 5005, Australia.

4. **A. V. Metcalfe**, School of Mathematical Sciences, University of
   Adelaide, South Australia 5005, Australia.

Chapter 4

Developments

4.1 Threshold Autoregressive Models for Directional Time Series

Threshold Autoregressive Models for Directional Time Series

Mahayaudin M. Mansor, Max E. Glonek, David A. Green, and Andrew V. Metcalfe

Abstract Many time series show directionality as plots against time and against time-to-go are qualitatively different. A stationary linear model with Gaussian noise is non-directional (reversible). Directionality can be emulated by introducing non-Gaussian errors or by using a nonlinear model. Established measures of directionality are reviewed and modified for time series that are symmetrical about the time axis. The sunspot time series is shown to be directional with relatively sharp increases. A threshold autoregressive model of order 2, TAR(2) is fitted to the sunspot series by (nonlinear) least squares and is shown to give an improved fit on autoregressive models. However, this model does not model closely the directionality, so a penalized least squares procedure was implemented. The penalty function included a squared difference of the discrepancy between observed and simulated directionality. The TAR(2) fitted by penalized least squares gave improved out-of-sample forecasts and more realistic simulations of extreme values.

Keywords Directional time series • Penalized least squares • Reversibility • Sunspot numbers • Threshold autoregressive models

1   Introduction

Directionality, defined as asymmetry in time [3], enables us to tell the difference between a sequence of observations plotted in time order (time series) and the sequence plotted in reverse time order (time-to-go). A clear example of directionality can be seen in the average yearly sunspot numbers 1700–2014 (Fig. 1). A time series model is reversible if, and only if, its serial properties are symmetric, with respect to time and time-to-go. A linear time series model with Gaussian errors (LGE model) is reversible, but a linear time series model with non-Gaussian errors is directional (LNGE model) [3]. Nonlinear time series models are also directional, whether or not the errors are Gaussian (NLGE and NLNGE models, respectively).

M.M. Mansor (✉) • M.E. Glonek • D.A. Green • A.V. Metcalfe
School of Mathematical Sciences, University of Adelaide, Adelaide, SA 5005, Australia
e-mail: mohdmahayaudin.mansor@adelaide.edu.au; max.glonek@adelaide.edu.au;
david.green@adelaide.edu.au; andrew.metcalfe@adelaide.edu.au
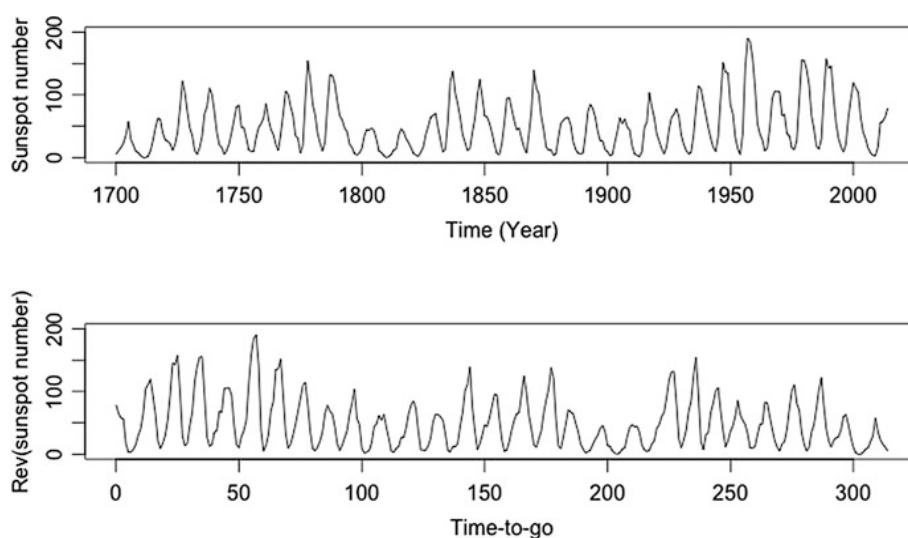
Fig. 1 Graphical inspection of directionality shows the sunspot observations rise more quickly than they fall in time order (*above*) and rise more slowly than they fall in reverse time order (*below*)

Directionality is important for forecasting because it indicates that models other than LGE should be considered [2, 3]. If the error distribution is better modeled as non-Gaussian, this will lead to more precise limits of prediction for forecasts. If nonlinear models provide a better fit than linear models, the forecast from nonlinear models will be more accurate. Modeling directionality also leads to more realistic simulation of extreme values.

The paper is arranged as follows. Section 2 describes well-established procedures for detecting directionality in time series, together with a modification for directional time series that are symmetrical about the time axis. In Sect. 3 we consider modeling directionality using TAR(2) models fitted by penalized least squares. In Sect. 4 we provide evidence of directionality in the sunspot series, discuss a model and simulation results. The model shows improved predictions and a more realistic distribution of maxima from a long simulation. A conclusion is given in Sect. 5.

2 Detecting Directionality

In general any trend or seasonality should be removed before investigating directionality in the stationary series. There are many possible directional features in time series, and these include, for example, sharp increases followed by slow recessions or slow increases followed by sharp decreases. In such cases the series is asymmetric with respect to time and also with respect to its mean or its median. In contrast a time series may exhibit both sharp increases and sharp decreases, followed by more

gradual returns to the median value. Such time series are asymmetric with respect to time but symmetric with respect to the median. Different statistics are appropriate for detecting directionality in series that are asymmetric or symmetric with respect to the median.

In this paper, we employ relatively simple and well-established tests [3] to detect directionality in time series: difference in linear quadratic lagged correlations; proportion of positive differences; skewness of differences; and tests based on comparisons of time from threshold to peak against time from peak to threshold. More recent tests are based on properties of Markov chains [1]; spectral estimation of kurtosis of differences in financial time series [9]; and the FeedbackTS package in R to detect time directionality occurring in specific fragments of time series [8].

2.1   Difference in Linear Quadratic Lagged Correlations

Directionality has the complementary concept of reversibility. Demonstrating evidence that a series is not reversible is another way of expressing that the series is directional. Following [3], a time series modeled by random variables $\{X_t\}$ for $t = 0, \pm 1, \pm 2, \ldots$ is reversible if the joint distribution of $X_t, X_{t+1}, \ldots, X_{t+r}$ and $X_{t+r}, X_{t+r-1}, \ldots, X_t$ is the same for all $r = 1, 2, \ldots$. In particular, a time series that is reversible has

$$\text{Corr}(X_t, X_{t+1}^2) = \text{Corr}(X_t^2, X_{t+1}). \tag{1}$$

A measure of directionality can be based on the difference in the sample estimates of these correlations. The non-dimensional measure used here is

$$\text{DLQC} = \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})(x_{t+1} - \bar{x})^2}{\left[\sum_{t=1}^{n}(x_t - \bar{x})^2\right]^{3/2}} - \frac{\sum_{t=1}^{n-1}(x_t - \bar{x})^2(x_{t+1} - \bar{x})}{\left[\sum_{t=1}^{n}(x_t - \bar{x})^2\right]^{3/2}}. \tag{2}$$

The rationale behind this statistic is as follows. Consider for example, there are sharp increases followed by slow recessions. Suppose the sharp increase occurs between x_t and x_{t+1}, then $(x_t - \bar{x})$ could be negative or positive but $(x_{t+1} - \bar{x})$ is very likely to be positive. It follows that $(x_t - \bar{x})(x_{t+1} - \bar{x})^2$ is negative or positive whereas $(x_t - \bar{x})^2(x_{t+1} - \bar{x})$ is positive and hence DLQC will tend to be negative. Both terms in DLQC are correlations, so bounds for the DLQC are $[-2, 2]$, but typical values in directional time series are smaller by two orders of magnitude. For the sunspot series, which exhibits clear directionality, with relatively sharp increases, DLQC is -0.06.

2.2 Methods Based on First Differences

More intuitive measures of directionality can be based on the distribution of lag one differences. For example, if there are sharp increases and slow recessions, there will be fewer large positive differences and more small negative differences. The distribution will be positively skewed. Let the observed time series be X_t and define the lag one first order differences[1] as

$$Y_t = X_t - X_{t-1} \quad \text{for} \quad t = 2, 3, \ldots, n. \tag{3}$$

2.2.1 Percentage of Positive Differences

The percentage of positive differences is

$$P^+ = \frac{\text{number of positive } Y_t}{\text{number of positive } Y_t + \text{ number of negative } Y_t} \times 100. \tag{4}$$

This formula excludes possible zero differences. If the time series is symmetric about the median and there are both sharp increases and sharp decreases, then there will tend to be both large positive and large negative differences, which tend to cancel out. Therefore, we adjust the measure to

$$P_{\text{abm}} = \frac{(P_{\text{above}}^+) + (P_{\text{below}}^-)}{(P_{\text{above}}^+) + (P_{\text{above}}^-) + (P_{\text{below}}^+) + (P_{\text{below}}^-)} \times 100, \tag{5}$$

where a difference is classified as above or below according to whether X_t is above or below the median. Also, P_{above}^+, and P_{below}^-, are the proportions of differences above, or below the median that are positive, or negative.

If a series is reversible, the expected values of the percentages in Eqs. (4) and (5) are 50 %. If differences are treated as independent, in a time series of length 1000, P_{abm} would need to be differ from 50 % by at least 3.2 % to be statistically significant at the 0.05 level.

2.2.2 Product Moment Skewness of Differences

A potentially more sensitive test for directionality is to consider the skewness of the distribution of differences [3] given by

$$\hat{\gamma} = \frac{\sum_{t=1}^{n}(y_t - \bar{y})^3/(n-1)}{\left[\sum_{t=1}^{n}(y_t - \bar{y})^2/(n-1)\right]^{3/2}}. \tag{6}$$

---

[1]In the following, "differences" refers to these lag one differences.

If a time series has both sharp increases and sharp decreases and is symmetric about the median we adapt the definition to be

$$\hat{\gamma}_{\text{abm}} = |\hat{\gamma}_{\text{above}}| + |\hat{\gamma}_{\text{below}}| \,. \tag{7}$$

Significant nonzero skewness of either $\hat{\gamma}$ or $\hat{\gamma}_{\text{abm}}$ is evidence of directionality.

2.3   *Threshold-Peak or Threshold-Trough Test*

Consider a threshold (H) set, in this investigation, at the upper quintile of the marginal distribution of the time series x_t. Suppose that $x_{j-1} < H$, $x_j > H$, and that x_t remains above H until $x_{j+k+1} < H$. Denote the time when x_t is greatest (peak value) for $j \leq t \leq (j + k)$ as $(j + p)$. Define the difference between time from threshold to peak and time from the peak to the threshold as

$$\text{DHPPH}_j = (k - p) - p. \tag{8}$$

A similar definition can be constructed for a threshold-trough test, using least values (troughs) of series of observations below the lower quintile (L). Denote the difference between time from threshold to trough and time from trough to threshold as DLTTL_j. Calculate DHPPH and DLTTL as the average of DHPPH_j and DLTTL_j respectively for all exceedances of H and excursions below L. The expected value of DHPPH and DLTTL is 0 for a reversible series.

2.4   *Evidence of Directionality*

In general, a directional time series will not show directionality on all of these statistics. Any one test being statistically significant at the α level, where $\alpha < 0.05$ say, is evidence of directionality. A conservative allowance for multiple testing would be to use a Bonferroni inequality and claim overall significance at less than an $m\alpha$ level where m is the number of tests.

3   Modeling Directionality

Mansor et al. [5] considered first order autoregressive processes AR(1) of the form

$$X_t = \alpha X_{t-1} + \epsilon_t \,, \tag{9}$$

and showed that the choice of non-Gaussian error distributions could lead to a variety of significant directional features. Furthermore, they demonstrated that

realizations from first order threshold autoregressive models TAR(1) with two thresholds (T_L, T_U) of the form

$$X_t = \begin{cases} \alpha_U X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T_U \\ \alpha_M X_{t-1} + \epsilon_t & \text{if } T_U < X_{t-1} < T_L \\ \alpha_L X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T_L \end{cases} \tag{10}$$

show substantial directionality even with Gaussian errors. They also found that the product moment skewness of differences ($\hat{\gamma}$) was generally the most effective statistic for detecting directionality. They subsequently fitted a second order threshold autoregressive model TAR(2) with one threshold (T) of the form

$$X_t = \begin{cases} \alpha_{1U} X_{t-1} + \alpha_{2U} X_{t-2} + \epsilon_t & \text{if } X_{t-1} > T \\ \alpha_{1L} X_{t-1} + \alpha_{2L} X_{t-2} + \epsilon_t & \text{if } X_{t-1} < T \end{cases} \tag{11}$$

to the first 200 values in the sunspot series, by nonlinear least squares. The TAR(2) gave some improvement over an AR(2) model for one-step ahead predictions of the remaining 115 values in the sunspot series that were not used in the fitting procedure. However they noted that there was scope for more realistic modeling of directionality.

Here we consider the strategy of using a penalized least squares procedure for the fitting of the TAR(2) model to the sunspot series. Initially, the objective function to be minimized was

$$\omega = \sum_{t=3}^{n} r_t^2 + \phi (\hat{\gamma}_{\text{simulated}} - \hat{\gamma}_{\text{observed}})^2, \tag{12}$$

where $\{r_t\}$ for $t = 3, \ldots, n$ are the residuals defined by

$$r_t = \begin{cases} x_t - \hat{\alpha}_{1U}(x_{t-1}) - \hat{\alpha}_{2U}(x_{t-2}) & \text{if } (x_{t-1}) > T \\ x_t - \hat{\alpha}_{1L}(x_{t-1}) - \hat{\alpha}_{2L}(x_{t-2}) & \text{if } (x_{t-1}) < T, \end{cases} \tag{13}$$

where $\{x_t\}$ for $t = 1, \ldots, n$ is the mean adjusted (to 0) time series, and $\hat{\gamma}_{\text{observed}}$ is the directionality calculated for the 200 year sunspot series.

For any candidate set of parameter values, the optimization routine has to determine not only the sum of squared errors, but also the directionality (γ). The directionality γ is not known as an algebraic function of the model parameters, so it is estimated by simulation of the TAR(2) model with the candidate parameters and resampled residuals. A simulation of length 2×10^5 was used to establish γ to a reasonable precision and this is referred to as $\hat{\gamma}_{\text{simulated}}$.

The R function *optim()* which uses the Nelder–Mead algorithm [6] was used to optimize the parameter values. The long simulation for every set of candidate parameters makes this a challenging optimization problem, but convergence was typically achieved within 30 min on a standard desktop computer. The sum of

squared residuals inevitably increases as ϕ increase from 0, but a substantial reduction in the difference between $\hat{\gamma}_{\text{simulated}}$ and $\hat{\gamma}_{\text{observed}}$ could be achieved with a relatively small increase in the sum of squared residuals.

However, the procedure was not found to be satisfactory because simulations with the optimized parameter values and resampled residuals were found to give lower marginal standard deviations than the standard deviation of the observed time series ($\hat{\sigma}_{\text{observed}}$) (the marginal standard deviation depending on the parameter values as well as the standard deviation of the error distribution). A lower marginal standard deviation would result in underestimation of the variability of extreme values, and lead to unrealistically narrow prediction intervals.

A solution is to include a requirement that the standard deviation of the fitted series should match that of the observed series in the optimization criterion. A standard deviation of the TAR(2) model ($\hat{\sigma}_{\text{simulated}}$) with candidate parameter values can conveniently be calculated along with the $\hat{\gamma}_{\text{simulated}}$. We modify (12) to

$$\omega = \sum_{t=3}^{n} r_t^2 + \phi_1(\hat{\gamma}_{\text{simulated}} - \hat{\gamma}_{\text{observed}})^2 + \phi_2(\hat{\sigma}_{\text{simulated}} - \hat{\sigma}_{\text{observed}})^2, \qquad (14)$$

where ϕ_1 and ϕ_2 are the weight given to mitigate the discrepancy in $\hat{\gamma}$ and $\hat{\sigma}$, respectively. The modification does not noticeably increase the run time. Detailed results are given in Sect. 4.

4   The Sunspot Series

We provide formal evidence of directionality in the average yearly sunspots (1700–2014) [7], and fit time series models to the first 200 points (1700–1900). We use the remaining 115 points (1901–2014) to compare the one-step-ahead forecast errors. We also compare the distribution of extreme values from the observed sunspots (1700–2014) and the simulated directional series using various error distributions.

4.1   Directionality in Sunspots

In Table 1 the DLQC, P^+ and $\hat{\gamma}$ all indicate directionality in the series and all have P-values of 0.00 (two-sided P-values calculated from a parametric bootstrap procedure) [4]. The P-value for P_{abm} and $\hat{\gamma}_{\text{abm}}$ is 0.28 and 0.79 respectively.

Table 1 Summary table of test statistics of directionality for the sunspot series (1700–2014)

Series	Length	Mean	sd	DLQC	P^+	P_{abm}	$\hat{\gamma}$	$\hat{\gamma}_{\text{abm}}$
Sunspot numbers	315	49.68	40.24	−0.0598	42.49 %	51.78 %	0.8555	1.5014

4.2 Threshold Autoregressive Model Fitted by Least Squares

We first consider AR(p) models of orders 1, 2, and 9. AR(9) corresponds to the lowest Akaike information criterion (AIC) for AR(p) models. The results are summarized in Table 6. The AR(2) model is a substantial improvement on the AR(1) model in terms of standard deviation of the errors. An AR(9) and ARIMA(2,0,1) give some further improvement on AR(2).

We consider the TAR(2) model in (11) with three different thresholds set at the 70 %, 80 %, and 90 % percentiles, respectively. The four parameters of the TAR(2) model, α_{1L}, α_{2L}, α_{1U}, and α_{2U}, are estimated by nonlinear least squares from the mean adjusted (to 0) time series. Of all the models considered, TAR(2)_90% with four estimated parameters (TAR(2)[LS]) is the best with an estimated standard deviation of the errors ($\hat{\sigma}_{\text{error}}$) of 13.94 (Table 2).

4.3 Threshold Autoregressive Model Fitted by Penalized Least Squares

We fit a TAR(2)_90% model by penalized least squares (TAR(2)[LSP]), finding the values of α_{1L}, α_{2L}, α_{1U}, and α_{2U} that minimize the objective function in (14). We determine a suitable value of ϕ_1 for fitting the TAR(2)[LSP] model to the first 200 sunspot numbers after fixing the value of ϕ_2 at 10^3 (for which $\hat{\sigma}_{\text{simulated}}$ is kept within 1 % of $\hat{\sigma}_{\text{observed}}$ which is 34.8). We compare all $\hat{\gamma}_{\text{simulated}}$ values modeled by TAR(2)[LSP] from a long simulation (length of 2×10^5) to the target skewness, $\hat{\gamma}_{\text{observed}}$ of 0.8344 in Table 3. A combination of $\phi_1 = 10^5$ and $\phi_2 = 10^3$ provides an

Table 2 Time series models for the sunspot series (1700–1900) compared by σ_{est}

AR model:	AR(1)	AR(2)	AR(9)
$\hat{\sigma}_{\text{error}}$	20.34	15.36	14.84
ARIMA model:	ARIMA(0,0,1)	ARIMA(1,0,1)	ARIMA(2,0,1)
$\hat{\sigma}_{\text{error}}$	21.80	16.60	14.72
TAR model, four parameters:	TAR(2)_70%	TAR(2)_80%	TAR(2)_90%
$\hat{\sigma}_{\text{error}}$	14.67	14.61	13.94

Table 3 Fitting TAR(2)[LSP]: $\hat{\gamma}_{\text{simulated}}$ and $\hat{\sigma}_{\text{error}}$ for selected ϕ_1 ($\phi_2 = 10^3$)

ϕ_1	$\hat{\gamma}_{\text{simulated}}$	$\hat{\sigma}_{\text{error}}$
0	0.1969	14.27
10^4	0.2623	14.04
10^5	0.6315	15.47
10^6	0.8103	16.00
10^{10}	0.8315	16.84
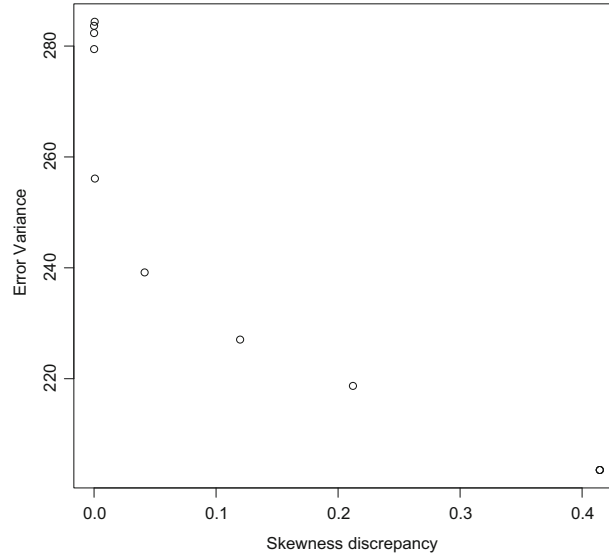
Fig. 2 Fitting TAR(2)[LSP] to the observed sunspots (1700–1900): trade-off between minimizing the sum of squared residuals, and minimizing the skewness discrepancy

improved approximation (0.6315) to the target directionality with a relatively small increase in $\hat{\sigma}_{\text{error}}$ of the TAR(2)[LSP] model.

We illustrate the relationship between $\hat{\sigma}^2_{\text{error}}$ and the squared difference between $\hat{\gamma}_{\text{simulated}}$ and $\hat{\gamma}_{\text{observed}}$ in Fig. 2. The $\hat{\sigma}_{\text{error}}$ is monotonically increasing with the increase of the ϕ_1 for $\phi_1 = 0, 10^1, 10^2, \ldots, 10^{10}$.

4.4 Details of Fitting AR(2), TAR(2)[LS], and TAR(2)[LSP]

The details of fitting AR(2), TAR(2)[LS], and TAR(2)[LSP] models are given in Tables 4, 5, and 6. The upper and the lower regimes of the TAR(2)[LS] and the TAR(2)[LSP] in Tables 5 and 6, respectively, are stable AR(2) processes which satisfy the requirements of the stationary-triangular region $\alpha_2 > -1$, $\alpha_1 + \alpha_2 < 1$ and $\alpha_1 - \alpha_2 > -1$ [2].

Table 4 Sample mean, two coefficients and $\hat{\sigma}_{\text{error}}$ of AR(2)

$\hat{\mu}$	$\hat{\alpha}_1$	$\hat{\alpha}_2$	$\hat{\sigma}_{\text{error}}$
44.11	1.3459	-0.6575	15.36

Table 5 Sample mean, four estimated parameters, and $\hat{\sigma}_{\text{error}}$ of TAR(2)[LS]

$\hat{\mu}$	T	α_{1L}	α_{2L}	α_{1U}	α_{2U}	$\hat{\sigma}_{\text{error}}$
44.11	51.79	1.5643	-0.8117	0.9978	-0.3158	13.94

Table 6 Sample mean, four estimated parameters, ϕ_1, ϕ_2, and $\hat{\sigma}_{\text{error}}$ of TAR(2)[LSP]

$\hat{\mu}$	T	α_{1L}	α_{2L}	α_{1U}	α_{2U}	ϕ_1	ϕ_2	$\hat{\sigma}_{\text{error}}$
44.11	51.79	1.2817	-0.4408	1.1429	-0.3626	10^5	10^3	15.47

Table 7 Test statistics of directionality in the simulated sunspots and the sunspot numbers (1700–1900)

Series	Length	Mean	sd	DLQC	P^+ (%)	P_{abm} (%)	$\hat{\gamma}$	$\hat{\gamma}_{\text{abm}}$
Sunspot numbers	200	44.11	34.76	-0.0609	43.94	48.48	0.8344	1.0646
TAR(2)[LS_G]	2×10^5	39.95	41.38	-0.0088	48.90	50.15	0.1444	0.2939
TAR(2)[LS_R]	2×10^5	44.38	39.25	-0.0173	48.26	49.05	0.2721	0.4878
TAR(2)[LSP_R]	2×10^5	45.25	35.18	-0.0266	45.50	49.18	0.6293	0.9501

4.5 *Simulation to Validate TAR(2)[LS] and TAR(2)[LSP]*

A simulation of 2×10^5 points from the TAR(2)[LS] model with Gaussian errors (TAR(2)[LS_G]) and with the resampled residuals (R) for TAR(2)[LS_R]; and TAR(2)[LSP_R] gave the statistics shown in Table 7.

The TAR(2)[LSP_R] gives the best fit to the first 200 observations in the sunspot series in terms of the statistics that were not included as criteria for fitting.

4.6 *Comparisons of One-Step-Ahead Predictions*

We compare the one-step-ahead forecasting performance of AR(2) with TAR(2)[LS] and TAR(2)[LSP] for the years 1901–2014 (Fig. 3). We define the forecasting performance by measuring the relative errors given by the following measures (Tables 8).

$$\text{Relative error } (E_{\text{rel}}) = \frac{(\text{actual} - \text{predicted})}{\text{actual}}, \tag{15}$$

$$\text{Absolute relative error } (|E_{\text{rel}}|) = \frac{|(\text{actual} - \text{predicted})|}{\text{actual}}. \tag{16}$$

The TAR(2)[LSP] model offers an improvement over the AR(2) and TAR(2)[LS] in terms one-step-ahead predictions.
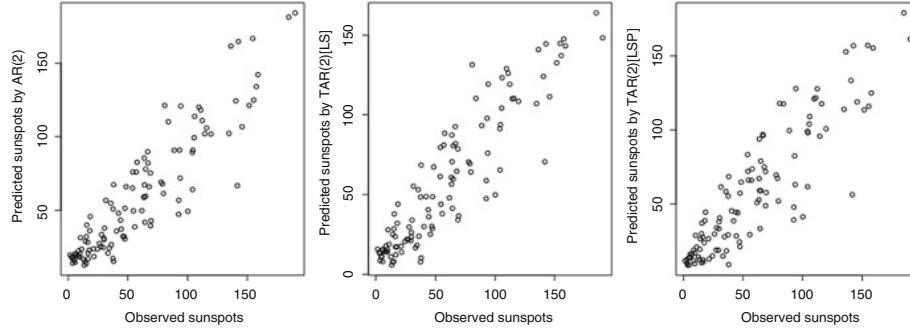
Fig. 3 Comparison of forecast values given by AR(2), TAR(2)[LS], and TAR(2)[LSP] at 90 % threshold

Table 8 Forecasting measures of predicted sunspots to the sunspot series from 1901 to 2014

Model	Mean(E_{rel})	sd(E_{rel})	Mean($|E_{rel}|$)	sd($|E_{rel}|$)
AR(2)	−0.6812	2.4777	0.8713	2.4169
TAR(2)[LS]	−0.4886	2.3014	0.7512	2.2289
TAR(2)[LSP]	−0.5015	2.1897	0.7394	2.1206

4.7   Comparisons of Distributions of 15-Year Extreme Values

We simulate 2×10^5 values using an AR(2) model with Gaussian errors (AR(2)_G), TAR(2)[LS_G], TAR(2)[LS_R], and TAR(2)[LSP_R] models. The upper and lower coefficients for TAR(2)[LS] and TAR(2)[LSP] are the optimized parameters in Tables 5 and 6 accordingly. We include another two different error distributions for TAR(2)[LSP] which are the back-to-back Weibull distribution, that is one Weibull distribution was fitted to the positive residuals and another to the absolute values of the negative residuals, (WD); and an Extreme Value Type 1 (Gumbel) distribution of minima (EV).

We refer TAR(2)[LSP] with WD and EV errors as TAR(2)[LSP_WD] and TAR(2)[LSP_EV], respectively. We calculate the extreme values for every 15 consecutive years in the simulated series of length 2×10^5, illustrate with boxplots (Fig. 4) and provide descriptive statistics (Table 9).

In general, TAR(2)[LSP] models simulate greater extreme values than TAR(2) [LS] and AR(2) models, as shown by inter-quartile range (IQR) and standard deviation (sd) in Table 9. Furthermore, 15-year extreme values from TAR(2)[LSP_DW] have the closest sd, and IQR, to the extreme values from the observed time series.

Table 9 Descriptive statistics of 15-year extreme values in the sunspot series (1700–2014) and the simulated series for each model with different residuals

Series	n	Median	Mean	Max	IQR	sd	Skewness
Sunspot numbers	21	111.0	112.20	190.20	53.90	37.89	0.1158
AR(2)_R	13,333	91.60	94.54	221.40	34.53	26.10	0.5903
TAR(2)[LS_G]	13,333	101.10	99.30	175.40	26.24	20.65	−0.2523
TAR(2)[LS_R]	13,333	102.70	102.00	195.10	27.58	21.94	0.0025
TAR(2)[LSP_R]	13,333	91.45	91.46	236.60	40.25	29.82	0.2110
TAR(2)[LSP_WD]	13,333	97.39	100.20	358.80	50.06	40.18	0.7913
TAR(2)[LSP_EV]	13,333	83.77	84.61	255.30	41.35	30.60	0.2973
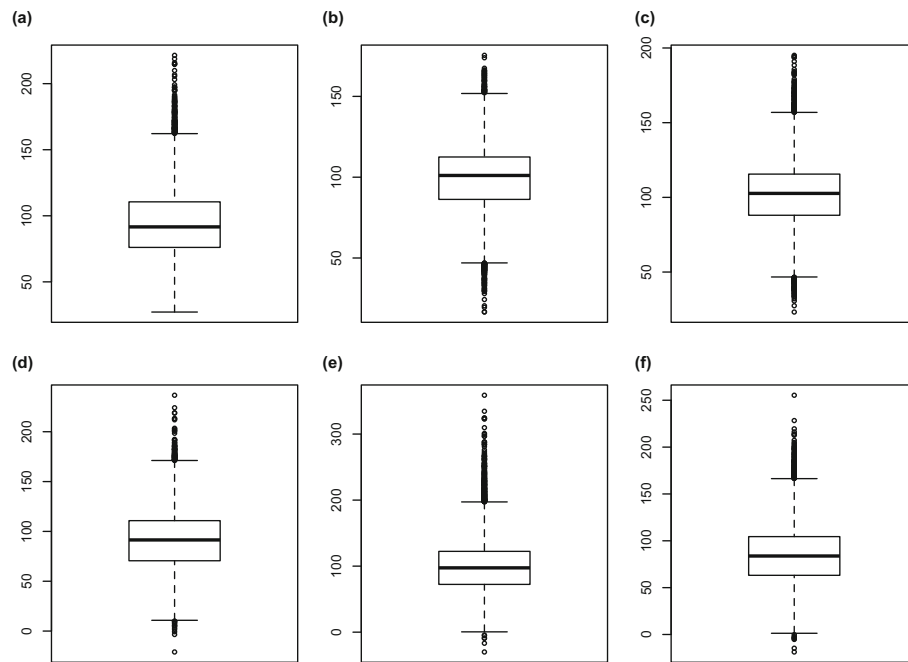
Fig. 4 Boxplot of 15-year extreme values in the simulated series for AR(2) and TAR(2) models. (**a**) AR(2)_R; (**b**) TAR(2)[LS_G]; (**c**) TAR(2)[LS_R]; (**d**) TAR(2)[LSP_R]; (**e**) TAR(2)[LS_WD]; (**f**) TAR(2)[LSP_EV]

5 Conclusion

There are many ways in which a time series can exhibit directionality, and different measures are needed to identify these different characteristics. TAR models provide a piecewise linear approximation to a wide range of nonlinear processes, and offer a versatile modeling strategy. The sunspot series shows clear directionality, a physical interpretation of which is given in [5]. We have shown that a nonlinear TAR(2)[LS]

model gives an improvement on, out of sample, one-step-ahead predictions made with an AR(2) model.

With the inclusion of the measure of directionality in the objective function (12) in the fitting procedures for the TAR(2)_90% model, we are able to reduce the discrepancy between the observed and the simulated directionality seen in the TAR(2)[LS model]. Furthermore, we have demonstrated that any consequential discrepancy in the marginal standard deviations in the fitted model may similarly be dealt with by the inclusion of the standard deviation term in the improved objective function (14). This TAR(2)[LSP] model yields improved one-step-ahead predictions for 115 out-of-sample values.

The use of resampled residuals in simulations of extreme values is unsatisfactory because the extreme errors in the simulation are restricted to the range of the residuals. In the case of the sunspots, back-to-back Weibull distributions provided good fit to the residuals and resulted in far more realistic simulations of 15-year extreme values.

In summary, we have modeled directionality in the sunspot series by explicitly using both the measure of directionality and standard deviation as fitting criteria. The explicit modeling of directionality has provided more accurate forecasting and more realistic modeling of extreme values.

References

1. Beare, B.K., Seo, J.: Time irreversible copula-based Markov models. Econ. Theory **30**, 1–38 (2012)
2. Chatfield, C.: The Analysis of Time Series: An Introduction, 6th edn., pp. 218–219, 223–224, 44–45. Chapman and Hall/CRC, London/Boca Raton (2004)
3. Lawrance, A.: Directionality and reversibility in time series. Int. Stat. Rev./Revue Internationale de Statistique **59**(1), 67–79 (1991)
4. Mansor, M.M., Green, D.A., Metcalfe, A.V.: Modelling and simulation of directional financial time series. In: Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015), pp. 1022–1028 (2015)
5. Mansor, M.M., Glonek, M.E., Green, D.A., Metcalfe, A.V.: Modelling directionality in stationary geophysical time series. In: Proceedings of the International Work-Conference on Time Series (ITISE 2015), pp. 755–766 (2015)
6. Nash, J.C.: On best practice optimization methods in R. J. Stat. Softw. **60**(2), 1–14 (2014)
7. Solar Influences Data Analysis Center, Sunspot Index and Long-term Solar Observations. http://www.sidc.be/silso (last accessed 17 October 2015)
8. Soubeyrand, S., Morris, C.E., Bigg, E.K.: Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data. Environ. Model Softw. **61**, 78–86 (2014)
9. Wild, P., Foster, J., Hinich, M.: Testing for non-linear and time irreversible probabilistic structure in high frequency financial time series data. J. R. Stat. Soc. A. Stat. Soc. **177**(3), 643–659 (2014)

4.2 Directionality and Volatility in Electroencephalogram Time Series

Statement of Authorship

Directionality and Volatility in Electroencephalogram Time Series

Published in *American Institute of Physics (AIP) Conf. Proc.* 1739, 2016.

Mahayaudin M. Mansor (Candidate)
Constructed statistical models, conducted analyses and interpreted results, wrote manuscript and acted as corresponding author.
Overall percentage: 80%

This paper reports on original research I conducted during the period of my Higher Degree by Research candidature and is not subject to any obligations or contractual agreements with a third party that would constrain its inclusion in this thesis. I am the primary author of this paper.

Signed: .. .. Date: ....*14 June 2017*....

By signing the Statement of Authorship, each co-author certifies that: (i) the candidate's stated contribution to the publication is accurate; (ii) permission is granted for the candidate in include the publication in the thesis; and (iii) the sum of all co-author contributions is equal to 100% less the candidate's stated contribution.

David A. Green
Supervised development of work and provided critical evaluation.

Signed: .. Date: *14/6/17*

Andrew V. Metcalfe
Supervised development of work and provided critical evaluation.

Signed: .... Date: *14/6/17*

Directionality and Volatility in Electroencephalogram Time Series

Mahayaudin M. Mansor[a)], David A. Green[b)] and Andrew V. Metcalfe[c)]

School of Mathematical Sciences, Level 6 & 7, Ingkarni Wardli Building,
North Terrace Campus, The University of Adelaide, SA 5005, Australia.

[a)]Corresponding author: mohdmahayaudin.mansor@adelaide.edu.au
[b)]david.green@adelaide.edu.au
[c)]andrew.metcalfe@adelaide.edu.au

Abstract. We compare time series of electroencephalograms (EEGs) from healthy volunteers with EEGs from subjects diagnosed with epilepsy. The EEG time series from the healthy group are recorded during awake state with their eyes open and eyes closed, and the records from subjects with epilepsy are taken from three different recording regions of pre-surgical diagnosis: hippocampal, epileptogenic and seizure zone. The comparisons for these 5 categories are in terms of deviations from linear time series models with constant variance Gaussian white noise error inputs. One feature investigated is directionality, and how this can be modelled by either non-linear threshold autoregressive models or non-Gaussian errors. A second feature is volatility, which is modelled by Generalized AutoRegressive Conditional Heteroskedasticity (GARCH) processes. Other features include the proportion of variability accounted for by time series models, and the skewness and the kurtosis of the residuals. The results suggest these comparisons may have diagnostic potential for epilepsy and provide early warning of seizures.

INTRODUCTION

The electroencephalogram (EEG) is obtained by monitoring electrical signals in areas of the brain. The comparison of EEG from healthy subjects with those from subjects who have epilepsy has the potential for improved diagnostics and therapeutics [1, 2]. Moreover the comparison may lead to early warning of forthcoming seizures, which would enable subjects to reach a safer environment before onset. Comparison of stationary time series generally requires consideration of more detailed characteristics than just the mean and autocovariance structure. Directionality often indicates non-linearity, which in itself may result from complex feedback systems that could be a cause of epilepsy. Volatility, which is defined as bursts of higher variability segments in the signal, is an indication of non-stationarity in the variance of the signal. Such volatility could be an indication of epilepsy.

Directionality or asymmetry in time can be seen in many stationary time series from various disciplines including medicine. A time series that can reasonably be considered as a realization of a first order stationary random process (that is any trend or seasonal effects have been identified and removed) is said to be directional if there is qualitative difference between observations plotted in chronological order and reverse time order (time-to-go). A clear example of directionality can be seen in the EEG time series (Fig. 2(d)). In time order there is a tendency for sharp increases followed by gradual recessions, whereas in time-to-go slow increases are followed by plummets.

In this paper, we investigate EEG time series from a group of five healthy volunteers (healthy group) and a group of five patients diagnosed with epilepsy (epilepsy group). For the healthy group, there were 100 pairs of EEG series recorded from 5 volunteers during a conscious relaxed state, using a standard multi-location placement of sensors. One of each pair of EEG series was recorded with the volunteer's eyes open (condition A) and this was followed by (or preceded by –the allocation being at random) a recording with the volunteer's eyes closed (condition B). For the epilepsy group, there were 100 triples of EEG series corresponding to three conditions C, D and E. The conditions were three different recording sites of hippocampal formation, epileptogenic and seizure region respectively [3]. The EEG series in each triple were recorded on the same subjects with epilepsy during the same session.

The paper is organized in five sections: comparing mean and standard deviation of EEG time series; detection of directionality and a novel method for modeling directionality, and the application of these techniques to the comparison of EEG from healthy subjects and the subjects with epilepsy; the use of a measure of volatility and models for volatility for distinguishing EEG from healthy subjects and subjects with epilepsy; a comparison of the residuals after fitting time series models to EEG from healthy subjects and subjects with epilepsy; and a conclusion.

EEG TIME SERIES

There are 100 time series of length 4097 for each of the conditions A up to E. The means of the means and standard deviations of these 100 time series are given in Table 1[1]. There is no significant difference between the means of the 100 series under condition A and the 100 series under condition B, eyes closed (P=0.12). However, the standard deviation of the series is higher under condition B (P≈.000). There is no significant difference between the means of the 100 series under conditions C, D and E (P=0.51), but there is a striking difference in standard deviations of EEG recorded at the different locations (P≈.000).

TABLE 1. Mean of the means and standard deviations of the 100 EEG time series for each condition. Also given are p-values for tests of equality of means and standard deviations of conditions within each group.

Group	Healthy volunteers		Epilepsy patients		
Condition	A	B	C	D	E
Mean	-6.26	-12.51	-8.88	-6.20	-4.75
P-value	0.116		0.505		
Standard deviation	40.73	61.11	50.83	65.62	306.61
P-value	0.000		0.000		

(a)   **(b)**
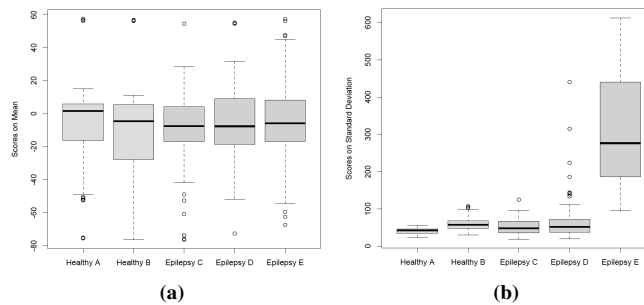
FIGURE 1. Boxplots of 100 EEG time series means (a) and standard deviations (b) by condition.

DIRECTIONALITY

A randomly selected time series for each condition is shown plotted in time order and against time-to-go in Fig. 2. Time series D_76 (76^{th} triple condition D) shows rapid increases and slow recessions about the spikes. It is harder to discern directionality in the other series and generally a formal quantitative test is required to quantify directionality.

---

[1]The p-values for the comparisons were obtained from a Bootstrap resampling procedure. The statistics considered for the sample of 100 pairs A,B were $\bar{d}/(s_d/\sqrt{100})$, and s_A/s_B respectively. The statistics considered for the 100 triples C,D,E were the F-ratios from randomized block analyses taking series means, and logarithms of series standard deviations, as the response.

(a) Healthy A_40   **(b)** Healthy B_98

(c) Epilepsy C_37   **(d)** Epilepsy D_76   **(e)** Epilepsy E_56
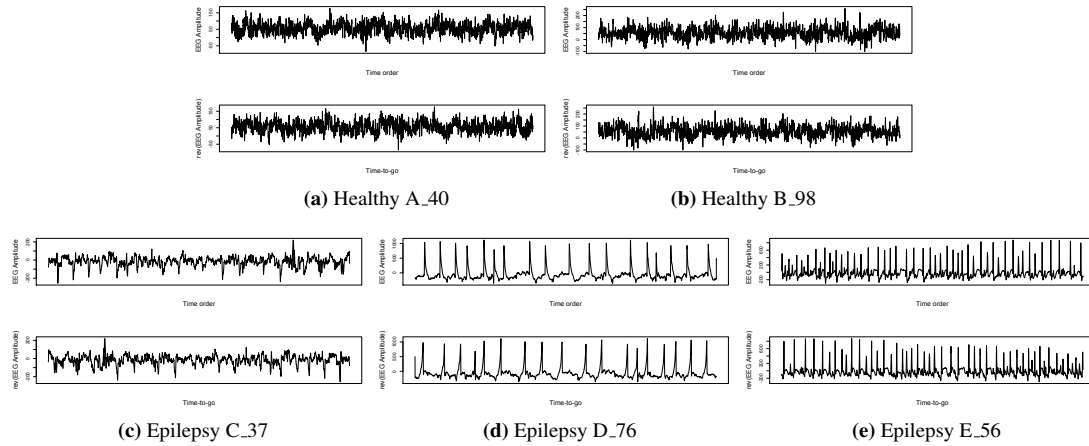
FIGURE 2. Selected EEG time series (length 4097) for the five conditions in time order (upper frames) and time-to-go (lower frames).

Detecting Directionality

Denote the time series model by $\{X_t\}$ for $t = 1, 2, ..., n$, and define the first differences, $Y_t = X_t - X_{t-1}$ for $t = 2, 3, ..., n$. If a process is reversible, that is it is not directional, the distribution of first differences will be symmetric. It follows that skewness of the first differences can be used as a measure of directionality [4]. In this paper, we employ the product moment skewness of first differences ($\hat{\gamma}_{dif}$), equation 1, to estimate the marginal skewness of the distribution of first order lag one differences [4]. EEG series with significant non-zero skewness of Y_t is evidence of directionality. This measure was chosen because it is generally the most intuitive and effective statistic for detecting directionality in time series from various disciplines (e.g.: [5, 6, 7]). Positive values of $\hat{\gamma}_{dif}$ correspond to a tendency for steep increases to be followed by slow decreases.

$$\hat{\gamma}_{dif} = \frac{\sum_{t=1}^{n}(y_t - \bar{y})^3/(n-1)}{[\sum_{t=1}^{n}(y_t - \bar{y})^2/(n-1)]^{3/2}}. \tag{1}$$

First, we measure directionality in the selected ECG series displayed in Fig. 2 then we determine the significance level of the statistic using a simulation procedure from a reversible process corresponding to the null hyphothesis of reversibility [6]. The results are given in Table 2. The *p*-order in the autoregressive models is determined by minimum Akaike information criterion (AIC) and the AR(AIC) models show substantial reduction in the standard deviation of the errors. A non-zero skewness coefficient with p-value less than 0.05 is said to be statistically directional at 5% significance level. Epilepsy series D_76 shows strong directionality in plots, and has the highest absolute value of directionality of 6.425 which is highly statistically significant. Epilepsy series C_37 possesses directionality of -0.6543 although very little evidence of directionality can be seen from plots. Epilepsy series E_56 is also seen to be statistically directional while healthy series A_40 and B_98 do not show significant directionality.

TABLE 2. Directionality, and change in standard deviation of errors after fitting AR model, for the series plotted in Fig. 2.

EEG series	Healthy A_40	Healthy B_98	Epilepsy C_37	Epilepsy D_76	Epilepsy E_56
Directionality, $\hat{\gamma}_{dif}$	0.060	-0.029	-0.654	6.425	-0.167
P-value	0.098	0.553	0.000	0.000	0.001
Standard deviation, s	33.92	46.98	52.71	223.65	160.20
{AR(AIC), $\hat{\sigma}_{est}$}	{35, 12.70}	{36, 9.32}	{36, 7.29}	{6, 28.43}	{23, 22.05}

Next, we measure directionality in 100 EEG time series for all conditions. The results are illustrated in Fig. 3. We conjecture that the lack of consistent sign for the directionalities that are high in absolute value is due to arbitrary polarity of the measuring procedure. This is suggested by the plots of C_37 and D_76 in Fig. 2 which are spikes down and spikes up respectively. In general the EEG from the epilepsy group C and D commonly show substantial directionality, and site E consistently shows directionality. In contrast the EEG from the healthy volunteers show little directionality but this may be partly due to the averaging of signals from multiple electrodes.

(a) Directionality   **(b)** Absolute value of directionality
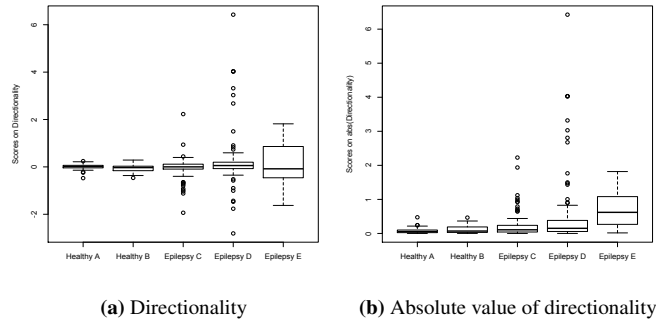
FIGURE 3. Boxplots of directionalities (a) and absolute value of directionalities (b) by condition (100 EEG time series for each condition).

Modelling Directionality

Directionality can be modelled using linear time series model with non-Gaussian errors [4] or using non-linear model with either Gaussian or non-Gaussian error distributions [5]. We fit an autoregressive model of order 2, AR(2), to the epilepy series D_76 which has the highest degree of directionality among all series (uppermost point in Fig. 3). The standard deviation of the AR(2) errors (Table 3) is only slightly higher than the standard deviation of the errors from the AR(6). The standard deviation of the errors is 13% of the marginal standard deviation of the time series (equivalent to AR(0)).

TABLE 3. Standard deviation of errors for AR models for series D_76.

Order	Epilepsy series D_76
0	$s = 223.65$
2	AR(2), $\hat{\sigma}_{error} = 29.48$
AIC	AR(6), $\hat{\sigma}_{error} = 28.43$

We now fit a threshold autoregressive model of order 2, TAR(2) with a threshold (T) set at the 90% percentile, equation 2, by minimizing a weighted sum of the sum of squared residuals and the squared discrepancy between observed and simulated directionality, equations 3 and 4 [6, 8] referred to as penalized least squares [LSP]. The fitted TAR(2) is unstable for $\phi = 0, 10^1, 10^2, 10^3$. One remedy would be to use constraint optimization with the constraints that $\alpha_{1U}, \alpha_{2U}, \alpha_{1L}, \alpha_{2L}$ satisfy $\alpha_2 > -1$, $\alpha_1 + \alpha_2 < 1$ and $\alpha_1 - \alpha_2 > -1$ [9]. However it turns out that increasing ϕ to $10^4, 10^5$ or 10^6 gives a stable model.

$$X_t = \begin{cases} \alpha_{1U}X_{t-1} + \alpha_{2U}X_{t-2} + \epsilon_t & \text{if } X_{t-1} > T \\ \alpha_{1L}X_{t-1} + \alpha_{2L}X_{t-2} + \epsilon_t & \text{if } X_{t-1} < T. \end{cases} \qquad (2)$$

We estimate α_{1U}, α_{2U}, α_{1L} and α_{2L} for TAR(2) using penalized least squares that includes the objective function of

$$\omega = \sum_{t=3}^{n} r_t^2 + \phi(\hat{\gamma}_{simulated} - \hat{\gamma}_{observed})^2, \tag{3}$$

where $\{r_t\}$ for $t = 3, ..., n$ are the residuals defined by

$$r_t = \begin{cases} x_t - \alpha_{1U}(x_{t-1}) - \alpha_{2U}(x_{t-2}) & \text{if } (x_{t-1}) > T \\ x_t - \alpha_{1L}(x_{t-1}) - \alpha_{2L}(x_{t-2}) & \text{if } (x_{t-1}) < T. \end{cases} \tag{4}$$

TABLE 4. Fitting TAR(2)[LSP].

Target	$\hat{\gamma}_{observed} = 6.425$		
Model	TAR(2)[LSP]		
ϕ	10^4	10^5	10^6
$\hat{\gamma}_{simulated}$	4.88	5.04	5.08
$\hat{\sigma}_{error}$	23.62	23.70	23.72

TABLE 5. Parameters of fitted AR(2) and TAR(2)[LSP].

Model	$\hat{\alpha}_1$	$\hat{\alpha}_2$	$\hat{\sigma}_{error}$			
AR(2)	1.675	-0.704	29.48			

Model	T	$\hat{\alpha}_{1L}$	$\hat{\alpha}_{2L}$	$\hat{\alpha}_{1U}$	$\hat{\alpha}_{2U}$	ϕ_1	$\hat{\sigma}_{error}$
TAR(2)[LSP]	215.76	1.683	-0.691	1.665	-0.694	10^6	23.72

TABLE 6. Comparison of observed time series and simulation of length 10^5 from the TAR(2)[LSP].

Series	Length	Mean	sd	$\hat{\gamma}_{dif}$
Epilepsy series D_76	4097	-51.76	223.65	6.425
TAR(2)[LSP]	10^5	-24.94	215.75	4.854

The TAR(2)[LSP] gives an improved fit on both AR(2) and AR(AIC) inasmuch as the standard deviation of the errors of 23.72, which is about 15% lower. Moreover, the constraint on directionality stabilized the fitted TAR(2) in this particular case.

VOLATILITY

Volatility in residuals after fitting a time series model is a tendency for bursts of high variability. It may be apparent in a plot of the residuals, but a more sensitive test is given by the autocorrelation function of the squared residuals. If there is no volatility 95% of autocorrelations of squared residuals will lie within $\pm\frac{2}{\sqrt{n}}$, where n is the length of the time series, so autocorrelations outside these limits provide evidence of volatility. Volatility is commonly modelled by a GARCH(1,1) process of the form

$$\sigma_t^2 = \omega + \alpha\,\epsilon_{t-1}^2 + \beta\,\sigma_{t-1}^2 \tag{5}$$

where ϵ_{t-1} is the error at time $t-1$ and σ_t^2 is the variance at time t. A single measure of volatility is given by $\frac{\omega}{1-\alpha-\beta}$, and hence the volatility can be estimated from the residuals by least squares. In this study a random sample of 10 time series for each condition was investigated. Volatility was found under conditions D and E but not under conditions A, B and C. The results are summarized below, and there is evidence that E is more volatile than D.

TABLE 7. Mean volatility and residuals standard deviation after fitting AR for 10 selected EEG time series within epilepsy group.

	Group	Epilepsy patients	
	Condition	D	E
	Volatility	10.59	48.90
	P-value	0.007	
Standard deviation of the residuals		10.99	33.96
	P-value	0.016	
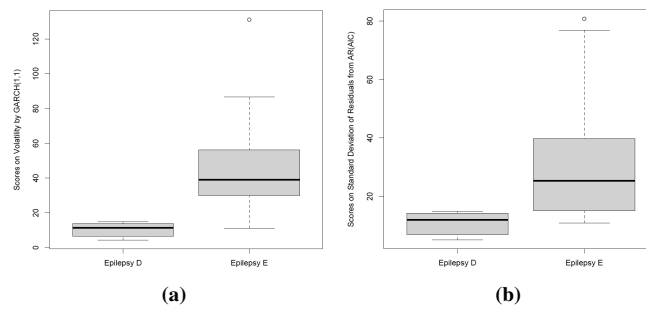
FIGURE 4. Boxplots of 10 EEG time series volatility values (a) and standard deviations of AR(AIC) residuals (b) by condition for epilepsy group.

CHARACTERISTICS OF RESIDUALS

Non-Gaussian errors impart some degree of directionality. We first compare the ratio of the mean standard deviation of the errors after fitting AR(AIC) to the mean marginal standard deviation of the time series. The smaller this ratio the more variability is explained by AR(AIC).

TABLE 8. Mean statistics of residuals from AR(AIC) for the 100 EEG time series for each condition.

	Group	Healthy volunteers		Epilepsy patients		
	Condition	A	B	C	D	E
Ratio of residual variability		0.23	0.18	0.15	0.14	0.13
	P-value	0.000			0.008	
Skewness of residuals		0.06	0.11	0.05	-0.07	-0.03
	P-value	0.397			0.861	
Kurtosis of residuals		5.66	20.78	13.28	19.81	20.32
	P-value	0.000			0.386	

This ratio depended on the condition and ranged from 0.23 for A to 0.13 for E. On average there was no evidence that the residuals were skewed. However, the average value of kurtosis was far higher than the expected value of 3 for a Gaussian distribution. This high value of kurtosis is a consequence of outlying values in the time series. The difference in average kurtosis between conditions A and B for healthy volunteers was significantly different.

(a) Proportion of residual variability **(b)** Skewness of residuals **(c)** Kurtosis of residuals
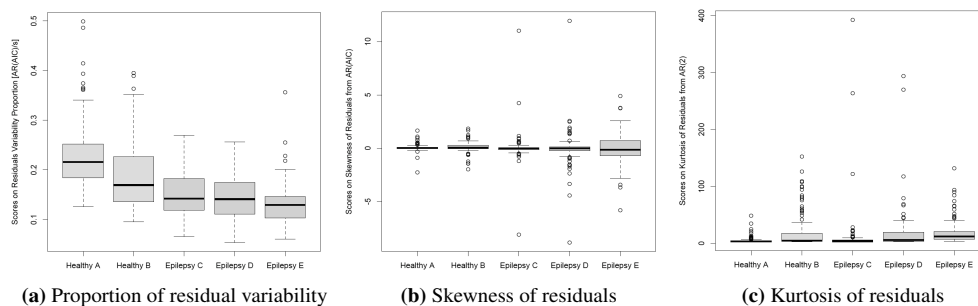
FIGURE 5. Boxplots of 100 EEG time series statistics for residuals from AR(AIC) by condition.

Unpaired T-test

Finally, we compare mean, standard deviation, directionality, ratio of residuals variability, skewness and kurtosis of residuals between condition A, and separately condition B, from healthy group with condition D from epilepsy group using unpaired t-test and the results are given in Table 9.

TABLE 9. Mean statistics of the 100 EEG time series of length 4097 for each group.

Comparison	Healthy A vs. Epilepsy D		Healthy B vs. Epilepsy D	
Mean	-6.26	-6.20	-12.51	-6.20
P-value		0.987		0.105
Standard deviation	40.73	65.62	61.11	65.62
P-value		0.000		0.458
Directionality	0.010	0.262	-0.071	0.262
P-value		0.035		0.006
Ratio of residuals variability	0.23	0.14	0.18	0.14
P-value		0.000		0.000
Skewness of residuals	0.06	-0.07	0.11	-0.07
P-value		0.463		0.315
Kurtosis of residuals	5.66	19.81	20.78	19.81
P-value		0.001		0.856

CONCLUSIONS

Although all brain signals were obtained using the same multi-channel amplifier system, the positions of electrodes differ between the healthy group and the epilepsy group. The brain signals for the healthy group (condition A and condition B) were taken from 19 electrodes (in accordance to the international electrode placement scheme), all over the volunteer's head, while EEG segments of the epilepsy group of the condition set C, D and E are recorded using the intracranial strip electrodes scheme for pre-surgical assessment of epilepsy patients from: the hippocampal formation zone (condition C); the epileptic site (condition D) of the brain during seizure free intervals; and epilepsy condition E series are from seizure zone [3]. The signals for the epileptic group under conditions D and E were remarkably different from the others. The D series were highly directional and the E series were highly variable. Further work may show that these findings have diagnostic potential. In particular it would be valuable to take readings from healthy subjects under conditions D and E.

Comparisons within the healthy group showed some statistically significant differences. The standard deviations of the time series under condition B tended to be higher than the standard deviations of time series under condition A. The 100 time series under condition A have on average negligible directionality (0.01), whereas the 100 time series in B have on average negative directionality (-0.07). The comparison between characteristics of residuals after fitting AR models to the healthy group showed that the AR model explains more of the variability in the time series under condition B; and the kurtosis of the residuals is higher under condition B. Comparisons within the epileptic group showed that there were statistically significant differences for: the standard deviation in the time series under conditions C, D and E; the variation of directionality for individual time series (Epilepsy C_37, Epilepsy C_76 and Epilepsy C_56); EEG time series under condition E are more volatile than D; and the ratio of residual variability. These differences may have potential as diagnostic or indication of onset, which would be highly valuable to people who are prone to epilepsy because they can remove themselves from a potentially dangerous situation.

ACKNOWLEDGMENTS

REFERENCES

[1] G. Petkov, M. Goodfellow, M. P. Richardson, and J. R. Terry, Frontiers in Neurology, **5**(261), 1–7 (2014)

[2] F. Miraglia, F. Vecchio, P. Bramanti, and P. M. Rossini, Clinical Neurophysiology, 1261–1268 (2016).

[3] R. G. Andrzejak, K. Lehnertz, F. Mormann, C. Rieke, P. David, and C. E. Elger, Physical Review E **64**(6; PART 1), 061907-1–061907-8 (2001).

[4] A. J. Lawrance, International Statistical Review/Revue Internationale de Statistique, JSTOR **59**(1), 67–79 (1991).

[5] M. M. Mansor, M. E. Glonek, D. A. Green, and A. V. Metcalfe, *Proceedings of the International Work-conference on Time Series (ITISE 2015)*, edited by O. Valenzuela, F. Rojas, H. Pomares, and I. Rojas (Granada, 2015), pp. 755–766.

[6] M. M. Mansor, D. A. Green, and A. V. Metcalfe, *Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)*, edited by T. Weber, M. J. McPhee, and R. S. Anderssen (Gold Coast, 2015), pp. 1022–1028.

[7] M. M. Mansor, F. L. Mohd. Isa, D. A. Green, and A. V. Metcalfe, *Proceedings of the 12th Engineering Mathematics and Applications Conference (EMAC 2015)*, edited by M. I. Nelson (ANZIAM-E Journal, Adelaide, 2016).

[8] M. M. Mansor, M. E. Glonek, D. A. Green, and A. V. Metcalfe, Springer, accepted: October 2015.

[9] C. Chatfield, *The Analysis of Time Series: An Introduction, Sixth Edition* (Chapman and Hall/CRC Publication, 2004).

4.3   Directionality in Time Series of Bank Log-return Share Prices

Directionality in Time Series
of Bank Log-return Share Prices

Mahayaudin M. Mansor,[*] David A. Green and Andrew V. Metcalfe

School of Mathematical Sciences, University of Adelaide, SA 5005, Australia

Abstract

The period 1999-2017 is split into three, relatively financially stable and two relatively financially unstable sub-periods. The directionality of the daily log-return time series for the five largest U.S. banks is investigated and found to be consistently high and positive during crisis periods than non-crisis periods. A moving directionality index is defined, and it is shown that moving directionality increases around economic crises and that it has potential as an early warning indicator of falling share prices. The performance of a portfolio, equally distributed over the five banks with a directionality investment criterion as a trading rule is investigated and performs consistently better than a rule based on volatility. Time series models that show directionality and volatility in realisations are fitted by least squares and penalized least squares. The best fitting model is used to evaluate the conditional value at risk (CVaR) for the directionality based trading rules.

Keywords: Reversibility, volatility, trading rule, investment simulation, penalized least squares, CVaR.

---

[*]Correspondence to: mohdmahayaudin.mansor@adelaide.edu.au

1   Introduction

Directionality is defined for stationary time series as asymmetry in time, and can be seen in many time series from various disciplines, including finance. Directionality is defined for realisations of stationary time series models in order to isolate the phenomenon from the effects of deterministic trends, stochastic trends as in a random walk, and seasonality which are often found in many business and financial time series. In practice, if a time series appears to have a trend or seasonal effects, then these characteristics should be identified and removed before considering directionality.

Directionality has a complementary concept of time reversibility (symmetry in time) (Lawrance 1991, Soubeyrand et al. 2014). A stationary random process $\{X(t)\}$ is said to be reversible if the joint probability between $\{X(t_1), ..., X(t_n)\}$ and $\{X(-t_1), ..., X(-t_n)\}$ is equal for every n, and every $t_1, ..., t_n$ (Weiss 1975). If a stationary random process is not reversible, then it is directional. There will, in principle, be qualitative differences between a realisation of a directional random process, a time series, plotted in time order and plotted in reverse time order (time-to-go) (Lawrance 1991).

Stationary linear time series models with Gaussian white noise are reversible, whereas non-linear time series models and linear models driven by non-Gaussian white noise are directional (e.g. Lawrance 1991). Therefore detecting directionality can indicate when non-linear time series models or at least non-Gaussian errors are appropriate to provide more accurate forecasts and more realistic generation of an ensemble of future scenarios. Threshold Autoregressive (TAR) models, which are based on piecewise linearization (Tong & Lim 1980), are also found to be useful for modelling directionality in time series and closer reproduction of directionality can be achieved by using a penalized least squares approach when fitting the TAR models (Mansor et al. 2016a,b,c, 2015b).

In some cases, directionality can be seen in stationary time series by comparing a plot against time order with a plot against time-to-go. For example, the United States (U.S.) civilian unemployment rate for the monthly series (Tsay 2005). However in the case of log-return time series of JP Morgan, it is difficult to discern directionality when comparing a plot of 315 log-returns (9 November 2004 to 8 February 2005) in time order with a plot of the 315 log-returns in time-to-go (Figure 1(a)). There is also no qualitative evidence of directionality in another sub-series of JP Morgan log-returns between January 2007 to December 2009 in Figure 1(b) where there are two major financial crises, the Global Financial Crisis (GFC) 2007-2008 and the subprime mortgage crisis from Dec 2007 to June 2009 (Allen & Christa 2013), happened in the U.S. during the period. Therefore we cannot rely on the graphical inspection of directionality alone and in general it is necessary to consider statistical measures to identify directionality.

In this paper, we focus on detecting directionality in log-returns time series of bank share prices because banks are central to business activity and the banking sector has experienced a series of severe financial and economic shocks which have affected banks' performance around the world. For example, banking crises cause bank distress followed by low credit lending activities that has an adverse effect on growth and leads to poor
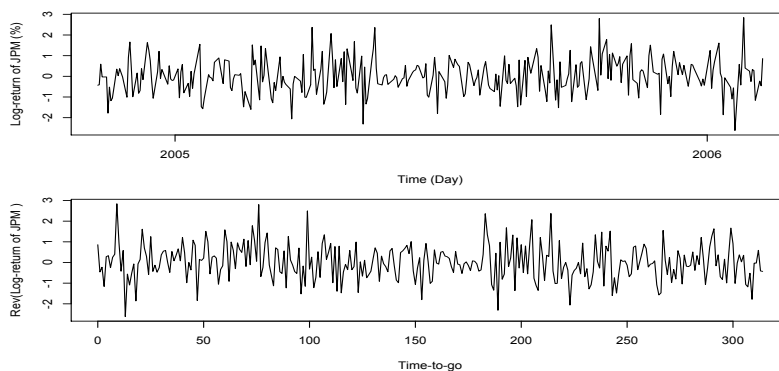
(a) JP Morgan log-return sub-series during stable period.
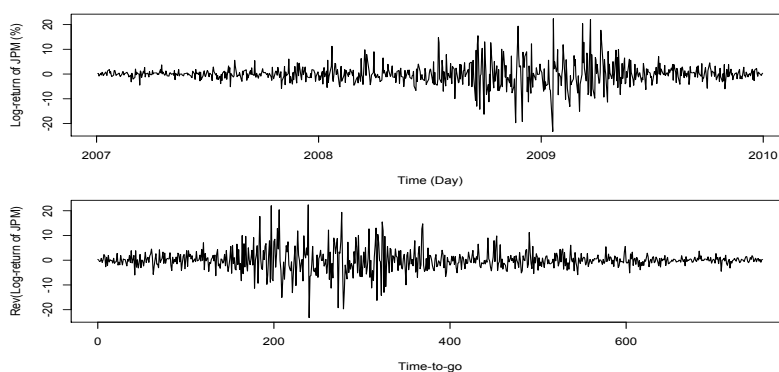
(b) JP Morgan log-return sub-series during financial crises.

Figure 1: Graphical inspection of directionality in sub-series of JP Morgan log-returns.

economic performance. These effects are particularly serious: in developing countries where alternative bank financing are more limited; in countries with less access to foreign finance; and where bank distress is more critical (Dell'Ariccia et al. 2008). Many banks had to undergo the financial stress test and received capital injections either from larger banks or massive government bailouts. In addition, Allen & Christa (2013) examined the effect of capital on the bank′s survival and market shares during two banking crises, three market crises and during normal times in the U.S. Their study conclude that capital, is the main buffer against the negative shocks, helps to increase survival rate and market shares for small banks at all times and for medium and large banks only during the banking crises. Consequently, detecting directionality in the log-return series probably reveal feedbacks following these extreme events. As we are interested to know the behaviour of this phenomenon during crises and non-crises periods.

3

In Section 2 we investigate directionality in the five largest banks in the U.S., and determine whether or not the directionality in the log-return series can plausibly be attributed to chance. In Section 3 we consider directionality as an indicator of financial crises in comparison to volatility. In Section 4, we explore the potential of directionality as a criterion for trading with a simulation of investment portfolio of shares from the five American banks. Then we consider a single asset investment for five non-US banks. JP Morgan Chase & Co. is the largest bank in the U.S. with total assets of US$2.35 trillion; the world's sixth largest bank by total assets and the world's most valuable bank by market capitalization (Keller & Chiglinsky Sept. 13, 2016). We investigate the JP Morgan time series in more detail in Section 5 where we fit time series models to the JP Morgan log-return time series including TAR model fitted by a penalized least squares. We examine the performance of directionality trading rules further by using CVaR for the JP Morgan share prices. Concluding remarks are given in Section 6.

2   Detecting Directionality

2.1   Data

The five largest banks in the U.S. by market capitalization for 2017 are JP Morgan Chase & Co., Wells Fargo & Co., Bank of America, Citibank Inc. and Goldman Sachs Group (relbanks 2017). The daily share prices of each bank are plotted in Figure 2(a, b, c, d and e), respectively. The common period for the five daily share price series is from 3 May 1999 to 17 February 2017, which is the longest common period since their initial public offering (IPO) in the New York Stock Exchange.

Brown & Warner (1985) highlight three issues concerning a study using daily share returns that include non-normality, non-synchronous trading and market model parameter estimation and variance estimation. However the non-normality of daily returns has no obvious impact on the study that is observed with the monthly data, autocorrelation and changes in variance conditional on daily returns can sometimes be advantageous (Brown & Warner 1985). In practice daily closing price is used as the closing position for share valuation as well as for the opening position for the next day valuation. In general it is useful to calculate historical share returns based on end-of-day market prices. Moreover there is considerable literature on forecasting and trading strategies that are based on daily observations (e.g. Zhong & Enke 2017, Broussard & Vaihekoski 2012, Zhai et al. 2007), which perform differently from those based on weekly or monthly data (e.g. Brown & Warner 1985, Lohpetch & Corne 2010).
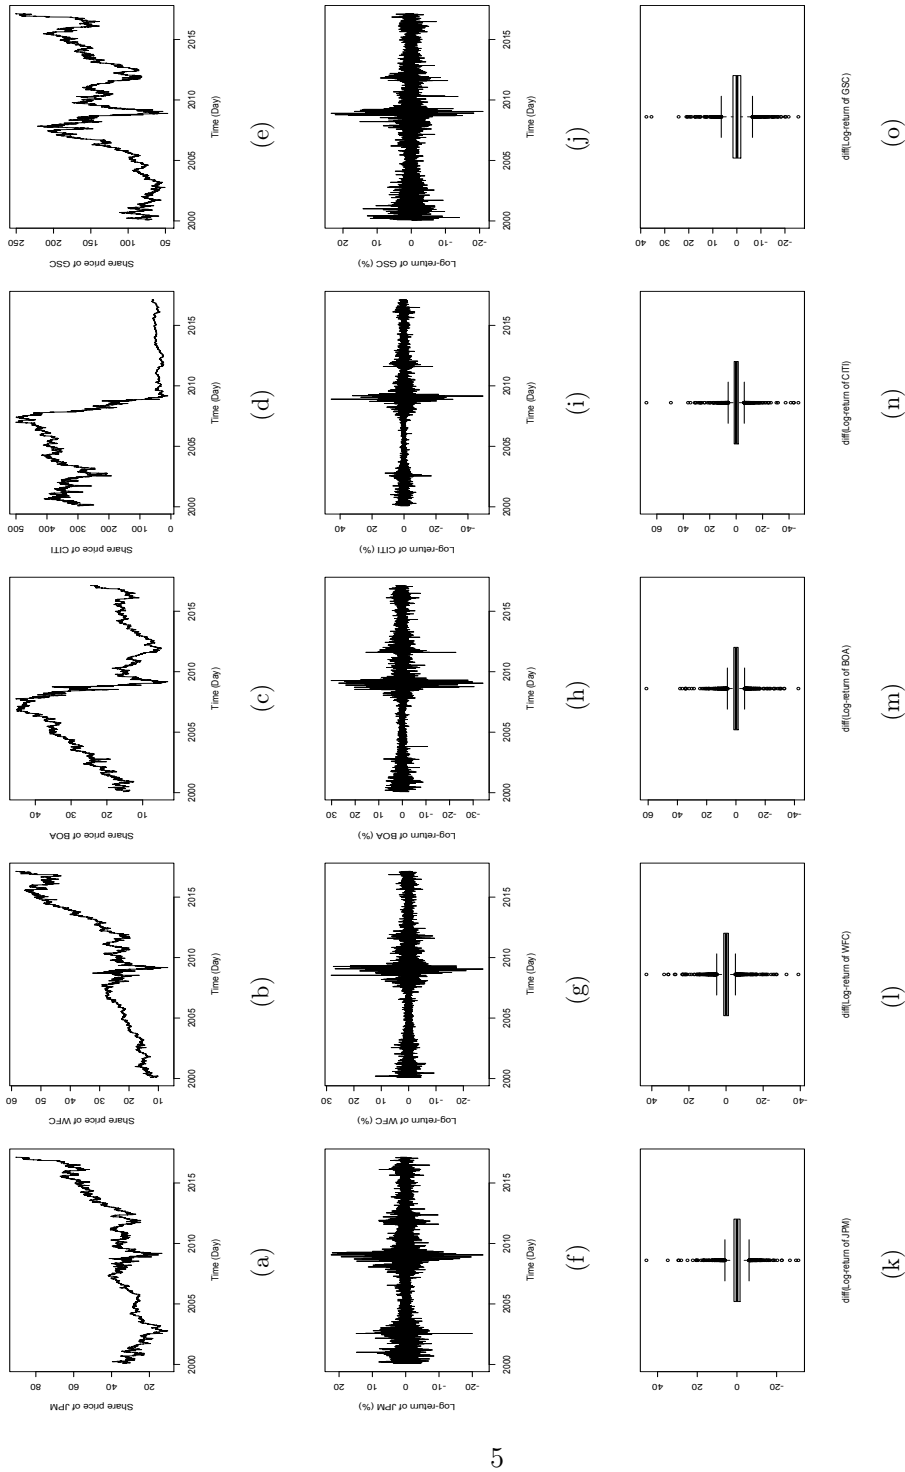
Figure 2: Daily share price, daily log-return and box plot of first differences of the log-returns of: JP Morgan (a, f and k); Wells Fargo (b, g and l); Bank of America (c, h and m); Citibank (d, i and n); and Goldman Sachs (e, j and o), respectively.

The remaining plots in Figure 2 are based on log-return of the series. The log-return time series, $\{x_t\}$, corresponding to the observed share price series for each bank, $\{s_t\}$, with length n is

$$x_t = ln \left\{ \frac{s_t}{s_{t-1}} \right\} \times 100 \qquad (1)$$

where $t = 2, 3, ...n$. The value of n is 4480 trading days in North America, where there are approximately 252 trading days per year (5 per week less 9 public holidays). This transformation will not affect the time reversibility of the process generating the original time series of last traded prices (Wild et al. 2014). These daily log-return series are considered as evenly spaced with no missing values. The five log-return series under consideration can reasonably be considered as realisations of stationary time series models. It is impractical to access directionality in any of the log-return series plots of JP Morgan (Figure 2(f)), Wells Fargo (Figure 2(g)), Bank of America (Figure 2(h)), Citibank (Figure 2(i)), or Goldman Sachs (Figure 2(j)) at this scale.

2.2   Detecting Directionality in log-returns

There are many ways in which a time series can exhibit directionality. Typically there can be a tendency for relatively rapid rises to a peak followed by slower recessions or vice-versa and this can be measured by the skewness of first differences, of the time series (Lawrance 1991). The rationale is that there will be fewer large positive differences and more small negatives differences leading to a positively skewed distribution. However both rapid rises and rapid falls with a slow return to the mean is sometimes apparent (Mansor et al. 2017b) as is localized directionality around extreme values (Soubeyrand et al. 2014).

Here we detect directionality in the log-return time series using the Product Moment Skewness of Differences given by:

$$\hat{\gamma}_d = \frac{\sum_{t=1}^{n}(y_t - \bar{y})^3/(n-1)}{[\sum_{t=1}^{n}(y_t - \bar{y})^2/(n-1)]^{3/2}} \qquad (2)$$

where the first differences (differences) of the log-return series are defined as

$$y_t = x_t - x_{t-1} \quad \text{for} \quad t = 3, 4, ...n. \qquad (3)$$

The $\hat{\gamma}_d$ is generally the most sensitive statistic for detecting directionality in financial time series (e.g. Mansor et al. 2015b, 2017). The statistical significance for the statistic is obtained by Monte Carlo simulations under the null hypothesis of a reversible process (Mansor et al. 2016c,a) and its sensitivity is quantified by the ratio of the observed directionality value to its standard error when sampling from an equivalent reversible process. The statistic has a mean of zero when calculated from a reversible process and if the ratio exceeds 2 the directionality is statistically significant at an approximate 0.05 level. This procedure ignores the sampling variability of the coefficient in the AR(p) but

Table 1: Directionality, and change in standard deviation of errors after fitting AR model, for the log-return time series.

Time series	JP Morgan	Wells Fargo	BankAmerica	Citibank	GoldmanSachs
Directionality, $\hat{\gamma}_d$	0.729	0.365	0.818	0.677	0.534
P-value	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Ratio	[23.2]	[12.1]	[26.3]	[20.4]	[16.6]
Std.Dev,s	2.556	2.463	3.032	3.222	2.399
$\{AR(p), \hat{\sigma}_{err}\}$	{36, 2.523}	{35, 2.392}	{35, 2.943}	{36, 3.163}	{34, 2.378}

Notes: The length of log-return time series from 1-Feb-2000 to 17-Feb-2017 is 4290; s is the marginal standard deviation; p, in the autoregressive process $AR(p)$ is the order of the AR model fitted using the Akaike information criterion (AIC); and $\hat{\sigma}_{err}$ is the estimated standard deviation of residuals in the $AR(p)$ model. (number) refers to P-value and [number] refers to the ratio of the test statistic to its standard deviation from Monte Carlo simulation.

this is negligible for long time series. If $\hat{\gamma}_d$ is statistically significantly different from zero, then we have evidence of skewness of differences and hence of directionality. A limitation of $\hat{\gamma}_d$ is that it is very sensitive to extreme outliers. However the box plots of the log-return differences in Table 2 (third row) while showing outlying values that are generally consistent with positive skewness, do not include any extreme outlying values. Results in Table 1 provide strong evidence of directionality in the log-return time series for all cases.

3 Directionality as an indicator of financial crises

3.1 Directionality and volatility during crisis and non-crisis periods

Allen & Christa (2013) identified seven financial crisis periods in the U.S. since 1990: two banking crises (the credit crunch (1990:Q1 to 1992:Q4) and the subprime lending crisis (2007:Q3 to 2009:Q4)); three market crises (stock market crash (1987:Q4), Russian debt crisis, and Long-Term Capital Management bailout (1998:Q3 to 1998:Q4)); together with dot.com bubble and September 11 attack on the World Trade Center (2000:Q2 to 2002:Q3).

In this paper we combined the dot.com bubble and the September 11 attack with the 2001-2002 recession (Wikipedia 2016*b*) as a sequence of major events that happened in the U.S between 2000 and 2002. We also combine the subprime mortgage crisis with the 2007-2008 GFC, with effects continuing into 2009, the collapse of Lehman Brothers in 2008 and the Great Recession of 2008-2012 (Wikipedia 2016*a*) as a second period of major shocks in the U.S. during 2007 to 2012. Our definition of financially stable and unstable periods is summarized in Table 2. However these periods do not consider any

major events that were specific to particular banks and had a direct impact on share price performance. For example, the effect of post-IPO for Goldman Sachs in 1999, the Bank of America's mergers and acquisitions (M&A) during 2004 to 2006 and the controversies affecting JP Morgan between 2002 and 2017 as well as its major M&A.

Table 2: Financially stable and unstable periods in the U.S. between May-1999 to Feb-2017.

Period	Status
May to Dec-1999	Stable
Jan-2000 to Dec-2002	Unstable
Jan-2003 to Dec-2006	Stable
Jan-2007 to Dec-2012	Unstable
Jan-2013 to Feb-2017	Stable

Summary statistics of the log-returns sub-time series of the JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs corresponding to the stable and unstable periods are given in Table 3(a), 3(b), 3(c), 3(d) and 3(e), respectively. Directionality in Table 3 is measured by the skewness of differences ($\hat{\gamma}_d$) of log-returns over the sub-periods. The volatility is defined by a standard deviation of log-returns over the sub-period:

$$sd = \sqrt{\frac{\sum (x_t - \bar{x}_t)^2}{m}} \tag{4}$$

where the summation is over the sub-period of length m and $\bar{x}_t$ is mean over the sub-period. The skewness is defined by:

$$skewness = \sqrt{\frac{\sum (x_t - \bar{x}_t)^3/m}{sd^3}} \tag{5}$$

and the kurtosis is defined by:

$$kurtosis = \sqrt{\frac{\sum (x_t - \bar{x}_t)^4/m}{sd^4}}. \tag{6}$$

The range is the difference between the highest and lowest log-returns over the sub-period.

The directionality is consistently positive during unstable periods and markedly higher during all of the unstable periods than it is in any of the stable periods. In contrast the volatility, given by standard deviation of log-returns is generally but not consistently higher during the unstable period, than it is during the stable period. For example, volatility during [2000 to 2002] unstable period is lower than [2003 to 2006] stable period for Wells Fargo, the Bank of America as well as for the Goldman Sachs.

Table 3: Summary statistics of the bank log-returns by stable and unstable sub-series.

Period	Mean	Range	Skewness	Kurtosis	Volatility	Directionality
(a) JP Morgan						
Stable	−0.031	13.3	0.31	3.2	2.48	−0.30
Unstable	−0.088	34.9	0.14	6.6	3.10	0.59
Stable	0.085	13.0	0.38	6.8	1.30	−0.11
Unstable	0.003	45.6	0.30	12.2	3.40	0.63
Stable	0.080	15.2	−0.04	6.1	1.32	0.31
(b) Wells Fargo						
Stable	−0.041	14.0	0.15	3.4	2.32	0.11
Unstable	0.028	21.5	0.22	6.1	1.98	0.46
Stable	0.054	7.0	0.13	4.2	0.88	0.04
Unstable	0.007	55.6	0.73	16.0	3.74	0.26
Stable	0.063	12.5	0.15	5.8	1.18	0.05
(c) Bank of America						
Stable	−0.204	13.7	0.27	3.3	2.49	0.03
Unstable	0.059	17.0	0.02	4.0	2.38	0.10
Stable	0.059	13.8	−1.51	19.1	0.95	−0.13
Unstable	−0.091	64.4	−0.21	15.4	4.59	0.64
Stable	0.075	14.5	−0.20	5.1	1.64	0.31
(d) Citibank						
Stable	0.067	10.7	0.31	2.8	2.25	0.08
Unstable	−0.008	29.0	−0.26	6.2	2.69	0.68
Stable	0.059	10.1	0.06	4.8	1.09	−0.04
Unstable	−0.168	95.1	−0.34	23.1	4.85	0.51
Stable	0.042	16.9	−0.33	6.2	1.58	0.09
(e) Goldman Sachs						
Stable	0.343	34.0	3.25	24.6	3.56	−1.75
Unstable	−0.041	28.2	0.27	5.0	3.01	0.15
Stable	0.110	10.4	0.12	3.9	1.38	0.01
Unstable	−0.025	44.5	0.33	13.6	3.06	0.66
Stable	0.070	13.1	−0.29	4.5	1.40	−0.05

9

Directionality is not consistently positive in stable periods and the negative directionality for Goldman Sachs during May to Dec-1999 is also noticeably large in absolute value. The other statistics do not show any consistent pattern between the stable and unstable periods. These results raise the question of whether monitoring directionality, rather than volatility, would give early warning of unstable periods. We investigate this in Section 3.2.

3.2   Moving directionality

A preliminary investigation based on a moving window of 9-month moving directionality (directionality in the log-returns of sub-series of length 189, which is 21 by 9, trading days in a year) over the full-length time series of log-returns. The formula used to calculate the 9-month moving directionality (MD) series is given in Equation (7),

$$\hat{\gamma}_{d,t} = \frac{\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta x}_t)^3/m}{[\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta x}_t)^2/m]^{3/2}}, \tag{7}$$

where $t = 189, ..., 4479$, $m = 189$, $\bar{\Delta x}_t = \sum_{i=0}^{188} \Delta x_{t-i}/m$ and x_t is the full-length time series of log-returns. Similarly for the volatility (σ), given by the marginal standard deviation of log-returns $\{x_t\}$, the 9-month moving volatility (MV) series is calculate using Equation (8).

$$\hat{\sigma}_i = \sqrt{\frac{\sum_{i=0}^{188}(x_{t-i} - \bar{x}_t)^2}{m}}. \tag{8}$$

where $\bar{x}_t = \sum_{i=0}^{188} x_{t-i}/m$.

The plots of the MD series and MV series together with the ratio between the MD and the MV series for the five banks are given in Figure 3. We use vertical lines in the plots to show the dates when the crises are reported in the news. For example, 15-Jan-2001 for the 2001 dotcom crash, 11-Sep-2001 for the 2001 Sept 11 attack on the World Trade Centre. The subprime mortgage crisis was first widely reported on 1-Jul-2007 and the collapse of Lehman Brothers was reported on 15-Sep-2008. No exact dates are found to indicate the recessions but the U.S. unemployment rate rose to 4.2% on 1-Feb-2001, the National Association of Securities Dealers Automated Quotations (NASDAQ) reached a 6-year low on 24-Sept-2002 and U.S. stock market fell dramatically on 1-Aug-2011 for the 2011 recession.

Directionality appears to be more affected by financial instability than the volatility and the moving directionality soars on two occasions during crises periods. In particular, the fourth vertical line of 24-Sept-2002 in Figure 3 (a, c and g), as well as for the second last vertical line of 14-Sept-2008 in Figure 3 (c and e). In contrast, volatility provides less effects for the fourth vertical line on Figure 3 (b and h), and for the second last line on Figure 3 (b, d and f). However, neither moving directionality nor moving volatility give early warning for Goldman Sachs. These obvious occurrences of directionality indicate the potential of directionality to provide early warning of financial crises which would be
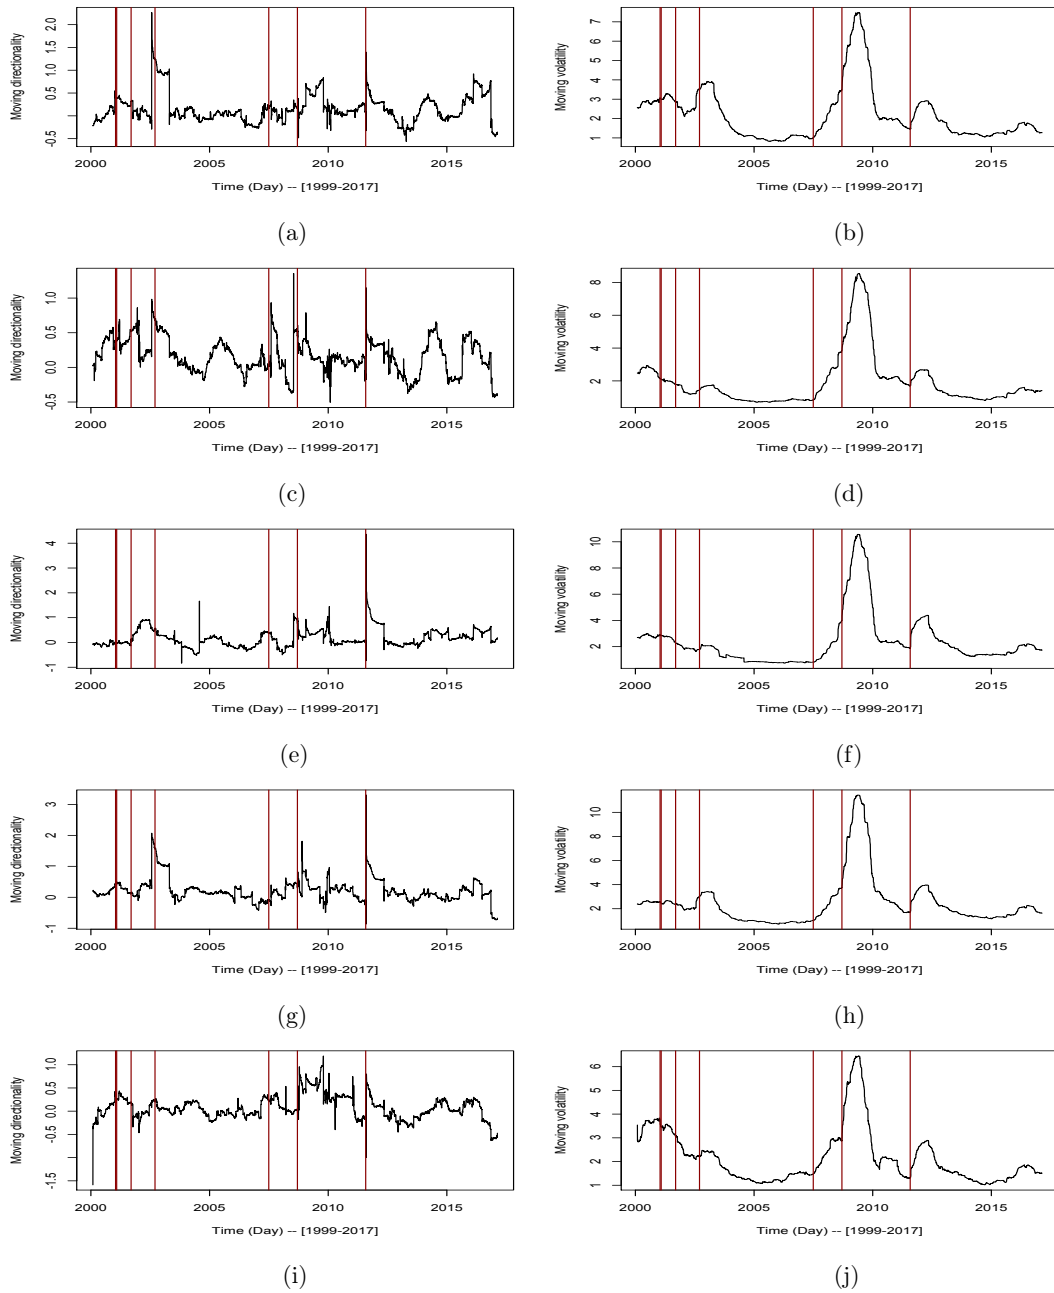
Figure 3: 9-month moving directionality and 9-month moving volatility: JP Morgan (a and b); Wells Fargo (c and d); Bank of America (e and f); Citibank (g and h); and Goldman Sachs (i and j), respectively.

11

useful to prompt investors to liquidate some part of the investment into less or low risk investment options such as mutual funds, certificates of deposit or money market funds, to minimize the potential loss. We have seen similar effects in the moving directionality at 3-months, 6-months and 12-months.

4 Investment simulation

The strategy is to monitor the directionality of the log-returns using 9-month moving directionality (MD) for each bank and to use this as an investment criterion with the aim of benefiting from market uncertainties. We consider a portfolio of shares of JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs. An investment capital of $500,000 is set aside to purchase up to $100,000 ($500,000/5) in each of the banks, in order to spread the risk. We investigated two trading rules for buying and selling the shares:

- Trading rule I:
    If $MD_i > T_i$, then we invest all the money allocated to $Bank_i$ in its shares. Next, if $MD_i \leq T_i$, then we sell the all shares in $Bank_i$; and

- Trading rule II:
    If $MD_i > T_i$, then we invest all the money allocated to $Bank_i$ in its shares. Next, if $MD_i \leq T_i/2$, then we sell the all shares in $Bank_i$

where T_i is set at the overall measure of directionality $\hat{\gamma}_d$ for $Bank_i$ (Table 1, third last column). Whilst the money allocated to $Bank_i$ is initially $100,000 this can increase due to profitable sales or interest payments, or decrease due to unprofitable sales, over time. We considered investing the capital for a period from 1-Feb-2000 to 17-Feb-2017. We checked MD value on a daily basis. We deposited the money in a bank paying 1% per annum (p.a) interest on the first day and buy shares when the MD_i values go above the threshold, hold, and then sell the shares when the MD_i values go below the threshold. The buying and selling processes continued and the revenue after each sell is deposited in the bank while waiting for the next buy. Results of this investment simulation for trading rule I and II are given in Table 4. The amount of money illustrated in Table 4 for each company is the value at the end of the investment period.

The compounded interest (R) is calculated using Equation 9:

$$A = P(1 + R)^{17} \tag{9}$$

where A is the final value of the portfolio and P is the initial investment capital. The strategy using directionality as a criterion in the trading rules provide capital appreciation at 2.43% compounded annually for trading rule I, and 2.15% compounded annually for trading rule II, over a 17-year investment tenure. For a comparison, we used the same

Table 4: Investment performance based on the directionality trading rules for a multi-assets portfolio of JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs from 1-Feb-2000 to 17-Feb-2017.

Equity	Threshold	Trading rule I	Trading rule II
JP Morgan	0.729	$198,261.70	$235,254.20
Wells Fargo	0.365	$84,638.22	$97,553.31
Bank of America	0.818	$227,486.60	$131,800.40
Citibank	0.677	$83,715.46	$59,732.46
Goldman Sachs	0.534	$158,431.10	$193,483.80
Initial investment capital		$500,000.00	$500,000.00
Final value of portfolio		$752,533.10	$717,824.20
Increase/decrease		40.88%	36.16%
Equivalent compounded interest		2.43% p.a	2.15% p.a

trading rules for volatility, where we monitored the volatility of the log-return (s) in Table 1 in the distribution of 9-month moving volatility (MV) for each company and the T_i is set at the overall volatility, given by s, for $Bank_i$ (Table 1, first column). The results are given in Table 5, but are not encouraging in this case, at least.

Table 5: Investment performance based on the volatility trading rules for the portfolio.

Equity	Threshold	Trading rule I	Trading rule II
JP Morgan	2.556	$93,403.22	$105,412.50
Wells Fargo	2.463	$182,653.40	$186,274.70
Bank of America	3.032	$62,901.37	$64,040.57
Citibank	3.222	$37,228.97	$57,351.92
Goldman Sachs	2.399	$109,892.60	$113,608.80
Initial investment capital		$500,000.00	$500,000.00
Final value of portfolio		$486,079.60	$526,688.50
Increase/decrease		−2.82%	5.20%
Equivalent compounded interest		−0.17% p.a	0.31% p.a

This simulation study is highly simplified but it suggests that monitoring directionality may have value as an aid to investment and merits further study. Instead of entering into the money market, we could also consider several other asset classes that are easy to liquidate with less impact on the investment value. We also noted that switching the investment from shares into other investment instruments or between shares involve with

multiple broker′s commissions and switching fees that are relevant to take into account when calculating the return on investment (ROI). Furthermore, there is a yearly dividend cash payout to investors who stay with the company during bullish or even bearish periods that can be reinvested into shares. The real ROI should consider the revenue minus the capital, deducting broker′s commissions, any switching fees and taxes on investment plus dividends.

Note that we buy shares when the MD value is more than the $\hat{\gamma}_d$ because from our observation, on many occasions the share price is low when the corresponding MD value is higher than the threshold value. In contrast, the share price is high when the corresponding MD value is lower than the threshold value. In particular for the case of JP Morgan and Goldman Sachs. This is consistent with the famous "buy low, sell high" strategy to take advantage of the market instability. Also note that we are lowering the threshold value by 50% for the trading rule II.

Table 6: Portfolio performance comparison between trading rules based on the directionality and the volatility for the 3-month, 6-month and 9-month series.

Series	Trading rule I	Trading rule II
3-month MD	2.04% p.a	3.31% p.a
3-month MV	0.77% p.a	0.97% p.a
6-month MD	3.00% p.a	2.48% p.a
6-month MV	1.05% p.a	0.51% p.a
9-month MD	2.43% p.a	2.15% p.a
9-month MV	−0.17% p.a	0.31% p.a

We applied the directionality and the volatility trading rules I and II for the moving directionality and for the moving volatility at 3-month and 6-month. The results, together with the 9-month MD, are given in Table 6. The portfolio performances are compared by the compounded annualised return (% p.a). Both directionality trading rules provide better results than the volatility rules in all cases. Next, we varied the threshold values by ±5% and ±10% for the directionality trading rule I and II. Their performances in Table 7 are consistently profitable. In this case, the 3-month MD provides better results than the others, and directionality trading rule II is generally performs better than rule I.

We applied the directionality trading rules I and II for five non-U.S. banks: Industrial & Commercial Bank of China (ICBC) (27-Oct-2006 to 20-Feb-2017); HSBC Holdings United Kingdom (HSBC U.K.) (13-Jul-1992 to 20-Feb-2017); Commonwealth Bank of Australia (CBA) (13-Sep-1991 to 20-Feb-2017); Royal Bank of Canada (RBC) (3-Jan-1975 to 17-Feb-2017); and Mitsubishi UFJ Financial Group of Japan (MUFG) (3-Apr-2001 to 20-Feb-2017). These banks are the largest bank in the country and are ranked

Table 7: Portfolio performance based on the directionality trading rules for 3-month, 6-month and 9-month MD.

Threshold	$T-10\%$	$T-5\%$	T	$T+5\%$	$T+10\%$
		Trading rule I			
3-month MD	3.31% p.a	2.85% p.a	2.04% p.a	2.17% p.a	2.30% p.a
6-month MD	2.89% p.a	3.03% p.a	3.00% p.a	1.98% p.a	3.19% p.a
9-month MD	3.08% p.a	3.04% p.a	2.43% p.a	2.02% p.a	1.64% p.a
		Trading rule II			
3-month MD	4.60% p.a	4.19% p.a	3.31% p.a	2.76% p.a	2.34% p.a
6-month MD	1.87% p.a	1.85% p.a	2.48% p.a	2.49% p.a	1.98% p.a
9-month MD	4.27% p.a	3.58% p.a	2.15% p.a	1.99% p.a	1.40% p.a

among world's largest banks in 2017 in terms of market capitalization (relbanks 2017). An investment capital of \$100,000 is set aside to purchase shares in each bank. Directionality in the log-returns of each bank is given by the $\hat{\gamma}_d$. Detailed results are illustrated in Table 8.

Table 8: Single asset investment performance based on the directionality trading rules for non-U.S. log-returns using 9-month MD.

Equity	Country	Ranking	Length	Directionality	Trading rule I	Trading rule II
ICBC	China	3	2352	0.593	4.02% p.a	2.80% p.a
HSBC	U.K.	6	6031	0.090	9.60% p.a	10.7% p.a
CBA	Australia	10	6241	0.046	8.38% p.a	8.54% p.a
RBC	Canada	11	10408	0.121	3.11% p.a	3.50% p.a
MUFG	Japan	14	3705	0.038	5.70% p.a	4.76% p.a

All the log-returns series in Table 8 are directional. The investment returns for all series show the potential of directionality trading rules for single asset investment. Nevertheless, there is a criterion for choosing an investment strategy between the directionality trading rule I and II, by using the conditional value at risk (CVaR) (Rockafellar & Uryasev 2002) and we investigate this in Section 5.4.

5   Modelling directionality for JP Morgan

The existence of directionality also suggests that stationary non-linear time series models or linear models driven by non-Gaussian white noise are appropriate for modelling the log-return series.

5.1   Model Fitting

We fitted the JP Morgan time series of log-returns from 1-Feb-2000 to 17-Feb-2017 with the first order autoregressive AR(1) model given in Equation 10, first order threshold autoregressive model TAR(1) with one threshold (T) given in Equation 11 and first order threshold autoregressive models TAR(1) with two thresholds (T_L, T_U) (TAR(1)2T) in Equation 12, where the $\epsilon_t \sim N(0, \sigma^2)$.

$$X_t = \phi_0 + \phi_1 X_{t-1} + \epsilon_t \ . \tag{10}$$

$$X_t = \begin{cases} \phi_{0U} + \phi_{1U} X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} > T \\ \phi_{0L} + \phi_{1L} X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} \leq T \end{cases} \ . \tag{11}$$

$$X_t = \begin{cases} \phi_{0U} + \phi_{1U} X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} > T_U \\ \phi_{0M} + \phi_{1M} X_{t-1} + \epsilon_t & \text{if} \quad T_L < X_{t-1} \leq T_U \\ \phi_{0L} + \phi_{1L} X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} \leq T_L \end{cases} \ . \tag{12}$$

Table 9: Time series models for the fitted JP Morgan log-returns compared by estimated standard deviation of errors $\hat{\sigma}_{err}$.

AR(0)	AR(1)	TAR(1)	TAR(1)2T	AR(36 by AIC)
2.5560	2.5483	2.5452	2.5345	2.5227

In Table 9, the marginal standard deviation of the fitted JP Morgan log-returns series is 2.5560, and the $\hat{\sigma}_{err}$ of AR(36 by AIC) is 2.5227. Although the models provide statistically significant improvements, the decrease in the estimated standard deviation of the errors is slight.

Furthermore, there is clear evidence of volatility in the log-return of JP Morgan given by the sample auto-correlogram function of the squared log-returns and this feature can be modelled using the Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) model by Nelson (1991). GARCH is an improved model in stochastic variance modelling to the Auto-Regressive Conditional Heteoskedastic (ARCH) model developed by Engle (1982). The models are given in Equation (13) and (14) respectively.

$$ARCH(q) \begin{cases} \sigma_t^2 = \omega + \sum_{j=1}^{q} \alpha_j \epsilon_{t-1}^2 \,, \\ \epsilon_t = \sigma_t z_t \,, \\ z_t \sim N(0,1). \end{cases} \tag{13}$$

$$GARCH(p,q) \begin{cases} \sigma_t^2 = \omega + \sum_{i=1}^{p} \beta_i \sigma_{t-1}^2 + \sum_{j=1}^{q} \alpha_j \epsilon_{t-1}^2 \,, \\ \epsilon_t = \sigma_t z_t \,, \\ z_t \sim N(0,1). \end{cases} \tag{14}$$

To ensure that the conditional variance σ_t^2 is a asymptotically stationary random sequence, the $\sum_{i=1}^{p} \beta_i + \sum_{j=1}^{q} \alpha_j$ in Equation 14, and $\sum_{j=1}^{q} \alpha_j$ in Equation 13, should not exceed 1 (e.g. Awartani & Corradi 2005).

The choice of values for p and q for the GARCH(p,q) model can be optimised but the GARCH(1,1) specification has proven to be an adequate representation for most financial time series (Lamoureux & Lastrapes 1990). Therefore, we fitted the residuals with GARCH(1,1) and compared its performance with ARCH(1). The AR(1)-ARCH(1), AR(1)-GARCH(1,1), TAR(1)-GARCH(1,1) and TAR (1)2T-GARCH(1,1) models can be expressed as follows:

AR(1)-ARCH(1):
$$\begin{cases} X_t = \phi_0 + \phi_1 X_{t-1} + \epsilon_t \,, \\ \epsilon_t = \sigma_t z_t \,, \\ \sigma_t^2 = \omega + \alpha_1 \epsilon_{t-1}^2 \,, \\ z_t \sim N(0,1). \end{cases} \tag{15}$$

AR(1)-GARCH(1,1):
$$\begin{cases} X_t = \phi_0 + \phi_1 X_{t-1} + \epsilon_t \,, \\ \epsilon_t = \sigma_t z_t \,, \\ \sigma_t^2 = \omega + \beta_1 \sigma_{t-1}^2 + \alpha_1 \epsilon_{t-1}^2 \,, \\ z_t \sim N(0,1). \end{cases} \tag{16}$$

TAR(1)-GARCH(1,1):
$$\begin{cases} X_t = \begin{cases} \phi_{0U} + \phi_{1U} X_{t-1} + \epsilon_t & \text{if } X_{t-1} > T \\ \phi_{0L} + \phi_{1L} X_{t-1} + \epsilon_t & \text{if } X_{t-1} \leq T \end{cases} \,, \\ \epsilon_t = \sigma_t z_t \,, \\ \sigma_t^2 = \omega + \beta_1 \sigma_{t-1}^2 + \alpha_1 \epsilon_{t-1}^2 \,, \\ z_t \sim N(0,1). \end{cases} \tag{17}$$

17

TAR(1)2T-GARCH(1,1):

$$
\begin{cases}
X_t = \begin{cases}
\phi_{0U} + \phi_{1U}X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} > T_U \\
\phi_{0M} + \phi_{1M}X_{t-1} + \epsilon_t & \text{if} \quad T_L < X_{t-1} \le T_U \\
\phi_{0L} + \phi_{1L}X_{t-1} + \epsilon_t & \text{if} \quad X_{t-1} \le T_L
\end{cases} , \\
\epsilon_t = \sigma_t z_t , \\
\sigma_t^2 = \omega + \beta_1 \sigma_{t-1}^2 + \alpha_1 \epsilon_{t-1}^2 , \\
z_t \sim N(0,1).
\end{cases}
\tag{18}
$$

Table 10: Estimated model coefficients for JP Morgan log-return series.

Model	Estimated parameters
AR(1)	$\hat{\phi}_0{=}0.025$, $\hat{\phi}_1{=}{-}0.079$
AR(1)-ARCH(1)	$\hat{\phi}_0{=}0.025$, $\hat{\phi}_1{=}{-}0.079$ $\hat{\omega}{=}3.398$ $\hat{\alpha}_1{=}0.535$
AR(1)-GARCH(1,1)	$\hat{\phi}_0{=}0.025$, $\hat{\phi}_1{=}{-}0.079$ $\hat{\omega}{=}0.026$, $\hat{\beta}_1{=}0.912$, $\hat{\alpha}_1{=}0.086$
TAR(1)-GARCH(1,1)	$\hat{\phi}_{0L}{=}{-}0.052$, $\hat{\phi}_{0U}{=}{-}0.009$ $\hat{\phi}_{1L}{=}{-}0.138$, $\hat{\phi}_{1U}{=}{-}0.028$ $\hat{\omega}{=}0.025$, $\hat{\beta}_1{=}0.915$, $\hat{\alpha}_1{=}0.083$
TAR(1)2T-GARCH(1,1)	$\hat{\phi}_{0L}{=}{-}2.384$, $\hat{\phi}_{0M}{=}0.027$, $\hat{\phi}_{0U}{=}{-}0.343$ $\hat{\phi}_{1L}{=}{-}0.463$, $\hat{\phi}_{1M}{=}{-}0.064$, $\hat{\phi}_{1U}{=}0.016$ $\hat{\omega}{=}0.025$, $\hat{\beta}_1{=}0.915$, $\hat{\alpha}_1{=}0.083$

The threshold, T, for TAR(1) was taken as the 0.80 quantile of the marginal distribution of the JP Morgan log-return time series. The choice was based on the minimum sum of squared residuals over the quantile values: 0.70, 0.75, 0.80, 0.85, 0.90 and 0.95. The T_L and T_U for TAR(1)2T were taken as 0.05 and 0.90 quantile respectively, based on the minimum sum of squared residuals over the same upper quantile values and 0.05, 0.10, 0.15, 0.20, 0.25, 0.30 for the lower quantiles. Constraints for stability for the TAR(1) is $-1 < \phi_{1L}, \phi_{1U} < 1$; and for the TAR(1)2T is $-1 < \phi_{1L}, \phi_{1M}, \phi_{1U} < 1$ (Chatfield 2004). All the estimated model parameters values in Table 10 satisfy the conditions for stationarity.

5.2   Model validation

The simulated series of length 10^5 by the AR(1)-GARCH(1,1) (Figure 4(c)) and by the TAR(1)2T-GARCH(1,1) (Figure 4(d)) appear to be far more realistic than the simulated

series given by AR(1)-ARCH(1) in (Figure 4(b)). The fluctuation in the variability of the log-returns (Figure 4(a)) is reasonably emulated by the GARCH model. A quantitative comparison including a measure of directionality is given in Table 11.

Table 11: Statistics in the fitted log-return of JP Morgan and in the simulated series.

Series	Length	Mean	Std.Dev	Directionality
JP Morgan	4290	0.024	2.56	0.73
AR(1)	10^5	0.017	2.57	0.01
AR(1)-ARCH(1)	10^5	0.018	2.70	-0.03
AR(1)-GARCH(1,1)	10^5	0.010	4.26	-0.06
TAR(1)	10^5	0.049	2.57	0.07
TAR(1)-GARCH(1,1)	10^5	0.050	4.03	0.10
TAR(1)2T-GARCH(1,1)	10^5	0.065	4.45	0.72

Directionality as measured by the skewness of differences $\hat{\gamma}_d$, is obtained in the simulated series as a consequences of the non-linear TAR(1) models together with a GARCH model for the errors, rather than a linear AR(1) model with Gaussian error. This is expected because the process is close to white noise (the autoregressive coefficient is close to zero) and white noise is a reversible process for any distribution of errors. The GARCH(1,1) model does emulate volatility (Figure 4(c and d) and better than the AR(1) (Figure 4(b)) for the simulations of length 10^5. It seems that 2 thresholds and GARCH are needed to model directionality but the standard deviation of the simulated series seems too high. In the next section, we investigate whether the use of penalized least squares can improve the fit in terms of standard deviation.

5.3 TAR model with penalized least square

There scope for more realistic modelling of directionality. This can be achieved by using a penalty for deviations from the observed directionality measures when fitting the TAR model. The strategy is to use directionality as a fitting criterion where the penalty function in the optimization criteria is a weighted sum of the error sum of squares and a measure of discrepancy between the observed directionality and the simulated directionality from long simulations (Mansor et al. 2016a). The objective function to be minimized is

$$\omega = \sum_{t=p+1}^{n} r_t^2 + \theta_1(\hat{\gamma}_{simulated} - \hat{\gamma}_{observed})^2 \tag{19}$$

where $\{r_t\}$ are the residuals from TAR model, p is the order of the autoregressive part of the model and θ_1 is the weight given to agreement between the $\hat{\gamma}_{simulated}$ and the target directionality value $\hat{\gamma}_{observed}$. The product moment skewness of differences ($\hat{\gamma}_d$) is used in Equation 19 as it is the most sensitive test for directionality.

19

(a) The fitted JP Morgan log-return series

(b) Simulated series by the AR(1)-ARCH(1) model

(c) Simulated series by the AR(1)-GARCH(1,1) model

(d) Simulated series by the TAR(1)2T-GARCH(1,1) model
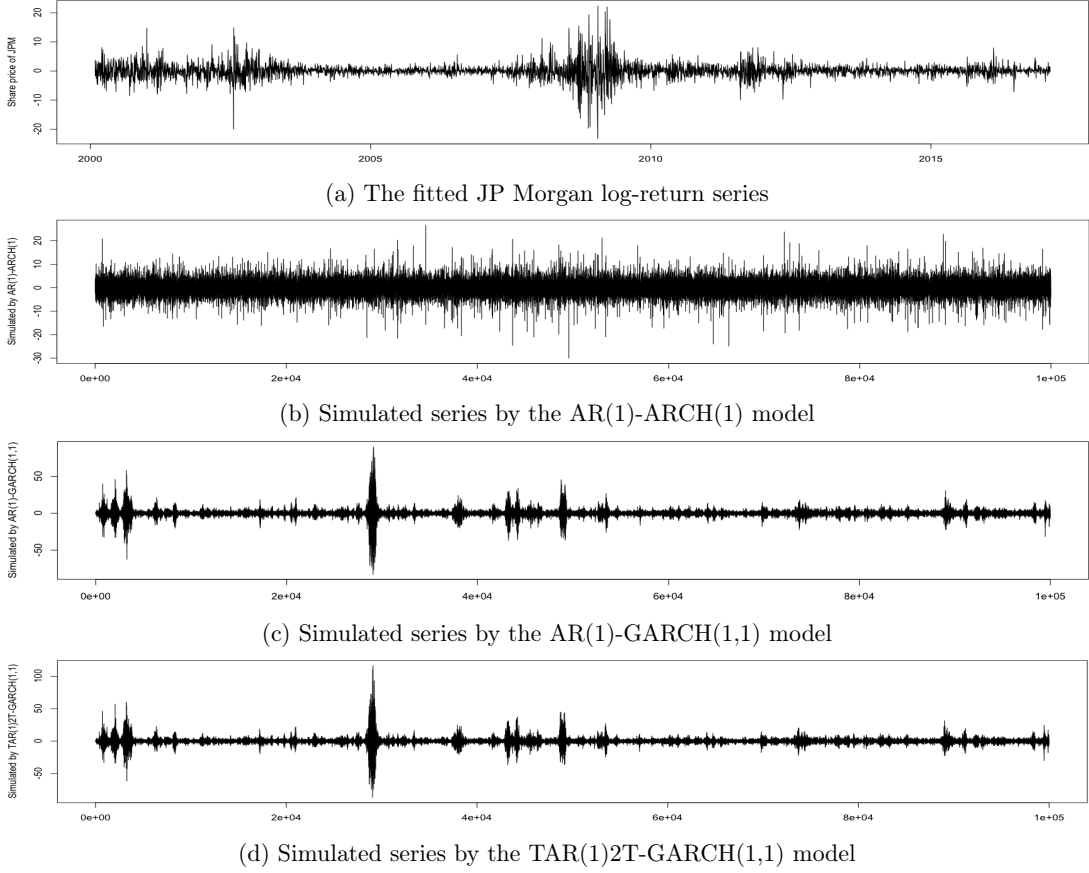
Figure 4: Qualitative comparison between series from the simulations and the observed log-return of JP Morgan.

Here, we demonstrate how the penalized least squares can be used in model fitting to mitigate directionality discrepancies for JP Morgan log-return series. We fit the first 75% of the series with the AR(1), TAR(1) and TAR(1)2T, and the target directionality ($\hat{\gamma}_{observed}$) for the fitted series is 0.6534. Next, we fit the TAR(1) and the TAR(1)2T models by the penalized least squares (PLS) procedure (TAR(1)[LSP] model and TAR(1)2T[LSP] model respectively) using the *R* function *optim()* which uses the Nelder-Mead algorithm (Nash 2014) to optimize the parameter values for the models. If the procedure is found to give lower marginal standard deviations $\hat{\sigma}_{simulated}$ than the standard deviation of the observed time series ($\hat{\sigma}_{observed}$), then modify Equation 19 to Equation 20:

$$\omega = \sum_{t=p+1}^{n} r_t^2 + \theta_1(\hat{\gamma}_{observed} - \hat{\gamma}_{simulated})^2 + \theta_2(\hat{\sigma}_{observed} - \hat{\sigma}_{simulated})^2 \qquad (20)$$

20

where θ_2 are the weight given to mitigate the discrepancy in $\hat{\sigma}$. Detailed results are given in Table 12 and Table 13. Note that TAR(1)[PLS_sd] is the TAR(1) model fitted by PLS in Equation 20.

Table 12: Fitting TAR models with PLS: the simulated directionality, $\hat{\gamma}_{d_{sim}}$, and $\hat{\sigma}_{err}$ for selected θ_1 values. For the TAR(1)[PLS_sd], the θ_2 is set at 10^3. The target $\hat{\gamma}_d$ is 0.6534; and the $\hat{\sigma}_{err}$ for TAR(1) is 2.8674, and for TAR(1)2T is 2.8658.

θ_1	TAR(1)[PLS] $\hat{\gamma}_{d_{sim}}$	$\hat{\sigma}_{err}$	TAR(1)2T[PLS] $\hat{\gamma}_{d_{sim}}$	$\hat{\sigma}_{err}$	TAR(1)[PLS_sd] $\hat{\gamma}_{d_{sim}}$	$\hat{\sigma}_{err}$
0	0.1431	2.8669	0.1609	2.8654	0.1449	2.8674
10^1	0.1434	2.8669	0.1612	2.8654	0.1450	2.8674
10^3	0.1787	2.8678	0.2219	2.8690	0.1456	2.8674
10^4	0.3374	2.8994	0.4165	2.9017	0.3718	2.9091
10^5	0.5776	3.0218	0.6344	2.9645	0.5603	3.0076
10^6	0.6461	3.0686	0.6485	2.9706	0.6529	3.0573
10^7	0.6520	3.0736	0.6534	2.9693	0.6518	3.0732
10^{10}	0.6534	3.0815	0.6534	2.9734	0.6534	3.0815

TAR models with PLS provide improved approximation to the target directionality of 0.6534. The $\hat{\gamma}_{d_{sim}}$ value increases as the value of ϕ_1 increase from 0 to 10^{10}. However the substantial reduction in the difference between the target directionality and the simulated directionality could be achieved with a relatively small increase in the sum of squared residuals (e.g. Mansor et al. 2016a).

Table 13: Model validation: comparing the statistics in the in-sample log-return of JP Morgan with the simulated JP Morgan series.

Series	θ_1	θ_2	Length	Mean	Std.Dev	Directionality
JP Morgan			3212	-0.012	2.882	0.6534
AR(1)-GARCH(1,1)			10^5	0.001	2.463	0.0659
TAR(1)-GARCH(1,1)			10^5	0.024	2.441	0.1429
TAR(1)2T-GARCH(1,1)			10^5	0.039	2.556	0.1608
TAR(1)2T[PLS]-GARCH(1,1)	10^2		10^5	0.043	2.556	0.1659
TAR(1)2T[PLS]-GARCH(1,1)	10^3		10^5	0.110	2.540	0.2219
TAR(1)[PLS]-GARCH(1,1)	10^4		10^5	0.092	2.613	0.3374
TAR(1)[PLS_sd]-GARCH(1,1)	10^4	10^3	10^5	0.110	2.641	0.3718

In Table 13, TAR(1)[PLS_sd] offers a potential improvement over TAR(1)[PLS] in particular for the simulated directionality and the simulated standard deviation. However

the improvement over the standard deviation is negligible with the $\hat{\sigma}_{err}$ of 2.91 compared to the overall results given by TAR(1)[PLS].
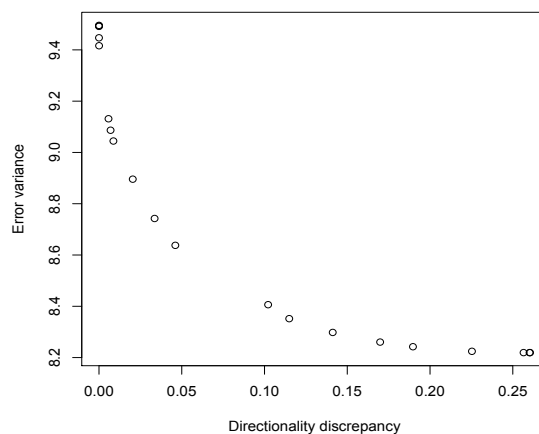
Figure 5: Fitting TAR(1)[PLS] to the in-sample JP Morgan log-returns: trade-off between minimizing the sum of squared residuals, and minimizing the directionality discrepancy.

We illustrate the relationship between the criteria in Equation 19 in Figure 5. The $\hat{\sigma}_{error}$ is monotonically increasing with the increase of the θ_1 for θ_1=0,10^1,10^2,...,10^{10} which corresponds to the values in Table 12 for TAR(1)[PLS].

5.4   Calculating CVaR

We simulate 5-year realisations of JP Morgan daily share prices using the TAR(1)2T-GARCH(1,1) model (Table 10, last row) for 1000 random occasions. In every occasion, we apply the directionality trading rule I and II, and we calculate the compounded annualised returns (% p.a) respectively. We calculate the CVaR(10%) (which is the mean of the valuations that are less than the lower 10% quantile) from a distribution of the simulated annualised returns. Results are given in Table 14, and the boxplot comparing the two trading rules are illustrated in Figure 6.

Table 14: CVaR, mean and standard deviation of a distribution of the simulated annualised returns based on the directionality trading rule I and II for JP Morgan.

Statistics	Trading rule I	Trading rule II
CVaR(10%)	−18.12% p.a	−18.19% p.a
Mean	3.03% p.a	3.33% p.a
Std.Dev	15.31	15.45

(a) Trading rule I
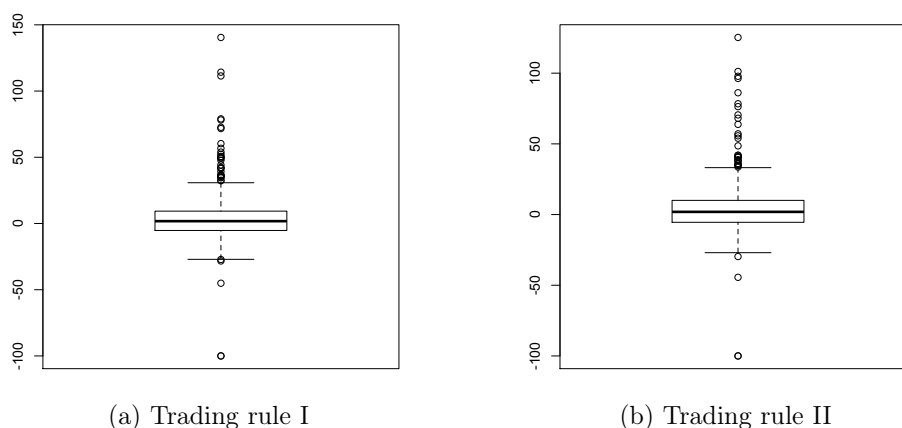
(b) Trading rule II

Figure 6: Boxplot of 5-year investment returns (% p.a) from the simulated series, by the TAR(1)2T-GARCH(1,1) model, based on investment simulations using directionality trading rule I and II with the 3-month moving directionality series.

Both directionality trading rules I and II were highly variable with CVaR(10%) of −18% p.a but the trading rule II had the higher mean value of 3.33.

6 Concluding remarks

We have investigated directionality measured by skewness of first differences in the time series of log-returns of five major U.S. banks. The study period was divided into five sub-periods which we defined as either relatively financially stable or financially unstable. The directionality was markedly high and positive during unstable periods, corresponding to a tendency for rapid increases in the log-returns followed by slower recessions. Although volatility also tends to increase during unstable periods, the relationship between directionality and volatility is not high and directionality appears to be more sensitive to the distinction between stable and unstable periods.

We defined a moving directionality over a 9 month period and found that it gave some early warning of an unstable period. This is quite plausible as it is likely that investors

and market traders would have been alerted to early signs of possible market instability and would have traded accordingly.

The incorporation of directionality into a share trading rules shows some promise and a more thorough appraisal requires stochastic simulation to investigate financial measures such as CVaR. We have fitted stochastic models for the log-returns based on threshold autoregressive process augmented with GARCH for the JP Morgan series. We have demonstrated a comparison of two directionality based trading rules for JP Morgan shares using the mean log-return and CVaR(10%) as criteria.

References

Allen, N. B. & Christa, H. S. B. (2013), 'How does capital affect bank performance during financial crises?', *Journal of Financial Economics* **109**(1), 146–176.

Awartani, B. M. & Corradi, V. (2005), 'Predicting the volatility of the {S&P}-500 stock index via {GARCH} models: the role of asymmetries', *International Journal of Forecasting* **21**(1), 167 – 183.

Broussard, J. P. & Vaihekoski, M. (2012), 'Profitability of pairs trading strategy in an illiquid market with multiple share classes', *Journal of International Financial Markets, Institutions and Money* **22**(5), 1188 – 1201.

Brown, S. J. & Warner, J. B. (1985), 'Using daily stock returns', *Journal of Financial Economics* **14**(1), 3 – 31.

Chatfield, C. (2004), *The Analysis of Time Series: An Introduction, Sixth Edition*, Chapman & Hall/CRC.

Dell'Ariccia, G., Detragiache, E. & Rajan, R. (2008), 'The real effect of banking crises', *Journal of Financial Intermediation* **17**(1), 89 – 112. Financial Contracting and Financial System Architecture.

Engle, R. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Keller, L. J. & Chiglinsky, K. (Sept. 13, 2016), 'Wells fargo eclipsed by {JP M}organ as world's most valuable bank', http://www.bloomberg.com (last accessed 5 October 2016).

Lamoureux, C. G. & Lastrapes, W. D. (1990), 'Persistence in variance, structural change, and the garch model', *Journal of Business & Economic Statistics* **8**(2), 225–234.

Lawrance, A. J. (1991), 'Directionality and reversibility in time series', *International Statistical Review/Revue Internationale de Statistique* **59**(1), 67–79.

Lohpetch, D. & Corne, D. (2010), *Outperforming Buy-and-Hold with Evolved Technical Trading Rules: Daily, Weekly and Monthly Trading*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 171–181.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2016a), *Time Series Analysis and Forecasting: Contributions to Statistics*, Springer International Publishing Switzerland, chapter Threshold Autoregressive Models for Directional Time Series.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2015b), Modelling and simulation of directional financial time series, *in* 'Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)', pp. 1022–1028.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2016b), 'Directionality and volatility in electroencephalogram time series', *American Institute of Physics (AIP) Conf. Proc.* **1739**, 020080:1–8.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2017), 'Directionality and volatility in high frequency time series', *High Frequency* submitted: 9 June 2017.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2017b), 'Detecting directionality in time series', *The American Statistician* submitted: 12 May 2017.

Mansor, M. M., Mohd. Isa, F. L., Green, D. A. & Metcalfe, A. V. (2016c), 'Modelling directionality for paleoclimatic time series', *ANZIAM J.* **57** (EMAC2015), C66–C81.

Nash, J. C. (2014), 'On best practice optimization methods in r', *Journal of Statistical Software* **60**(2), 1 – 14.

Nelson, D. B. (1991), 'Conditional heteroskedasticity in asset returns: A new approach', *Econometrica* **59**, 347–370.

relbanks (2017), 'World's largest banks 2017', http://www.relbanks.com/worlds-top-banks/market-cap (last accessed 24 February 2017).

Rockafellar, R. T. & Uryasev, S. (2002), 'Conditional value-at-risk for general loss distributions', *Journal of Banking & Finance* **26**(7), 1443–1471.

Soubeyrand, S., Morris, C. E. & Bigg, E. K. (2014), 'Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data', *Environmental Modelling & Software* **61**, 78–86.

Tong, H. & Lim, K. S. (1980), 'Threshold autoregression, limit cycles and cyclical data', *Journal of the Royal Statistical Society. Series B (Methodological)* **42**(03), 245–292.

Tsay, R. S. (2005), *Analysis of Financial Time Series, Second Edition*, Wiley Series in Probability and Statistics, John Wiley & Sons Inc.

Weiss, G. (1975), 'Time-reversibility of linear stochastic processes', *Journal of Applied Probability* **12**(4), 831–836.

Wikipedia (2016*a*), 'Financial crisis of 2007–2008', https://en.wikipedia.org (last accessed 11 November 2016).

Wikipedia (2016*b*), 'List of economic crises', https://en.wikipedia.org (last accessed 5 October 2016).

Wikipedia (2016*c*), 'Subprime mortgage crisis', https://en.wikipedia.org (last accessed 11 November 2016).

Wikipedia (2017), 'Trading day', https://en.wikipedia.org (last accessed 25 March 2017).

Wild, P., Foster, J. & Hinich, M. J. (2014), 'Testing for non-linear and time irreversible probabilistic structure in high frequency financial time series data', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **177**(3), 643–659.

Zhai, Y., Hsu, A. & Halgamuge, S. K. (2007), *Combining News and Technical Indicators in Daily Stock Price Trends Prediction*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1087–1096.

Zhong, X. & Enke, D. (2017), 'Forecasting daily stock market return using dimensionality reduction', *Expert Systems with Applications* **67**, 126 – 139.

4.4 Directionality and Volatility in High Frequency Time Series

Directionality and Volatility
in High Frequency Time Series

Mahayaudin M. Mansor,* David A. Green and Andrew V. Metcalfe

School of Mathematical Sciences, University of Adelaide, SA 5005, Australia

Abstract

We provide empirical evidence of directionality in high frequency multivariate time series of the five largest U.S banks between 1999 to 2017. The directionality is more apparent during crisis periods than during non-crisis periods, and it has only a low association with volatility. We use directionality and volatility as a regime-switching criterion between two-regime threshold vector autoregressive (TVAR) models for forecasting share prices. We compare the forecasting performances using mean relative error squared, and a weighted-average of the forecasting error, with weights based on the estimated conditional variance, for individual model components and as a group. We have demonstrated that moving directionality can provide early warning of increased volatility and crisis periods, and has potential for improving one-step ahead forecasts using TVAR(1) models.

Keywords: Multivariate forecasting, reversibility, irreversibility, TVAR-GARCH model, regime-switching criterion, bank share prices.

---

*Correspondence to: mohdmahayaudin.mansor@adelaide.edu.au

1 INTRODUCTION

High frequency financial time series can refer to observations recorded daily or sub-daily, at time intervals that can be as small as seconds, for transactions in security markets (e.g. Cont 2001, Yan et al. 2003, Chavez-Demoulin et al. 2005, Chavez-Demoulin & McGill 2012). Analysing such data presents new challenges because high frequency time series have some unique characteristics that are not seen at lower frequencies (Tsay 2005). These characteristics include, but are not limited to, non-synchronous trading, bid-ask spread, duration models, and models associated with volatility. The study of high frequency time series is also important for various issues related to trading process and market micro-structure. For example, it allows researchers to compare the efficiency of different trading system in price discovery; to study the dynamics of bid and ask quotes of a particular security market; to examine the order dynamics; and to determine the market liquidity provider (Tsay 2005).

We analyse day-to-day closing share prices for the five largest U.S. banks over a 17 year period, and intraday records down to a one-minute sampling interval for the same banks. We focus on detecting a characteristic called directionality in daily and intraday log-return time series. If a time series is directional it is in principal possible to discern qualitative differences between the time series plotted in time order and the time series plotted in reverse time order (time-to-go) in long records. Trends and asymmetric seasonal effects, which are often found in financial time series, will impart directionality. But, trend and seasonal effects are typically modelled as deterministic processes, and are estimated and removed at the start of the analysis. Stochastic trends are a feature of non-stationary models, such as the random walk, which is a common model for financial time series, can be removed by differencing. So, it is more useful to consider directionality in the context of first order stationary time series models so as to distinguish it from trends and seasonality. We do not require second order stationary as financial time series often exhibit volatility, periods of increased variability, which can be modelled with Generalized Auto-Regressive Conditional Heteroskedasticity (GARCH) models.

In this paper we refer to the following definition of stationary. If a first order stationary random process $X(t)$ has the property that $\{X(t_1), ..., X(t_n)\}$ and $\{X(-t_1), ..., X(-t_n)\}$ have the same joint probability distribution for every n, and every $t_1, ..., t_n$, then the process is reversible (symmetry in time) (Weiss 1975). A stationary random process without this property is not reversible, and is said to be directional.

In some cases appropriate plotting of the time series, considered as a realization of an underlying random process, can show a lack of symmetry which is evidence of directionality (Lawrance 1991). An example is the yearly sunspot series (1700 to 2014) shown in Figure 1. The series has sharp increases to the peaks followed by slow recessions to the troughs, when plotted in time order, and in contrast it has slow increases before the peaks followed by plummets to the troughs when plotted against time-to-go. The five most distinctive asymmetric peaks are highlighted with circles in Figure 1, and a similar

2

shape can be seen in most of the other peaks. Other examples include: fragmented time series of rainfall (Soubeyrand et al. 2014); EEG signals from subjects diagnosed with epilepsy (Mansor et al. 2016b); the North Greenland Ice Core Project (NGRIP) series and the Vostok ice core series (Mansor et al. 2016c); and the United States (U.S.) monthly civilian unemployment rate series (Tsay 2005).
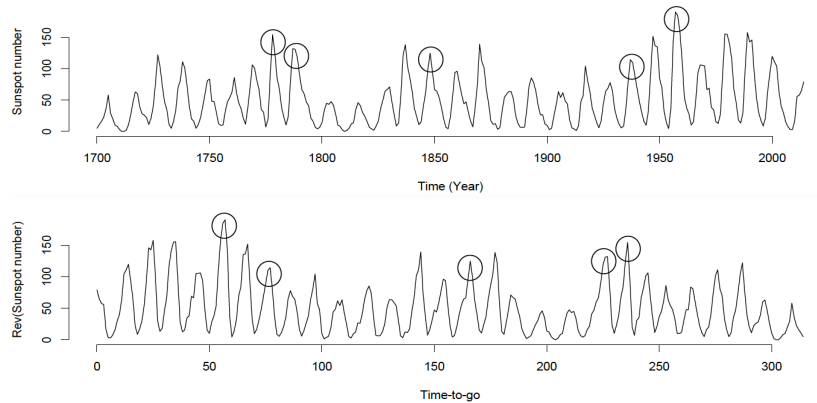
Figure 1: Sunspot series plotted in time order [above] and against time-to-go [below]. There is a clearly visible tendency for rapid rises and relatively slow recessions in time order, whereas relatively slow recoveries and sharp reductions are seen in the time-to-go plot.

However, graphical inspection for directionality between observations plotted in time order and the observations plotted in time-to-go is impractical for high frequency data. This is because the period is too long to visually detect relatively fine scale features if the entire record is plotted. Short excerpts could separate sample points sufficiently to see evidence of directionality, but directionality can occur at different time scales or intermittently and be missed in short excerpts. For example, the daily log-return (logarithm of ratio of today's price to yesterday's price) sub-series of 315 days for JP Morgan in Figure 2(a and b) do not show any clear directionality, although there is statistical evidence of directionality in the entire series. Moreover, it is difficult to discern directionality in any of the plots of daily log-returns over the entire period of JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs (Figure 3(b, h, n), Figure 4(b and h), respectively), despite convincing statistical evidence of the phenomenon.

High frequency time series are often analyzed in the frequency domain. In particular, spectral analysis is useful for identifying dynamics that are masked by noise. However, spectral analysis is based on fitting sinusoidal functions and cannot identify directionality. Wavelet analysis provides a frequency decomposition that changes over time but we cannot discern any evidence of directionality in wavelet decompositions of time series that we have analysed.

3
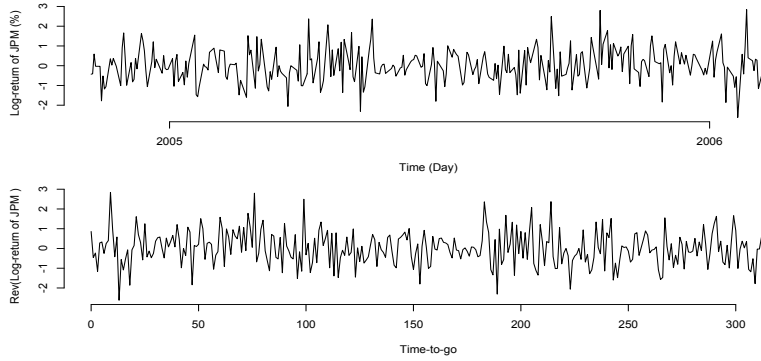
(a) JP Morgan sub-series of 315 log-returns between 9 Nov. 2004 and 8 Feb. 2005 during low frequency period.
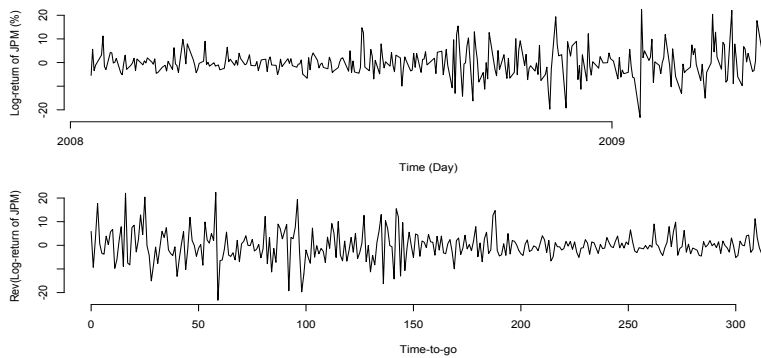
(b) JP Morgan sub-series of 315 log-returns between 15 Jan. 2008 and 15 Apr. 2009 during high frequency period.

Figure 2: Sub-series of JP Morgan log-returns during crisis (a) and non-crisis (b) periods, plotted in time order [above] and in time-to-go [below]. There is no clearly visible directionality in either case.

Directionality can be a consequence of feedback in a system following extreme events. Occasional extreme events induce feedbacks that locally disturb the symmetric behavior of a time series (Soubeyrand et al. 2014). For example, the link between extreme values and directionality has been discussed by: Soubeyrand et al. (2014) in the context of intense rainfall; Mansor et al. (2015) in the context of vigorous magnetic fields of the Sun which are apparent in sunspot series; and Mansor et al. (2016c) in the context of extreme paleoclimatic events on ice-core time series. In the case of finance, the world economy has often experienced a series of severe financial and economic shocks. A recent

4

example is the subprime mortgage crisis in 2007 that lead to a severe financial crisis from 2008 to 2009 which affected the performance of financial institutions in the U.S. (e.g. Allen & Christa 2013). Directionality has been shown to be more marked during crisis periods than it is in non-crisis periods (Figure 5, second row). Furthermore, detecting directionality in the log-return series may identify feedback behaviour following extreme events and provide early warning of decreases in share prices.

Directionality can be incorporated into linear time series models, typically autoregressive moving average ARMA(p,q) models, by introducing non-Gaussian errors. Also non-linear models such as bilinear, random coefficient, and threshold are inherently directional (Lawrance 1991). In principle, if the distribution of the errors is better modelled as a non-Gaussian process, then this will lead to more accurate forecasts. Also, if non-linear models such as threshold autoregressive (TAR) models (Tong & Lim 1980) provide better fit than autoregressive (AR) linear models, then the forecast of these non-linear models will be more accurate. For example, Mansor et al. (2016a) demonstrate that a TAR model gives improved forecasting results on an AR model, and back-to-back Weibull distributions provide more realistic simulations of 15-year extreme values for directional time series of sunspots.

In general, TAR models provide a piecewise linear approximation to a wide range of non-linear processes. TAR models are appropriate for financial time series with directionality component because these models use threshold space to capture several non-linear characteristics which are commonly observed in practise as asymmetry in declining and rising patterns (directionality) of a process (Tsay 2005). These ideas generalize to multivariate time series models that can show directionality in simulations, through either non-Gaussian errors or non-linearity or both.

In this paper, we consider a threshold vector autoregressive (TVAR) model for a portfolio of daily log-returns of five American shares, and we propose a strategy for using directionality as a regime-switching criterion between two-regime TVAR models in the model fittings. We investigate the performance of this innovation by comparing daily share prices forecasting performances of the TVAR models with a vector autoregressive (VAR) model. Another prominent feature of many time series of share returns is volatility, and we investigate the relationship between directionality and volatility. Volatility can be modelled by a GARCH model (Nelson 1991). The non-Gaussian errors for multivariate time series modelled as a vector of GARCH processes.

The remainder of this paper is organized into 4 sections. Section 2 describes the daily data in more detail. Section 3 provides empirical evidence of directionality in the log-returns, and investigate the association between directionality and volatility. Here, we also consider detecting directionality in 1-minute log-returns for the same banks. In Section 4, we compare the multivariate prediction models used for in-sample and out-of-sample one-step ahead daily share prices and volatility forecasts. Concluding remarks are given in Section 5.

2 RESEARCH DATA

We consider the daily share prices of the five largest, by market capitalization for 2017, banking institutions in the U.S. (relbanks 2017), and details are given in Table 1.

Table 1: Periods of time series of daily share prices of the five largest banks in the U.S.

Bank	Ticker	Market Cap ($\approx$US$b)	World Ranking	Indentifier i	Start	Finish
JP Morgan	JPM US	300	1	1	28/07/1980	17/02/2017
Wells Fargo	WFC US	277	2	2	28/07/1980	17/02/2017
Bank of America	BAC US	229	4	3	28/07/1980	17/02/2017
Citibank Inc.	C US	160	7	4	29/10/1986	17/02/2017
Goldman Sachs	GS US	97	12	5	03/05/1999	17/02/2017

The common period for the five daily share price series is from 3 May 1999 to 17 February 2017, which is the longest common period since their initial public offering (IPO) in the New York Stock Exchange. We focus on the U.S. stock exchange market because it is the world's largest stock exchange by market capitalization, and the U.S. economy is the largest in the world in terms of gross domestic product (GDP).

The share prices plot for each bank is given in Figure 3(JP Morgan, Wells Fargo, Bank of America) and 4(Citibank, Goldman Sachs). The daily share price plots appear to be plausibly modelled as random walks, and there is no evidence against the null hypothesis of a unit root, using the Augmented Dickey-Fuller test, to support this claim. These series show some similarities, in particular a major upward trend after 2005 followed by a sharp downward trend beginning 2007 and persisting towards the end of 2009. The decline seems particularly noticeable for the Bank of America, Citibank and Goldman Sachs. Major recoveries have been made by JP Morgan, Wells Fargo and Goldman Sachs after 2012, despite some short-term downward trends, and these banks are currently at their highest points ever. After the plummet, Bank of America shows some improvements and is recovering slowly while Citibank is moving in sideways trend for about 7 years.

The remaining plots in Figure 3 and 4 are based on log-return of the series. The log-return time series, $\{x_{t,i}\}$, corresponding to the observed share price series for each bank i (Table 1, column 5), $\{y_{t,i}\}$, with length n is

$$x_{t,i} = ln\left\{\frac{y_{t,i}}{y_{t-1,i}}\right\} \times 100 \tag{1}$$

where $t = 2, 3, ...n$. The value of n is 4480 trading days in North America, where there are approximately 252 trading days per year (5 per week less nine public holidays) (e.g. Huang et al. 2017).

6
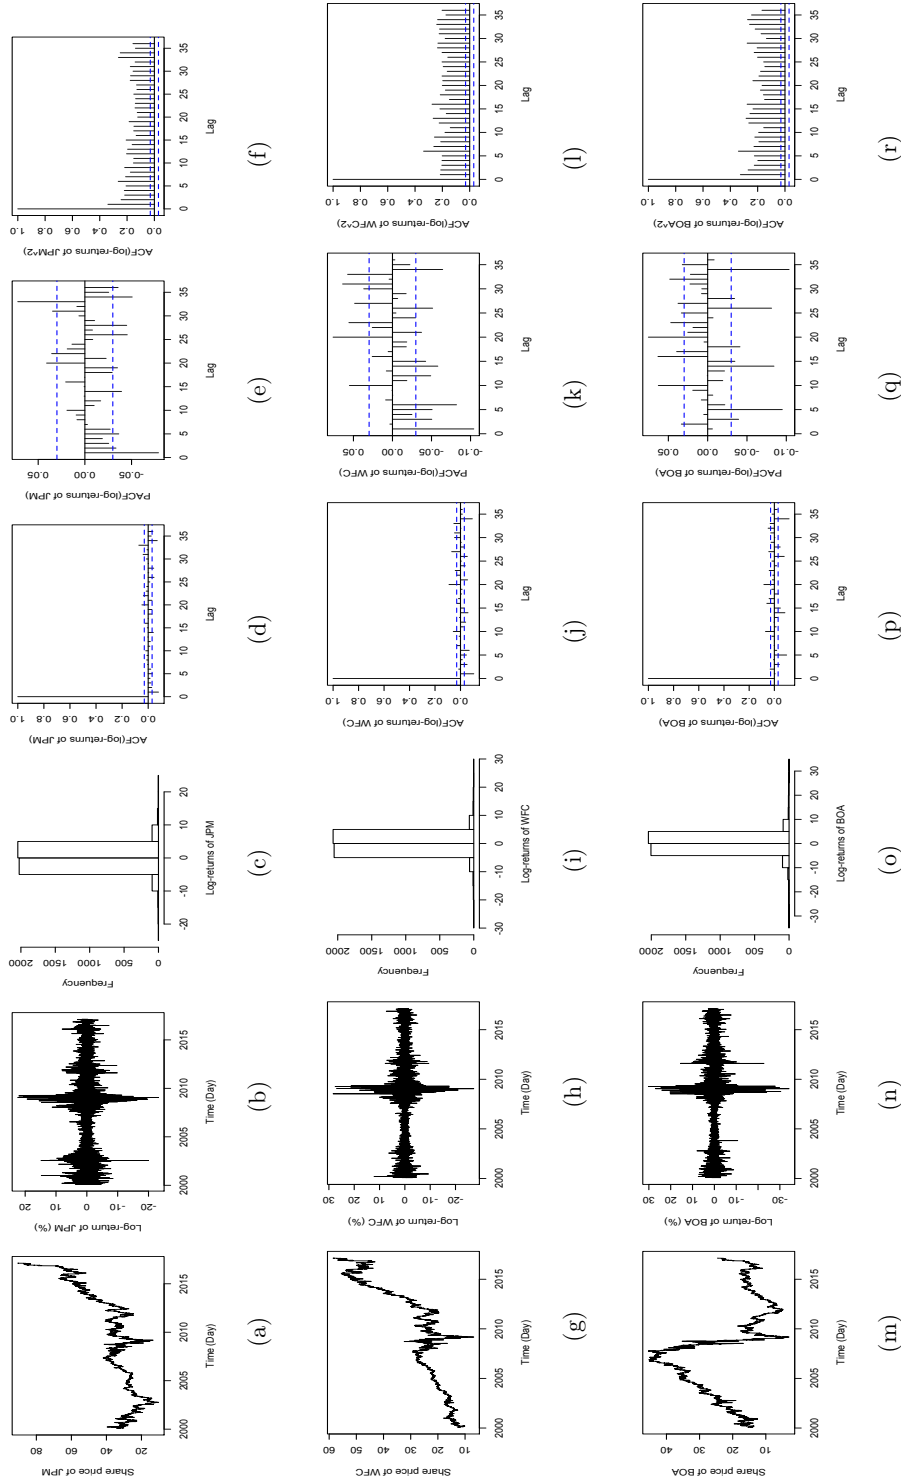
Figure 3: Daily share price, daily log-return, histogram of the log-returns, auto-correlogram function (ACF) of the log-returns, partial ACF (PACF) of the log-returns and ACF of the squared log-returns of: JP Morgan (a) to (f); Wells Fargo (g) to (l); and Bank of America (m) to (r), respectively.
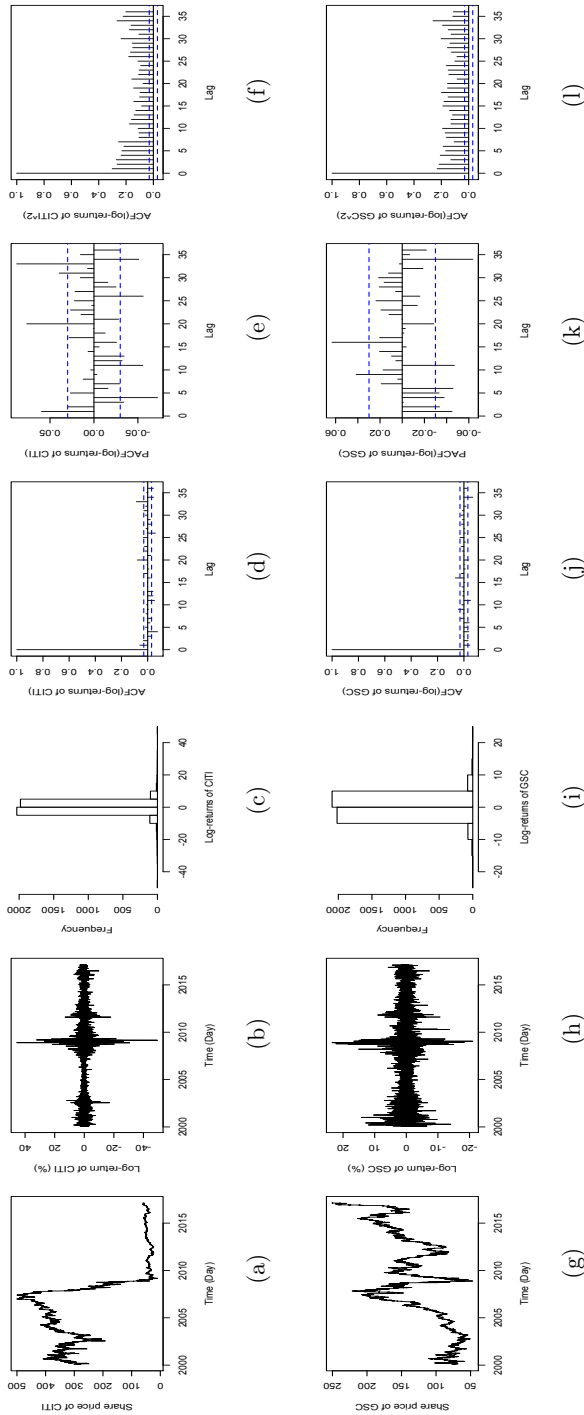
Figure 4: Daily share price, daily log-return, histogram of the log-returns, ACF of the log-returns, PACF of the log-returns and ACF of the squared log-returns of: Citibank (a) to (f); and Goldman Sachs (g) to (l), respectively.

The daily log-return series are considered as evenly spaced with no missing values. The five series can reasonably be considered as realisations of stationary time series models, and Table 2 provides summary statistics.

Table 2: Summary statistics of the log-return time series.

Time series	Mean	Std.Dev	SE mean	Skewness	Kurtosis	Minimum	Maximum
JP Morgan	0.024	2.556	0.039	0.27	13.2	−23.2	22.4
Wells Fargo	0.036	2.467	0.038	0.90	27.2	−27.2	28.3
Bank of America	0.011	3.032	0.046	−0.35	25.9	−34.2	30.2
Citibank	−0.037	3.222	0.049	−0.53	39.3	−49.5	45.6
Goldman Sachs	0.027	2.399	0.037	0.28	12.0	−21.0	23.5

The standard error of the means are larger than the means themselves so there is no statistical evidence that the mean is non-zero, over this period at least. However all but Citibank have positive estimate of the mean log-return, representing shares that have increased in value over this period. Citibank share value dropped from above \$500 to below \$1.00 in 2008 following the subprime mortgage and the global financial crisis.

Table 3: Cross-correlation matrix of the log-return series.

Time series	JPMorgan	WellsFargo	B.America	Citibank	GoldmanS
JP Morgan	**1.00**	.76	.76	.73	.70
Wells Fargo	.76	**1.00**	.81	.71	.60
Bank of America	.76	.81	**1.00**	.79	.63
Citibank	.73	.71	.79	**1.00**	.64
Goldman Sachs	.70	.59	.63	.64	**1.00**

The log-return series with the highest volatility, given by the marginal standard deviation of log-returns, is the Citibank log-returns, and the range $(45.6 - (−49.5))$ is also the highest. The histograms in Figure 3 and 4 (column 3) show that the log-return distributions are highly kurtotic with some degree of skewness. The magnitude of skewness and the kurtosis are given in Table 2. These high values of kurtosis indicate that the distributions tend to have extreme outliers.

Similarity between two log-return series at lag-zero, concurrent effect, is given by the correlation coefficients in the cross-correlation matrix (CCM) in Table 3 (all are statistically significant, P<0.01). The strength of linear dependence between log-returns series is generally strong positive (r >0.7), but Goldman Sachs has slightly lower values of correlation with the other banks. Co-integration test of Phillips-Ouliaris for the five series concludes that the multiple series are co-integrated (P<0.01).

3 DIRECTIONALITY

3.1 Detecting Directionality

There are many ways in which time series can exhibit directionality and in general formal statistical measures are required to quantify the phenomenon. Apart from slow increases followed by plummets, or the reverse, time series can exhibit both sharp rises and sharp falls followed by a slow recovery to the mean level, and also intermittent directionality as a consequence of feedback to counteract extreme inputs to the system. We investigate directionality in the log-return series using a suite of directionality indicators. The suite consists of two statistics based on the original series, $\{x_t\}$, and four statistics based on the first differences of the series, defined as

$$\Delta x_t = x_t - x_{t-1}, \tag{2}$$

for $t = 3, 4, ..n$.

The first statistic calculated for the log-return series itself is the differences in linear quadratic lagged correlation

$$DLQC = \frac{\sum_{t=2}^{n-1}(x_t - \bar{x})(x_{t+1} - \bar{x})^2}{[\sum_{t=2}^{n}(x_t - \bar{x})^2]^{3/2}} - \frac{\sum_{t=2}^{n-1}(x_t - \bar{x})^2(x_{t+1} - \bar{x})}{[\sum_{t=2}^{n}(x_t - \bar{x})^2]^{3/2}}. \tag{3}$$

Sharp increases, followed by slow recessions, will tend to give negative DLQC because, for example, below average values followed by high values will give negative contributions to the first term and positive contributions to the second.

The second statistic calculated for the log-return series is based on the mean of the five observations following a peak less the mean of the five observations preceding the peak (POT), where peaks are defined as peaks above some threshold value (in this case the 80^{th} percentile). The rational is that the mean over a slow recession from a peak will be higher than the mean over a rapid rise to a peak and POT will be positive. The statistic is defined by

$$t_{paired\ statistic} = \frac{\bar{d}}{(s_d/\sqrt{k})} \tag{4}$$

where $\bar{d}$ is the mean of differences between the mean of h observations before a peak and the mean of h observations after the peak, s_d is the standard deviation of the differences and k is the number of paired samples, for a paired comparison.

The other four statistics are based on the Δx_t: the proportion of positive Δx_t relative to the sum of positive and negative Δx_t to allow for 0 difference (P_d^+); the skewness of the distribution of Δx_t ($\hat{\gamma}_d$); the L-skewness of the distribution of Δx_t (LSK); and the skewness of absolute differences of $\Delta x_{t,i}$ from the mean ($\hat{\gamma}_{dab}$). The rational for all these statistics is that sharp increases followed by a slow recession will tend to give smaller proportion of the relatively large positive differences and a distribution of differences that is positively skewed.

10

Table 4: Summary table of the suite of directionality indicators in the sunspots and the log-return time series.

Method based on	original series		first differences			
Time series	$DLQC$	POT	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
Sunspots	-0.060	4.832	42.5%	0.856	0.149	0.665
n=315, s=40.24	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
AR(9), $\hat{\sigma}_{err}$=15.54	[5.57]	[4.38]	[4.25]	[6.64]	[6.15]	[5.84]
(a) JP Morgan	-0.770	3.371	49.4%	0.729	0.018	0.042
n=4290, s=2.556	(0.00)	(0.00)	(0.14)	(0.00)	(0.00)	(0.19)
AR(36), $\hat{\sigma}_{err}$=2.523	[23.2]	[3.35]	[1.43]	[23.2]	[3.10]	[1.30]
(b) Wells Fargo	-0.399	2.060	49.2%	0.365	0.008	0.316
n=4290, s=2.463	(0.00)	(0.04)	(0.04)	(0.00)	(0.16)	(0.00)
AR(35), $\hat{\sigma}_{err}$=2.392	[12.2]	[2.09]	[1.99]	[12.1]	[1.42]	[10.2]
(c) Bank of America	-0.778	4.051	48.2%	0.818	0.017	0.891
n=4290, s=3.032	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
AR(35), $\hat{\sigma}_{err}$=2.943	[26.3]	[4.06]	[3.97]	[26.3]	[2.95]	[28.7]
(d) Citibank	-0.582	2.237	48.7%	0.677	0.022	0.301
n=4290, s=3.222	(0.00)	(0.03)	(0.00)	(0.00)	(0.00)	(0.00)
AR(36), $\hat{\sigma}_{err}$=3.163	[20.4]	[2.25]	[3.04]	[20.4]	[3.60]	[9.63]
(e) Golman Sachs	-0.538	2.849	48.9%	0.534	0.013	0.342
n=4290, s=2.399	(0.00)	(0.00)	(0.01)	(0.00)	(0.03)	(0.00)
AR(34), $\hat{\sigma}_{err}$=2.378	[16.6]	[2.85]	[2.36]	[16.6]	[2.23]	[11.0]

Notes: n is the length of the log-return time series from 1-Feb-2000 to 17-Feb-2017; s is the marginal standard deviation; $\hat{\sigma}_{err}$ is the estimated standard deviation of residuals in the AR(p by AIC) model; (number) refers to P-value; and [number] refers to the ratio of the directionality indicator to its standard deviation from Monte-Carlo simulation.

The significance levels of the directionality indicators are obtained by Monte Carlo simulation, implemented with the following procedure. An autoregressive model of order p, AR(p), is fitted to the time series of log-returns. The order of the AR(p) is selected as that which minimizes Akaike information criterion (AIC). One thousand realizations, of the length of the observed time series, are generated from the fitted AR(p) using Gaussian white noise errors with a standard deviation equal to the standard deviation of the residuals. The directional indicators are calculated for each realization and so their sampling distribution, under the null hypothesis of a reversible process, is obtained empirically. The P-value associated with a directional indicator calculated for the observed time series is equal to the proportion of directional indicators, calculated for the simulated time series, that are larger in absolute magnitude (Mansor et al. 2016c,a). This procedure

ignores the sampling variability of the coefficient in the AR(p) but this is negligible for long time series. The results of the statistical tests are summarized in Table 4.

All the indicators, for the sunspots as well as the five bank log-returns, are consistent with sharp increases followed by relatively slow recessions and are highly statistically significant. In this application, the skewness of differences ($\hat{\gamma}_d$) appears to be the most sensitive statistic for detecting directionality, in terms of its value relative to its standard error under the null hypothesis of reversibility, and it is used throughout the remainder of the paper. There is indisputable evidence of directionality, over the study period at least, but the time series are long and statistical significance does not imply practical significance. The practical implications are investigated in the following sections.

3.2 Directionality in moving averages of daily log-returns and sampled log-returns

In this section we consider directionality in: moving averages (MA) of 2, 5, and 10-day daily log-returns; and log-returns sampled every 2, 5 and 10 days. The values of the skewness of differences statistic ($\hat{\gamma}_d$) are shown in Table 5, with the 1-day directionality shown for comparison.

Table 5: Directionality in the daily log-returns, 2-day, 5-day and 10-day log-returns.

Series	1-day	2-day	5-day	10-day
JP Morgan MA	0.729	0.465	0.222	-0.146
JP Morgan sampled	0.729	0.843	0.501	0.632
Wells Fargo MA	0.365	0.382	0.233	0.187
Wells Fargo sampled	0.365	0.851	1.231	-0.546
Bank of America MA	0.818	0.350	0.226	-0.116
Bank of America sampled	0.818	0.615	1.362	0.502
Citibank MA	0.677	1.482	0.788	0.304
Citibank sampled	0.677	1.938	1.195	0.378
Goldman Sach MA	0.534	0.614	0.517	0.133
Goldman Sach sampled	0.534	0.752	0.857	0.365

With a few partial exceptions directionality decreases with the length of the moving average. This is expected as averaging reduces the impact of extreme events. In contrast the sampling period has a less predictable effect. The values of directionality for Wells Fargo, which had the lowest 1-day directionality, in particular vary widely.

12

3.3   Moving directionality and volatility

The objective of this section is to investigate whether directionality is a consistent feature of a time series or a transient phenomenon in response to shocks to the system as suggested by Soubeyrand et al. (2014). If it is a transient phenomenon, then it may be linked to periods of increased variability (volatility). We use a 9-month moving directionality (MD), 21 days by 9 months is 189 trading days, which is defined in terms of skewness of the first differences in the log-returns.

The formula used to calculate the MD series is given in Equation (5),

$$\hat{\gamma}_{d,t} = \frac{\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta}x_t)^3/m}{[\sum_{i=0}^{188}(\Delta x_{t-i} - \bar{\Delta}x_t)^2/m]^{3/2}}, \tag{5}$$

where $t = 189, ..., 4479$, $m = 189$, $\bar{\Delta}x_t = \sum_{i=0}^{188}\Delta x_{t-i}/m$ and x_t is the full-length time series of log-returns. Similarly for the volatility (σ), given by the marginal standard deviation of log-returns $\{x_t\}$, the 9-month moving volatility (MV) series is calculate using Equation (6).

$$\hat{\sigma}_i = \sqrt{\frac{\sum_{i=0}^{188}(x_{t-i} - \bar{x}_t)^2}{m}}. \tag{6}$$

where $\bar{x}_t = \sum_{i=0}^{188} x_{t-i}/m$. The plots of the MD series and the MV series for the five banks, together with the log-return series are given in Figure 5 with vertical lines corresponding to the date of the crises being reported in the news. For example, 15-Jan-2001 for the 2001 dotcom crash, 11-Sep-2001 for the 9/11 attack on the World Trade Centre, the subprime mortgage crisis was first widely reported on 1-Jul-2007 and the collapse of Lehman Brothers was reported on 15-Sep-2008. No exact date is found to indicate the recessions, but the U.S. unemployment rate rose to 4.2% on 1-Feb-2001, the National Association of Securities Dealers Automated Quotations (NASDAQ) reached a 6-year low on 24-Sept-2002 and U.S. stock market fell dramatically on 1-Aug-2011 for the 2011 recession.

The moving directionality appears to be more responsive to imminent crisis periods than the moving volatility, in particular for the fourth vertical line of 24-Sept-2002 in Figure 5 ((f), (g) and (i)) as well as for the second last vertical line of 14-Sept-2008 in Figure 5((g) and (h)), where directionality is found to peak before the crises. Volatility provides some indication for the fourth vertical line only on Figure 5 ((k) and (n)), and on Figure 5 ((k), (l) and (m)) for the second last line. However, neither moving directionality nor moving volatility give early warning for Goldman Sachs. Nevertheless, the plots show some potential for directionality as an indicator of impending financial crisis.
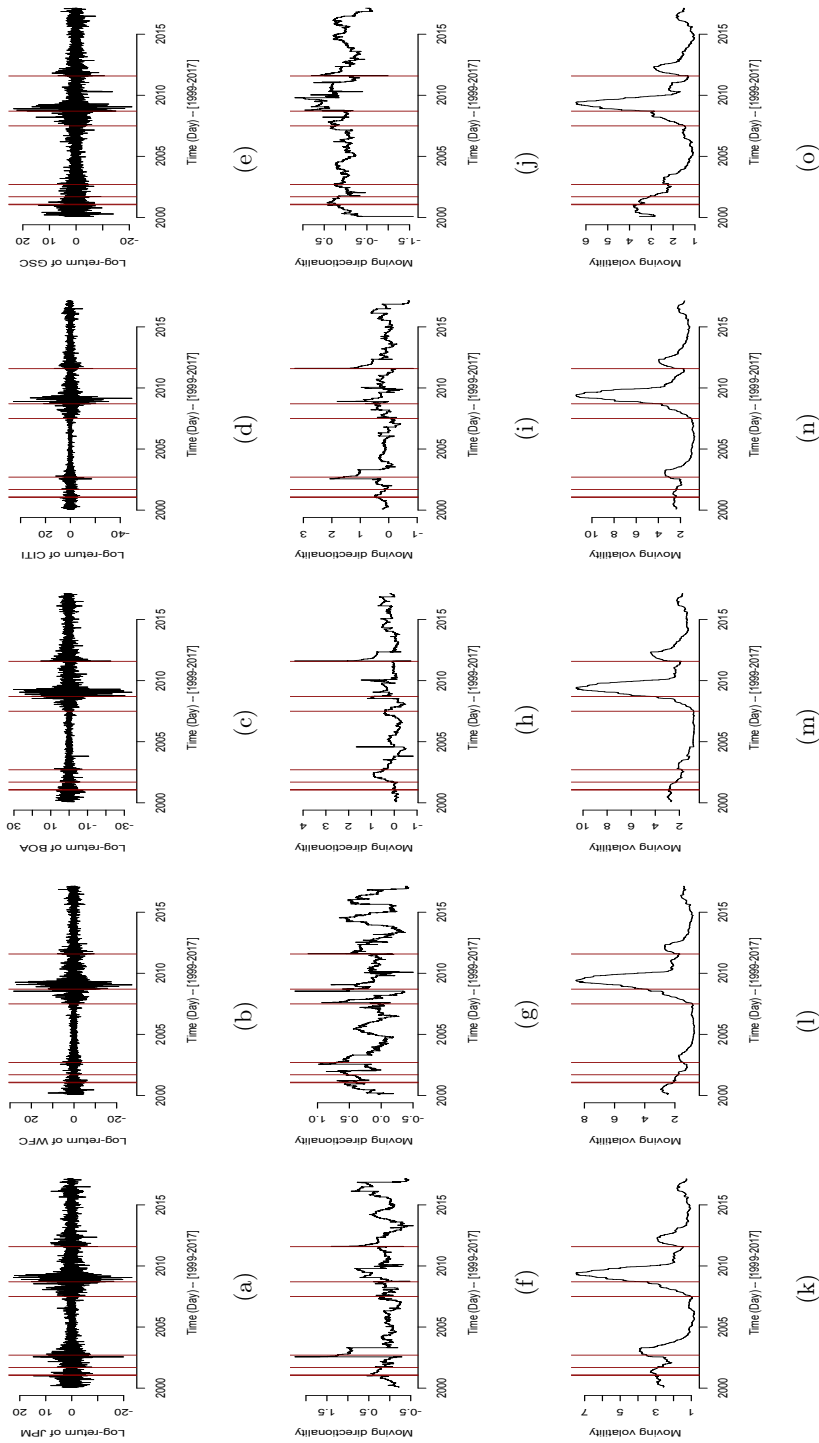
Figure 5: Daily log-return, 9-month moving directionality and 9-month moving volatility of: JP Morgan (a, f and k); Wells Fargo (b, g and l); Bank of America (c, h and m); Citibank (d, j and n); and Goldman Sachs (e, j and o), respectively.
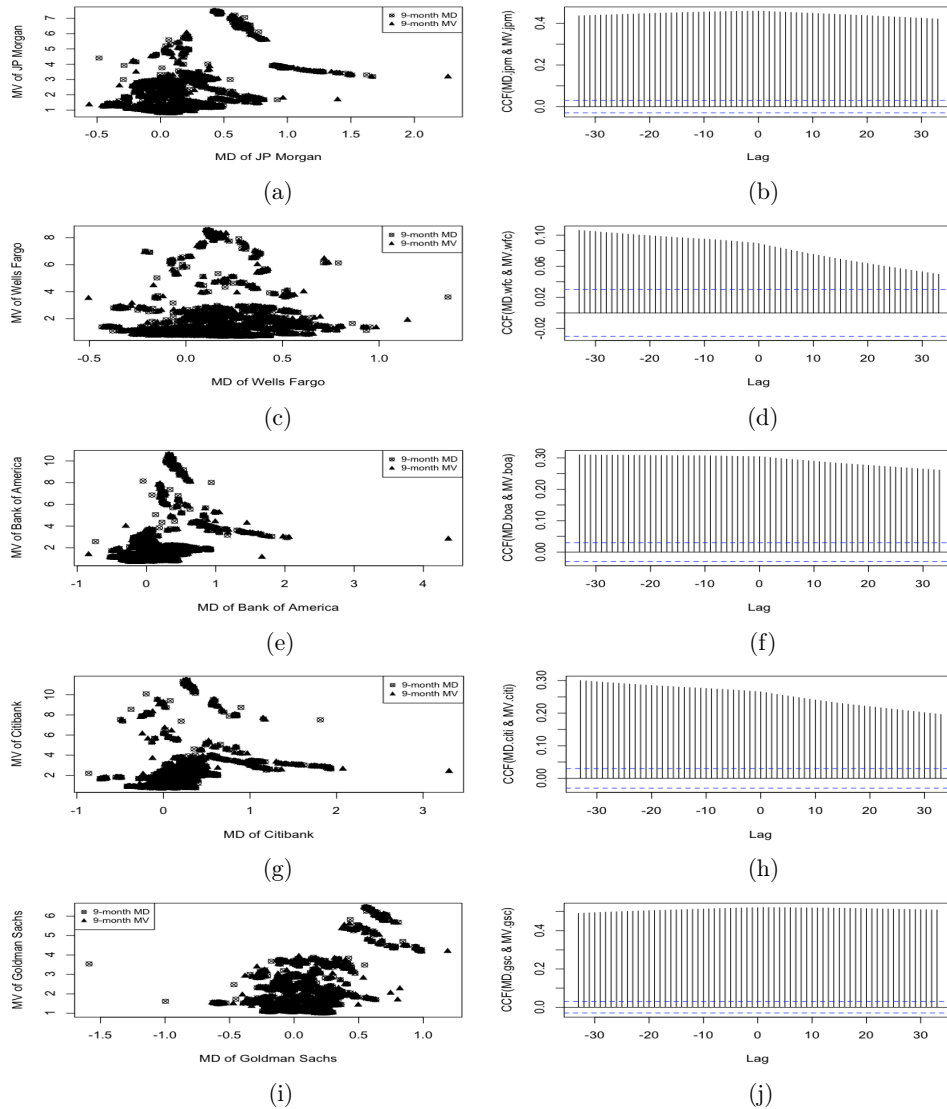
Figure 6: Scatter plot and cross-correlation function (CCF) show the relationship between the 9-month moving directionality and the 9-month moving volatility of: JP Morgan (a and b); Wells Fargo (c and d); Bank of America (e and f); Citibank (g and h); and Goldman Sachs (i and j), respectively.

We show the relationship between directionality and volatility for all five log-return series using scatter plot of the 9-month moving directionality $\{MD\}$ versus the 9-month moving volatility $\{MV\}$ in Figure 6 (first column). There is no obvious linear pattern

15

from the plots, and the strength of association between MD and MV series for JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs is 0.46, 0.09, 0.30, 0.27 and 0.52, respectively. The cross-correlation coefficients are low which indicates that directionality provides additional information to volatility. Furthermore, directionality seems to lead the occurrence of volatility, at least in the case of Wells Fargo, Bank of America and the Citibank given by the CCF in Figure 6(d, f and h), respectively.

3.4 Analyzing minute data

We have time series sampled at 1-minute intervals over a period of: 26 years for JP Morgan, Wells Fargo and Bank of America; 23 years for Citibank; and 18 years for Goldman Sachs. The 1-minute trading data is provided by Tickdatamarket. We also aggregated the time series to provide 3-minute and 6-minute records. The summary statistics including the length and proportion of zeroes (PZ) of the 1-minute time series are given in Table 6. The 1-minute data is a very long time series and has a substantial PZ corresponds to the trading period. The sampling distributions of the statistics we calculate are affected by the extraordinary high kurtosis of the 1-minute log-return data (order of 1000) arising from outliers and exacerbated by the peak at 0 attributable to a large proportion of zeroes.

Table 6: Summary statistics of the 1-minute log-return time series.

Ticker	Length	PZ	Mean	Std.Dev	SE.mean	Skew	Kurtosis	Min	Max
JPM US	2192338	27.6%	1.5×10^{-4}	0.149	1.0×10^{-4}	0.50	481.9	-15.7	17.2
WFC US	2048441	34.0%	1.5×10^{-4}	0.147	1.0×10^{-4}	4.88	1190.0	-13.9	28.1
BAC US	2159226	38.3%	5.9×10^{-5}	0.154	1.0×10^{-4}	3.69	1447.3	-18.0	21.8
C US	2124731	27.4%	0.00	0.173	1.2×10^{-4}	9.70	4611.7	-38.3	46.5
GS US	1730944	9.6%	6.2×10^{-5}	0.125	9.5×10^{-5}	4.66	2013.7	-15.3	27.5

The 1-minute time series of log-returns and the corresponding spectra, with span 10^4 seconds, are shown in Figure 7. Periods of volatility can be seen in the time series plot. The spectra all show a clear decrease as the frequency becomes lower which is characterise of high frequency processes such as the autoregressive order 1 model with a negative coefficient. If an AR(1) model is fitted to the time series of log-returns, then the coefficients were -0.10, -0.11, -0.07, -0.07 and -0.02 for JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs, respectively. This is consistent with a slight tendency for over correction, inasmuch as decreases immediately followed by increases are more frequent than consecutive decreases. Similarly for increases.

Results from the suite of directionality indicators for minutes data of each bank are given in Table 7. Many of the directionality measures appear to vary widely. However the POT and the P_d^+ are generally consistent in detecting directionality for the minutes
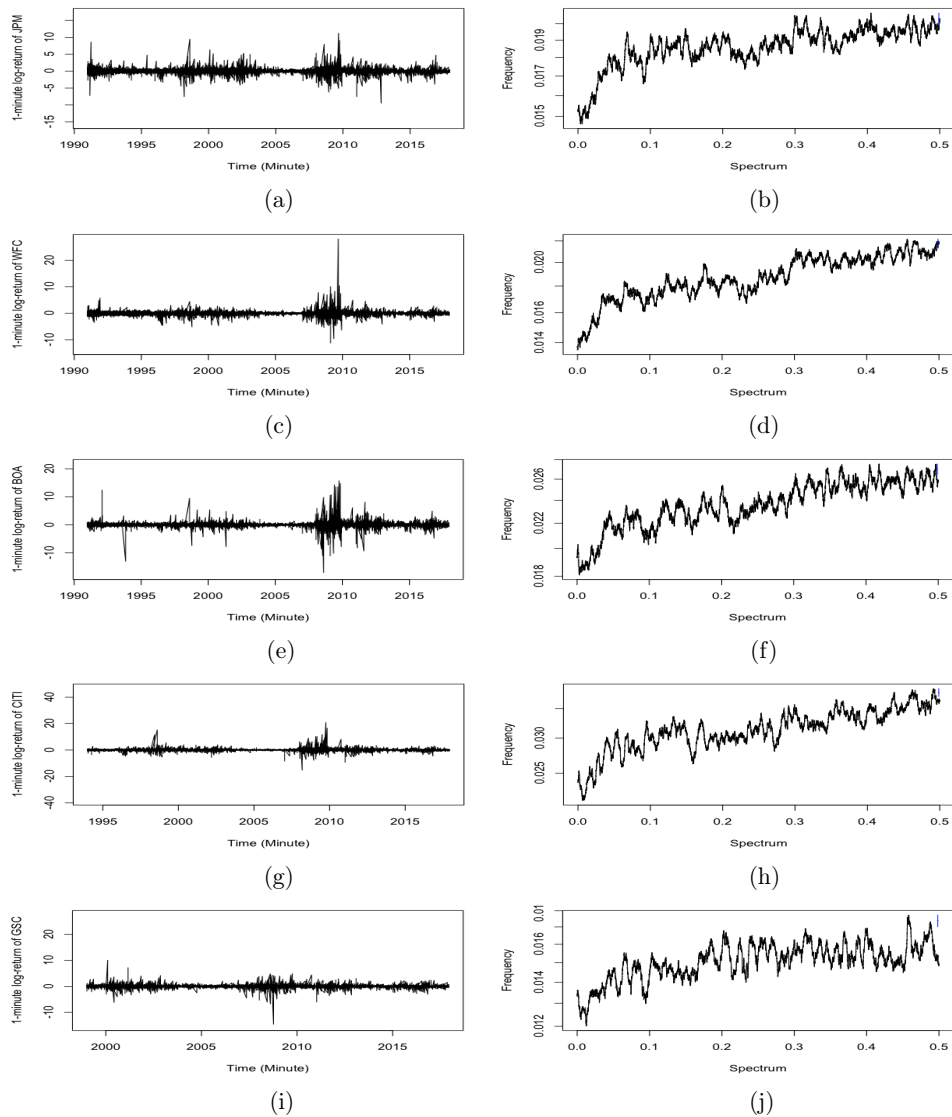
Figure 7: 1-minute log-return and its smoothed spectrum: JP Morgan (a and b); Wells Fargo (c and d); Bank of America (e and f); Citibank (g and h); and Goldman Sachs (i and j), respectively.

data. In the case of log-returns for U.S. banks end of the day values seem more useful for monitoring directionality.

Table 7: Directionality indicators in the 1-minute, 3-minute and 6-minute log-return time series.

Method based on	original series		first differences			
Time series	$DLQC$	POT	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
(a) JP Morgan						
1-minute	-0.621	35.66	46.34%	0.571	0.002	0.278
3-minute	-0.582	10.42	48.21%	0.785	0.001	0.909
6-minute	-0.411	5.183	49.15%	0.599	0.002	0.545
(b) Wells Fargo						
1-minute	-0.044	44.55	45.91%	0.040	0.004	0.736
3-minute	-0.694	10.59	47.51%	0.973	0.003	1.920
6-minute	0.327	3.178	48.92%	-0.480	0.002	-0.569
(c) Bank of America						
1-minute	-0.091	52.10	45.97%	0.087	0.003	1.698
3-minute	-0.215	14.31	47.14%	0.311	0.003	0.836
6-minute	-0.182	6.111	48.60%	0.283	0.004	-0.285
(d) Citibank						
1-minute	0.470	40.38	46.56%	-0.449	0.005	1.359
3-minute	0.192	11.23	48.39%	-0.283	0.002	1.032
6-minute	1.587	3.722	49.38%	-2.379	0.001	-2.148
(e) Goldman Sachs						
1-minute	-0.034	12.90	48.20%	0.035	0.001	0.687
3-minute	0.325	5.926	49.38%	-0.461	0.001	1.336
6-minute	-0.384	5.398	49.78%	0.585	-0.001	0.354

4 MULTIVARIATE FORECASTING

4.1 Multivariate time series models for forecasts

A vector autoregressive of order p, VAR(p), model is a basic linear multivariate time series model and its general form is:

$$\boldsymbol{X}_t = \varPhi_0 + \varPhi_1 \boldsymbol{X}_{t-1} + ... + \varPhi_p \boldsymbol{X}_{t-p} + \boldsymbol{\mathcal{E}}_t, \tag{7}$$

where $\boldsymbol{X}_t$ is a i-dimensional vector of i^{th} univariate time series modelled by random variables, $\{X_{t,1}, ..., X_{t,I}\}$, for $t = 190, 191, ..., n$ and $i = 1, 2, ..., I$, $\varPhi$ is a $I \times I$ matrix of the model coefficients, and $\boldsymbol{\varepsilon}_t$ is a sequence of serially uncorrelated multivariate random vectors with mean zero and non-negative definite covariance matrix $\Sigma_{\boldsymbol{\varepsilon}}$. The delay of 189 days at the beginning of the series is to allow for moving directionality to be used as a

threshold.

In this section, we fit five daily log-return series of JP Morgan $\{x_{t,1}\}$, Wells Fargo $\{x_{t,2}\}$, Bank of America $\{x_{t,3}\}$, Citibank $\{x_{t,4}\}$ and Golman Sachs $\{x_{t,5}\}$ with the VAR process of order one, VAR(1) in Equation (8).

$$\boldsymbol{X}_t = \Phi_0 + \Phi_1 \boldsymbol{X}_{t-1} + \boldsymbol{\mathcal{E}}_t. \tag{8}$$

The ACF plot in Figure 3(d, j), Figure 4(d and j) has significant spike at lag 1, and this is confirmed by the PACF plot in Figure 3(e, k), Figure 4(e and k). We note that there are significant higher-order autocorrelations at the larger number of lags for the five log-return series but for simplicity reason we use the VAR(1) as a reasonable first approximation for the five dimensional case. The VAR(1) model for five-component case consists of the following five equations:

$$
\begin{aligned}
x_{t,1} &= \phi_{10} + \phi_{11}x_{t-1,1} + \phi_{12}x_{t-1,2} + \phi_{13}x_{t-1,3} + \phi_{14}x_{t-1,4} + \phi_{15}x_{t-1,5} + \varepsilon_{t,1}, \\
x_{t,2} &= \phi_{20} + \phi_{21}x_{t-1,1} + \phi_{22}x_{t-1,2} + \phi_{23}x_{t-1,3} + \phi_{24}x_{t-1,4} + \phi_{25}x_{t-1,5} + \varepsilon_{t,2}, \\
x_{t,3} &= \phi_{30} + \phi_{31}x_{t-1,1} + \phi_{32}x_{t-1,2} + \phi_{33}x_{t-1,3} + \phi_{34}x_{t-1,4} + \phi_{35}x_{t-1,5} + \varepsilon_{t,3}, \\
x_{t,4} &= \phi_{40} + \phi_{41}x_{t-1,1} + \phi_{42}x_{t-1,2} + \phi_{43}x_{t-1,3} + \phi_{44}x_{t-1,4} + \phi_{45}x_{t-1,5} + \varepsilon_{t,4}, \\
x_{t,5} &= \phi_{50} + \phi_{51}x_{t-1,1} + \phi_{52}x_{t-1,2} + \phi_{53}x_{t-1,3} + \phi_{54}x_{t-1,4} + \phi_{55}x_{t-1,5} + \varepsilon_{t,5}.
\end{aligned} \tag{9}
$$

In addition, the five log-return series provide empirical evidence of directionality. Non-Gaussian errors will impart some directionality but the effect is slight when the value of the coefficients in the VAR(1) model are close to 0. Threshold vector autoregressive models (TVAR), provide a piecewise linear approximation to a wide range of non-linear processes (Tong & Lim 1980), can contribute more to the directionality. A five dimensional two-regime TVAR model of order 1 (TVAR(1)) can be expressed as follows:

$$
x_{t,1} = \begin{cases} \phi_{10}^{U} + \phi_{11}^{U}x_{t-1,1} + \phi_{12}^{U}x_{t-1,2} + \phi_{13}^{U}x_{t-1,3} + \phi_{14}^{U}x_{t-1,4} + \phi_{15}^{U}x_{t-1,5} + \varepsilon_{t,1}^{U} \text{ if } x_{t-1,1} > T_1 \\ \phi_{10}^{L} + \phi_{11}^{L}x_{t-1,1} + \phi_{12}^{L}x_{t-1,2} + \phi_{13}^{L}x_{t-1,3} + \phi_{14}^{L}x_{t-1,4} + \phi_{15}^{L}x_{t-1,5} + \varepsilon_{t,1}^{L} \text{ if } x_{t-1,1} \leq T_1 \end{cases} ,
$$

$$
x_{t,2} = \begin{cases} \phi_{20}^{U} + \phi_{21}^{U}x_{t-1,1} + \phi_{22}^{U}x_{t-1,2} + \phi_{23}^{U}x_{t-1,3} + \phi_{24}^{U}x_{t-1,4} + \phi_{25}^{U}x_{t-1,5} + \varepsilon_{t,2}^{U} \text{ if } x_{t-1,2} > T_2 \\ \phi_{20}^{L} + \phi_{21}^{L}x_{t-1,1} + \phi_{22}^{L}x_{t-1,2} + \phi_{23}^{L}x_{t-1,3} + \phi_{24}^{L}x_{t-1,4} + \phi_{25}^{L}x_{t-1,5} + \varepsilon_{t,2}^{L} \text{ if } x_{t-1,2} \leq T_2 \end{cases} ,
$$

$$
\begin{aligned}
x_{t,3} &= \cdots \qquad\qquad \cdots \qquad\qquad \cdots \qquad\qquad \cdots \qquad , \\
x_{t,4} &= \cdots \qquad\qquad \cdots \qquad\qquad \cdots \qquad\qquad \cdots \qquad ,
\end{aligned}
$$

$$
x_{t,5} = \begin{cases} \phi_{50}^{U} + \phi_{51}^{U}x_{t-1,1} + \phi_{52}^{U}x_{t-1,2} + \phi_{53}^{U}x_{t-1,3} + \phi_{54}^{U}x_{t-1,4} + \phi_{55}^{U}x_{t-1,5} + \varepsilon_{t,5}^{U} \text{ if } x_{t-1,5} > T_5 \\ \phi_{50}^{L} + \phi_{51}^{L}x_{t-1,1} + \phi_{52}^{L}x_{t-1,2} + \phi_{53}^{L}x_{t-1,3} + \phi_{54}^{L}x_{t-1,4} + \phi_{55}^{L}x_{t-1,5} + \varepsilon_{t,5}^{L} \text{ if } x_{t-1,5} \leq T_5 \end{cases} ,
$$

$$\tag{10}$$

where $\{x_{t-1,1}\}$, $\{x_{t-1,2}\}$, $\{x_{t-1,3}\}$, $\{x_{t-1,4}\}$ and $\{x_{t-1,5}\}$ are the lag-1 log-return series of JP Morgan, Wells Fargo, Bank of America, Citibank and Goldman Sachs sequentially. The values of the T_i were taken as the 0.95 quantiles of the marginal distributions of the

threshold variables. Therefore, T_1 is the 0.95 quantile of $\{x_{t-1,1}\}$ for JP Morgan, T_2 is the 0.95 quantile of $\{x_{t-1,2}\}$ for Wells Fargo and so forth. The choice was based on minimum sum of squared residuals over the quantile values of 0.50, 0.55, 0.60, 0.65, 0.70, 0.75, 0.80, 0.85, 0.90 and 0.95.

In the previous section, we showed that directionality becomes more apparent during crisis periods than it is in non-crisis periods for all five log-return series. This characteristic can be modelled by fitting separate sets of coefficients to the TVAR(1) model during periods of high and low directionality. We refer to this as a two-regime TVAR(1) model where the regime is determined by whether or not moving directionality exceeds an external threshold (TVAR(1)MD). The TVAR(1)MD model for the five-variate case is given in Equation (11).

$$
x_{t,1} = \begin{cases} \begin{cases} \phi_{10}^{U1} + \phi_{11}^{U1}x_{t-1,1} + \phi_{12}^{U1}x_{t-1,2} + ... + \phi_{15}^{U1}x_{t-1,5} + \varepsilon_{t,1}^{U1} \text{ if } x_{t-1,1} > T_1 \\ \phi_{10}^{L1} + \phi_{11}^{L1}x_{t-1,1} + \phi_{12}^{L1}x_{t-1,2} + ... + \phi_{15}^{L1}x_{t-1,5} + \varepsilon_{t,1}^{L1} \text{ if } x_{t-1,1} \leq T_1 \end{cases} \text{ if } MD_1 > T_{D1} \\ \begin{cases} \phi_{10}^{U2} + \phi_{11}^{U2}x_{t-1,1} + \phi_{12}^{U2}x_{t-1,2} + ... + \phi_{15}^{U2}x_{t-1,5} + \varepsilon_{t,1}^{U2} \text{ if } x_{t-1,1} > T_1 \\ \phi_{10}^{L2} + \phi_{11}^{L2}x_{t-1,1} + \phi_{12}^{L2}x_{t-1,2} + ... + \phi_{15}^{L2}x_{t-1,5} + \varepsilon_{t,1}^{L2} \text{ if } x_{t-1,1} \leq T_1 \end{cases} \text{ if } MD_1 \leq T_{D1} \end{cases},
$$

$$x_{t,2} = \cdots \qquad \cdots \qquad \cdots \qquad \cdots \quad ,$$

$$x_{t,3} = \cdots \qquad \cdots \qquad \cdots \qquad \cdots \quad ,$$

$$x_{t,4} = \cdots \qquad \cdots \qquad \cdots \qquad \cdots \quad ,$$

$$
x_{t,5} = \begin{cases} \begin{cases} \phi_{50}^{U1} + \phi_{51}^{U1}x_{t-1,1} + \phi_{52}^{U1}x_{t-1,2} + ... + \phi_{55}^{U1}x_{t-1,5} + \varepsilon_{t,5}^{U1} \text{ if } x_{t-1,5} > T_5 \\ \phi_{50}^{L1} + \phi_{51}^{L1}x_{t-1,1} + \phi_{52}^{L1}x_{t-1,2} + ... + \phi_{55}^{L1}x_{t-1,5} + \varepsilon_{t,5}^{L1} \text{ if } x_{t-1,5} \leq T_5 \end{cases} \text{ if } MD_5 > T_{D5} \\ \begin{cases} \phi_{50}^{U2} + \phi_{51}^{U2}x_{t-1,1} + \phi_{52}^{U2}x_{t-1,2} + ... + \phi_{55}^{U2}x_{t-1,5} + \varepsilon_{t,5}^{U2} \text{ if } x_{t-1,5} > T_5 \\ \phi_{50}^{L2} + \phi_{51}^{L2}x_{t-1,1} + \phi_{52}^{L2}x_{t-1,2} + ... + \phi_{25}^{L2}x_{t-1,5} + \varepsilon_{t,5}^{L2} \text{ if } x_{t-1,5} \leq T_5 \end{cases} \text{ if } MD_5 \leq T_{D5} \end{cases}.
$$

$$(11)$$

The external threshold variable MD_i is the 9-month moving directionality and T_{Di} is taken as the 0.80 quantile of the marginal distributions for i from 1 to 5, for JP Morgan up to Goldman Sachs respectively. The internal thresholds for both regimes are taken as the 0.95 quantile of the marginal distributions of the internal threshold variable for each series.

For comparison, we explore the possibility of using 9-month moving volatility $\{MV\}$ in Section 3.3 as the external threshold variable for Equation (11) (TVAR(1)MV model). Another possibility is to use a combination of the 9-month $\{MD\}$ and the 9-month $\{MV\}$, called the moving threshold indicator (MTI), as an external threshold variable for the TVAR(1) in Equation (11) (TVAR(1)MTI model). The MTI is the standardized 9-month MD plus the standardized 9-month MV. The internal thresholds for both regimes are again taken as the 0.95 quantile of the marginal distributions of the internal threshold variable for each series.

To check for stationarity, we use a characteristic function polyroot in R to calculate zeroes or roots of polynomials using the determinant of the estimated model coefficients (Cowpertwait & Metcalfe 2009) in Equation (9), (10) and (11). All estimated model parameters values for the VAR(1), TVAR(1), TVAR(1)MD, TVAR(1)MV and TVAR(1)MTI satisfy the conditions for stationarity since all roots exceed unity in its absolute term.

4.2   Performance of share prices forecasts

The directionality that is apparent in the time series is a consequence of non-linearity or non-Gaussian errors or both. Given that the time series are reasonably modelled as autoregressive series of order 1 with a coefficient value close to zero the effect of non-Gaussian errors is slight and introducing non-linearity by allowing for a two regime TVAR(1) model contributes more to the directionality. Given that the TVAR(1) improves the fit in terms of directionality, in this section our aim is to investigate whether the TVAR(1) also improved one-step ahead forecasts.

We provide formal evidences of directionality in the daily log-return series, and we fit the VAR(1), TVAR(1), TVAR(1)MD, TVAR(1)MV and TVAR(1)MTI models to the five log-return series. We make one-step ahead forecasts of share prices ($y_{t+1|t,i}$) (a) in-sample and (b) 70:30% out-of-sample. For (b) we fit the models to the first 70% log-returns ($\{x_{t,i}\}$) and use these models for one-step ahead forecasts for the remaining 30% share prices. The models are not re-fitted using previously forecast values when forecasting further. One-step ahead share prices forecasts $\hat{y}_{t+1|t,i}$ are obtained by the inverse of Equation (1).

$$\hat{y}_{t+1|t,i} = \exp\left\{\frac{\hat{x}_{t+1|t,i}}{100}\right\} \times y_{t,i} \tag{12}$$

where $\hat{x}_{t+1|t,i}$ is the one-step ahead log-returns forecasts and $y_{t,i}$ is the observed stock prices at time t.

We measure the forecasting performances of the VAR(1), TVAR(1), TVAR(1)MD, TVAR(1)MV and TVAR(1)MTI models for case (a) and (b) using the mean relative error squared (MRES) in Equation (13) for individual component $MRES_i$ in each time series models and as a group $MRES$.

$$\mathcal{E}_{t+1,i} = y_{t+1,i} - \hat{y}_{t+1|t,i}, \tag{13}$$

$$MRES_i = \frac{1}{T-1}\sum_{t=1}^{T-1}\left(\frac{\mathcal{E}_{t+1,i}}{y_{t+1,i}}\right), \tag{14}$$

$$MRES = \frac{1}{I}\sum_{i=1}^{I}(MRES_i), \tag{15}$$

where $\mathcal{E}_{t+1,i}$ is the one-step ahead forecast error, $y_{t+1,i}$ is the observed share price with length T and I is the number of banks, in this case is five. Detailed results are given in Table 8.

Based on group performances, the TVAR(1) models perform slightly better than the VAR(1) model for in-sample and out-of-sample. The TVAR(1)MD provides the least group forecasting errors for case (a) and the TVAR(1) for case (b). In general, the TVAR(1)MD model performs slightly better than the VAR(1), TVAR(1), and TVAR(1)MTI models at the individual level. For example the in-sample case of JP Morgan, Citibank and Goldman Sachs, and for the out-of-sample JP Morgan.

Table 8: Forecasting performance by MRES.

Model	VAR(1)	TVAR(1)	TVAR(1)MD	TVAR(1)MV	TVAR(1)MTI
(a) In-sample one-step ahead forecast					
$MRES_{jpm}$	6.43×10^{-4}	6.43×10^{-4}	6.39×10^{-4}	6.39×10^{-4}	6.39×10^{-4}
$MRES_{wfc}$	5.83×10^{-4}	5.82×10^{-4}	5.79×10^{-4}	5.75×10^{-4}	5.74×10^{-4}
$MRES_{boa}$	9.31×10^{-4}	9.16×10^{-4}	9.03×10^{-4}	9.02×10^{-4}	8.97×10^{-4}
$MRES_{citi}$	1.06×10^{-3}	1.05×10^{-3}	1.01×10^{-3}	1.03×10^{-3}	1.03×10^{-3}
$MRES_{gms}$	5.70×10^{-4}	5.67×10^{-4}	5.61×10^{-4}	5.63×10^{-4}	5.62×10^{-4}
$MRES$	7.57×10^{-4}	7.51×10^{-4}	7.39×10^{-4}	7.41×10^{-4}	7.39×10^{-4}
(b) 70:30% out-of-sample one-step ahead forecast					
$MRES_{jpm}$	2.09×10^{-4}	2.10×10^{-4}	2.09×10^{-4}	2.14×10^{-4}	2.09×10^{-4}
$MRES_{wfc}$	1.55×10^{-4}	1.56×10^{-4}	1.62×10^{-4}	1.56×10^{-4}	1.58×10^{-4}
$MRES_{boa}$	3.33×10^{-4}	3.35×10^{-4}	3.42×10^{-4}	3.45×10^{-4}	3.37×10^{-4}
$MRES_{citi}$	5.06×10^{-4}	3.04×10^{-4}	3.14×10^{-4}	3.09×10^{-4}	3.12×10^{-4}
$MRES_{gms}$	2.21×10^{-4}	2.23×10^{-4}	2.22×10^{-4}	2.25×10^{-4}	2.23×10^{-4}
$MRES$	2.85×10^{-4}	2.46×10^{-4}	2.50×10^{-4}	2.50×10^{-4}	2.48×10^{-4}

4.3 Performance of share price forecasts adjusted for volatility

There is clear evidence of volatility in the five log-return time series given by the ACF of the squared log-returns in Figure 3(f, l, r), Figure 4(f and l), respectively. This feature can be modelled as GARCH process and its general form for univariate residuals is given in Equation (16).

$$GARCH(p,q) \begin{cases} h_t = \omega + \sum_{a=1}^{p} \beta_i h_{t-1} + \sum_{b=1}^{q} \alpha_j \varepsilon_{t-1}^2 , \\ \varepsilon_t = \sqrt{h_t}\, z_t , \\ z_t \sim N(0,1), \end{cases} \qquad (16)$$

where h_t is the conditional variance of random disturbance that depend linearly on the past behavior of itself and the squared random error ε_{t-1}^2. We ensure that the h_t is a asymptotically stationary random sequence as the sum of β and α should not exceed 1 (e.g. Awartani & Corradi 2005). The choice of values for order p and q for the GARCH(p,q) model can be optimised but GARCH(1,1) specification has been found to be an adequate representation for most financial time series (Lamoureux & Lastrapes 1990). We generalize the univariate volatility model in Equation (16) for modelling volatility processes of multiple log-returns using GARCH(1,1) (e.g. Tsay 2005) given in Equation (17).

$$GARCH(1,1)_i \begin{cases} h_{t,i} = \omega_i + \beta_{1,i}h_{t-1} + \alpha_{1,i}\varepsilon_{t-1,i}^2 \, , \\ \varepsilon_{t,i} = \sqrt{h_{t,i}} \, z_{t,i} \, , \\ z_{t,i} \sim N(0,1). \end{cases} \tag{17}$$

We fit the GARCH(1,1) model to every component residuals of the VAR(1), TVAR(1), TVAR(1)MD, TVAR(1)MV and TVAR(1)MTI models, where their $\{\varepsilon_{t,i}\} \sim N(0, h_t)$ and $i = 1, 2, ...5$. The component residuals are $\{\varepsilon_{t,1}\}$ for the JP Morgan, $\{\varepsilon_{t,2}\}$ for the Wells Fargo, $\{\varepsilon_{t,3}\}$ for the Bank of America, $\{\varepsilon_{t,4}\}$ for the Citibank and $\{\varepsilon_{t,5}\}$ for the Goldman Sachs. All the component residuals are either serially uncorrelated or with minor lower order serial correlations, but serially dependent. In addition, there is evidence for ARCH effects in each of the component residuals given by the Lagrange Multiplier (LM) test (P < 0.01), to support the claim, for all models.

We illustrate the before and after effects of GARCH(1,1) on the component residuals of the VAR(1) using ACF and Q-Q plot in Figure 8. The ACF and Q-Q plot of the component residuals squared of the VAR(1) (Figure 8(first and second column), respectively) show the need of GARCH(1,1) process. The Q-Q plot of the component residuals squared of the VAR(1)-GARCH(1,1) in Figure 8(last column) is reasonably straight and the ACF in Figure 8(third column) suggests no significant serial correlations and serially independent, and is approximately normally distributed.

Next, we make: (a) in-sample one-step ahead volatility forecasts $\hat{h}_{t+1|t,i}$, and (b) 70:30% out-of-sample one-step ahead volatility forecasts $\hat{h}_{t+1|t,i}$ where we fit the GARCH (1,1) to the first 70% of the component residuals and estimated $\hat{h}_{t+1|t,i}$ for the 30% $\hat{h}_{t+1|t,i}$.

We propose a forecasting measure to incorporate the performance of $\hat{h}_{t+1|t,i}$ by the GARCH(1,1) model with $\hat{y}_{t+1|t,i}$ by the VAR(1), TVAR(1), TVAR(1)MD, TVAR(1)MV and TVAR(1)MIT models. The measure is the mean weighted of the forecasting error to the estimated variance $\hat{h}_{t+1|t,i}$ (MWES). The rational is that a large forecasting error is less misleading if it is accompanied by a large estimated of the variance, and hence wide prediction intervals.

$$MWES_i = \frac{\sum_{t=1}^{T-1}[\mathcal{E}_{t+1,i} \times \frac{1}{\hat{h}_{t+1|t,i}}]}{\sum_{t=1}^{T-1}\frac{1}{\hat{h}_{t+1|t,i}}}, \tag{18}$$
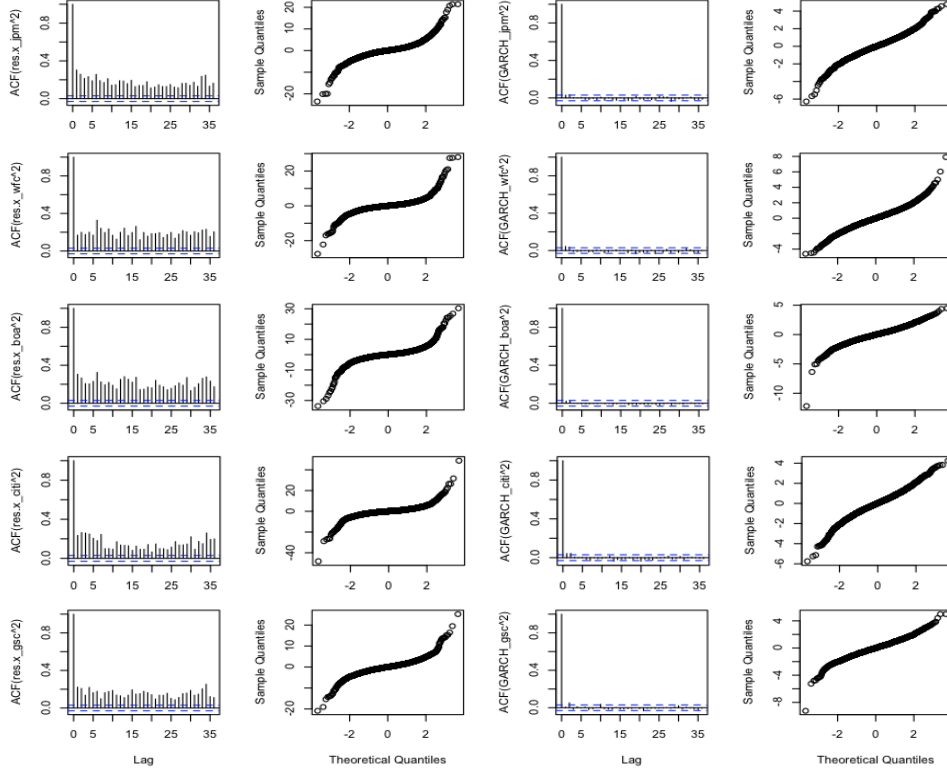
Figure 8: ACF and quantile-to-quantile (Q-Q) plots of the residuals squared of: the VAR(1) (column 1 and 2) and the VAR(1)-GARCH(1,1) (column 3 and 4), for the model components of JP Morgan (row 1), Wells Fargo (row 2), Bank of America (row 3), Citibank (row 4) and Goldman Sachs (row 5).

$$MWES = \frac{1}{I} \sum_{i=1}^{I} (MWES_i), \tag{19}$$

where $\mathcal{E}_{t+1,i}$ is the one-step ahead forecasting error given in Equation (13), T is the length of the observed share prices $\hat{y}_{t+1|t,i}$, and I is the number of banks. The $MWES_i$ in Equation (18) is the individual performance of each model component and the $MWES$ in Equation (19) is the overall forecasting performance. Results for the in-sample and 70:30% out-of-sample one-step ahead forecast using these measures are given in Table 10.

In the case of group performance for (a), the TVAR(1)MD-GARCH(1,1) model offers improvement on the VAR(1), TVAR(1) and TVAR(1)MV models. The TVAR(1) with MTI external threshold variable-GARCH(1,1) model performs better in most individual cases and provides further reduction in $MWES$. The VAR(1)-GARCH(1,1) model is better at the individual level for the out-of-sample case. The TVAR(1), TVAR(1)MD and

24

TVAR(1)MTI models with GARCH(1,1) perform equally better for the overall forecasting performance for this case.

Table 9: Forecasting performance by MWES.

Model	VAR(1)-GARCH(1,1)	TVAR(1)-GARCH(1,1)	TVAR(1)MD-GARCH(1,1)	TVAR(1)MV-GARCH(1,1)	TVAR(1)MTI-GARCH(1,1)
(a) In-sample one-step ahead forecast					
$MWES_{jpm}$	2.62×10^{-1}	2.64×10^{-1}	2.66×10^{-1}	2.62×10^{-1}	2.62×10^{-1}
$MWES_{wfc}$	1.03×10^{-1}	1.03×10^{-1}	1.03×10^{-1}	1.02×10^{-1}	1.02×10^{-1}
$MWES_{boa}$	9.23×10^{-2}	9.40×10^{-2}	9.37×10^{-2}	9.08×10^{-2}	9.11×10^{-2}
$MWES_{citi}$	1.16×10^{1}	1.16×10^{1}	1.13×10^{1}	1.11×10^{1}	1.11×10^{1}
$MWES_{gms}$	3.99	4.01	4.04	4.00	3.99
$MWES$	3.21	3.22	3.15	3.12	3.11
(b) 70:30% out-of-sample one-step ahead forecast					
$MWES_{jpm}$	3.89×10^{-1}	3.96×10^{-1}	4.00×10^{-1}	3.98×10^{-1}	3.91×10^{-1}
$MWES_{wfc}$	1.79×10^{-1}	1.81×10^{-1}	1.84×10^{-1}	1.86×10^{-1}	1.88×10^{-1}
$MWES_{boa}$	4.34×10^{-2}	4.39×10^{-2}	4.39×10^{-2}	4.33×10^{-2}	4.35×10^{-2}
$MWES_{citi}$	6.56×10^{-1}	3.92×10^{-1}	3.91×10^{-1}	3.87×10^{-1}	3.85×10^{-1}
$MWES_{gms}$	4.23	4.28	4.27	4.33	4.29
$MWES$	1.10	1.06	1.06	1.07	1.06

5   CONCLUSION

We have found strong evidence of directionality in the daily log-return time series, derived from closing share prices, for the five largest U.S. banks. Moreover, directionality appears to be a transient phenomenon, possibly as a consequence of feedback in the system following severe financial shocks which initiate crisis periods. In particular, the sign of moving directionality can change. Nevertheless, it persists for sufficiently long periods to be highly statistically significant. These findings suggest that the most informative strategy for monitoring directionality is to calculate it over a moving window, and we refer to this as moving directionality. We have demonstrated that moving directionality becomes more noticeable during crisis periods than it is during non-crisis periods.

The volatility of the daily log-returns series also increases during crisis periods. However, we have demonstrated that moving directionality becomes more noticeable during crisis periods than it is during non-crisis periods, and that it is somewhat more sensitive to crisis periods than moving volatility. The association between moving directionality,

which is partly related to third order properties of the time series such as skewness of the error distribution, and moving volatility, a second order property, is weak. In same cases at least, directionality tends to lead increased volatility and so can provide early warning of increased volatility. We conclude that directionality is a potentially useful feature of high frequency financial time series that is distinct from volatility. For example, monitoring directionality on a daily basis could alert investors to sell some of their shares and re-invest in lower risk options.

We also investigate time series of log-returns sampled at 1-minute intervals. It seems that the extraordinary high kurtosis of the 1-minute series affects the consistency of most of the directionality measures considered for the daily log-return series. The main exceptions were the POT and P_d^+ measures, and we will investigate these in future work. We conclude that end of the day values are more useful for monitoring directionality in the case of log-returns for the five U.S. banks.

We used multivariate autoregressive models of order one, VAR(1), and first-order non-linear threshold autoregressive, TVAR(1), with non-Gaussian errors, modelled by GARCH(1,1) processes, models for a portfolio of daily log-returns of the five U.S. banks to forecast the share prices. The VAR(1) models emulate very slight directionality through non-Gaussian errors, but the non-linear TVAR(1) models are better at reproducing the directionality seen in the time series. We considered the direct use of moving directionality as a regime-switching criterion between two-regime TVAR models, and compared this strategy with volatility regime-switching and with a combination of both directionality and volatility for the regime switching. The forecasting performance of the various models was compared in terms of root mean squared errors $MRSE$ and squared errors weighted inversely by the estimated variance of the prediction $MWES$. The TVAR(1) model offers slight but consistent improvement on the VAR(1) model. Further improvements are made in $MRSE$ and $MWES$ through the moving directionality threshold variable MD and the moving threshold indicator MTI, respectively for the in-sample one-step ahead forecasting. But, in the case of 70-30% out-of-sample forecast, the TVAR(1) thresholding on the past value was the best. It is possible that dynamic updating of the TAR(1)MD parameters might give an improved performance.

References

Allen, N. B. & Christa, H. S. B. (2013), 'How does capital affect bank performance during financial crises?', *Journal of Financial Economics* **109**(1), 146–176.

Awartani, B. M. & Corradi, V. (2005), 'Predicting the volatility of the {S&P}-500 stock index via {GARCH} models: the role of asymmetries', *International Journal of Forecasting* **21**(1), 167 – 183.

Chavez-Demoulin, V., Davison, A. C. & McNeil, A. J. (2005), 'A point process approach to value-at-risk estimation', *Quantitative Finance* **5**(2), 227–234.

Chavez-Demoulin, V. & McGill, J. (2012), 'High-frequency financial data modeling using hawkes processes', *Journal of Banking & Finance* **36**(12), 3415–3426.

Cont, R. (2001), 'Empirical properties of asset returns: stylized facts and statistical issues'.

Cowpertwait, P. S. P. & Metcalfe, A. V. (2009), *Introductory Time Series with R*, 1st edn, Springer Publishing Company, Incorporated.

Huang, X., Rui, Y., Shen, J. & Tian, Y. (2017), 'Us class action lawsuits targeting foreign firms: The country spillover effect', *Journal of Corporate Finance* .

Lamoureux, C. G. & Lastrapes, W. D. (1990), 'Persistence in variance, structural change, and the garch model', *Journal of Business & Economic Statistics* **8**(2), 225–234.

Lawrance, A. J. (1991), 'Directionality and reversibility in time series', *International Statistical Review/Revue Internationale de Statistique* **59**(1), 67–79.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2015), Modelling directionality in stationary geophysical time series, *in* 'Proceedings of the International work-conference on Time Series (ITISE 2015)', pp. 755–766.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2016a), *Time Series Analysis and Forecasting: Contributions to Statistics*, Springer International Publishing Switzerland, chapter Threshold Autoregressive Models for Directional Time Series.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2016b), 'Directionality and volatility in electroencephalogram time series', *American Institute of Physics (AIP) Conf. Proc.* **1739**, 020080:1–8.

Mansor, M. M., Mohd. Isa, F. L., Green, D. A. & Metcalfe, A. V. (2016c), 'Modelling directionality for paleoclimatic time series', *ANZIAM J.* **57** (EMAC2015), C66–C81.

Nelson, D. B. (1991), 'Conditional heteroskedasticity in asset returns: A new approach', *Econometrica* **59**, 347–370.

relbanks (2017), 'World's largest banks 2017', http://www.relbanks.com/worlds-top-banks/market-cap (last accessed 24 February 2017).

Chapter 5

Conclusion

5.1   Concluding Remarks

The findings from all the papers included in this thesis are summarized into four themes representing the overall significance of this study to the body of time series analysis and forecasting.

Theme I: Detecting directionality in time series.

The first stage in a time series analysis is usually to identify and then remove any trend and seasonal effects before fitting stationary time series models. It is straightforward to check for directionality after removing any trend and seasonal effects, and before fitting stationary models. If directionality is detected, then this indicates that the errors would be better modelled by a non-Gaussian distribution or that a non-linear model should be considered or both.

Detecting directionality in the following time series, collected from natural and anthropogenic processes from business, environmental science, finance and medicine, is considered in this study. Further information about these time series are summarised in Appendix B.

1. Sunspot numbers
2. deseasonalized Brisbane rainfall
3. Southern Oscillation Index
4. detrended and deseasonalized Vostok ice-core
5. deseasonalized NGRIP ice-core
6. EEG records from healthy subjects
7. EEG records from subjects diagnosed with epilepsy
8. detrended and deseasonalized AUS visitor arrivals
9. AUDUSD exchange rate
10. GBPUSD exchange rate
11. AUS 2-year bond yield
12. U.S. unemployment rate
13. JP Morgan log-returns
14. Wells Fargo log-returns
15. Bank of America log-returns
16. Citibank log-returns
17. Goldman Sachs log-returns

There are many ways in which a time series can exhibit directionality, and different quantitative tests detect different forms of directionality. In this study, seven statistics comprising three methods based on the original series and four methods based on the first differences are considered. The methods based on the original series are Difference in Linear Quadratic Lagged Correlations (DLQC), Markov Chain Detailed Balances (MCDB) and Peaks Over Threshold Test (POT). The first differences based methods are Percentage of Positive Differences (P_d^+), Product Moment Skewness of Differences ($\hat{\gamma}_d$), L-skewness of Differences (LSK) and Product Moment Skewness of Differences of Absolute Values about the Mean ($\hat{\gamma}_{dab}$). DLQC, $\hat{\gamma}_d$ and P_d^+ are the statistics found in the literature and the others are the alternative directionality measures proposed in this study. These seven statistics are referred to as a suite of directionality statistics. Directionality in time series is measured using the suite of directionality tests, and estimate their standard errors and compare the performance of the associated tests using Monte-Carlo procedures. The aim is to determine which tests are more powerful at detecting specific forms of directionality in time series.

The suite of directionality statistics distinguished four categories of directional characteristics in time series:

- Category 1: *A time series with rapid rises followed by slow recessions, or slow increases followed by fast recessions.*

  Eight time series fall into this category: the sunspots; the detrended and deseasonalized Vostok; deseasonalized NGRIP; the EEG (epilepsy) records; the AUDUSD exchange rate; the GBPUSD exchange rate; the AUS 2-year bond yield; and the U.S. unemployment rate. This form of directionality is best detected by the methods based on the first differences. The estimator $\hat{\gamma}_d$ appears to be the best for detecting directionality.

- Category 2: *A time series with rapid rises above the mean and rapid recessions below the mean (or rapid returns to the mean).*

  The statistic $\hat{\gamma}_{dab}$ was introduced to detect this category of directionality. However, the DLQC and POT statistics can also detect directionality if it is marked. The Southern Oscillation Index series and the deseasonalized Brisbane rainfall appear to be in this category.

- Category 3: *A time series with directionality or asymmetric patterns above a threshold.*

  Example of time series for this category are: (i) the time series in Category 1, except the exchange rates; and (ii) the time series in Category 2 because POT ignores information below the threshold. This form of directionality is best detectable by the POT but sensitive to the chosen averaging parameter between the mean of number observations before a peak and the mean of number observations following the peak.

- Category 4: *Intermittent directionality.*

  Directionality seems to be intermittent in some time series, share log-returns in particular. This can be monitored by calculating a moving directionality index. This could be based on any of the measures but we have focused on a moving form of the product moment skewness of differences that appears to work well with shorter series.

  In this study, we discovered that directionality is a feature of high frequency time series, particularly in the log-returns of U.S. bank shares. The suite of directionality tests is needed for high frequency time series because the graphical inspection of directionality became impractical for a long time series, and directionality may not be apparent in short sub-series of the log-returns. Directionality was found to be a transient phenomenon and to vary with classification of the period into financially stable and unstable sub-periods. This may be a consequence of intermittent feedback in the system following extreme events of severe financial shocks. The directionality is also quite sensitive to sampling intervals, similar to the way volatility works. Moreover, we demonstrated in this study that the directionality appeared to be more marked, consistently high and positive – more affected by financial shocks, during crisis periods than it is during non-crisis periods.

Through this study, we also developed a general procedure to detect directionality in univariate time series, involving the suite of directionality tests. The proposed procedure is given in Figure 1, in which a general process for detecting directional time series is summarised.

In conclusion, because certain tests are more sensitive at detecting specific forms of directionality in univariate time series, we recommend the suite of seven statistics remains as part of the general procedure.
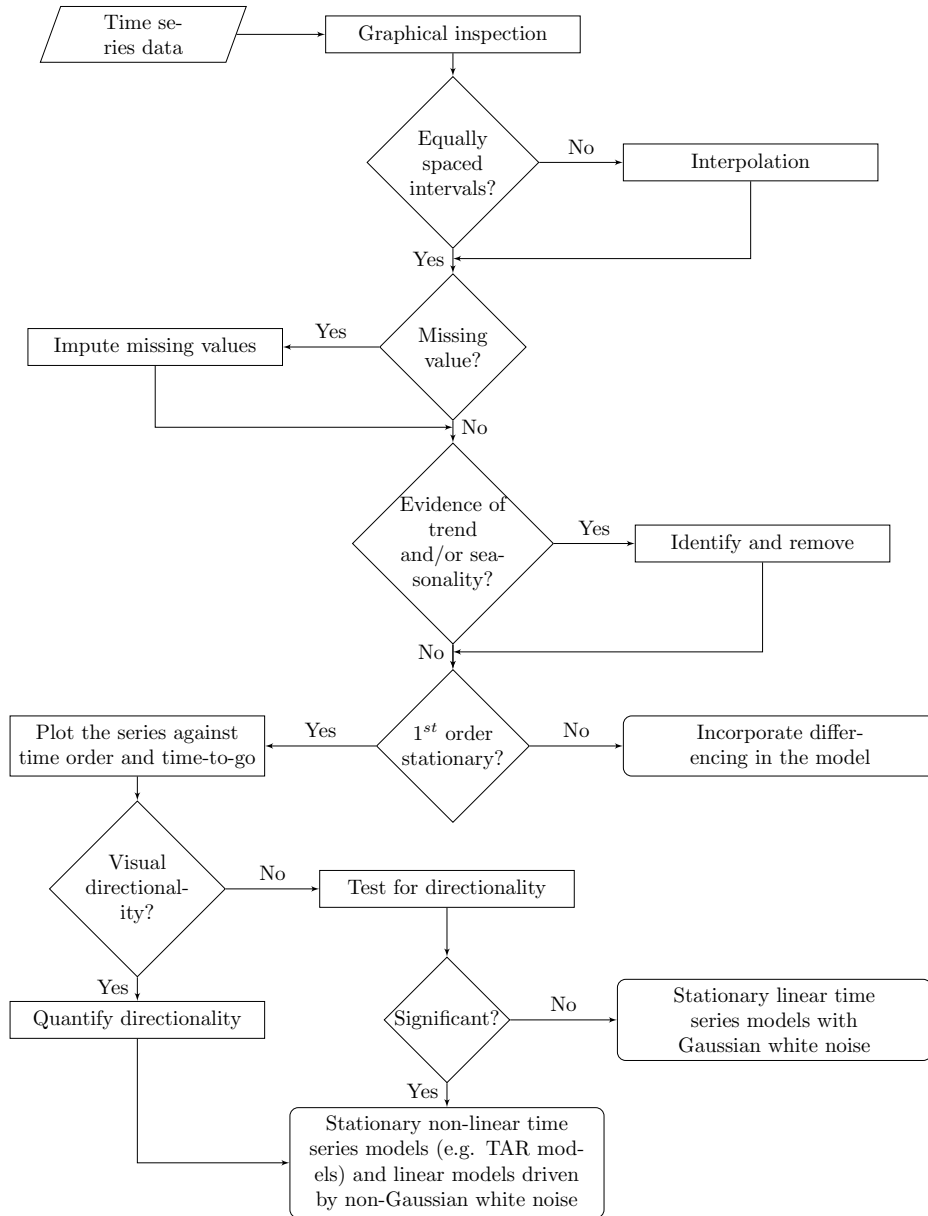
Figure 1: A general procedure for detecting directionality in time series.

It is noted in this study that in some cases there are physical reasons to expect directionality in time series. For example, stream flow will increase rapidly following a storm because of the immediate overland run-off, but will decrease slowly due to the rain infiltrating the catchment and augmenting the ground water flow to the river. Once directionality has been detected in a time series, it can provide insight into the physical processes underlying the time series.

A brief discussion about the link between directionality and extreme events follows. Vigorous magnetic fields of the sun for the sunspots and extreme paleoclimatic events in ice-core time series is suggested. For example, sunspots are dark regions of intense magnetic field that are associated with solar flares and coronal mass ejections. They are relatively cool because their intense magnetic fields inhibit the rise of heat from the solar interior, but are relatively bright at higher frequencies of radiation. Sunspots can appear alone, or in close connection to other sunspots making an active region. The energy that sunspots lack in heat is compensated for by the energy of the magnetic field. The magnetic fields rise above the surface and remain strong, while the rest of the sun has weak overlying magnetic fields. The strong magnetic fields form into loops that confine solar plasma and heat it to extreme temperatures in excess of 1 million K. It is highly plausible that such extreme electromagnetic events will promote feedback mechanisms (National Oceanic and Atmospheric Administration).

Theme II: Applications from the detection of directionality.

(a) Application to Finance

In relation to the nature of directionality in the high frequency log-returns, a moving directionality index over a 9 month period is defined for the five U.S banks, to monitor the behaviour of directionality during crisis and non-crisis periods, in comparison to moving volatility index. We show in the study that the moving directionality and the moving volatility are loosely correlated, and directionality is leading the occurrence of volatility in the case of Wells Fargo, Bank of America and Citibank. From this, we conclude that monitoring directionality could provide early warning of increased volatility and an early warning indicator of falling share prices. As a result, directionality may have potential as criterion for a stop-loss order, to sell shares when the price reaches a certain level, in equity trading to limit the investment losses.

The strategy is to monitor the directionality of the log-returns using 9-month moving directionality for each bank and to use this as an investment criterion with the aim of benefiting from market uncertainties. The proposed directionality trading rules for buying and selling the shares are:

- If $MD_i > T_i$, then we invest all the money allocated to $Bank_i$ in its shares.

- If $MD_i \leq T_i$, then we sell the all shares in $Bank_i$; and

where MD_i is the moving directionality, T_i is set at the overall measure of directionality for $Bank_i$.

The rationale is to buy shares when the MD value is more than the observed directionality because our observation is that on many occasions the share price is low when the corresponding MD value is higher than the threshold value. In contrast, the share price is high when the corresponding MD value is lower than the threshold value (particularly for JP Morgan and Goldman Sachs). This is consistent with the famous "buy low, sell high" strategy to take advantage of the market instability. In the simulation, directionality trading rules perform consistently better than volatility trading rules for a portfolio of the five U.S. bank shares. While this investment simulation is highly-simplified, it suggests that monitoring directionality may have value as an aid to investment and merits further study.

(b) Application to Medicine

Detecting directionality in electroencephalogram (EEG) time series provides potential physical interpretation of directionality found in the time series. Directionality was found in the EEG signals sampled from subjects diagnosed with epilepsy but not in the EEG records sampled from volunteers with no diagnosis of epilepsy. This finding suggests that a directional EEG time series may not be a good sign in humans. Although this interpretation is highly speculative, the possibility that directionality might provide or contribute towards early warning of an epileptic seizure is sufficiently important to warrant further investigation. This would be very meaningful to those people prone to epilepsy because they could take precautionary measures to mitigate untoward events.

Theme III: Time series models that explicitly emulate direction-ality.

There are many possible models for time series. As with any mathematical modelling there is no correct model, however, a good model will give a close match to the observed statistic. There are at least two reasons for fitting time series models: (i) to make short term forecasts; and (ii) for simulation studies.

Detecting directionality is useful because it indicates when non-linear time series models, or at least non-Gaussian errors, are appropriate, leading to more accurate forecasts and more realistic scenarios (Lawrance 1991) than can otherwise be made. In this study, we investigated the modelling aspects of directionality using threshold autoregressive (TAR) models with Gaussian, non-Gaussian, or resampled residuals from the fitted models for the errors, and autoregressive (AR) models with non-Gaussian errors to improve the fit in terms of directionality statistic found in the studies.

The goodness-of-fit between time series models used in this study was measured by the standard deviation of the errors. In general, TAR models were found to be satisfactory for the directional time series based on the least squares error criterion, and substantially simpler relative to AR model ruled by the Akaike information criterion (AIC). All the estimated model parameters were satisfied by the constraints for stationarity.

We considered Beta and t-distributions for the symmetric non-Gaussian errors in the study. Excess kurtosis of the chosen distributions need to be quite extreme to see substantial directionality for a first order AR processes. For the asymmetric non-Gaussian errors, we used exponential, Gumbel, and back-to-back Weibull distributions on AR and TAR models to investigate directionality through simulations. We demonstrated in the study that the choice of non-Gaussian error distributions could lead to a variety of significant directional features.

TAR models fitted by the penalized least squares (PLS) strategy produced improved results in terms of reproduction of directionally for the time series, with relatively small increase (or small decrease) in the sum of squared residuals. We demonstrated that TAR models with PLS strategy are reproducible models for modelling directionality in the time series of: sunspots; GBPUSD exchange rate; Australian two-year bond yield rate; U.S. unemployment rate; EEG signals from a patient diagnosed with epilepsy; log-

returns of JP Morgan; and ice-core records of Vostok and NGRIP.

When the PLS procedure with directionality fitting criterion was found to give lower marginal standard deviation than the observed standard deviation, then the standard deviation of the fitted series can be included as part of the optimization criteria. Although this modification did not noticeably increase the runtime, the optimization problem has become more challenging, particularly when determining the optimum combination of the criterion's weights for the routine to minimise not only the discrepancies in directionality and standard deviation, but also the sum of squared errors.

Theme IV: Applications from modelling directionality.

(a) Application for testing investment strategy

In general, TAR models with suitable non-Gaussian errors will provide realistic simulation of directional processes. We have shown that TAR models, fitted by using least squares and the PLS strategy, with non-Gaussian residuals, modelled by GARCH model, for the errors closely emulates the target directionality for JP Morgan log-returns. The TAR model with the best fit in term of directionality is the used in the stochastic simulations to calculate the conditional value at risk (CVaR) from a distribution of the simulated annualised returns for choosing an investment strategy in the investment simulations.

(b) Application to univariate forecasting for the sunspots

In this study. we demonstrated that the TAR model fitted by the PLS strategy has significantly reduced the target discrepancies between the observed and the simulated directionality. This has yielded more accurate forecasting and more realistic simulations of extreme values for the directional time series of sunspots.

(c) Application to multivariate forecasting for the log-returns

Some empirical evidence of directionality in high frequency multivariate time series of log-returns of the five largest U.S banks for the prediction of unstable periods has been presented in the study, showing that the directionality has only a low association with volatility which could provide additional information to volatility.

Multivariate autoregressive models of order one (VAR(1)) and first-order non-linear threshold autoregressive (TVAR(1)) models with non-Gaussian errors modelled by GARCH(1,1) processes for a portfolio of daily log-returns of the five U.S. banks to forecast the share prices were also considered in the study. While the VAR(1) models emulate very slight directionality through non-Gaussian errors, the non-linear TVAR(1) models are better at reproducing the directionality seen in the time series as threshold models provide a piecewise linear approximation to a wide range of non-linear processes (Tong & Lim, 1980).

We demonstrated that directionality becomes more apparent during crisis periods than it is in non-crisis periods for all five log-return series. This characteristic can be modelled by fitting separate sets of coefficients to the TVAR(1) model during periods of high and low directionality. We refer to this as a two-regime TVAR(1) model where the regime is determined by whether or not moving directionality (MD) exceeds an external threshold (TVAR(1)MD). The direct use of moving directionality as a regime-switching criterion between two-regime TVAR models is then compared with moving volatility (MV) regime-switching, and with a combination of both moving directionality and volatility (moving threshold indicator MTI) for the regime switching.

The forecasting performance of the various models was compared in terms of root mean squared errors $MRSE$ and squared errors weighted inversely by the estimated variance of the prediction $MWES$. The TVAR(1) model offers slight, yet consistent improvement on the VAR(1) model. Further improvements are made in $MRSE$ and $MWES$ through the moving directionality threshold variable MD and the moving threshold indicator MTI respectively, for the in-sample one-step ahead forecasting. However, in the case of 70-30% out-of-sample forecast, the TVAR(1) thresholding on the past value was the best. It is possible that dynamic updating of the TAR(1)MD parameters might give an improved performance.

In conclusion, the time series models incorporating directionality or associated with directionality provide promising results in terms of prediction limits and more realistic simulations than the models that do not. This finding suggests that directionality could be used as a criterion for choosing stationary non-linear time series models, particularly TAR model and TVAR models for univariate and multivariate directional time series, respectively.

5.2 Further Work

Going forward, I plan to develop an R package for the suite of directionality statistics including the Monte-Carlo procedures to calculate the statistical significance of directionality. This package would benefit R users with similar interests.

Once evidence of directionality has been established, a suitable time series model needs to be identified. There are no correct time series models and, in contrast to linear Gaussian models, a limitless set of possibilities ready for discovery. We have considered in this study TAR and AR models with a range of non-Gaussian errors including shocks (outlier), volatility and resampled residuals. There is additional scope for developing guidelines for choosing models for directional time series.

The physical explanations of directionality using medicine time series should be continued, for example, by assessing directionality in blood glucose and blood insulin time series from healthy subjects and comparing the findings with subjects diagnosed with diabetes or obesity. Another possibility is to compare directionality in blood glucose records after plasmapheresis and steroid treatments. These investigations could lead to new knowledge for therapeutic or diagnostic purposes. I am also keen to continue the work on epilepsy. The brain signals for the epilepsy group were taken from recordings using the intracranial strip electrodes scheme for pre-surgical assessment of epilepsy patients from the hippocampal formation zone, the epileptic site of the brain during seizure free intervals, and the seizure zone. The EEG samples set at the epileptogenic and seizure zones of the subject's brain were remarkably different, and the epileptogenic series were highly directional and the series from seizure zone were highly variable. Further work may show that these findings have diagnostic potential. In particular, it would be valuable to take EEG readings from healthy subjects from the same epileptic and seizure sites of the brain. Also, a larger sample size than used in our study would help generalising the findings.

There are many other possible extensions to this study in the future, for example, to investigate directionality in non-linear music production or to explore the inclusion of directionality as an updating function into the Holt-Winters method alongside with the three smoothing functions for level, trend and seasonality. I am also keen to consider real-world applications of direc-

tionality, for example, the use of directionality in trading rules for buying and selling shares. Although the financial simulations in our study showed possible potential, substantial additional work is required to give the findings greater practical implication. I could consider a portfolio of shares of sector mix, for example, an investment portfolio of shares from financial, telecommunication, technology and transportation sectors. I could also compare the performance of directionality trading rules to these sector individually. I have noted that the investment simulation is highly simplified. Instead of entering into the money market while waiting for the next buy, I could also consider several other asset classes that are easy to liquidate with less impact on the investment value. I also noted that switching the investment from shares into other investment instruments or between shares involves multiple brokers' commissions and switching fees that are relevant when calculating the return on investment (ROI). Furthermore, there is a yearly dividend cash payout to investors electing to stay with the company during bullish or even bearish periods that can be reinvested into shares. The real ROI should consider the revenue minus the capital, deducting brokers' commissions, any switching fees and taxes on investment plus dividends.

Appendix A

Artificial Time Series

Testing for Directionality in Artificial Time Series

We have considered four deterministic series of length 600, with and without Gaussian noise added. We used the suite of statistics to detect directionality in these artificial time series, based on the general procedure for detecting directionality in univariate time series illustrated in Paper 1.

i. Deterministic Series

(a) Symmetric (control case)

(b) Sawtooth profile

(c) Rotating-shark

(d) Repeating-chirp

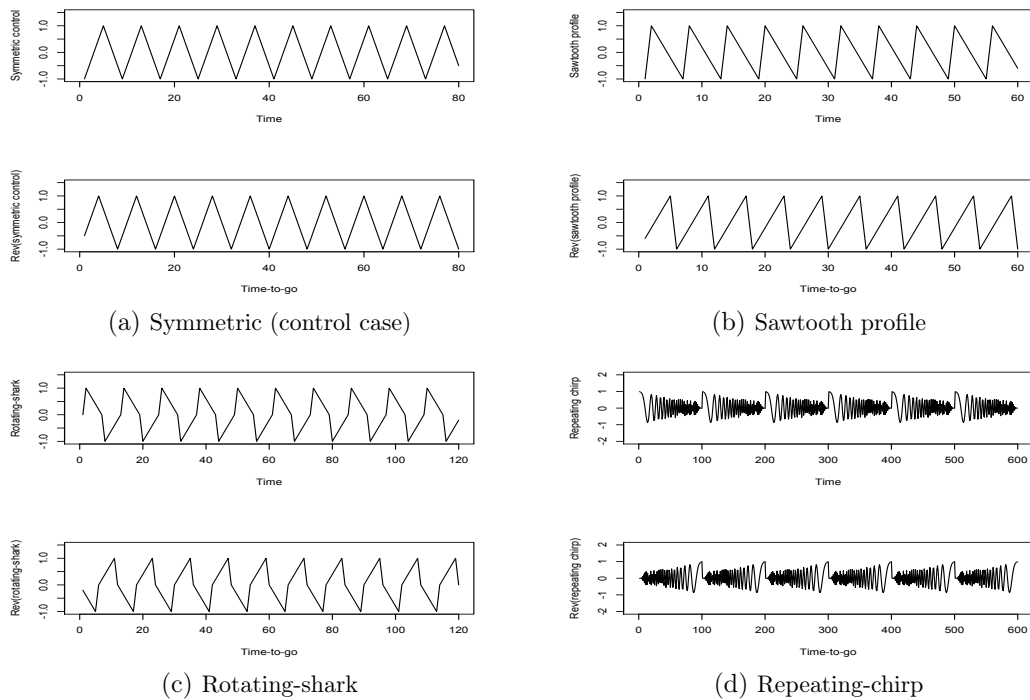Figure 1: Deterministic series.

The deterministic stationary time series in Figure 1 are described as follows.

- A symmetric triangular pattern (-1, -0.5, 0, 0.5, 1, 0.5, 0, -0.5) repeated 75 times (Figure 1(a)).

- A sawtooth profile, sharp increase and slow decrease, defined by the sequence (-1, 1, 0.6, 0.2, -0.2, -0.6) repeated 100 times. This is asymmetric about the median value of 0 (Figure 1(b)).

- A profile that we refer to as a rotating-shark fin (rotating-shark) that has both sharp increases and sharp decreases followed by slow returns to the mean value of 0. It is defined by the sequence (0, 1, 0.8, 0.6, 0.4, 0.2, 0, -1, -0.8, -0.6, -0.4, -0.2) repeated 50 times. This is symmetric about the mean value of 0 (Figure 1(c)).

- A repeating chirp which has a clear directionality while being symmetric about the mean. It is defined as $e^{-t/100}cos(2\pi f_t t)$ repeated 6 times, where $t = 1, 2, ..., 100$ and $f_t = \frac{t}{4(max(t)-min(t))} + 1$ (Figure 1(d)).

Although these series have a repeating pattern, the period is considered unknown to distinguish it from seasonality, and the aim is to investigate the performance of measures when a time series has known asymmetry. We apply the tests of directionality to these series, to verify directionality, and the results are given in Table 1 together with the length n and the standard deviation s of the series.

Table 1: Test statistics of directionality for deterministic series.

Method based on	The original series					The first differences			
Artificial series	DLQC	MCDB	POT(2)	POT(3)	POT(5)	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
(a) Symmetric (control) n=600, s=0.61	0.00	0.00	$\frac{0}{0}$ {	$\frac{0}{0}$ T=0.50	$\frac{0}{0}$ }	50.1%	0.00	0.00	0.00
(b) Sawtooth profile n=600, s=0.68	-1.33	0.67	$\frac{1.2}{0}$ {	$\frac{0.8}{0}$ T=0.60	$\frac{0}{0}$ }	16.7%	1.78	0.67	0.00
(c) Rotating-shark n=600, s=0.61	0.00	0.25	$\frac{0.8}{0}$ {	$\frac{0.8}{0}$ T=0.60	$\frac{0.8}{0}$ }	49.9%	0.00	0.00	1.78
(d) Repeating-chirp n=600, s=0.47	0.00	0.07	2.31 {	1.92 T=0.49	2.07 }	47.9%	-0.01	0.00	0.38

Note: For the POT, T refers to threshold value at 80-th percentile in the original series. The result of the POT test is presented as (mean difference)/(standard error of mean differences), to avoid NaN as a consequence of division by 0.

- *Control.* The results for the symmetric control are as expected and there is no indication of directionality.

- *Sawtooth.* DLQC, MCDB, P_d^+, $\hat{\gamma}_d$ and LSK all indicate directionality. The POT method also detects directionality with averaging of 2 or 3 observations before and after the peak but not with an averaging of 5 suggesting that this statistic may be sensitive to the averaging parameter used. The remaining measure $\hat{\gamma}_{dab}$, designed to

2

detect directionality in time series models that are symmetric about the median, is 0 as expected for the sawtooth.

- *Rotating-shark.* The signal is symmetric about the time series median and only MCDB, POT and $\hat{\gamma}_{dab}$, which was designed for this case, detect the directionality.

- *Repeating-chirp.* This signal has a clear repeating pattern and the increasing frequency through each chirp defines a clear direction. The directionality measures for the repeating chirp are sensitive to the sampling interval, but the POT, P_d^+ and $\hat{\gamma}_{dab}$ do detect directionality with the time series of length 600.

ii. Deterministic Series with Added Gaussian White Noise

(a) Symmetric (control) with AWGN   (b) Sawtooth profile with AWGN

(c) Rotating-shark with AWGN   (d) Repeating-chirp with AWGN
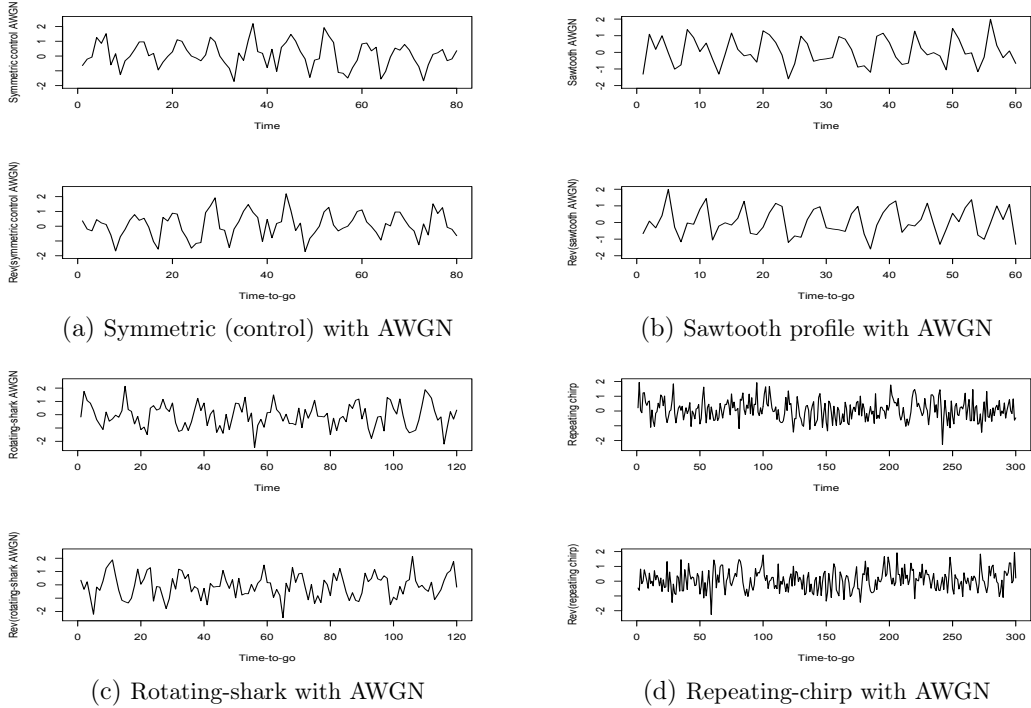
Figure 2: Deterministic series with AWGN.

Next, we create another four artificial time series that are more realistic by including additive Gaussian white noise (AGWN), with mean 0 and standard deviation 0.5, to the deterministic series in the previous section. The deterministic series with AGWN are plotted in time order (upper sub-figure) and in reverse time order (lower) in Figure 2.

Table 2: Summary table of test statistics of directionality and p-values for deterministic series with added Gaussian white noise (AGWN) with mean 0 and standard deviation 0.5.

Method	Based on the original series					Based on the first differences			
Artificial series	DLQC	MCDB	POT(2)	POT(3)	POT(5)	P_d^+	$\hat{\gamma}_d$	LSK	$\hat{\gamma}_{dab}$
(a) Symmetric (control) with AGWN $n{=}600,\ s{=}0.77$	-0.03	0.05 {0.05}	0.26 {	-0.80 $T{=}0.69$	-0.57 }	49.4% {0.00%}	0.06	0.01	0.01
AR(21), $\hat{\sigma}_{err}{=}0.54$	(0.55) [0.55]	(0.32) [0.39]	(0.80) [0.26]	(0.49) [0.72]	(0.59) [0.57]	(0.58) [0.48]	(0.50) [0.66]	(0.40) [0.83]	(0.96) [0.06]
(b) Sawtooth profile with AGWN $n{=}600,\ s{=}0.84$	-0.77	0.33 {0.06}	11.7 {	5.77 $T{=}0.78$	-0.18 }	42.1% {0.00%}	0.90	0.19	0.04
AR(24), $\hat{\sigma}_{err}{=}0.57$	(0.00) [5.27]	(0.00) [10.5]	(0.00) [6.25]	(0.00) [4.66]	(0.85) [0.18]	(0.00) [4.83]	(0.00) [5.52]	(0.00) [6.33]	(0.67) [0.41]
(c) Rotating-shark with AGWN $n{=}600,\ s{=}0.78$	-0.06	0.03 {0.04}	-0.39 {	1.91 $T{=}0.65$	2.78 }	48.4% {0.00%}	0.14	0.02	-0.07
AR(24), $\hat{\sigma}_{err}{=}0.55$	(0.18) [1.27]	(0.64) [0.43]	(0.69) [0.40]	(0.07) [1.81]	(0.02) [2.36]	(0.18) [1.33]	(0.12) [1.55]	(0.27) [1.13]	(0.42) [0.78]
Rotating-shark with 0.50×AGWN $n{=}600,\ s{=}0.65$	-0.02	0.10 {0.04}	6.91 {	7.55 $T{=}0.60$	11.3 }	50.4% {0.00%}	0.10	0.002	0.53
AR(24), $\hat{\sigma}_{err}{=}0.30$	(0.40) [0.66]	(0.01) [2.89]	(0.00) [5.68]	(0.00) [5.33]	(0.00) [7.04]	(0.65) [0.36]	(0.22) [1.23]	(0.91) [0.11]	(0.00) [3.88]
(d) Repeating-chirp with 0.50×AGWN $n{=}600,\ s{=}0.55$	0.02	0.04 {0.05}	-0.84 {	0.01 $T{=}0.54$	0.11 }	50.3% {0.00%}	-0.04	-0.01	0.08
AR(1), $\hat{\sigma}_{err}{=}0.52$	(0.62) [0.48]	(0.51) [0.12]	(0.41) [0.82]	(0.99) [0.01]	(0.93) [0.11]	(0.79) [0.20]	(0.67) [0.44]	(0.74) [0.34]	(0.36) [0.91]

Notes for Table 2:
1. In the series description: n is the length of the time series; s is the marginal standard deviation; p, in AR(p), is the order of the AR model fitted using AIC; and $\hat{\sigma}_{err}$ is the estimated standard deviation of residuals in the AR(p) model.
2. (number) refers to p-value, [number] refers to the ratio of the test statistic to its standard deviation and {number} in MCDB column is the mean of the test statistic from Monte-Carlo procedure. {number} in P_d^+ refers to proportion of zero in the distribution of differences in original series. T in POT columns refers to threshold value at 80-th percentile in the original series. POT(2), POT(3) and POT(5) are the POT with averaging of 2, 3 and 5 observations before and after the peak.

For the sawtooth profile (Figure 2(b)), it is possible to detect the relatively rapid increases from the time and time-to-go plots. However, it is hard to discern any differences from the plots in the other cases. Results from formal statistical tests are given in Table 2.

- *Control AGWN*. As expected, none of these statistics detects directionality for the symmetric (control) with AWGN.

- *Sawtooth AGWN*. The first six statistics all detect the directionality, except POT(5) which indicates that an averaging parameter of 5 in POT is not suitable for this case. As expected, $\hat{\gamma}_{abm}$ does not detect directionality.

- *Rotating-shark AGWN*. The series is only detected statistically directional by POT(5), although MCDB and $\hat{\gamma}_{dab}$ also showed substantial directionality in the deterministic signal. However, we can see similar performance by the MCDB, POT(2), POT(3), POT(5) and $\hat{\gamma}_{dab}$ when the signal to noise ratio is doubled (rotating-shark with half amplitude AGWN).

- *Repeating-chirp with half amplitude AGWN*. The repeating pattern is obscured by the AWGN, and directionality is not detected by any of the statistics at the 5% level of significance. However, a wavelet analysis is generally suitable for detecting a repeating deterministic pattern within which there is a change of frequency. Figure 3 is a Haar wavelet analysis of the first 2^9 points in the time series (512 out of 600), drawn using the Haar wavelet option in the package wavethresh in the $\boldsymbol{R}$ software. The first five chirps appear with the frequency increasing from 4 to 8 as time increases. The beginning of the sixth chirp is also apparent.
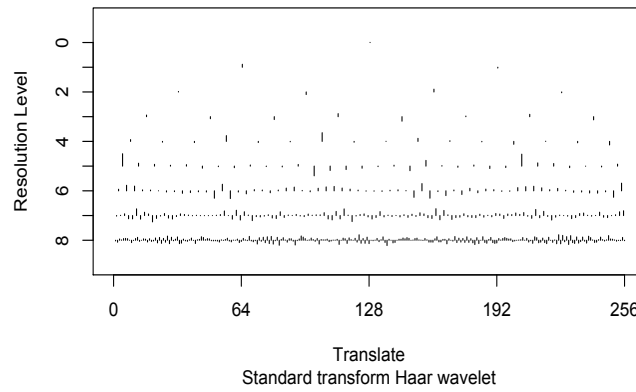
Figure 3: Haar discrete wavelet coefficients plot for repeating-chirp with half amplitude AGWN.

We can draw some tentative conclusions from these simulation results. They suggest that $\hat{\gamma}_d$ is preferable to P_d^+ and that LSK offers little advantages over $\hat{\gamma}_d$. The $\hat{\gamma}_{dab}$ has potential to detect directionality in a series that is symmetric about the time series median, when the other methods based of the first differences do not. A useful set of statistics for detecting directionality might be restricted to MCDB, POT, $\hat{\gamma}_d$ and $\hat{\gamma}_{dab}$. Our results showed that, none of the tests provided statistically significant evidence of directionality in the symmetric (control) series.

Appendix B

Time Series Data

A List of Time Series Used in the Study

Time series	Time interval	Process	Discipline	Directionality	Category	Source
Sunspot numbers	Daily	Natural	Environmental Sc.	Yes	1 & 3	SILSO, Belgium
ds Brisbane rainfall	Yearly	Natural	Environmental Sc.	Yes	2 & 3	Australian Bureau of Meteorology
Southern Oscillation Index	Monthly	Natural	Environmental Sc.	Yes	2 & 3	Australian Bureau of Meteorology
dtds Vostok ice-core	50-year	Natural	Environmental Sc.	Yes	1 & 3	Petit et al., 1999
ds NGRIP ice-core	50-year	Natural	Environmental Sc.	Yes	1 & 3	NGRIP members
EEG (normal subjects)	173.61 Hz	Natural	Medicine	No	NA	Klinik für Epileptologie
EEG (epilepsy patients)	173.61 Hz	Natural	Medicine	Yes	1 & 3	Klinik für Epileptologie
dtds AUS visitor arrivals	Monthly	Anthropogenic	Business	No	NA	Australian Bureau of Statistics
AUD-USD exchange rate	Daily	Anthropogenic	Finance	Yes	1	Bloomberg
GBP-USD exchange rate	Daily, quarterly	Anthropogenic	Finance	Yes	1	Bloomberg
AUS 2-year bond yield	Monthly	Anthropogenic	Finance	Yes	1 & 3	Reserve Bank of Australia
U.S. unemployment rate	Yearly	Anthropogenic	Finance	Yes	1 & 3	U.S. Bureau of Labor Statistics
JP Morgan log-returns	Daily, minutes	Anthropogenic	Finance	Yes	intermittent	Bloomberg, Tickdatamarket
Wells Fargo log-returns	Daily, minutes	Anthropogenic	Finance	Yes	intermittent	Bloomberg, Tickdatamarket
Bank of America log-returns	Daily, minutes	Anthropogenic	Finance	Yes	intermittent	Bloomberg, Tickdatamarket
Citibank log-returns	Daily, minutes	Anthropogenic	Finance	Yes	intermittent	Bloomberg, Tickdatamarket
Goldman Sachs log-returns	Daily, minutes	Anthropogenic	Finance	Yes	intermittent	Bloomberg, Tickdatamarket

Appendix C

R Code

The Suite of Directionality Tests

The following algorithms assume an evenly spaced intervals time series that contains no missing values. Any evidence of trend and/or seasonality in a time series should be identified and removed, which can be regarded as a realization of a stationary random process, before considering directionality.

i. Visual Inspection

```
### load a time series in the row vector form
data=read.table{utils}
### the series in time order
data.ts=ts(data, frequency=1, start=the time of the first point)
### the series in reverse time order (time-to-go)
data.reverse.ts=ts(rev(data), frequency=1, start=0))
### plot the series to discern directionality in plots
par(mfrow=c(2,1))
plot(as.ts(data.ts), xlab="Time", ylab="the series")
plot(as.ts(data.reverse.ts), xlab="Time-to-go", ylab="Rev(the series)")
```

ii. Difference in Linear Quadratic Lagged Correlations (DLQC)

```
### define x as the original series
x=data
### start DLQC
mx=x-mean(x)
sum12=sum21=0
for(i in 1:(length(x)-1))
{
sum12=sum12+mx[i]*mx[i+1]^2
sum21=sum21+mx[i]^2*mx[i+1]
}
DLQC=(sum12/length(x)-sum21/length(x))/sd(x)^(3)
print(c("DLQC",round(DLQC, digits=4)))
### end DLQC
```

iii. Markov Chain Detailed Balances (MCDB)

```
### define x as the original series
x=data
### start MCDB
### calculate first quantile, Q1
Q1=sort(x)[floor(0.25*length(x))]
### calculate second quantile, Q2
Q2=sort(x)[floor(0.50*length(x))]
```

```
### calculate third quantile, Q3
Q3=sort(x)[floor(0.75*length(x))]
### divide a time series into a four state Markov Chain and
### label each observation according to its state
x.lab=NULL
for(i in 1:length(x))
{
### label x as 4 if x>Q3 (state 4)
if(x[i]>Q3){x.lab=c(x.lab,4)}
### label x as 3 if Q3>=x>Q2 (state 3)
else if(x[i]>Q2){x.lab=c(x.lab,3)}
### label x as 2 if Q2>=x>Q1 (state 2)
else if(x[i]>Q1){x.lab=c(x.lab,2)}
### label x as 1 if otherwise (state 1)
else{x.lab=c(x.lab,1)}
}
### define all possible pairs of transition, e.g. m13 is a transition
### between a point of state 1 and the next point of state 3, in sequence
m11=m12=m13=m14=0
m21=m22=m23=m24=0
m31=m32=m33=m34=0
m41=m42=m43=m44=0
### count no. of transitions for each pair
for(t in 1:(length(x)-1))
{
if(x.lab[t]==1 && x.lab[t+1]==1){m11=m11+1}
else if(x.lab[t]==1 && x.lab[t+1]==2){m12=m12+1}
else if(x.lab[t]==1 && x.lab[t+1]==3){m13=m13+1}
else if(x.lab[t]==1 && x.lab[t+1]==4){m14=m14+1}
else if(x.lab[t]==2 && x.lab[t+1]==1){m21=m21+1}
else if(x.lab[t]==2 && x.lab[t+1]==2){m22=m22+1}
else if(x.lab[t]==2 && x.lab[t+1]==3){m23=m23+1}
else if(x.lab[t]==2 && x.lab[t+1]==4){m24=m24+1}
else if(x.lab[t]==3 && x.lab[t+1]==1){m31=m31+1}
else if(x.lab[t]==3 && x.lab[t+1]==2){m32=m32+1}
else if(x.lab[t]==3 && x.lab[t+1]==3){m33=m33+1}
else if(x.lab[t]==3 && x.lab[t+1]==4){m34=m34+1}
else if(x.lab[t]==4 && x.lab[t+1]==1){m41=m41+1}
else if(x.lab[t]==4 && x.lab[t+1]==2){m42=m42+1}
else if(x.lab[t]==4 && x.lab[t+1]==3){m43=m43+1}
else {m44=m44+1}
}
### calculate total no. of transitions for each row (R1:R4)
```

```
R1=m11+m12+m13+m14
R2=m21+m22+m23+m24
R3=m31+m32+m33+m34
R4=m41+m42+m43+m44
### Markov chain (MC) matrix consists transition probability for each pair
Row1=c(m11/R1, m12/R1, m13/R1, m14/R1)
Row2=c(m21/R2, m22/R2, m23/R2, m24/R2)
Row3=c(m31/R3, m32/R3, m33/R3, m34/R3)
Row4=c(m41/R4, m42/R4, m43/R4, m44/R4)
MC=rbind(Row1,Row2,Row3,Row4)
### calculate steady-state for MC
### note that (2^n)-step ahead varies for different series
A=MC%*%MC
B=A%*%A
C=B%*%B
.  .    .
.  .    .
L=K%*%K
M=L%*%L
### check steady state status for MC, use matrix M and L to verify
check=sum(round(sum(M[1:length(MC)]-L[1:length(MC)]), digits=4))
### extract pi from the diagonal elements of the state state MC
pi1=M[1,1]
pi2=M[2,2]
pi3=M[3,3]
pi4=M[4,4]
### calculate detailed balance
if(check==0)
{
MCDB=abs(pi1*MC[1,2]-pi2*MC[2,1]) + abs(pi1*MC[1,3]-pi3*MC[3,1]) +
abs(pi1*MC[1,4]-pi4*MC[4,1]) + abs(pi2*MC[2,3]-pi3*MC[3,2]) +
abs(pi2*MC[2,4]-pi4*MC[4,2]) + abs(pi3*MC[3,4]-pi4*MC[4,3])
print(c("MCDB",round(MCDB, digits=4)))
}else{stop("Diagonal entries are not in steady state,increase n-step ahead")}
### end MCDB
```

iv. Peaks Over Threshold Test (POT)

```
### define x as the original series
x=data
### start POT
### set the threshold level (TL), e.g. 80%
TL=0.80
### compute threshold value (T)
```

```
T=sort(x)[floor(TL*length(x))]
### set h as the observations before and after a peak e.g. 5
h=5
### AT is set to be a vector of observations above T
AT=x[x>T]
### LT is the location of AT in the data set
LT=which(x>T)
### identify location of candidate for peaks that are the maximum
### value in independent excursions above T
Lmax=0
Lm=0
for(i in 1:length(LT))
{
if(length(LT) != 0){
Lmax[i]=which(x==max(x,na.rm=TRUE))[1]
M=Lmax[i]
if(length(which(!is.na(x[abs(M-h):abs(M+h)])))==(2*h+1)){
x=replace(x[1:length(data)],abs(M-h):abs(M+h),NaN)
Lm[i]=M
}else{
x=replace(x[1:length(data)],M,NaN)
Lm[i]=NaN
}
AT=x[x>T]
LT=which(x>T)
}}
### remove NaN in Lm, a vector of locations of candidate for peaks
Loc=Lm[which(!is.na(Lm))]
### ensure the location holds value above threshold
LP=Loc[which(data[Loc[1:length(Loc)]]>T)]
### peaks from the data set
PEAK=data[LP[1:length(LP)]]
### calcuate the differences between average before and after peak
d=ma=mb=0
for(j in 1:length(LP))
{
#mean before the peak
mb[j]=mean(data[abs(LP[j]-h):abs(LP[j]-1)],na.rm=TRUE)
#mean after the peak
ma[j]=mean(data[abs(LP[j]+1):abs(LP[j]+h)],na.rm=TRUE)
d[j]=ma[j]-mb[j]
}
t_paired=mean(d)/sqrt(var(d)/length(d))
```

```
print(c("threshold value", T))
print(c("POT(h)",round(t_paired, digits=4)))
### end POT
```

v. Distribution of First Differences(y)

```
### define x as the original series
x=data
### Compute the lag one first order differences (y)
### x_0 a sequence of {x2,x3,...,xn}
x_0=x[2:length(x)]
### x_1 a sequence of {x1,x2,...,xn-1}
x_1=x[1:(length(x)-1)]
y=(x_0)-(x_1)
```

vi. Percentage of Positive Differences (P_d^+)

```
### start P^{+}_d
w=0
pos=length(which(y>w))
neg=length(which(y<(-w)))
PP=pos/(pos+neg)
print(c("P^{+}_d",round(PP, digits=4)))
### end P^{+}_d
```

vii. Product Moment Skewness of Differences ($\hat{\gamma}_d$)

```
### start gamma.hat_d
Sbar=mean(y)
S=0
for(k in 1:length(y))
{
S[k]=(y[k]-Sbar)^3
}
gamma.hat_d=(sum(S)/length(y))/sd(y)^3
print(c("gamma.hat_d",round(gamma.hat_d, digits=4)))
### end gamma.hat_d
```

viii. Product Moment Skewness of Differences for Symmetrical Time Series about the Median ($\hat{\gamma}_{dab}$)

```
x=data
### median corrected data
w=abs(x-median(x))
### calculate distribution of differences for w (y)
```

```
w_0=w[2:length(w)]
w_1=w[1:(length(w)-1)]
yw=(w_0)-(w_1)
### start gamma.hat_dab
Mbar=mean(yw)
M=0
for(m in 1:length(yw))
{
M[m]=(yw[m]-Mbar)^3
}
gamma.hat_dab=(sum(M)/length(yw))/sd(yw)^3
print(c("gamma.hat_dab",round(gamma.hat_dab, digits=4)))
### end gamma.hat_dab
```

ix. L-skewness of Differences (LSK)

```
### start LSK
### calculate sample probability weighted moments
### ranked y in ascending order
sy=sort(y)
sumb0=sumb1=sumb2=0
for(j in 1:length(y)){sumb0=sumb0+sy[j]}
for(j in 2:length(y)){sumb1=sumb1+sy[j]*((j-1)/(length(y)-1))}
for(j in 3:length(y)){
sumb2=sumb2+sy[j]*((j-1)*(j-2))/((length(y)-1)*(length(y)-2))
}
### second order measures the spread or scale of the distribution
L2=2*(sumb1/length(y))-(sumb0/length(y))
### third Order measures of skewness
L3=6*(sumb2/length(y))-6*(sumb1/length(y))+(sumb0/length(y))
### estimate the dimensionless analogues of skewness (LSK)
LSK=L3/L2
print(c("LSK",round(LSK, digits=4)))
### end LSK
```

(b) *R* Code for the Penalized Least Squares Fit

An example of a TAR(1) model, with a threshold T set at the 80% percentile, fitted by PLS

```
### load a time series in the row vector form

Y=read.table{utils}

### initial parameters
### target is the observed directionality
### alpha is the estimated coefficient of AR(1)
### phi1 is the weight to mitigate directionality discrepancy

N=10^3
phi1=10^5
TH=0.8
target=0.61
alpha=par1=par2=0.91
n=length(Y)

y=Y[2:n]
x=Y[1:(n-1)]
dat=data.frame(x, y)
T=sort(y)[floor(TH*length(y))]

### start PLS

TAR1=function(N,v1,v2){
z=rep(0,n)
for(i in 2:n){
z[i]= ifelse(x[i-1]<T,(v1),(v2))*x[i-1]
}
RES=rep(0,n)
y=c(0,y)
for(i in 1:n){
RES[i]= y[i]-z[i]
}
RES.s=sample(RES, N, replace=TRUE)
zp=rep(0,N)
zpd=rep(0,(N-1))
for(i in 2:N){
zp[i]= ifelse(zp[i-1]<T,(v1),(v2))*zp[i-1] + RES.s[i]
```

```
zpd[i-1]=zp[i]-zp[i-1]
}

sk.TAR=(sum((zpd-mean(zpd))^3)/(N-1))/sd(zpd)^3
skmt=((target-sk.TAR)^2)
return(list(A=sk.TAR,B=skmt))
}

min.RSS=function(data,par){
with(data,(sum(((ifelse(x<T,(par[1]),(par[2]))*x)-y)^2) +
phi1*TAR1(N,(par[1]),(par[2]))$B))
}
OP=optim(par=c((par1),(par2)),min.RSS,data=dat)

### The estimated TAR(1) parameters

L_alpha1=OP$par[1]
U_alpha1=OP$par[2]

### end PLS
```

Bibliography

Allen, N. B. & Christa, H. S. B. (2013), 'How does capital affect bank performance during financial crises?', *Journal of Financial Economics* **109**(1), 146–176.

Andrzejak, R. G., Lehnertz, K., Mormann, F., Rieke, C., David, P. & Elger, C. E. (2001), 'Indications of nonlinear deterministic and finite-dimensional structures in time series of brain electrical activity: Dependence on recording region and brain state', *Physical Review E* **64**(061907), 1–8.

Awartani, B. M. & Corradi, V. (2005), 'Predicting the volatility of the {S&P}-500 stock index via {GARCH} models: the role of asymmetries', *International Journal of Forecasting* **21**(1), 167 – 183.

Bajpai, P. (2017), 'The world's top 10 economies', http://www.investopedia.com/ articles/investing/022415/worlds-top-10-economies.asp (last accessed 24 February 2017).

Beare, B. K. & Seo, J. (2014), 'Time irreversible copula-based markov models', *Econometric Theory* **30**(5), 923–960.

Broussard, J. P. & Vaihekoski, M. (2012), 'Profitability of pairs trading strategy in an illiquid market with multiple share classes', *Journal of International Financial Markets, Institutions and Money* **22**(5), 1188 – 1201.

Brown, S. J. & Warner, J. B. (1985), 'Using daily stock returns', *Journal of Financial Economics* **14**(1), 3 – 31.

Chatfield, C. (2004), *The Analysis of Time Series: An Introduction, Sixth Edition*, Chapman & Hall/CRC.

Chavez-Demoulin, V., Davison, A. C. & McNeil, A. J. (2005), 'A point process approach to value-at-risk estimation', *Quantitative Finance* **5**(2), 227–234.

Chavez-Demoulin, V. & McGill, J. (2012), 'High-frequency financial data modeling using hawkes processes', *Journal of Banking & Finance* **36**(12), 3415–3426.

Cont, R. (2001), 'Empirical properties of asset returns: stylized facts and statistical issues'.

Cowpertwait, P. S. P. & Metcalfe, A. V. (2009), *Introductory Time Series with R*, 1st edn, Springer Publishing Company, Incorporated.

Cox, D. R. (1981), 'Statistical analysis of time series: Some recent developments', *Scandinavian Journal of Statistics* **8**, 93–115.

Dell'Ariccia, G., Detragiache, E. & Rajan, R. (2008), 'The real effect of banking crises', *Journal of Financial Intermediation* **17**(1), 89 – 112. Financial Contracting and Financial System Architecture.

Engle, R. (1982), 'Autoregressive conditional heteroskedasticity with estimates of the variance of united kingdom inflation', *Econometrica* **50**, 987–1007.

Finance, G. (2017), 'Google finance: Stock quotes & charts, financial news, currency conversions, or track your portfolio', https://www.google.com/finance (last accessed 11 April 2017).

Granger, C. W. J. & Andersen, A. (1978), 'On the invertibility of time series models', *Stochastic Processes and their Applications* **8**(1), 87–92.

Hosking, J. R. M. (1990), 'L-moments: analysis and estimation of distributions using linear combinations of order statistics', *Journal of the Royal Statistical Society. Series B (Methodological)* **52**(1), 105–124.

Huang, X., Rui, Y., Shen, J. & Tian, Y. (2017), 'Us class action lawsuits targeting foreign firms: The country spillover effect', *Journal of Corporate Finance* .

Keller, L. J. & Chiglinsky, K. (Sept. 13, 2016), 'Wells fargo eclipsed by {JP M}organ as world's most valuable bank', http://www.bloomberg.com (last accessed 5 October 2016).

Kroese, D. P., Taimre, T. & Botev, Z. I. (2011), *Handbook of Monte Carlo Methods*, John Wiley & Sons.

Lamoureux, C. G. & Lastrapes, W. D. (1990), 'Persistence in variance, structural change, and the garch model', *Journal of Business & Economic Statistics* **8**(2), 225–234.

Lawrance, A. J. (1991), 'Directionality and reversibility in time series', *International Statistical Review/Revue Internationale de Statistique* **59**(1), 67–79.

Lohpetch, D. & Corne, D. (2010), *Outperforming Buy-and-Hold with Evolved Technical Trading Rules: Daily, Weekly and Monthly Trading*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 171–181.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2015a), Modelling directionality in stationary geophysical time series, *in* 'Proceedings of the International work-conference on Time Series (ITISE 2015)', pp. 755–766.

Mansor, M. M., Glonek, M. E., Green, D. A. & Metcalfe, A. V. (2016a), *Time Series Analysis and Forecasting: Contributions to Statistics*, Springer International Publishing Switzerland, chapter Threshold Autoregressive Models for Directional Time Series.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2015b), Modelling and simulation of directional financial time series, *in* 'Proceedings of the 21st International Congress on Modelling and Simulation (MODSIM 2015)', pp. 1022–1028.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2016b), 'Directionality and volatility in electroencephalogram time series', *American Institute of Physics (AIP) Conf. Proc.* **1739**, 020080:1–8.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2017a), 'Detecting directionality in time series', *The American Statistician* submitted: 12 May 2017.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2017b), 'Directionality in time series of bank log-return share prices', *Journal of Business & Economic Statistics* submitted: 14 June 2017.

Mansor, M. M., Green, D. A. & Metcalfe, A. V. (2018), 'Directionality and volatility in high-frequency time series', *High Freq.* **00**, 1–17.

Mansor, M. M., Mohd. Isa, F. L., Green, D. A. & Metcalfe, A. V. (2016c), 'Modelling directionality for paleoclimatic time series', *ANZIAM J.* **57** (EMAC2015), C66–C81.

Martens, P. C., Nandy, D. & Munoz-Jaramillo, A. (2011), Meridional surface flows and the recent extended solar minimum, *in* 'Bulletin of the American Astronomical Society', Vol. 1, p. 1705.

Miraglia, F., Vecchio, F., Bramanti, P. & Rossini, P. M. (2016), 'Eeg characteristics in "eyes-open" versus "eyes-closed" conditions: Small-world network architecture in healthy aging and age-related brain degeneration', *Clinical Neurophysiology* **127**(2), 1261–1268.

Nandy, D., Muñoz-Jaramillo, A. & Martens, P. C. (2011), 'The unusual minimum of sunspot cycle 23 caused by meridional plasma flow variations', *Nature* **471**(7336), 80–82.

NASA (2017), 'The sunspot cycle', https://solarscience.msfc.nasa.gov (last accessed 17 March 2017).

Nash, J. C. (2014), 'On best practice optimization methods in r', *Journal of Statistical Software* **60**(2), 1 – 14.

Nelson, D. B. (1991), 'Conditional heteroskedasticity in asset returns: A new approach', *Econometrica* **59**, 347–370.

NGRIP_members, Andersen, K. K., Azuma, N., Barnola, J.-M., Bigler, M., Biscaye, P., Caillon, N., Chappellaz, J., Clausen, H. B., Dahl-Jensen, D., Fischer, H., Flückiger, J., Fritzsche, D., Fujii, Y., Goto-Azuma, K., Grønvold, K., Gundestrup, N. S., Hansson, M., Huber, C., Hvidberg, C. S., Johnsen, S. J., Jonsell, U., Jouzel, J., Kipfstuhl, S., Landais, A., Leuenberger, M., Lorrain, R., Masson-Delmotte, V., Miller, H., Motoyama, H., Narita, H., Popp, T., Rasmussen, S. O., Raynaud, D., Röthlisberger, R., Ruth, U., Samyn, D., Schwander, J., Shoji, H., Siggard-Andersen, M.-L., Steffensen, J. P., Stocker, T., Sveinbjörnsdottir, A. E., Svensson, A., Takata, M., Tison, J.-L., Thorsteinsson, T., Watanabe, O., Wilhelms, F. & White, J. (2004), 'High-resolution record of northern hemisphere climate extending into the last interglacial period', *Nature* **431**, 147–151.

NOAA (2015), 'National oceanic & atmospheric administration, earth system research laboratory', https://esrl.noaa.gov (last accessed 17 October 2015).

Petit, J. R., Jouzel, J., Raynaud, D., Barkov, N. I., Barnola, J. M., Basile, I., Bender, M., Chappellaz, J., Davis, M., Delaygue, G., Delmotte, M., Kotlyakov, V. M., Legrand, M., Lipenkov, V. Y., Lorius, C., PEpin, L., Ritz, C., Saltzman, E. & Stievenard, M. (1999), 'Climate and atmospheric history of the past 420,000 years from the vostok ice core, antarctica', *Nature* **399**(6735), 429–436.

Petkov, G., Goodfellow, M., Richardson, M. P. & Terry, J. R. (2014), 'A critical role for network structure in seizure onset: A computational modeling approach', *Frontiers in Neurology* **5**, 261.

Priestley, M. B. (1980), 'State-dependent models: a general approach to nonlinear time series analysis', *Journal of Time Series Analysis* **1**(1), 47–71.

Rao, T. S. (1981), 'On the theory of bilinear time series models', *Journal of the Royal Statistical Society. Series B (Methodological)* **43**(2), 244–255.

R_Core_Team (2013), *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria.

relbanks (2017), 'World's largest banks 2017', http://www.relbanks.com/worlds-top-banks/market-cap (last accessed 24 February 2017).

Rockafellar, R. T. & Uryasev, S. (2002), 'Conditional value-at-risk for general loss distributions', *Journal of Banking & Finance* **26**(7), 1443–1471.

Sankarasubramanian, A. & Srinivasan, K. (1999), 'Investigation and comparison of sampling properties of {L}-moments and conventional moments', *Journal of Hydrology* **218**(1), 13–34.

Sharifdoost, M., Mahmoodi, S. & Pasha, E. (2009), 'A statistical test for time reversibility of stationary finite state markov chains', *Applied Mathematical Sciences* **52**, 2563–2574.

SILSO (2015), 'Solar influences data analysis center, sunspot index and long-term solar observations', https://www.sidc.be/silso (last accessed 17 October 2015).

Soubeyrand, S., Morris, C. E. & Bigg, E. K. (2014), 'Analysis of fragmented time directionality in time series to elucidate feedbacks in climate data', *Environmental Modelling & Software* **61**, 78–86.

Tong, H. & Lim, K. S. (1980), 'Threshold autoregression, limit cycles and cyclical data', *Journal of the Royal Statistical Society. Series B (Methodological)* **42**(03), 245–292.

Tsay, R. S. (2005), *Analysis of Financial Time Series, Second Edition*, Wiley Series in Probability and Statistics, John Wiley & Sons Inc.

van Prehn, J., Vincken, K. L., Sprinkhuizen, S. M., Viergever, M. A., van Keulen, J. W., van Herwaarden, J. A., Moll, F. L. & Bartels, L. W. (2009), 'Aortic pulsatile distention in young healthy volunteers is asymmetric: Analysis with {ECG}-gated {MRI}', *European Journal of Vascular and Endovascular Surgery* **37**(2), 168 – 174.

Weiss, G. (1975), 'Time-reversibility of linear stochastic processes', *Journal of Applied Probability* **12**(4), 831–836.

Wikipedia (2016*a*), 'Financial crisis of 2007–2008', https://en.wikipedia.org (last accessed 11 November 2016).

Wikipedia (2016*b*), 'List of economic crises', https://en.wikipedia.org (last accessed 5 October 2016).

Wikipedia (2016*c*), 'Subprime mortgage crisis', https://en.wikipedia.org (last accessed 11 November 2016).

Wikipedia (2017), 'Trading day', https://en.wikipedia.org (last accessed 25 March 2017).

Wild, P., Foster, J. & Hinich, M. J. (2014), 'Testing for non-linear and time irreversible probabilistic structure in high frequency financial time series data', *Journal of the Royal Statistical Society: Series A (Statistics in Society)* **177**(3), 643–659.

Yan, B., Zivot, E. et al. (2003), Analysis of high-frequency financial data with s-plus, Technical report, UWEC-2005-03. URL http://ideas. repec. org/p/udb/wpaper/uwec-2005-03. html.

Zhai, Y., Hsu, A. & Halgamuge, S. K. (2007), *Combining News and Technical Indicators in Daily Stock Price Trends Prediction*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 1087–1096.

Zhong, X. & Enke, D. (2017), 'Forecasting daily stock market return using dimensionality reduction', *Expert Systems with Applications* **67**, 126 – 139.