



Provided by the author(s) and NUI Galway in accordance with publisher policies. Please cite the published version when available.

Title	A complex systems approach to financial market analysis: Nonlinearity, regime shifts and early warning indicators
Author(s)	Metadjer, Naoise
Publication Date	2017-08-16
Item record	http://hdl.handle.net/10379/7131

Downloaded 2021-05-12T09:37:42Z

Some rights reserved. For more information, please see the item record link above.



NATIONAL UNIVERSITY OF IRELAND GALWAY

**A complex systems approach to financial
market analysis:**

Nonlinearity, regime shifts and early warning indicators

Author:

Naoise METADJER

Supervisor:

Dr. Srinivas RAGHAVENDRA

*A thesis submitted in fulfillment of the requirements
for the degree of Doctor of Philosophy*

in the

Discipline of Economics,
J.E. Cairnes School of Business and Economics

August 2017

Contents

Declaration of Authorship	v
Abstract	vii
Acknowledgements	ix
1 Introduction	1
1.1 The Linear Stochastic Modelling Framework	1
1.2 The Stylised Facts of Financial Time Series	3
1.2.1 Aggregate Normality	6
1.2.2 Power Laws and Long Memory of Returns	8
1.2.3 Volatility Clustering	10
1.2.4 Time Varying Cross Correlation of Returns	13
1.3 Motivation and Research Question	14
1.4 Structure of Thesis	15
1.5 Thesis outputs	17
1.5.1 Papers	17
1.5.2 Conferences	18
2 Complex Systems and Financial Markets	19
2.1 Introduction	19
2.2 Complex Systems	20
2.3 Economic complexity	22
2.3.1 Decision making under uncertainty	22
2.3.2 Endogenous dynamics	25
2.3.3 Phase transitions	27
2.3.4 Power laws and scale free networks	29
2.4 Early warning signals	31
2.5 Conclusion	34
3 Nonlinear Dynamical Systems and Bifurcation Theory	35
3.1 Introduction	35
3.2 Nonlinear Dynamical Systems Theory	35
3.2.1 Dynamical Systems	35
3.2.2 Chaos and Strange Attractors	38
3.2.3 Phase Space Reconstruction	39
3.2.4 Correlation Dimension	41
3.3 Bifurcation Theory	42
3.3.1 Fixed Points and Bifurcations	42
3.3.2 Hysteresis	45
3.4 Conclusion	45

4	Nonlinearity in Stock and Bond Markets	47
4.1	Introduction	47
4.2	Methodology	51
4.2.1	BDS Test	51
4.2.2	Surrogate Data Analysis	53
4.2.3	Invariant Statistics	56
4.3	Data	60
4.4	Results	63
4.4.1	BDS Test	63
4.4.2	Surrogate Data Analysis	66
4.5	Conclusion	74
5	Critical Transitions in Financial Markets	77
5.1	Introduction	77
5.2	Theoretical Framework: Bifurcation Theory	81
5.2.1	Critical Slowing Down	85
5.3	Empirical Methodology	87
5.4	Data	90
5.5	Results	91
5.5.1	Sensitivity to Parameter Choice	96
5.5.2	Comparative Analysis	97
5.5.3	False Positive Early Warning Signals	101
5.6	Conclusion	106
6	Network Topology and Systemic Risk	109
6.1	Introduction	109
6.2	Literature Review	112
6.2.1	Systemic Risk	112
6.2.2	The Correlation Network Approach	113
6.3	Data	116
6.4	Methodology	119
6.4.1	Minimum Spanning Trees	119
6.4.2	Network Metrics	122
6.4.3	Time Parameter	124
6.5	Results	126
6.5.1	S&P 500 Static Analysis	126
6.5.2	Dynamic Network Metrics	131
6.6	Conclusions	140
7	Conclusions	143
7.1	Overview of Thesis and Contributions	143
7.2	Limitations	149
7.3	Future Research	150
A	Appendix Chapter 4	153
	Bibliography	157

Declaration of Authorship

I, Naoise METADJER, declare that this thesis titled, "A complex systems approach to financial market analysis: *Nonlinearity, regime shifts and early warning indicators*" and the work presented in it are entirely my own. I confirm that no part of this thesis has previously been submitted for a degree or any other qualification at this or any other University.

Signed:

Date:

Abstract

The history of financial markets over the past century points to the stylised fact that markets build up to a peak and then crash. Many of the standard methods for risk estimation and modelling of financial time series rely on the linear stochastic modelling framework. Under this approach, the interaction of market participants are assumed to be independent and, when taken on aggregate, cancel each other out.

In order to better capture the build-up of risk in the financial system, the methods applied in this thesis allow for endogenous dynamical behaviour, caused by the complex interaction of market participants as they react and adapt to the trends and patterns they create at an aggregate level. Endogenous dynamics can lead to the build-up of instabilities in a complex system, pushing it closer to a critical threshold. When the critical point is reached, the system may abruptly switch between alternate equilibria. In effect, a regime shift occurs in the system.

We investigate whether we can detect certain universal features of complex systems approaching a regime shift to develop early warning indicators of financial crises. Focusing on sovereign bond and stock market time series in the periods leading up to financial crises, the thesis proposes a number of potential indicators of risk building up in the financial system. In particular, we present evidence of nonlinear dependence structures, critical slowing down, and changing network topology in financial markets in the periods preceding financial crises.

Taken on aggregate, the results presented in this thesis indicate that moving beyond the linear stochastic framework and applying methods that can capture complex interactions of market participants, and the resultant emergent patterns they create, adds significant value in the development of early warning signals for financial crises.

Acknowledgements

I would like to express thanks to my supervisor, Dr. Srinivas Raghavendra, whose expert knowledge of the field, combined with strong guidance, patience and insights strongly contributed to the development of my research.

I am would also like to express my gratitude to Prof. John McHale, Prof. Alan Aherne, Mr. Cian Twomey, Dr. Petri Piironen, Dr. Aidan Kane, Dr. Gerard Turley, Mrs Claire Noone, Mrs Imelda Howley, and all the staff of the JE Cairnes School of Business and Economics.

I thank all the doctoral students, past and present, for the many fond memories I have of my time in NUI Galway. The people whose friendship I have enjoyed are too long to list, but I truly appreciated the debates over lunch, social outings to the college bar, games of soccer, camping and general fun that I experienced during my time in NUIG. A particular thanks to Kitty Moloney, Brendan Walsh and Stephen O'Neill for help, guidance and feedback on my research throughout the process.

Thanks to my parents, Kay and Jack, as well as to all my brothers, sisters, friends and colleagues whose encouragement and support was always appreciated.

Finally, a special thanks to my wife Nicola who has supported me every step of the way. I really couldn't have finished this without your love, patience and encouragement. Thanks for always keeping a smile on my face and a song in my heart.

This project was made possible by funding from the Irish Research Council and The Hardiman Research Scholarships, as well as write-up bursaries from the Central Bank of Ireland and the J.E. Carines Business School of Business and Economics.

Chapter 1

Introduction

1.1 The Linear Stochastic Modelling Framework

The recent global financial crisis and Eurozone sovereign debt crisis highlighted deficiencies in common risk and economic modelling techniques which were prevalent amongst market participants. The root causes of the global financial crisis have been well documented¹. Amongst the factors which led to the increased fragility of the financial system were: low interest rates leading to booms in housing and asset markets; deregulation of financial markets and institutions; financial innovation, whose risks were misunderstood by market participants, increasing the size, complexity, opacity and interconnectedness of the global financial system; and increased leverage funded by short term debt. The U.S. subprime crisis set into effect a cycle of mark to market losses and fire-sale of securitised products, driven by similar trading strategies followed by financial institutions across the globe. The opacity of balance sheets meant lenders were uncertain of the ability of counterparties to repay their debt and short term funding sources dried up. A credit crunch ensued with severe spill-over effects for the real economy.

The decade leading up to the financial crisis saw a huge increase in the use of mathematical risk and pricing models. Physicists, engineers, mathematicians and computer scientist were employed in large numbers as quantitative analysts or "quants". There was a belief that as banks' returns were being boosted by larger balance sheets, financed by increased leverage, risks were reduced by advances in

¹For a review see Taylor (2009), Mishkin (2010) and Karmin and Pounder DeMarco (2010).

risk management (Haldane, 2009). However, the mathematical rigour of the models concealed the weaknesses of the assumptions underlying them (Colander et al., 2009).

In order to facilitate modelling, quants relied on the hypothesis that most market data follow a stochastic process, which depends only on past observations of itself and other market data (Danielsson, 2002). This hypothesis assumes that there are a very large number of market participants and that on aggregate their actions are random and cannot influence the market. This assumption, however, falls foul of the Lucas (1976) critique: that changes in policy will systematically alter the structure of econometric models. With advances in financial risk modelling and, in particular, the rise in popularity of Value at Risk (VaR) modelling frameworks, financial institutions began following similar risk management strategies. Under the Basel II capital accord, VaR was placed as the preferred model for market risk modelling. While institutions still had the flexibility to devise the precise nature of their models, a 99% 10 day VaR was set as standard (BCBS, 2004).

The implications of these developments were not fully appreciated. During times of increased financial market volatility VaR limits would be breached and portfolios rebalanced, potentially by large numbers of banks and investment firms simultaneously. Resultant positive feedback loops would exacerbate market movements, with declines in asset prices of magnitudes deemed virtually impossible under the assumptions of the models. Thus, the very act of modelling market risk changes the distributional properties of risk (Danielsson, 2002). This flawed assumption contributed to the "black swan" events outlined in (Taleb, 2007).

The individualistic approach to risk management, which takes the behaviour of market participants as given (Colander et al., 2009), failed to recognise the complex nature of the financial system. Interactions between economic agents at a micro level can lead to emergent phenomena at a macro level that are exceedingly difficult to predict (Solomon and Golo, 2014). In this sense, financial institutions and regulators alike were guilty of the "fallacy of composition". They implemented policies based on the assumption that the financial system is safe if each financial institution is safe. This faulty assumption is implicitly applied in a stochastic risk modelling or pricing framework characterised by the central limit theorem, where the law of large

numbers means that interactions cancel each other out on average and lead to a smooth bell curve (Zigrand, 2014). During the period 1998 to 2007, which Haldane (2009) refers to as the golden decade, this approach to risk management ignored the build-up of systemic risk caused by the complex interaction of agents.

Another key assumption that was inherent in risk modelling during the golden decade was, that the basic statistical properties of financial data are broadly the same during stable periods as they are during crises (Danielsson, 2002). Risk modellers failed to take into account that financial and economic systems are characterised by non-stationarities. The distribution of outcomes during crises display fatter tails, skewness, and increased correlation across asset classes and markets when compared with stable periods. The implications of the assumption of stationarity was most acutely felt in the pricing of credit derivatives and mortgage backed securities. Simulation methods were used, calibrated using historical data on the distribution of losses. The illusion of control provided by mathematical modelling meant that financial institutions and ratings agencies believed that market data from the golden decade was representative of potential future losses. Neglecting to incorporate the possibility of non-stationarities meant that the risks associated with these products were seriously underestimated. In the words of Reinhart and Rogoff (2009), they believed that “this time is different”.

1.2 The Stylised Facts of Financial Time Series

The narrative for the failure to detect the build of systemic risk in the months and years preceding the financial crisis has its roots in linear stochastic risk modelling methodologies. In financial time series analysis, it is commonly assumed that financial asset prices follow a random walk process, where the current price is equal to the price in the previous time period plus a shock variable. A drift variable is also often included to capture the tendency for prices to trend upwards for prolonged periods of time. This can be represented as follows:

$$P_t = \mu + P_{t-1} + \epsilon_t, \quad (1.1)$$

where P_t is the price at time t , μ is the drift term and ϵ_t is a independently and identically distributed (IID) error term, with mean 0 and variance σ^2 . This means that changes in price are uncorrelated over time (Campbell, Lo, and MacKinlay, 1997). As well as being IID, the error term is commonly assumed to be normally distributed. The random walk process described above is non-stationary in both the mean and variance. Using the assumption of the random walk leads to a non-zero probability that the price will be negative, therefore we use the log of P_t :

$$\ln(P_t) = \mu + \ln(P_{t-1}) + \epsilon_t \quad (1.2)$$

or,

$$r_t = \mu + \epsilon_t, \quad (1.3)$$

where r_t is the log-returns, and approximates the continuously compounded rate of return. For sufficiently small Δt , the log-returns are approximately equal to the simple returns for a time series. Equation 1.3 implies that the log-returns are normally distributed, with mean 0 and variance σ^2 (Campbell, Lo, and MacKinlay, 1997).

One important implication of IID error terms is that of stationarity. A stochastic process is stationary if, for every time increment h , the joint distribution $(x_{t_1}, x_{t_2}, \dots, x_{t_N})$ is the same as the joint distribution $(x_{t_1+h}, x_{t_2+h}, \dots, x_{t_N+h})$ (Wooldridge, 2012). However, a somewhat weaker assumption of covariance stationarity generally suffices for statistical inference and risk modelling. A covariance stationary stochastic process has a constant mean, constant variance, and the covariance between x_t and x_{t+h} depends only on t and not on h .

While a linear stochastic framework provides a parsimonious and analytically tractable methodology for modelling financial markets, a number of stylised facts² have emerged which indicate that such a framework may not be appropriate for

²Stylised facts are those based on empirical observation of assets across multiple markets, regions and time periods.

financial market analysis in all circumstances (see Cont (2001) and Jondeau, Poon, and Rockinger (2007) for an overview). A number of these facts are described below.

1. Financial returns exhibit insignificant serial correlation except at high, intra-day time frequencies.
2. The unconditional returns distribution has fatter tails than expected from a normal distribution, with power law or Pareto-like tails, and a tail index between 2 and 5. This indicates that the unconditional distribution does not follow a normal distribution.
3. The unconditional distribution of returns is negatively skewed with an asymmetry between gains and losses. This means that large draw-downs occur more frequently than large gains.
4. Financial returns exhibit intermittency, or irregular bursts of activity. This intermittency is largely related to volatility clustering (serial correlation in volatility).
5. After correcting for volatility clustering, the conditional distribution of returns continues to exhibit fat tails and negative skewness.
6. The autocorrelation function of absolute returns exhibits a power law decay, indicating long range dependence.
7. Leverage effects are present in financial returns, meaning that negative returns increase volatility more than positive returns.
8. Returns exhibit time varying cross-correlation, with increases in correlation detected during times of increased volatility.

In this section we present a number of the most salient stylised facts of financial time series, briefly discussing their implications for financial modelling. The purpose of the section is to highlight the shortcomings of the linear stochastic framework for modern risk modelling. While we focus on stock market returns, due to the availability of high frequency data, the nature of the stylised facts are such that they apply across multiple asset classes.

1.2.1 Aggregate Normality

A normal distribution has the useful property that it is entirely described by its first two moments, mean and variance. The probability density function for the normal distribution is given as follows,

$$P(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-(x-\mu)^2/(2\sigma^2)}, \quad (1.4)$$

where μ is the mean of the distribution and σ is the standard deviation, or square root of the variance. When presented with an empirical time series of financial returns, it is often useful to calculate the first four moments of the empirical distribution.

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^T r_t, \quad (1.5)$$

where $\hat{\mu}$ is the sample mean and is a measure of the central tendency of the empirical distribution.

$$\hat{\sigma}^2 = \frac{1}{T-1} \sum_{t=1}^T (r_t - \hat{\mu})^2, \quad (1.6)$$

where $\hat{\sigma}^2$ is the sample variance and is a measure of dispersion or spread of the empirical distribution.

$$\hat{s} = \frac{1}{T} \sum_{t=1}^T \frac{(r_t - \hat{\mu})^3}{\hat{\sigma}^2}, \quad (1.7)$$

where \hat{s} is the sample skewness and is a measure of asymmetry of the distribution. Skewness is zero for a normal distribution due to asymptotic symmetry.

$$\hat{k} = \frac{1}{T} \sum_{t=1}^T \frac{(r_t - \hat{\mu})^4}{\hat{\sigma}^2}, \quad (1.8)$$

where \hat{k} is the sample kurtosis and is a measure of the fatness of tails of the empirical distribution. The kurtosis of a normal distribution is 3 so we generally express the statistic in terms of excess kurtosis ($\hat{k} - 3$).

One of the stylised facts of financial data is that of aggregate normality. This result is highlighted in Figure 1.1, which displays the standardised distribution of log-returns of Exxon Mobil stock prices (XOM) for time increments of 5 minutes to 1 month. Standardisation of the distribution involves subtracting $\hat{\mu}$ and then dividing by $\hat{\sigma}$. The standard normal distribution is included in Figure 1.1 for comparative purposes. It is clear from the returns distribution that higher frequency returns do not follow a normal distribution, with a leptokurtic distribution³ being evident. Table 1.1 provides the empirical moments of the returns' distributions as well as the Jarque-Bera test statistic, which is calculated as follows:

$$JB = n \left[\frac{\hat{s}^2}{6} + \frac{(\hat{k} - 3)^2}{24} \right], \quad (1.9)$$

where JB is the Jarque-Bera test statistic and has an asymptotic chi-square distribution with two degrees of freedom. A test statistic of 0 indicates normally distributed data. The skewness and kurtosis for all return frequencies display non-normal properties, with negative skewness and fat tails. However, the distribution gets closer to the normal distribution as we lower the frequency of the returns. The implication of this result is that statistics and risk models estimated using the assumption of normally distributed returns will underestimate the fraction of returns in the left tail of the returns distribution.

TABLE 1.1: Moments and Jarque-Bera test Statistic of XOM returns distribution (Jan 1999–Dec 2011)

Series	Mean	Std. Dev.	Skewness	Kurtosis	JB Test Statistic
5 min	0.00	.002	-6.8	774.94	6,327,523,690
20 min	0.00	0.004	-3.92	238.50	153,844,208
1hour	0.00	0.007	-2.13	94.08	8,392,039
1 day	0.00	0.017	-0.65	19.76	53,604
1 week	0.00	0.034	-1.05	5.83	1,090
1 month	0.00	0.058	-0.38	2.85	59
Std. Norm.	0.00	1.000	0.00	0.00	0

³A leptokurtic distribution is one with a higher peak and fatter tails than the normal distribution.

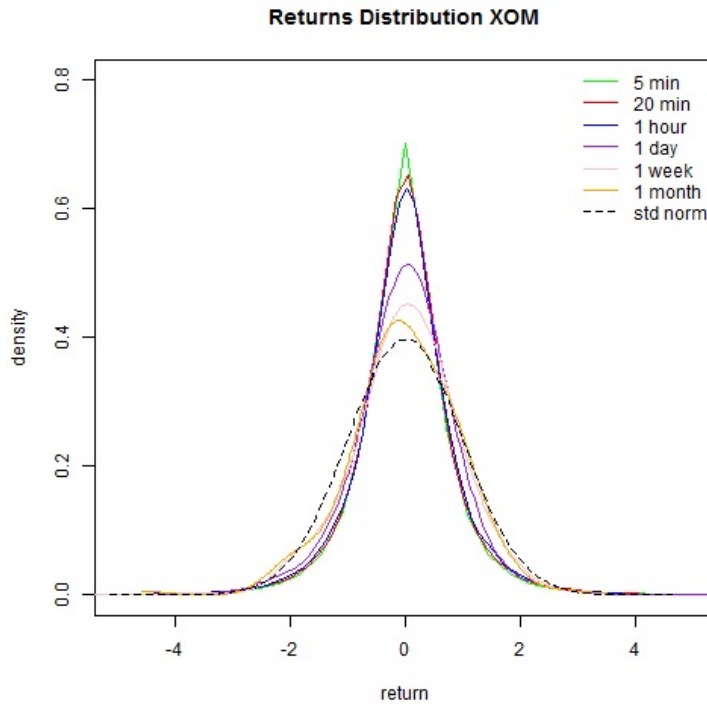


FIGURE 1.1: Distribution of returns at increasing levels of aggregation for XOM (Jan 1999–Dec 2011). The standard normal distribution is represented by dotted line. Distributions are fitted using the R bkde function.

1.2.2 Power Laws and Long Memory of Returns

A number of empirical studies have found consistent evidence across a large number of securities and asset classes, that financial returns display a power law distribution with a tail index between 3 and 5⁴. This means that the tails of financial returns can be approximately modelled according to the cumulative density function,

$$P(x) = 1 - \left(\frac{x}{x_{min}} \right)^{-\alpha+1}, \quad (1.10)$$

where α in the above equation is the tail index and x_{min} denotes the first point in the power law scaling region (Gillespie, 2015). For efficient market theories, which assume that stock price movements are driven by stochastic shocks related to news events, this would mean that such shocks would have to be power law distributed with $\alpha \approx 3$ (Gabaix et al., 2007). The tail index is inversely related to the thickness of the tails of the distribution. The Gaussian distribution returns $\alpha = \infty$ (Cont, 2001).

⁴see Cont (2001) and Gabaix et al. (2007) for a useful discussion.

1.2. The Stylised Facts of Financial Time Series

In Figure 1.2, we demonstrate the power law tail distribution for daily log-returns of XOM. The tail index, $\alpha \approx 3.5$, is within the range specified above and is much lower than we would expect from normally distributed returns. As discussed above, such a finding has significant implications for risk models, such as the parametric Value at Risk model, which often rely on the Gaussian assumption. Efforts have been made to improve risk modelling techniques by explicitly modelling the tails of the distribution using Extreme Value Theory (EVT), amongst other methods. Such methods are an improvement on linear parametric models, which ignore the fatness of the tails of returns distributions. However, they often suffer from limitations, such as the assumption of IID tail observations and the stationarity of the distribution of returns. An in-depth discussion of EVT is beyond the scope of this chapter. For a good overview of the benefits and limitations of EVT see Diebold, Schuermann, and Stroughair (2000).

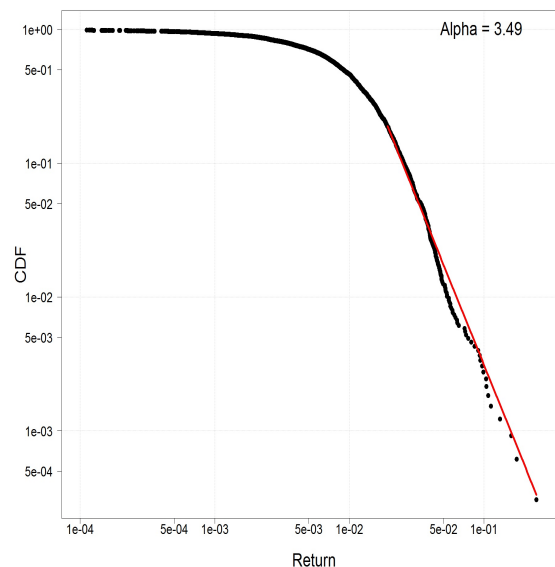


FIGURE 1.2: Power law distribution of absolute daily returns of XOM (Jan 1999–Dec 2011)

Power law distributions are also present in the autocorrelation function (ACF) of absolute returns, with a time decay exponent of between 0.2 and 0.4 (Cont, 2001). This indicates a long memory property in absolute returns, where the ACF decays according to $t^{-\beta}$. We fit a power law distribution to the ACF of XOM absolute daily log-returns (Figure 1.3), finding that the ACF decays according to $t^{-0.7}$, which is outside the range indicated in Cont (2001) but still indicates a long memory

dependence structure in daily returns. This long memory property violates the assumption of IID returns and, if properly exploited, can improve risk models and forecasts. Absolute returns, alongside squared returns, are common proxies for the variance of a financial time series.

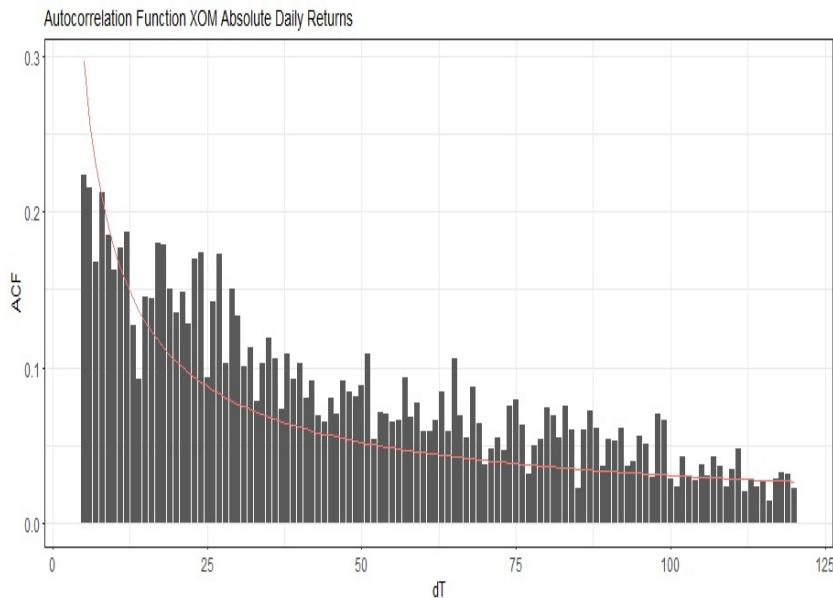


FIGURE 1.3: Power law distribution of the ACF of absolute daily returns for XOM (Jan 1999–Dec 2011). The ACF decays according to a power law with an exponent of -0.7

1.2.3 Volatility Clustering

Conditional heteroskedasticity, or time varying volatility, is a ubiquitous feature of financial returns. Figure 1.4 demonstrates a prevalence for financial returns to go through periods of high and low volatility. This feature of financial returns led to the introduction of stochastic volatility modelling techniques, such as the AutoRegressive Conditional Heteroskedasticity (ARCH) model by Engle (1982). The ARCH model was extended to the Generalised ARCH (GARCH) model by Bollerslev (1986). Essentially, the GARCH model captures serial dependence in volatility by modelling it as an AutoRegressive Moving Average (ARMA) process. Serial dependence in volatility can help explain the long memory dependence structure found in both absolute and squared returns, and violates the assumption of constant unconditional variance. The GARCH(1,1) model, which is commonly found to be sufficient to

1.2. The Stylised Facts of Financial Time Series

capture the effects of volatility clustering in financial time series, is outlined below.

$$r_t = \mu_t + \epsilon_t \quad (1.11)$$

$$\epsilon_t = \sigma_t z_t \quad (1.12)$$

$$z_t \sim N(0, 1) \quad (1.13)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta \sigma_{t-1}^2 \quad (1.14)$$

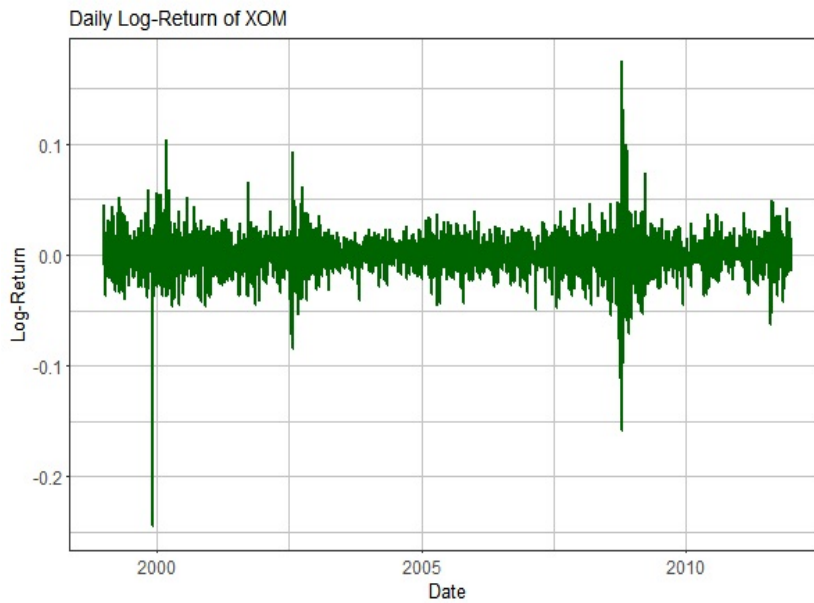


FIGURE 1.4: Daily Log>Returns of XOM (Jan 1999 - Dec 2011)

Equation 1.11 is called the mean equation. We can see it is very similar to Equation 1.3 for log-returns. The key difference is that the drift term, μ_t , is now time dependent. In fact, the GARCH specification allows the mean equation to be modelled as an ARMA process, in order to capture serial dependence that is unrelated to volatility clustering. In the GARCH model, the error terms are no longer IID. The term ϵ_t is modelled as the product of an IID normally distributed stochastic process with mean 0 and standard deviation 1, and a conditional variance term. The

conditional variance term is then modelled in equation 1.14 as a function of past squared error terms (a proxy for volatility) and lagged model variance (σ_{t-1}^2). The GARCH model can have p lags of the squared error term and q lags of the model variance term. However, as discussed above, the GARCH(1,1) specification is often sufficient to capture serial dependence in the variance of financial time series.

In Figure 1.5, we display the ACF of the absolute values of the residuals from a GARCH(1,1) model estimated from XOM daily log-returns. The long-memory dependence structure that is visible in the log-returns (Figure 1.3) has been removed. This indicates that the long memory component of log-returns is closely related to serial dependence in volatility, and may be successfully exploited using stochastic volatility modelling frameworks. However, referring to the descriptive statistics in Table 1.2, we find strong evidence of fat tails in the GARCH(1,1) residuals. Furthermore, our tests for dependence focus on the ACF, which is a linear statistic. Nonlinear test statistics may find evidence of nonlinear dependence structure in the residuals. We investigate this further in Chapter 4.

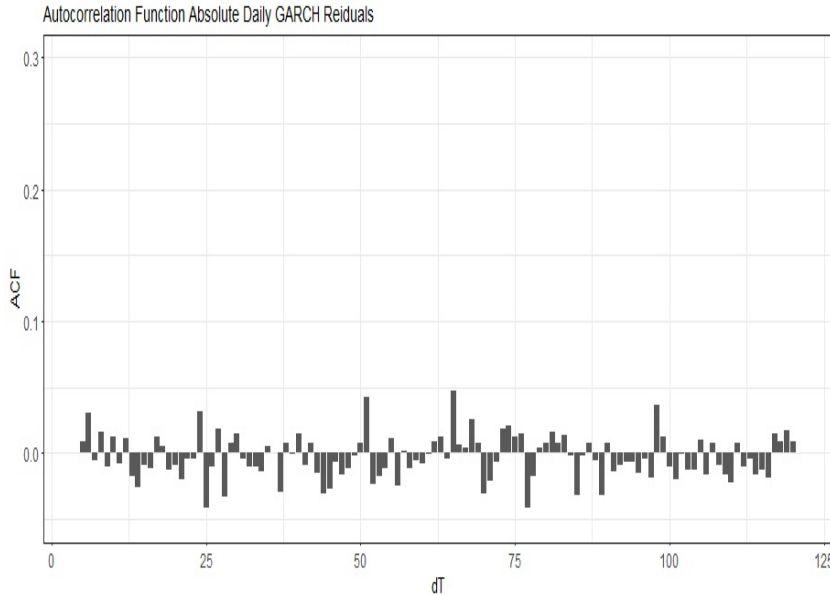


FIGURE 1.5: Power law distribution of the ACF of absolute daily GARCH(1,1) residuals XOM (Jan 1999 - Dec 2011)

TABLE 1.2: Moments and Jarque-Bera test Statistic of XOM Daily GARCH residuals (Jan 1999 - Dec 2011)

Series	Mean	Std. Dev.	Skewness	Kurtosis	J.B. Test Stat
1 day GARCH	-0.01	1.000	-1.11	14.21	28,274

1.2.4 Time Varying Cross Correlation of Returns

The final stylised fact of financial time series that we will discuss is time varying cross correlation of returns. Portfolio risk models frequently calibrate the covariance matrix of asset returns to historical data. If used to manage the expected risk and return of the portfolio, this assumes that correlation coefficients calculated from historical data are representative of future dependence structures between the assets. This assumption falls foul of the fact that the asset covariances are not stationary. Figure 1.6 displays the mean correlation coefficient for the log-returns of five global equity indices. There is a large degree of volatility in the series over time. In particular, we can detect an increase in the correlation coefficient from about 2005 until the onset of the subprime crisis in early 2007. Another sharp increase in correlation occurs during the increased volatility surrounding the Lehmann Brothers crash of the 15th of September 2008. This highlights the dangers of blindly assuming that risk estimates that are calibrated to historical data are representative of future market movements. Prices across all asset classes tend to move together during severe crises, even if they are uncorrelated or weakly correlated during normal times. Thus, risk models which assume stationarity of the covariance matrix will underestimate losses during crisis periods.

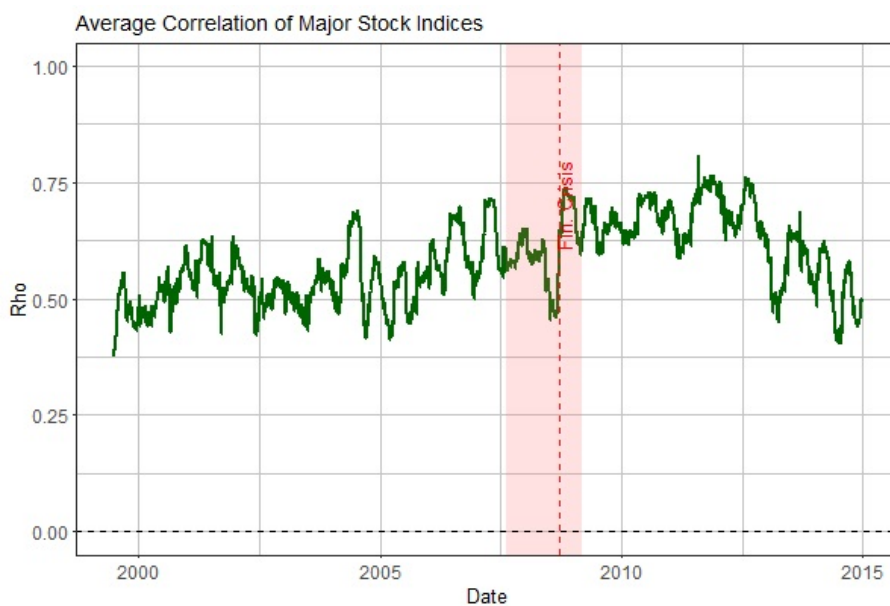


FIGURE 1.6: Mean correlation coefficient for the MSCI Emerging Market, FTSE North America, MSCI Asia Pacific, MSCI BRIC and MSCI European equity indices log-returns

1.3 Motivation and Research Question

The sections above outline a number of empirical features of financial time series that are common across different regions, asset classes and time periods. These features violate the assumptions of common financial risk and forecasting models, and can lead to an under-estimation of risk. When quantitative models are followed blindly, without regard for flaws in the underlying assumptions, extreme market movements can cause portfolio managers to rebalance their portfolios in large numbers. Taken on aggregate, such actions by large numbers of market participants can lead to dangerous positive feedback loops and severe price declines, with spill-over effects to the real-economy. The above example highlights one of the key weaknesses of the linear stochastic modelling approach: that the actions of market participants are assumed to be independent. Thus, the framework does not allow for endogenous dynamics, which can help explain a number of the stylised facts of financial time series, i.e. large asset price movements, nonlinear dependence structures, and non-stationary covariance⁵.

In order to better understand the build-up of risks over time, we must treat financial markets as complex systems with elements that react and adapt to the results of their own actions and the actions of others. This approach is necessary to inform more reliable and realistic risk modelling techniques, and to generate consistent early warning signals for financial crises. Complex systems have been shown to generate endogenous, non-equilibrium, nonlinear dynamics, such as abrupt switching between alternate states. Such dynamics may explain the tendency for financial markets to build-up to a peak and then crash. Furthermore, studies of complex systems have found common patterns as the system approaches a regime shift. Examples of these patterns include nonlinear dependence structures, the emergence of power law distributions, critical slowing down and log-periodic oscillations.

The central research question of this thesis is: Can we exploit certain universal features of systems approaching a regime shift to detect the build-up of instabilities in financial markets and generate consistent, reliable early warning indicators of

⁵Chapter 2 outlines a body of literature which discusses the evidence for and implications of endogenous dynamics in economic systems.

financial crises? This central research question is in turn broken down into three research objectives, which are addressed in Chapter 4 to Chapter 6 of this thesis as follows:

1. To determine if nonlinear dependence structures are present in financial time series in the lead-up to financial crises.
2. To investigate if early warning signals of regime shifts were present prior to the Eurozone sovereign debt crisis and, if so, whether or not these signals can be used as robust indicators of financial crises in general.
3. To examine the usefulness of measures of time-varying correlation network topology as indicators of systemic risk in equity markets.

1.4 Structure of Thesis

The structure of this thesis is as follows:

Chapter 2 discusses complex systems theory as an alternate framework to the economic and financial market analysis that is generally applied in mainstream economic theory. The chapter outlines the main characteristics of complex systems and considers the behavioural finance roots of endogenous economic dynamics. We outline how systems with multiple heterogeneous interacting elements can lead to positive feedback loops and sustained out of equilibrium dynamics. We also discuss how sequential learning and adaption, a prevalence for imitation, and technological innovation can lead to phase transitions and emergent phenomena, such as power law distributions, in economic and financial systems. Finally, we review literature related to early warning signals for financial market crashes, with roots in emergent phenomena observed in financial markets and other complex systems.

Chapter 3 outlines the theory of nonlinear dynamical systems and presents a number of the key concepts that are used throughout the thesis, including phase spaces, fixed points, chaos and bifurcations. We introduce a number of key tools that are used in the study of nonlinear dynamical systems and define the methods of phase space reconstruction, including: estimation of the correlation dimension, the false nearest

neighbours algorithm and methods for choosing an appropriate time lag. Finally, we outline a number of the “normal forms” of bifurcations, which lead to the creation, destruction or stability switching of fixed points.

Chapter 4 presents the results of a number of tests for nonlinearity, which are applied to the returns of stock market indices and sovereign bond yields. We focus on the periods preceding the global financial crisis and Eurozone sovereign bond crisis respectively. Phase transitions, or bifurcations, are a result of strong nonlinear responses in a system, and cannot occur in a system which is completely described by a set of linear equations. Thus, finding evidence of nonlinearity in financial time series is a key first step to developing an understanding of financial crises through the lens of bifurcations. We find mixed evidence for the presence of nonlinear dependence in the data, focusing on the BDS test and surrogate data analysis applied to standardised GARCH residuals. Stronger evidence exists for nonlinear dependence in sovereign bond yield returns than for stock returns, although the results are sensitive to the methodology used.

Chapter 5 presents empirical evidence that sovereign bond markets may have undergone a phase transition between alternate equilibria during the Eurozone debt crisis. We present evidence of a phenomenon called “critical slowing down”, that theory predicts should precede such transitions. We examine the statistical properties of sovereign bond yield data for trends which have been shown to precede catastrophic regime shifts between alternate steady states in many real world dynamical systems. Our results indicate that the critical transitions approach may provide an alternate method to study financial market crashes. However, the phenomenon of critical slowing down, estimated from the statistical properties of the financial data, may provide false positive indications of impending crashes if used in isolation as an early warning signal for stock markets.

Chapter 6 investigates changing correlation network topology as a methodology to detect the build-up of instabilities in stock markets. The results of Chapter 4 and Chapter 5 indicate that stock markets may need to be analysed at a lower level of aggregation than the market index. Chapter 6 uses minimum spanning tree analysis on the correlation networks of individual component stock returns for four global

stock market indices. We assess a number of indicators of interconnectedness and spread of the network, and propose topological shrinking of the network as an indicator of systemic risk. This dynamical behaviour indicates an increased likelihood for a shock to propagate through the financial system as the crisis period is approached. The topological shrinking of the minimum spanning tree in the lead-up to the crisis indicates an increase in the systematic component of equity returns. We also note an apparent relationship between financial sector centrality and sectoral modularity in markets where the financial sector is the largest component of the index.

Chapter 7 presents concluding remarks and a discussion of the findings of the thesis. Future work which may follow from this thesis and some limitations of the thesis are also included.

1.5 Thesis outputs

1.5.1 Papers

The outputs from this thesis are discussed below. Regarding journal submission, three papers will be submitted to academic journals and are based on edited versions of each of the empirical chapters of this thesis. In addition to journal submissions, details on conference presentations are also provided. Chapter 4, titled “Nonlinearity in Stock and Bond Markets” will be submitted to *Studies in Nonlinear Dynamics and Econometrics*. Chapter 5, titled “Critical Transitions in Financial Markets” will be submitted to *PloS one*. Finally Chapter 6, titled “Network Topology and Systemic Risk” will be submitted to *Journal of Network Theory in Finance*. In addition to these three papers, one NUI Galway working paper was submitted based on the work completed during this thesis. This is titled “Critical transitions in Eurozone sovereign bond markets”.

1.5.2 Conferences

- PhD Student Seminars, NUI Galway, “Nonlinear Dependence in Financial Time Series”, Paper presentation, May 2012.
- Inter-Departmental Brown Bag Seminar, NUI Galway, Early Warning Signals for Regime Shifts in Financial Markets”. Paper presentation, February 2013.
- Irish Research Council Annual Symposium, Dublin Convention Centre, “Critical Transitions in Eurozone Sovereign Bond Markets”. Poster presentation, September 2013.
- Irish Society of New Economists 10th Annual Conference, NUI Maynooth, “Critical Transitions in Eurozone Sovereign Bond Markets”. Paper presentation, September 2013.
- Ways of Seeing International Conference, NUI Galway, “Dynamics of Regime Shifts in Financial Markets”. Paper presentation, April 2014
- Irish Economic Association Annual Conference, Dublin Convention Centre, “Interconnectedness as a measure of systemic risk potential in the S&P 500”. Paper presentation, May 2015.
- Irish Economic Association Annual Conference, NUI Galway, “Critical Transitions in Eurozone Sovereign Bond Markets”, Paper presentation, May 2016.
- Financial Risk and Network Theory Annual Conference, Cambridge Centre for Risk Studies, “Interconnectedness as a measure of systemic risk potential in the S&P 500”. Paper presentation, September 2016.

Chapter 2

Complex Systems and Financial Markets

2.1 Introduction

In this chapter, we outline a body of literature that provides an alternate framework for economic and financial market analysis to that generally applied in mainstream economic theory. With roots in the theories of dynamical systems, bifurcations, chaos and networks, complex systems theory is the study of systems with multiple heterogeneous, interacting elements, whose interdependent behaviour leads to emergent phenomena at different levels of aggregation¹. In particular, we consider systems where learning and adaptation are important, and where individual elements react and adapt to the emergent patterns they create at the aggregate level. Such behaviour often leads to positive feedback loops, such as information cascades, which generate sustained out of equilibrium dynamics and have been used to explain bubble dynamics in financial market.

The purpose of this chapter is threefold. Firstly, to define complex systems and provide a broad overview of a number of common characteristics that have been detected in complex systems in many fields of study. Secondly, to review theory and evidence of emergent phenomena in financial markets. Finally, to discuss the implications of this framework for providing early warning signals of financial crises.

¹We define emergent phenomena as ones in which the interaction of system elements at a micro level leads to regular patterns, or order, at the meso and macro level.

Thus, we aim to provide both a theoretical and empirical grounding for the methods applied in this thesis.

The remainder of the chapter is structured as follows: in Section 2.2, we define complex systems and discuss their main characteristics. In Section 2.3, we discuss the behavioural finance roots of endogenously driven economic dynamics². In this section we also outline how sequential learning and adaptation, a prevalence for imitation, and technological innovation, can lead to non-equilibrium dynamics, phase transitions, and power law distributions. In Section 2.4, we discuss the implications of the complex systems literature for the development of potential early warning signals for financial market crashes. Finally, in Section 2.5, we provide some concluding remarks.

2.2 Complex Systems

Complex systems are ones in which the interaction of many agents leads to emergent phenomena. In other words, the behaviour and interaction of agents at a micro level leads to global behaviour which is very different from its origins (Miller and Page, 2009). As Anderson (1972) stated in a seminal paper on the subject of emergent phenomena, “Psychology is not applied biology, nor is biology applied chemistry”. Anderson’s argument is essentially that complex systems comprise hierarchies, or different levels of aggregation. At each level of analysis new concepts, laws and generalisations are needed, due to the inability start with the fundamental laws of behaviour of individual elements (agents) and model the macro behaviour of a system. Even if we know the laws which govern the interaction of agents, it is still not trivial to determine the properties of the system as a whole (Simon, 1962). Keynes (1936) recognised this issue for economic systems, highlighting the mismatch between individual actions and aggregate outcomes.

In this thesis, we are interested in a subset of complex systems in which evolutionary learning and adaptation play a key role. Such systems have been termed Complex

²Endogenous dynamics are ones which come from the operation of the system itself, rather than from external forcing

Adaptive Systems (CAS), and have the defining feature that the individual elements of the system interact with each other, but also adapt or react to the patterns they create at a macro level (Arthur, 1999). Mitchell (2009) provides a general definition of a CAS as one in which large networks of components, with no central control and simple rules of operation, give rise to complex collective behaviour, sophisticated information processing, and adaptation via learning or evolution. This continuum of reaction and adaptation leads to a system that is constantly changing over time, and to endogenously driven dynamical behaviour. As a result, CAS are probabilistic not deterministic, difficult to predict and control, driven by nonlinear network connections between elements of the system, and are characterised by feedback loops between the elements of the system and the patterns they create (Farmer et al., 2012). Fontana (2010) provides a more granular outline of the features of CAS as systems which comprise many diverse parts, have irreversible histories, and exhibit nonlinear dynamics. This is due to system components operating on different temporal and spatial scales. It is this nonlinearity that leads to unique behaviour at different levels of aggregation. Furthermore, the system components react adaptively to change in a manner which increases their probability of survival. CAS can maintain themselves out of equilibrium and tend to self-organise to a critical state, often called Self-Organised Criticality (SOC). In this state, the system is vulnerable to cascades and phase transitions following relatively minor shocks. However, it is also believed that the system is in a state of maximum adaptability and information flow at the point of SOC (Eidelson, 1997).

The general definition and features of complex systems discussed above leads us to a number of further, more specific, definitions of complex systems. The first definition is that of dynamic complexity, defined in Day (1994) as a system whose dynamics do not lead asymptotically to a fixed point, a limit cycle, or to explosive behaviour. Nonlinearity is a necessary but insufficient condition for dynamic complexity (Holt, Rosser Jr, and Colander, 2011). The predominant tools used in the analysis of dynamic complexity are nonlinear dynamical systems theory, chaos theory and bifurcation theory. This is due to the observation that many dynamically complex systems exhibit sensitivity to initial conditions (chaotic dynamics) and critical

dependence on parameters (bifurcations).

The second definition is that of computational complexity. In general terms, computational complexity relates to the number of bits of information or, more generally, the computing resources required to solve a problem, or complete a task. For example, given a variable which represents the state of a system, computational complexity could describe the amount of computing resources which would be required to recreate the dynamics of the state variable. A strict definition of a computationally complex system is: if the system is computable then it is not complex (Fontana, 2010). Mitchell (2009) provides a number of potential measures of computational complexity. Firstly, she suggests complexity as size, e.g. the number of system components. Secondly, complexity as entropy, i.e. the level of randomness in the system. Thirdly, complexity as algorithmic complexity, which is shortest computer programme possible to generate the information contained in the system.

The third definition of complexity is that of connective complexity. Connective complexity is the study of the relationships that exist within the system, with a focus on understanding the forces that maintain order and those that lead to disorder (Fontana, 2010). Network analysis and Agent Based Modelling (ABM) are key tools in understanding connective complexity and, in particular, examining the emergent phenomena which occur at the different levels of aggregation.

2.3 Economic complexity

2.3.1 Decision making under uncertainty

The mainstream approach to financial economics, and macroeconomics in general, is one where the law of large numbers and the central limit theorem implies that the heterogeneous behaviour of a large number of independent agents washes out on average, leading to a smooth bell curve (Miller and Page, 2009). It is the assumption of the independence of agent behaviour that allows economics to circumnavigate the fallacy of composition, highlighted by Keynes (1936), and to build theories and models based on a rational, utility maximising, representative agent. For example, in

models of demand and supply, economists generally assume that the interaction of buyers and sellers leads to an equilibrium price. They do not take into account that the subjective decisions of buyers (sellers) can affect the outcomes and actions of other buyers (sellers), and lead to non-equilibrium dynamics. The Efficient Market Hypothesis (EMH) (Malkiel and Fama, 1970), which forms the basis of much of modern financial theory, assumes that agents form independent, rational expectations of prices based on information which is available at the time. This means that new information is incorporated into prices immediately and, at a given point in time, the best prediction of the future price of an asset is the current price. All dynamical behaviour is driven by a stochastic news process.

If economic agents are constrained in their rationality, and if their actions are not independent, if they interact, learn and adapt this introduces the possibility of behaviour which is very different from the mean. Simon (1957) proposes that, in reality, agents may be boundedly rational, particularly when faced with uncertain outcomes. They may have access to limited information, or search costs may be too high to obtain all information required. Furthermore, boundedly rational individuals have limited cognitive ability. Even if they have sufficient information, they may not be able to work out the best course of action to optimize their outcomes. Constrained cognitive ability is directly related to the computational complexity of the economic system and financial markets. As a result, Simon (2000) proposes that behaviour of people in the real world is as much determined by their inner memories and processes as it is by the outer environment to which they are exposed.

Tversky and Kahneman (1975) argue that rules of thumb, or heuristics, are pervasive when individuals are faced with decision making under uncertainty. The authors also highlight a number of biases which occur as a result, and lead to sub-optimal outcomes. A particular heuristic which is prevalent in individuals is representativeness, where individuals faced with determining whether dataset A is generated by model B will examine the degree to which the characteristics of A match the key characteristics of B . This can lead to over-estimation of the probability that A comes from B if the characteristics of A are similar to B , even if the probability of B is small. Representativeness also leads to the bias of sample size neglect. For

example, an individual may make inferences that A comes from B due to similar characteristics, despite the fact that the sample size of A is not sufficiently large to make such claims. This bias can lead to momentum effects in asset prices if traders see that asset prices have risen and assume that this is representative of the fact that prices will go up in the future. Representativeness can also be linked to availability bias. When faced with an uncertain situation, individuals will search their memory for relevant information in order to categorise the situation. More recent or more extraordinary events are more likely to be evoked and found relevant, even if the probability of their occurrence is low.

Another heuristic, which is explained in Tversky and Kahneman (1975), is anchoring. Individuals who are asked to estimate a quantity will typically form an initial estimate and adjust away from it based on the feedback they receive, or from observing the estimates of others. This leads to the bias of over-anchoring, where individuals will under-adjust away from their initial estimate. Over-anchoring can lead to the underestimation of a damaging event, despite new information coming to light that increases the objective probability of that event. The effect of over-anchoring is magnified when combined with the bias of overconfidence. This means that individuals are overconfident in their judgements despite being poorly calibrated when it comes to assigning probabilities (Barberis and Thaler, 2003).

Shiller (1995) argues that human society has an evolutionary advantage in reacting collectively to information, and in establishing collective memory, common assumptions and conventions. This social force towards conformity increases with the size of the group acting in a similar manner (Easley and Kleinberg, 2010). The bias towards conformity can lead to herding and speculative bubbles in financial markets. As discussed in Shiller (2002), major portfolio allocation decisions tend to be based on common consensus, rather than individual analysis. This is due to the complexity of the task of synthesizing all available information and contrary opinions, while making a subjective decision as to the correct action. Therefore, speculative bubbles are generally supported by some superficially plausible theory as to why asset price increases are supported by fundamentals. A common error amongst investors is that everyone assumes that someone else has carefully analysed the theory Shiller (2002).

The biases discussed above can lead to speculative bubbles due to herding, and feedback between asset prices and investor actions. However, biases and feedback loops can also lead to negative asset price spirals. Prechter Jr (2001) argue that many biases in individuals relate to evolved behaviours to avoid negative results and attain positive results. These biases can lead to impulsive behaviour, which may not be supported by fundamentals. Feedback loops exist between impulsive or subconscious mental activity and stress. For a trader, poor results can lead to stress and related impulsive behaviour, which generates further poor results and stress. Since all market participants share the same environment, the combination of similar stress triggers can lead to global patterns, such as asset fire-sales and liquidity spirals during financial crises.

2.3.2 Endogenous dynamics

Constrained rationality, and the fundamental uncertainty of the outcome of actions, is compounded by the presence of other boundedly rational market participants whose combined actions will affect the outcome of a particular investment strategy. Participants must form subjective beliefs about the subjective beliefs of others, and will continuously adapt their strategies based on the success of their actions (Arthur, 2013). A key factor in this adaptive process is that decisions of participants do not occur simultaneously (Fontana, 2010). Each adaptation and action cascades in the system, as the choices of participants affects the performance of others. This leads to endogenously generated fluctuations as agents explore, learn and adapt, which Arthur (1999) likens to a constant Brownian motion.

Technological innovation provides a second source of endogenous fluctuations. Arthur (2013) argues that the force of technological disruption is self-perpetuating. The author gives the example of the computer, whose development demanded further technologies, such as computer languages and data storage. Moreover, each breakthrough makes further breakthroughs possible. Therefore, Arthur (2013) claims that the economy is in a permanent state of disruption. Solomon and Golo (2014) further outline the disruptive effects of technological or financial innovation, stating that the introduction of new technology causes some large firms to shrink and other

smaller firms to grow. Sectors that are more capable of exploiting new technologies will start to grow, while others will decline. As a result, even when the macro economy appears to be stable, there are continuous, non-equilibrium fluctuations at both the micro level of interacting agents and the meso level of expanding, and declining, firms and sectors.

Feedback loops are a type of endogenous dynamical behaviour that have been used to explain bubbles in financial markets, and can fundamentally alter the dynamics of a system (Miller and Page, 2009). Negative feedback loops, whose forces cause a system to converge to a stable equilibrium, are familiar territory in economic theory. The laws of supply and demand states that the interaction of buyers and sellers will dampen the effect of a price shock, and cause the market to converge quickly to an equilibrium price and quantity. While negative feedback loops lead to stability in a system, positive feedback loops amplify changes and lead to instability. The result of the presence of both positive and negative feedback loops is complex dynamical behaviour (Arthur, 1999). The implication for financial markets is that the system does not blindly revert to equilibrium, nor explode and die (Zigrand, 2014).

An example of feedback loops in financial markets is provided by herding due to imitation. Where traders have independent but imperfect information, they may rationally discount their private information if they observe a large number of other traders following a different strategy. In effect, they make inferences about the information held by others from their actions (Easley and Kleinberg, 2010). The information effect increases where large numbers of traders are following the alternate strategy, and where the system is very complex. In such a case, the probability that a trader will follow the herd increases. This can lead to asset prices increasing far beyond their fundamental value, as more and more traders follow the herd. Positive feedback loops related to herding due to imitation, often referred to as information cascades, are fundamentally fragile because the cascade can easily collapse when new public information comes to light (Easley and Kleinberg, 2010).

Herding may also occur in systems where agents are assumed to be rational but have incomplete information. Rational observational learning involves individuals making rational Bayesian inferences from the behaviour of others (Hirshleifer and

Hong Teoh, 2003). Moreover, herding may also be caused by strong payoff externalities. Examples of payoff externalities are given by the Diamond and Dybvig (1983) bank run model, or in mutual fund runs outlined in Chen, Goldstein, and Jiang (2010). In Chen, Goldstein, and Jiang (2010), if investors in a fund believe that a redemption shock will lead to a fire-sale of assets into an illiquid market, and that the resulting losses may trigger further redemptions, they have the opportunity to limit their losses by redeeming their shares. A “first mover advantage” occurs because a significant portion of the costs related to redemptions are borne by the remaining investors in the fund, and not by those who have cashed out. Taken on aggregate, the first mover advantage scenario can lead to self-fulfilling cycles related to positive feedback loops between investors’ decisions to redeem and asset price decreases. These examples show that, even with fully rational agents, the interdependence of investor actions and the structure of financial contracts can lead to endogenous dynamics and the possibility for multiple equilibria.

2.3.3 Phase transitions

Sustained out of equilibrium dynamics can lead to bouts of volatility, or temporary instability, and cause the system to evolve towards a critical state (Eidelson, 1997). In other words, the build-up of instabilities in the system reaches a critical point and a relatively small perturbation can cause abrupt, discontinuous changes to the state of the system. These shifts, also known as regime shifts, phase transitions, critical transitions or bifurcations, are caused by the critical dependence of the state of the system on parameters. In general, emergent phenomena do not occur in complex systems until some parameter reaches a critical level. The parameter in question generally relates to the intensity of adjustment (adaptation) of individuals, or the degree of connectivity in the system (Arthur, 2013). Strogatz (1994) argue that, if we examine a dynamical system in phase space and the topological structure of the system changes as a parameter is varied, a bifurcation has occurred. Bifurcations cause changes in the number or stability of fixed points, closed orbits, or saddle points in the system.

Catastrophe theory, developed by Thom (1972) and Zeeman (1977) from the more general bifurcation theory, takes the view that the world is mostly smooth and stable, but is subject to discontinuous change when a key parameter reaches a critical point (Rosser, 2000). In a catastrophic bifurcation, as the parameter is varied, the system can switch from a state with a single equilibrium to multiple equilibria. When the parameter reaches a critical point, the state of the system will collapse to a new equilibrium. The empirical features of financial crises led Zeeman (1974) to develop a catastrophe model of stock markets, where the percentage of the market held by chartist investors over fundamentalists acts as the critical parameter. This model was initially heavily criticised, due to the qualitative nature of its arguments and the violation of certain mathematical assumptions (Sussmann and Zahler, 1978). However, in recent years there has been increased interest in the catastrophic bifurcation as an explanation of financial crises. Barunik and Kukacka (2015), for example, fit a catastrophe model to U.S. stock market data over the period 1984–2010. They find that the catastrophe model outperforms linear or logistic models in many time periods. However, the performance of the model between 2003 and 2009 cannot be distinguished from, or underperforms, the logistic model, indicating that a catastrophe model may not explain the discontinuous change experienced in the U.S. stock market during the recent global financial crisis.

The catastrophic bifurcation is only one type of phase transition that can cause structural change in complex systems³. Simulation models of financial markets with heterogeneous agent types, such as fundamental value traders and trend followers, can generate phase transitions from equilibrium to complexity to chaos, or from equilibrium to complexity to multiple equilibria (Arthur, 2013). These models are generally called Agent Based Models (ABMs), and often apply relatively simple behavioural rules. ABMs have been developed to incorporate learning and adaptation, with roots in bounded rationality of Simon (1957), biases and heuristics of Tversky and Kahneman (1975), and insights into investor behaviour from surveys by Frankel and Froot (1990) and Taylor and Allen (1992). The critical parameters in many ABMs relate to the switching intensity between strategies, risk aversion, and

³see Chapter 3 for a more in-depth discussion of bifurcation theory, and an overview of a number of transitions which can occur in dynamical systems.

the deviation of asset prices from fundamentals (Brock and Hommes, 1998; Gaunersdorfer, 2000; Hommes, 2001; Chiarella and He, 2002; Brock, Hommes, and Wagener, 2005). ABMs of financial markets generally incorporate a stochastic news component; however, it is the endogenous dynamics which drive the system towards the critical point, with stochastic shocks merely triggering the phase transition.

2.3.4 Power laws and scale free networks

Power law distributions are ubiquitous features of CAS and are also empirical features of many economic systems. For example, power law distributions have been found in wealth and income distributions (Atkinson and Piketty, 2007), firm size (Gabaix and Landier, 2008), and financial market fluctuations (Stanley et al., 2002). Furthermore, power laws can be found in the network topology of CAS. Scale-free networks are ones which have a large number of nodes with few connections (low degree nodes) and a very small number of nodes with many connections (high degree nodes). Scale-free networks are self-similar⁴, have a power law degree distribution and are an emergent property of CAS. They tend to emerge in CAS due to their resilience to the removal of a random node in the network and thus, are robust to random node failure. However, if a high degree node is deleted it can have serious effects on the stability of the system (Mitchell, 2009). Scale-free networks are generally seen as a consequence of two mechanisms: firstly, networks grow continuously by adding nodes, and secondly, nodes attach preferentially to the most connected existing nodes (Albert and Barabási, 2002). Studies of financial network topology has found evidence of scale-free networks in inter-bank markets (Boss et al., 2004; Iori et al., 2008; Cajueiro and Tabak, 2008), derivatives (Peltonen, Scheicher, and Vuillemeys, 2014; Kenny, Killeen, and Moloney, 2016) and credit networks (Fujiwara et al., 2009; De Masi, 2009).

Solomon and Golo (2014) assert that power law distributions emerge in systems, such as economic systems, if the elements of the system follow an autocatalytic process. In other words, the rate at which the elements of the system grow is proportional to its

⁴Self-similarity in dynamical systems is related to strange attractors or fractals, where if we zoom in on any part of an attractor, it will have features which are topologically or statistically similar to the whole.

current state. For example, if the rate of accumulation of wealth is proportional to the level of wealth held by an individual. Autocatalytic growth is also called preferential attachment in network terminology (Easley and Kleinberg, 2010) and, more generally, proportionate growth (Sornette, 2006). However, Schweitzer et al. (2009) state that the presence of a power law distribution does not necessarily imply proportionate growth, because it is only one of many types of processes which can lead to power law distributions⁵. In general, systems with strong nonlinear dynamics, or stochastic multiplicative amplification, appear to be characterised by power law distributions (Sornette, 2006).

The emergence of power law distributions in financial time series was first noted in Mandelbrot (1963). Mandelbrot proposed that the distribution of many financial returns had much longer tails than the Gaussian distribution, and could be better modelled by the Paretian (power law) distribution. Since Mandelbrot's seminal paper, a large body of literature has emerged that has analysed the emergence of power laws in financial markets. One apparently universal feature of financial market fluctuations is that the distribution of stock market returns follows a Levy distribution in which the tails of the distribution converge to a power law, with an exponent approximately equal to -3 (Lux and Alfarano, 2016). This feature of stock market returns appears to be universal across economies and time periods, with a relatively constant scaling exponent (Stanley et al., 2002).

This inverse cubic power law distribution holds when analysing both company level returns and stock market indices (Liu, Gopikrishnan, and Stanley, 1999; Gopikrishnan et al., 1999). Universal power law scaling can also be found in the temporal dependence structure of stock price volatility. Ding, Granger, and Engle (1993) found evidence of power law decay in the autocorrelation structure of the absolute value of stock market returns, a common proxy for volatility. A scaling exponent of -0.2 to -0.3 appears to be universal across all returns time series. Evidence of power law distributions is also provided for trading volume in international stock markets (Gopikrishnan et al., 2000; Plerou and Stanley, 2007), and bid-ask spreads (Plerou, Gopikrishnan, and Stanley, 2005).

⁵see Gabaix (2009) for discussion of a number of alternate power law generating mechanisms.

Lux and Alfarano (2016) review a number of economic models that are capable of generating power law distributions, with exponents which are similar to those empirically observed in financial data. Examples include percolation, interacting agent and discrete choice models. A key feature of these models is that the emergence of the power laws are a result of endogenously driven dynamics. However, the results are highly dependent on modelling inputs, such as the number of agents and the noise level. Researchers do not yet fully understand why power laws emerge in financial markets; however, their presence and universal nature points to complex, endogenous dynamics caused by interacting market participants (Lux and Alfarano, 2016).

2.4 Early warning signals

Sensitivity to initial fluctuations and randomness is a feature of self-organising behaviour in CAS (Easley and Kleinberg, 2010). In the initial phases, for example after a phase transition from an organised to a disorganised state, it is not clear which system elements will become most important in terms of connections or size. In other words, there are multiple possible outcomes or basins of attraction. This feature is called path dependence, and is related to the non-ergodicity of complex systems, i.e. where a shock at one point in time affects the long-run state of the system (Durlauf, 2005). The implication for economic and financial systems is that it may not be possible to extrapolate new behaviour from past crises, since the new collective organization is generally completely different from the previous one (Schweitzer et al., 2009).

Despite the inability to extrapolate from past behaviour, and the unpredictability of potential outcomes, it may be possible to leverage certain universal features of systems approaching a critical point to detect the build-up of instabilities in financial markets. A series of papers by Didier Sornette, Anders Johansen and collaborators found log periodic power laws as fingerprints of critical behaviour, such as ruptures of spherical pressurised metal tubes (Johansen and Sornette, 2000), earthquakes (Johansen, Saleur, and Sornette, 2000; Huang, Saleur, and Sornette,

2000), and financial crashes (Sornette, Johansen, and Bouchaud, 1996; Johansen, Ledoit, and Sornette, 2000). Johansen, Sornette, and Ledoit (1999) observe power law acceleration in stock market prices, decorated by log periodic oscillation, in the build-up to financial crises. The authors propose a model where the emergence of power law behaviour is due to self-similarity of markets approaching a critical point. This self-similarity is driven by the imitative behaviour of traders, with log periodic oscillations caused by complex power law exponents. It is self-similarity that allows local imitative behaviour to cascade to the global environment (Johansen, Sornette, and Ledoit, 1999). This model is based on the proposal that the main concepts underlying financial market dynamics are imitation, herding, self-organised co-operation, and positive feedback loops (Sornette, 2014). Examples of successful prior-prediction of peaks in asset markets using the log periodic power law framework can be found for the U.K. real estate market (Zhou and Sornette, 2003), the U.S. real estate market (Zhou and Sornette, 2006), oil prices (Sornette, Woodard, and Zhou, 2009) and the Chinese Stock market (Jiang et al., 2010)⁶.

While log periodic power laws have their roots in mechanical ruptures, another potential early warning signal for phase transitions has recently gained popularity in ecological sciences. Critical Slowing Down (CSD) has been found as a universal property of systems approaching a catastrophic bifurcation (Wissel, 1984). CSD has been detected prior to regime shifts in real world complex systems that are believed to have passed through a catastrophic bifurcation, such as climate systems (Dakos et al., 2008; Lenton et al., 2012) and ecological systems (Rietkerk et al., 2004; Carpenter et al., 2011). Recently, evidence has been provided that CSD occurs prior to non-catastrophic bifurcations (Boettiger and Hastings, 2012; Kéfi et al., 2013), meaning that the measure could be used more generally as an indicator of change in a complex system. Recent efforts to detect CSD prior to financial crises have produced mixed results. Guttal et al. (2016) find no evidence of CSD prior to stock market crashes in the U.S., German and U.K. stock markets over the past century. The authors find evidence of rising variability as a universal feature of markets approaching a crash, suggesting a stochastic transition as the likely cause of abrupt

⁶For a summary of additional evidence of advanced prediction of bubbles and crashes see (Sornette et al., 2013).

transitions in financial markets. However, Tan and Cheong (2014) find strong evidence of CSD in the U.S. housing market in the lead-up to the sub-prime crisis, with weaker evidence for the 1997–1998 Asian financial crisis and the 2000–2001 technology bubble.

Another potential method for detecting impending regime shifts in complex systems is to examine changes in network topology and interconnectedness. May (1972) found that large complex systems may suddenly become unstable if the average level of connectedness, or strength of the connections, in the system reaches a critical level. Moreover, certain network structures are more stable than others to disturbances. As discussed above, disassortative networks, such as those which exhibit power law degree distributions, can be more stable to random perturbations (May, Levin, and Sugihara, 2008). Thus, changes in network topology can lead to an unstable system, and can push the system closer to a critical point (Schweitzer et al., 2009). A number of topological phase transitions have been noted in complex networks, such as the emergence of giant connected components, scale-free networks, critical levels of connectedness for propagation of epidemics, and condensation transitions⁷. Therefore, dynamic analysis of measures related to network structure, degree distribution and connectedness may provide indicators of impending regime shifts. Haldane (2009) proposes that financial networks may exhibit a robust but fragile characteristic, where connections may help dissipate shocks throughout the financial system. However, over a certain level of connectivity the shocks may propagate through-out the system and cause contagion. While empirical evidence for phase transitions in financial networks is limited, a number of simulation papers provide frameworks for phase transitions. The Gai and Kapadia (2010) and Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) models of the interbank market provide recent examples. Both studies find connectivity and concentration to be key variables in the propagation of shocks. Wiliński et al. (2013) provide some evidence of transitions from a hierarchal scale free structure to a superstar structure prior to the global financial crisis, using dynamic minimum spanning tree analysis of the Frankfurt Stock Exchange.

⁷see Dorogovtsev, Goltsev, and Mendes (2008) for an overview.

2.5 Conclusion

In this chapter we have outlined a framework for the study of economic and financial markets as complex systems by presenting a body of empirical and theoretical research. We have discussed how the interdependent actions of boundedly rational agents can lead to nonlinear, non-equilibrium dynamics, related to endogenously driven fluctuations and feedback loops. This dynamical behaviour can push a system towards a critical point, where changes in underlying parameters can lead to structural change in system. Emergent phenomena, such as power law distributions and scale free networks, which are universal features of complex systems, have been shown to occur in many economic and financial systems. These emergent phenomena provide evidence and motivation for the study of financial markets as complex dynamical systems, and financial crises as critical phenomena.

This chapter proposes a number of potential avenues for the research of early warning signals of financial crises, with roots in complex dynamical behaviour. Research into log periodic power laws, and related empirical evidence, is well developed for financial market crashes. However, there is scope for empirical research into the presence of CSD prior to crashes in financial markets. There is also a dearth of empirical research into early warning signals of phase transitions in financial networks. Therefore, the research conducted in the remainder of this thesis is motivated by the destructive nature of the recent financial crisis, the empirical and theoretical evidence of complex economic dynamics presented in this chapter, and the scarcity of reliable early warning signals for the build-up of financial instabilities.

Chapter 3

Nonlinear Dynamical Systems and Bifurcation Theory

3.1 Introduction

This chapter provides a brief overview of a number of the key concepts and methodologies discussed in the remainder of this thesis. In Sections 3.2.1 and 3.2.2 we provide an overview of the key features of dynamical systems, including trajectories, fixed points, attractors and chaos. Furthermore, in Section 3.2.3 we provide an introduction to some of the tools of nonlinear dynamical systems theory that are used in this thesis, such as phase space reconstruction, the correlation dimension and the false nearest neighbours algorithm. Finally, in Section 3.3 we discuss bifurcations as the primary method by which fixed points in a nonlinear dynamical system are created, destroyed or switch stability. In section 3.3 we also provide examples of the normal forms of the key classes of bifurcations.

3.2 Nonlinear Dynamical Systems Theory

3.2.1 Dynamical Systems

A dynamical system can be defined in terms of a phase space, M , and a map or law of motion, f . Where the phase space, $M \in \mathbb{R}^d$, and x^d is a point in d -dimensional phase space that defines the state of the system at a particular point in time. The

evolution rule, f^t , is usually represented by a d -dimensional map, or a set of ordinary differential equations, that maps $x^d \mapsto x_t^d$ (Kantz and Schreiber, 2004). If there is a definite rule that maps x^d to a unique point, x_t^d , the system is deterministic (Cvitanovic et al., 2005). We can then use the notation

$$x_{n+1}^d = F(x_n^d), \quad n \in \mathbb{Z} \quad (3.1)$$

for a discrete time, and for continuous time

$$\frac{d}{dt}x^d(t) = f(x^d(t)), \quad t \in \mathbb{R} \quad (3.2)$$

A trajectory is the set of points in M which traces out the evolution of x^d to x_t^d . In a fully deterministic system, where the f^t is independent of time, and is Lipschitz continuous, M will be smooth and for almost any initial value a unique trajectory will exist¹. A key implication of the uniqueness of trajectories is that trajectories will not intersect. If they did then, for a given starting point there would be more than one solution, causing the uniqueness of trajectories to be violated (Strogatz, 1994).

When examining a nonlinear dynamical system, we are generally interested in the existence, number and stability of fixed points. For a continuous system, this can be done by analytically solving the systems of ordinary differential equations that describe its dynamical behaviour, obtaining the eigenvalues and eigenvectors. We can tell a lot about the stability of a system by examining its eigenvalues to classify the fixed points. For example, in a 2-dimensional system where both eigenvalues have positive real parts, the fixed point in the system will be unstable and will repel trajectories away from that point. Where both eigenvalues have negative real parts, the fixed point will be stable and attract trajectories towards that point. If there is one positive and one negative attractor, the trajectories along the eigenvector corresponding to negative eigenvalue will be attracted towards and trajectories along the positive eigenvalue will be repelled away from the fixed point. The situation where one eigenvalue is positive and one negative is called a saddle point. Finally,

¹Cvitanovic et al. (2005) state that there is a subset of points, for example tips of wedges and cusps, for which trajectories are not defined. Hence the use of *almost any initial value* in this statement.

where one or more eigenvalue is imaginary we can get centres and spirals (Strogatz, 1994).

Figure 3.1 provides an example of a simple nonlinear model that generates stable fixed points, unstable fixed points and saddle points. The model is the Lotka-Volterra model of competition between two species, and is reproduced from the specification outlined in (Strogatz, 1994).

$$\begin{aligned} \dot{x} &= 3x - x^2 - 2xy \\ \dot{y} &= 2y - xy - y^2 \end{aligned} \tag{3.3}$$

In the phase portrait in Figure 3.1, the dark lines represent trajectories and the light coloured arrows represent the direction and velocity of trajectories starting from a particular point in phase space. Stable fixed points are represented by a red node, whereas unstable fixed points and saddle nodes are represented by a white node. There are two stable fixed points, or attractors, in the system at $(0, 2)$ and $(3, 0)$. There is an unstable fixed point, or repeller, at $(0, 0)$ and finally a saddle point at $(1, 1)$. This example demonstrates that multiple outcomes can occur in even simple 2-dimensional nonlinear systems. The sphere of influence of the stable fixed points in the system is called the basin of attraction. This means that a point starting in the basin of attraction of a stable fixed point will converge asymptotically to that fixed point.

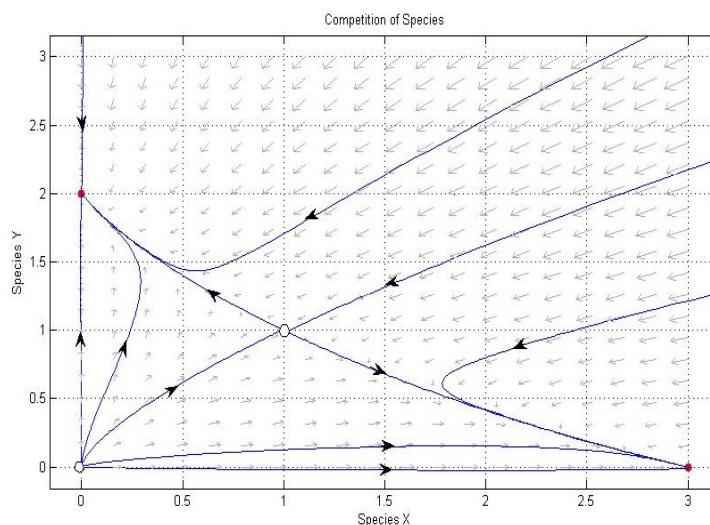


FIGURE 3.1: Model of competing species.

3.2.2 Chaos and Strange Attractors

The evolution of deterministic systems is given by a trajectory in the state space forming an attractor: a geometrical shape which exists in a subset of the state space (Papana, 2009). When the dimension of a nonlinear dynamical system is ≥ 3 , very complicated geometrical objects can occur in phase space. Such attractors are called strange attractor, have a non-integer dimension, and are found in chaotic systems (Kantz and Schreiber, 2004). Strange attractors occur due to repeated stretching and folding, are bounded to a subset of phase space, and generally have a fractal geometry. Fractals are complex geometric shapes which exhibit some degree of self-similarity, i.e. if we zoom in on any part of a fractal it will have features which are topologically or statistically similar to the whole.

Chaotic systems are particularly interesting as they are both deterministic and unpredictable, due to sensitive dependence on initial conditions. A system is said to be chaotic when tiny deviations in initial conditions can lead to exponential divergence in trajectories. One of the simplest chaotic attractors is given by the Rössler (1976) system of equations.

$$\begin{aligned}\dot{x} &= -y - z \\ \dot{y} &= x + ay \\ \dot{z} &= b + z(x - c)\end{aligned}\tag{3.4}$$

The Rossler system is plotted in Figure 3.2 in 2-dimensions (left pane) and 3-dimensions (right pane) with $a = b = 0.2$ and $c = 8$, which gives a chaotic, strange attractor. As stated above, strange attractors are a result of stretching and folding. In 2-dimensions we can see the stretching as the trajectories spiral outwards. We have to go to the third dimension to see the folding back towards the initial starting point. It is this folding in the third dimension that ensures the uniqueness or non-intersection of trajectories (Strogatz, 1994). We can see that the Rossler attractor is bounded, due to folding, such that it is confined to a subset of 3-dimensional phase space.

However, as is characteristic for chaotic systems, we get exponential divergence of initially close trajectories.

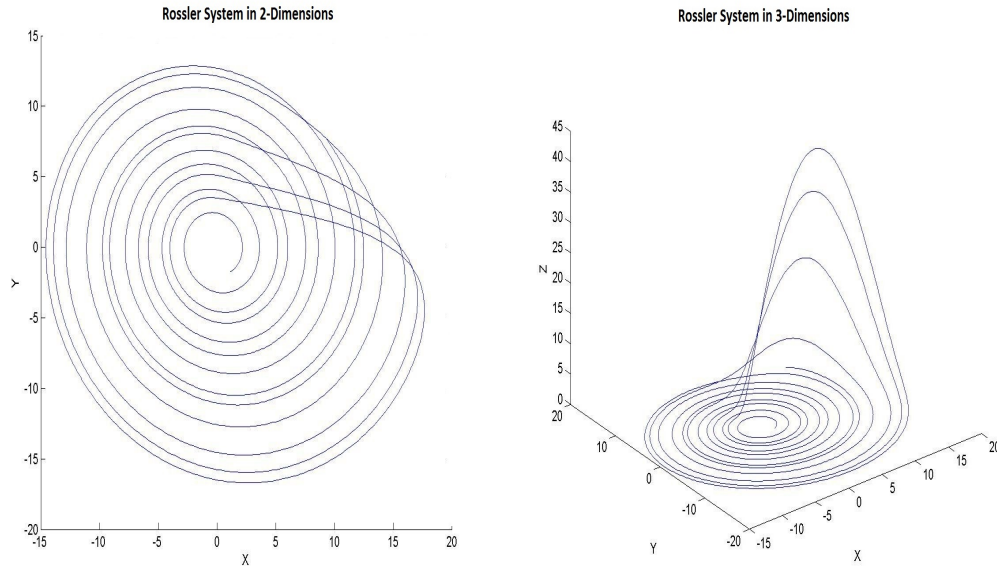


FIGURE 3.2: Rossler strange attractor in 2 and 3-dimensions.

In a chaotic system, the properly averaged exponent of deviation, called the Lyapunov exponent, quantifies the strength of chaos. While a number of different Lyapunov exponents can be defined for a dynamical system, the most important is the maximal Lyapunov exponent, which is defined by Kantz and Schreiber (2004) as follows:

Let two points in phase space, x_{n1}^d and x_{n2}^d whose initial distance $\|x_{n1}^d - x_{n2}^d\| = \delta_0 \ll 1$ and whose distance Δn steps into the future is given by $\delta_{\Delta n} = \|x_{n1+\Delta n}^d - x_{n2+\Delta n}^d\|$. The maximal Lyapunov exponent, λ , is defined as

$$\delta_{\Delta n} \simeq \delta_0 e^{\lambda \Delta n} \quad (3.5)$$

A positive value for λ indicates exponential divergence of initially close trajectories and a chaotic dynamical system.

3.2.3 Phase Space Reconstruction

A key challenge in nonlinear time series analysis is to understand the deterministic system that is generating the time series when we have no knowledge of the underlying system of equations. To do so, the observations must be converted into

states space values. This is completed by means of phase space reconstruction by the method of delays (Kantz and Schreiber, 2004). A reconstruction of the attractor underlying the time series, x_t , is given by the time delay embedding

$$x_t \mapsto (x_t, x_{t-\tau}, \dots, x_{t-de+\tau}), \quad (3.6)$$

$$= s_t \quad (3.7)$$

where de is the embedding dimension and τ is the time lag. The Takens (1981) embedding theorem states that, as long as the embedding dimension is greater than $2D + 1$, the embedding will preserve the dynamical properties of the underlying attractor, where D is the dimension of the true state space. However, for the estimation of topological invariants, such as the correlation dimension, Sauer, Yorke, and Casdagli (1991) generalise Takens' result and find that an embedding dimension greater than twice the box counting (fractal) dimension is sufficient. Sauer, Yorke, and Casdagli (1991) define the box-counting dimension of an attractor, A , as

$$D_F = \lim_{r \rightarrow \infty} \frac{\log N(r)}{-\log r}, \quad (3.8)$$

where $N(r)$ is the number of cubes of size r which intersect the attractor, A .

In practice a standard method for choosing an appropriate embedding dimension, which will provide a faithful reconstruction of the attractor, is provided by the *false nearest neighbours* (FNN) algorithm (Kennel, Brown, and Abarbanel, 1992). The idea behind the algorithm is, for a time series x_t with embedding $s_t(de)$, to gradually increase the embedding dimension de and count the number of nearest neighbours which are within threshold ϵ of each other at each dimension. Given an embedding dimension that is too small to provide a faithful reconstruction, the topological structure of the underlying attractor will not be preserved (Papana, 2009). In this case the reconstructed vector will have (false) neighbours which would not be within distance ϵ of each other in a higher dimension. A sufficient embedding dimension is

provided by the value of de for which the fraction of neighbours that are false is sufficiently small.

The next challenge is to find a good estimate of the time lag τ . Kantz and Schreiber (2004) state that choosing too small a time lag can lead to an embedding vector where the successive elements are strongly correlated, whereas too large a time lag can lead to successive elements which are almost independent. Two alternate methods for choosing the time lag are outlined in Kantz and Schreiber (2004). The first is the first zero of the autocorrelation function, which the authors find provides a good trade-off between too large and too small a time lag. The second method is the first minimum of the time delayed mutual information, (Fraser and Swinney, 1986).

3.2.4 Correlation Dimension

For embedded time series s_t with embedding dimension de , the correlation dimension is a measure of how the points are distributed within the embedding space. In order to estimate the correlation dimension, the correlation sum is defined for s_t as the fraction of all possible pairs of points which are closer than a given distance ϵ . The correlation sum is given by the equation

$$C(\epsilon) = \frac{2}{N(N-1)} \sum_{i=1}^N \sum_{j=1}^N \Theta(\|s_i - s_j\| < \epsilon), \quad (3.9)$$

where $\Theta(X) = 1$ if and only if X is true and $\Theta(X) = 0$ in all other cases. $\|\cdot\|$ denotes the supremum. The correlation sum is the discrete time equivalent of the correlation integral and estimates the probability that, for dimension de , two randomly chosen points are closer than ϵ . In the limit where the number of data points goes to infinity and ϵ is small, the correlation sum increases according to a power law as ϵ is increased i.e.

$$C(\epsilon) \propto \epsilon^{C^D}, \quad (3.10)$$

where C_D is the correlation dimension and can be defined by

$$C_D = \lim_{\epsilon \rightarrow \infty} \lim_{N \rightarrow \infty} \frac{\log C(\epsilon)}{\log \epsilon} \quad (3.11)$$

According to Small (2005), noise has a correlation dimension equal to the embedding dimension. Therefore, if we calculate the correlation dimension of an IID stochastic process at any embedding dimension (de), C_D will equal de . A correlation dimension lower than the embedding dimension is an indication that the underlying system is not dominated by noise and there is some underlying deterministic structure. The standard method for estimation of the correlation dimension is set out in Grassberger and Procaccia (1983). One plots $\log C(\epsilon)$ against $\log \epsilon$ and calculates the slope of the graph as $\epsilon \rightarrow \infty$. Due to the finite nature of many time series of interest, a scaling region must be picked where ϵ is moderately small.

Creedy and Martin (1994) state that there are numerous problems with the application of the correlation dimension to time series which are finite and noisy. If a deterministic process has a high dimension, the relationship between C_D and de for moderate sized de will be similar for deterministic and stochastic processes. Furthermore, for near unit root stochastic processes the C_D can be mistaken for that of a deterministic process. Finally, if the data has a low signal to noise ratio then C_D may be inaccurate. Therefore, for a finite dataset it is not possible to confirm with certainty whether or not an underlying stochastic process has an infinite correlation dimension.

3.3 Bifurcation Theory

3.3.1 Fixed Points and Bifurcations

In this section we provide an introduction to bifurcation theory, with examples of simple bifurcations which can occur even in one-dimensional nonlinear systems. The objective of the section is to introduce bifurcations as the primary mechanism by which fixed points are created, destroyed or switch stability in nonlinear dynamical

3.3. Bifurcation Theory

TABLE 3.1: Examples of simple bifurcations in 1-dimension

Bifurcation	Normal form	Equilibria as r passes threshold
Saddle node	$\dot{x} = r + x^2$	No fixed points \rightarrow Two Stable Creation of fixed points.
Transcritical	$\dot{x} = rx - x^2$	One Stable and One Unstable \rightarrow One Unstable and One Stable Stability switching of fixed points.
Supercritical Pitchfork	$\dot{x} = rx - x^3$	One Stable \rightarrow Two Stable and One Unstable Both creation and stability switching of fixed points.

systems. Bifurcations are the qualitative changes to a dynamical system induced by changes in some key parameters of the system. In other words, the dynamical system undergoes a qualitative change when it passes through the bifurcation point, caused by changes in some key parameters (Strogatz, 1994). Bifurcations that involve changes in the existence, number, or stability of fixed points are known as local bifurcation. Global bifurcations involve changes in the entire orbit or trajectory of a system. Furthermore, bifurcations can be classified as continuous when eigenvalues in the system become stable or unstable, or discontinuous when eigenvalues appear or disappear (Sprott, 2003).

In Table 3.1 we describe three common local bifurcations that occur in one-dimensional continuous nonlinear systems. The normal form for the saddle node bifurcation is $\dot{x} = r + x^2$, where x represents the state of the system and r is the bifurcation parameter. A fixed point in the system occurs where $\dot{x} = 0$. The saddle node bifurcation is the basic mechanism by which fixed points are created and destroyed in a dynamical system. For $r < 0$ there are no fixed points in the system; however, as we increase r to $r = 0$ a half stable fixed point appears. If we further increase r to $r > 0$, the half stable fixed point splits into one stable and one unstable fixed point. The bottom panel of Figure 3.3, illustrates the existence, position and stability of fixed points in the saddle node bifurcation equation as r is varied. The saddle node bifurcation is a discontinuous bifurcation as, once the bifurcation point is crossed from either direction, the fixed points (and associated eigenvalues) appear or disappear. $\dot{x} = r + x^2$ is called the normal form of the saddle node bifurcation. It is so called as even much more complicated and higher dimensional systems will behave similarly when we examine the system close to the fixed point as a saddle

node bifurcation occurs.

The next bifurcation of interest is the transcritical bifurcation. It is one in which the fixed points in the system are not created nor destroyed, rather the stability of fixed points switch once the bifurcation point is reached. The normal form of the transcritical bifurcation is $\dot{x} = rx - x^2$. It is the primary mechanism for changes in the stability of fixed points in a system (Strogatz, 1994). We can see from the bifurcation diagram in the lower panel of Figure 3.3 that the stability of fixed points switches as r is varied. As r goes from $r < 0$ to $r > 0$, the stable fixed point at $x = 0$ becomes unstable and the unstable fixed point which approaches 0 along the diagonal becomes stable.

Pitchfork bifurcations occur in systems which have symmetry. Strogatz (1994) gives the example of load bearing beam. When the load is small, the beam is stable in the vertical position. After the load breaches the beam's load bearing capacity, the vertical beam becomes unstable and may buckle to the left or the right. Therefore, the stable (vertical) fixed point has switched from stable to unstable and two new fixed points (left and right buckled positions) are created. The lower panel of Figure 3.3 shows the supercritical pitchfork bifurcation. As r moves from $r < 0$ to $r > 0$, the stable fixed point at $x = 0$ becomes unstable and two new stable fixed points are created.

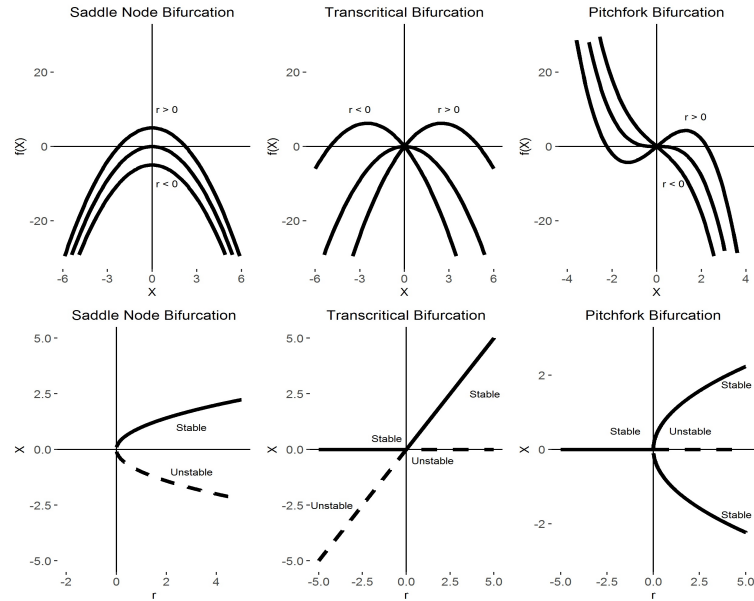


FIGURE 3.3: Saddle node, transcritical and supercritical pitchfork bifurcations. The top panel represents X versus \dot{X} with $r = 5$, $r = -5$ and $r = 0$. The bottom panel presents the bifurcation diagram of fixed points.

3.3.2 Hysteresis

In the preceding section we described the saddle node bifurcation as discontinuous, due to the appearance or disappearance of eigenvalues in the system. The saddle node bifurcation can also be described as explosive as it lacks hysteresis, (Sprott, 2003). Hysteresis means that the bifurcation will occur for different values of the parameter once the direction of the parameter is reversed. In Figure 3.3 we can see that the saddle node bifurcation occurs at $r = 0$ regardless of the direction of the parameter. An example of a transition in which hysteresis occurs is given by adding an imperfection parameter to the supercritical pitchfork bifurcation. We can see from Equation 3.12 below that when the imperfection parameter (h) is equal to zero we are left with a supercritical pitchfork bifurcation.

$$\dot{x} = h + rx - x^3 \tag{3.12}$$

Where the value of r is positive, varying h causes the system to switch between alternate equilibria. When $h = 0$ there are two stable fixed points in the system at $x = \pm\sqrt{r}$ and an unstable fixed point at the $x = 0$. If we increase (decrease) h , the stable fixed point to the right (left) of the origin begins to converge with the unstable

fixed point. If h passes a critical value, the system will collapse to the left (right) hand fixed point and the system will have only a single stable equilibrium. Reducing h back below the critical threshold is not sufficient to reverse the transition. This lack of reversibility at the same parameter value for which the bifurcation occurred is known as hysteresis. In Chapter 5 we discuss the theory of discontinuous bifurcations with hysteresis in more detail, providing examples and outlining potential applications in explaining discontinuous regime shifts in real world dynamical systems.

3.4 Conclusion

This chapter has outlined a number the key concepts and tools of nonlinear dynamical systems theory and bifurcation theory, which will be referred to throughout this thesis. Of particular importance in the following chapters will be the concepts of nonlinear systems, chaos, phase space reconstruction, discontinuous bifurcations and hysteresis.

Chapter 4

Nonlinearity in Stock and Bond Markets

4.1 Introduction

The objective of this chapter is to test financial time series for evidence of nonlinear dependence in the lead-up to financial crises. Specifically, we focus on the daily log-returns of stock indices in the period preceding the 2007-2009 global financial crisis, and sovereign bond yields in the lead-up to the Eurozone sovereign bond crisis, which began in earnest in 2010. In Chapter 2 we discussed how complex systems are characterised by nonlinear network connections and how nonlinearity is a necessary condition for dynamic complexity. A key motivation for the approach taken in this chapter is the statement of Dakos et al. (2012a), that critical transitions (bifurcation) are triggered by strong nonlinear responses in a system. Thus, evidence of nonlinearity should be present in a system approaching a critical transition. Finding evidence of nonlinear dependence in the stock and bond market data would open the possibility that such crises can be explained as a bifurcation or regime shift in the system. When a regime shift occurs in a dynamical system, the system switches abruptly between alternate steady states (equilibria) and displays qualitatively different behaviour to the original state ¹.

As discussed in Chapter 1, the dominant methods in the analysis of financial time series are based on the linear paradigm, which states that small causes lead to small

¹For more details on regime shifts in financial markets, refer to Chapter 5 of this thesis.

effects. Linear equations can only lead to exponentially decaying or growing periodic oscillations, with all irregular behaviour caused by stochastic shocks (Kantz and Schreiber, 2004). Data which originates from a nonlinear data generating process can generate much richer types of behaviour, such as sudden bursts of volatility and large asset price movements (Hsieh, 1991). Moreover, nonlinear chaotic time series can behave in a manner that is similar to stochastic processes, but are governed by nonlinear deterministic behaviour. Such systems can be predictable in the short run but not the long run, due to sensitivity to initial conditions (Barnett, Medio, and Serletis, 2015)². While financial time series undeniably have stochastic components, the behaviour of asset prices during crises are qualitatively similar to bifurcations in real world dynamical systems, such as ecosystems and climate systems³.

The question of whether financial data is the result of a nonlinear or low dimensional chaotic dynamical system has long been explored by academics and market practitioners. Campbell, Lo, and MacKinlay (1997) state that interaction between agents in the market, the price discovery process and the dynamics which generate macroeconomic fluctuations are all nonlinear. While financial time series data may be the product of a nonlinear data generating process and may exhibit nonlinear dependence, a number of challenges exist in detecting such dynamical behaviour. Firstly, the strength of a nonlinear signal in the data may be weakened by the ability of traders to perceive and trade against these patterns (LeBaron, 1994). Secondly, multiple interactions between market participants and rapid economic fluctuations result in high dimensionality in financial time series (Kyrtsov and Serletis, 2006). Such high dimensionality can lead to nonlinear chaotic time series which are, for all practical purposes, indistinguishable from random data (Hsieh, 1991). In fact, pseudo-random number generators often use chaotic systems to generate random looking numbers. Thirdly, financial time series are finite and noisy, which poses a challenge in the estimation of statistics that are used to detect the signature of chaotic

²In this chapter, when we speak of nonlinear systems and nonlinear time series we are referring to data generated by a system of nonlinear equations. These can be either completely deterministic or stochastic nonlinear equations. In such systems we may be able to detect nonlinear dependence, i.e. where the current observation of a time series depends on past observations and the functional form of that dependence is nonlinear. Refer to Chapter 3 for more details on nonlinear dynamical systems.

³Chapter 5 discusses a number of studies which provide examples of bifurcations in real world dynamical systems.

dynamics, such as the correlation dimension and Lyapunov exponent (Kyrtsov and Serletis, 2006)⁴. However, if financial time series are not too complex, and the signal to noise ratio is sufficient, short term predictability may be exploitable and nonlinear models may be used to improve forecastability.

Despite the challenges, many empirical studies have presented evidence for and against nonlinear structure in financial time series. Brock and Sayers (1988) use the Grassberger and Procaccia (1983) (henceforth GP) correlation dimension⁵ to detect chaotic dynamics in the residuals of an AR model applied to macroeconomic data, finding that, once linear serial correlation is removed, no chaotic dynamics are detected. Scheinkman and LeBaron (1989) use the GP algorithm on the residuals of an ARCH model applied to U.S. stock returns, finding evidence of low dimensional chaos. The authors find a correlation dimension lower than the theoretical $2 \log_{10} N$ Ruelle (1989) limit, indicating the presence of low dimensional chaotic dynamics. Ruelle (1989) however, states that the result is likely due to remaining smoothness in the data following filtering using AR and ARCH models. Haniias, Curtis, and Thallasinos (2007) apply the GP methodology directly to stock prices of the Athens stock exchange using a large time lag to reduce dependence in the data, finding evidence of low dimensional chaos.

There are a number of problems in applying the GP algorithm to financial data, in particular the difficulty of picking an appropriate scaling region and the lack of an asymptotic distribution for statistical inference⁶. As a result, Brock, Dechert, and Scheinkman (1987) (henceforth BDS) test statistic was developed, which introduced an asymptotic distribution theory for statistics based on the correlation integral. The null hypothesis of the test is that the data being tested is IID. The BDS test is generally applied to financial data after an appropriate model is fitted to remove any serial correlation and volatility clustering in the data. If the model is correctly specified, the resulting residuals should be IID. Thus, when applied to the standardised residuals of a well specified model, rejection of the null hypothesis

⁴As discussed in Chapter 3, the correlation dimension is a measure of the degree to which an attractor fills up phase space for a given dimension, whereas the Lyapunov exponent is a measure of sensitivity to initial conditions. Both are used as tests for deterministic chaos.

⁵The correlation dimension is outlined in Equations 3.9–3.11 of Chapter 3

⁶See Ruelle (1989) and Kyrtsov and Serletis (2006) for an overview.

provide indirect evidence of nonlinear dependence in the data. Evidence of nonlinear dependence has been found using the BDS test in stock markets (Scheinkman and LeBaron, 1989; Hsieh, 1991), macroeconomic data (Brock and Sayers, 1988), derivatives data (Moloney and Raghavendra, 2011) and exchange rates (Brooks, 1996; Kyrtsov and Serletis, 2006). Hsieh (1991) finds that, while the BDS test rejects the null of IID for GARCH and EGARCH standardised residuals of stock returns, a more generalised variance model removes the remaining structure in the data and the BDS test accepts the null of IID residuals⁷.

The BDS test provides a powerful statistic and lends itself effectively to the examination of the adequacy of models that assume IID residuals. It is not a test for non-linearity or chaos, and offers no information on the type of dependence detected in the residuals. Surrogate data analysis (Theiler et al., 1992) can provide information in this regard by computing an invariant statistic⁸ on the data being tested. This statistic is then compared to the same invariant statistic computed on Monte Carlo generated time series that are consistent with some null hypothesis regarding the data but otherwise random. For example, if the tester believes that the data is generated by a nonlinear dynamical system he/she will use a linear, stochastic model to generate the surrogate data (Soofi et al., 2014). Thus, the null hypothesis is that the data is generated by a linear, stochastic system. Rejection of the null occurs if the invariant statistic computed on the original data falls outside the distribution of statistics calculated on the surrogate data.

A number of algorithms have been developed for generating data consistent with different null hypotheses (Theiler et al., 1992; Schreiber and Schmitz, 1996; Kugiumtzis, 2002a; Schreiber, 1998; Small, Yu, and Harrison, 2001). A body of literature exists that applies surrogate data tests for nonlinearity to financial time series. Evidence of nonlinear structure can be found for exchange rates (Soofi and Galka, 2003; Zhang, Soofi, and Wang, 2011) and for stock markets (Small and Tse, 2003); however, evidence of deterministic chaos is limited. The lack of evidence for

⁷Hsieh (1991) models the current standard deviation as a function of lagged standard deviation but, unlike the ARCH and GARCH specifications, includes an innovation variable in the variance equation.

⁸Small (2005) defines an invariant statistic as a quantity which describes aspects of the dynamical behaviour of a system and does not depend on the coordinate system, i.e. the invariant calculated on the original system will be the same as that calculated on the delay embedding reconstruction.

deterministic chaos may be due to the unavailability of methods that can separate chaotic dynamics from exogenous noise (Das and Das, 2007).

In this chapter, we apply a number of tests for nonlinearity to the returns of stock market indices and sovereign bond yields in the periods preceding the global financial crisis and Eurozone sovereign bond crisis respectively. We find mixed evidence for the presence of nonlinear dependence in the data, focusing on the BDS test and surrogate data analysis applied to standardised GARCH residuals. Stronger evidence exists for nonlinear dependence in sovereign bond yield returns than for stock returns; however, the results are sensitive to the methodology used. Section 4.2 explains the BDS methodology in more detail and provides details of the algorithms used to generate the surrogate data sets, as well as the invariant statistics used for hypothesis testing. Section 4.3 describes the stock index and sovereign bond samples. In Section 4.4 the results of the BDS tests and the surrogate data analysis are presented. Finally, in Section 4.5 we conclude with a brief discussion of the results and their implications.

4.2 Methodology

4.2.1 BDS Test

The BDS test is a test for nonlinearity based on the correlation integral. The null hypothesis of the test is that the time series under analysis is IID. It was originally designed as a test for deterministic chaos, but also as a test of goodness of fit when applied to the residuals of a model (Brock et al., 1996). The BDS test statistic uses the fact that the correlation integral is a measure of the probability that two points in embedding space are closer than a given distance (ϵ). The correlation integral, $C_{de}(\epsilon)$, is the continuous time version of the correlation sum described in Equation 3.9 of Chapter 3. If the points are independent of each other, then the probability of two points being closer than ϵ in de dimensions will be the probability that they are closer

than ϵ in one dimension to the power of de , or

$$C_{de}(\epsilon) = [C_1(\epsilon)]^{de} \quad (4.1)$$

The standardized BDS statistic is estimated as

$$BDS = \frac{\sqrt{n}(C_{de}(\epsilon) - [C_1(\epsilon)]^{de})}{\sigma_{de}(\epsilon)}, \quad (4.2)$$

where de is the embedding dimension of the phase space reconstruction and $C_{de}(\epsilon)$ is the correlation integral.

No alternate hypothesis is specified for the BDS statistic; therefore, rejecting the null hypothesis indicates that the times series in question is not IID but does not infer anything about the type of dependence that remains in the residuals. A key benefit of the BDS statistic is that it is asymptotically normally distributed with $N(0, 1)$.

Therefore, it is possible to specify critical values and calculate p-values given a large enough sample (Brock et al., 1996).

As we are using the BDS test for financial time series, we must first apply a GARCH type model to correct for volatility clustering and apply the test to the standardised residuals. Lima (1996) states that in order for the BDS statistic to be asymptotically distributed according to a standard normal distribution, when applied to model residuals, the model in question must be a linear additive model, or be capable of being cast in this format. Therefore, in order for GARCH residuals to be compatible, we must use the log-squared standardised residuals (Caporale et al., 2005).

Where the scaling value is large and volatility clustering is present, the power of the BDS test is reduced (Moloney and Raghavendra, 2011). The standard method for carrying out the BDS test is to apply the test over a range of embedding dimensions, and over a range of scaling values (generally ϵ/σ between 0.5 and 2) (Brock and Sayers, 1988). According to Kanzler (1999), asymptotic convergence of the BDS statistic to the standard normal distribution, as well as the independence of the probability of falsely rejecting the null hypothesis from the choice of de and ϵ , may not hold for small samples. The author carries out a Monte Carlo study where he

calculates critical values for the BDS statistic over different embedding dimension, scaling values and sample sizes. In this chapter we use the Kanzler (1999) critical values for hypothesis testing.

4.2.2 Surrogate Data Analysis

The purpose of surrogate data analysis, as described in Theiler et al. (1992), is to test a null hypothesis regarding the data generating process that governs the dynamics of the time series of interest. A number of surrogate data sets are generated, which preserve the statistical properties of the original data that are relevant to the null hypothesis but are otherwise random. For example, if our null hypothesis is that of IID data, any random shuffle of our original time series would provide a surrogate that is consistent with this hypothesis.

The primary application of surrogate data analysis in the literature is to test for evidence of nonlinearity. In such tests, surrogate data sets are generated which preserve the linear dependence structure of the data. An invariant statistic is then computed on the surrogate data sets and also on the original data. The invariant statistic should be able to capture nonlinear deterministic behaviour present in the original data that does not exist in the surrogates. Finally, the distribution of the statistics is examined. If the statistic calculated on the original data does not fall within the distribution of statistics calculated on the surrogates we can reject the null hypothesis.

A number of different methods have been set out in the literature to generate surrogates consistent with the hypothesis that the original data is generated by a linear stochastic process. Amongst the most popular are algorithms based on the Fourier transform. The Wiener Khinchin theorem states that the Fourier transform of the autocorrelation function of a signal is equal to its power spectrum⁹ (Kantz and Schreiber, 2004). Therefore, we can create a surrogate time series with the same autocorrelation function as the one we are testing by multiplying the Fourier

⁹The power spectrum is defined as the squared modulus of the continuous Fourier transform and gives the square of the amplitude by which a frequency contributes to a signal. The discrete series version of the power spectrum is called the periodogram.

transform of the data by random phases and then transforming it back to the time domain (Schreiber and Schmitz, 2000).

Another hypothesis that can be tested using Fourier based surrogates is that deviations from the normal distribution are caused by a static, invertible, measurement function. In other words, the data is consistent with a nonlinear transformation of linear filtered noise. Methods such as the Amplitude Adjusted Fourier Transform (AAFT) (Theiler et al., 1992) are used to test this hypothesis. However, the AAFT algorithm has been criticised as the power spectra of the original and AAFT surrogate can diverge when the sample size is small, and also because the inversion of the measurement function is not equal the true inversion. This will turn into a bias in the power spectrum, causing it to be flatter than that of the original data, and can lead to false rejection of the null hypothesis (Schreiber and Schmitz, 2000).

Schreiber and Schmitz (1996) propose an alternative iterative method (IAAFT) which will have power spectra practically indistinguishable from the original time series.

Kugiumtzis (2008) states that for AAFT surrogates, if the measurement functions is non-monotonic, the surrogates cannot match the linear structure of the original time series. Moreover, a discrepancy occurs in IAAFT due to the iterative process used in estimation, which starts from a flat spectrum and finishes with the same accuracy of approximation for each of the surrogates. Kugiumtzis (2008) states that this discrepancy, when combined with small variance of the spectra of the surrogates, causes a bias in favour of rejection of the null hypothesis. Kugiumtzis (2002a) proposes a method for the generation of surrogate data sets that does not suffer from the biases of the AAFT and IAAFT algorithms, called Statistically Transformed Autoregressive Process (STAP). The STAP method generates surrogate data sets, Z , from the original data, X , which are realisations of a scalar linear stochastic process with the same autocorrelation structure, $\rho_X = \rho_Z$ and marginal cdf $\Phi_X = \Phi_Z$ (Kugiumtzis, 2002a). The method is different from the AAFT and IAAFT approaches as it does not approximate the sample periodogram, instead approximating the sample autocorrelation and building a proper autoregressive model in order to generate surrogate data (Kugiumtzis, 2008).

Small and Tse (2003) demonstrate that surrogate data sets consistent with the

hypothesis of a nonlinear transformation of linearly filtered noise are not able to capture GARCH effects in financial time series. They find that surrogates generated using the IAAFT algorithm cause rejection of the null hypothesis when the methods are applied to daily stock market, gold and foreign exchange returns. Therefore, if we use IAAFT or STAP surrogate methods to test for nonlinearity in financial time series, we must do so using standardised GARCH residuals. However, Small and Tse (2003) highlight two problems testing model residuals for nonlinearity. Firstly, when testing the model residuals, you are testing the hypothesis that a specific model generated the data. Secondly, pre-whitening the data using a model can either mask chaotic dynamics or introduce spurious determinism. Therefore, it is preferable to apply surrogate data testing directly to the original data and to use surrogate generating algorithms that can capture the hypothesis being tested.

Small, Yu, and Harrison (2001) propose a methodology based on local-linear modelling approaches, Pseudo-Periodic Surrogates (PPS), which can mimic minor nonlinearity in the original data but will not reproduce the dynamics entirely. Small and Tse (2003) state that this method will generate surrogates which reproduce large scale dynamics, such as trends or persistence, but should not capture deterministic chaos or higher order nonlinearity. As such, rejection of the null will imply that the data may not be modelled by local linear models, autoregressive models or state dependent noise processes such as GARCH. The Small, Yu, and Harrison (2001) algorithm is outlined briefly below

Step 1 is to construct a time delay embedding

$\{x_t\}_{t=1}^N \mapsto (x_t, x_{t-1}, \dots, x_{t-de+1}) = \{s_t\}_{t=1}^{N-de+1}$, calling A the reconstructed attractor.

In step 2, select a random point in A as an initial condition, $z_1 \in A$ and let $i = 1$.

In step 3 choose a nearest neighbour of z_i , $s_j \in A$ according to the probability

distribution

$$Prob(s_i = s_t) \propto \exp\left(-\frac{\|s_t - z_i\|}{\rho}\right), \quad (4.3)$$

where ρ is the noise parameter. In step 4 set $z_{i+1} = s_{i+1}$ and if $i < N$ repeat step 3.

The surrogate time series is the first scalar component of z_t . The choice of noise parameter is critical in determining the level of nonlinearity preserved in the

surrogates. Small and Tse (2003) suggest choosing ρ such that the number of short segments where the surrogate and original data sets are identical is maximised¹⁰.

4.2.3 Invariant Statistics

According to Schreiber and Schmitz (1997), the invariant statistics calculated on the original data should be sensitive to the kind of nonlinearity present in the data, and it should be possible to estimate this value with low variance. The measures must be powerful in discriminating between linear dynamics and weak nonlinear signatures (Schreiber and Schmitz, 2000). We avoid estimation of the GP correlation dimension and the Lyapunov exponent due to the issues with estimation for short, noisy time series discussed in Section 4.1. In this chapter we use a number of statistics which have been shown to be robust in detecting (weakly) nonlinear dynamics (Schreiber and Schmitz, 1997; Small and Tse, 2003; Small, 2005; Zhao et al., 2006; Kugiumtzis and Tsimpiris, 2010). These statistics are

1. Cumulative Mutual Information
2. Algorithmic Complexity
3. Nonlinear Prediction Error

Cumulative Mutual Information: Mutual information between two time series X and Y , $I(X; Y)$, is described in Cover and Thomas (1991) as the reduction in the uncertainty of X due to the knowledge of the value of Y . When used contemporaneously between two variables, mutual information is a measure of correlation and can detect both linear and nonlinear relationships. Mutual information can be expressed as

$$I(X; Y) = \sum_{x \in X} \sum_{y \in Y} p(x, y) \log \frac{p(x, y)}{p(x)p(y)}, \quad (4.4)$$

where $p(x, y)$ is the joint distribution of X and Y and $p(x)p(y)$ is the product of the individual probability mass functions of X and Y respectively.

¹⁰The authors propose a segment of length 2 is appropriate for financial time series.

Mutual Information is derived from the concept of entropy, $H(X)$, which is a measure of uncertainty in a random variable X .

$$H(X) = - \sum_{x \in X} p(x) \log p(x) \quad (4.5)$$

This expression for entropy is called Shannon entropy (Shannon, 1948). The relationship between entropy and mutual information is given in Cover and Thomas (1991) as

$$I(X, Y) = H(X) - H(X | Y) = H(Y) - H(Y | X) \quad (4.6)$$

Where $H(X | Y)$ is the conditional entropy of X given Y and is given by the equation

$$H(X | Y) = - \sum_{x \in X} \sum_{y \in Y} p(x, y) \log p(x | y), \quad (4.7)$$

where $p(x | y)$ is the conditional probability of x given y . When applied between a time series and its own lagged values, Mutual information gives a measure of linear and nonlinear serial dependence. Thus, if mutual information calculated on the original data is statistically different from mutual information calculated on the surrogate data, there is evidence of nonlinear serial dependence in the original data. As proposed in Kugiumtzis and Tsimpiris (2010), we use cumulative mutual information which sums the magnitude of the mutual information values up to a given maximum lag. A maximum lag of 5 trading days is deemed to be sufficient due to the well documented fast decay of correlation dynamics in financial returns. In this chapter, we estimate mutual information using equi-probable binning.

Algorithmic Complexity: Algorithmic complexity is a measure which allows testing for deterministic dynamics in surrogate data analysis (Zhao et al., 2006).

When applying the measure to a financial time series data, one must first convert the high precision data to a sequence of elements from a finite, small set. In this chapter we convert the data into a binary sequence where the data falls into the binning $\{0, 1\}$ with equal probability. After the conversion, we are left with a binary series which is the same length as the original time series. Algorithmic complexity is then calculated as the number of sequences one observes in the symbolic sequence as a fraction of the

maximum possible number of sequences, where the maximum is equivalent to the number of sequences observed in a random time series (Small, 2005).

The Lempel-Ziv algorithm (Lempel and Ziv, 1976), which forms the basis of many file compression methods such as WinZip, is used to calculate the complexity of a time series. The methodology is outlined in Zhao et al. (2006) and described as follows:

For a sequence, S , of length n , where $S = (s_1, s_2, \dots, s_n)$. For a binary encoding system, each s_i is one of 2 symbols, $s_i \in A = [0, 1]$. $c(n)$ is the counter of unique sub-sequences in the sequence S . P and Q are two sub-strings of S ; PQ is the concatenation of P and Q and $PQ\pi$ is the concatenation of P and Q with the last symbol deleted. $v(PQ\pi)$ is the set of all sub-strings of $PQ\pi$. The procedure for calculating algorithmic complexity is as follows.

1. Initialise $c(n) = 1, P = s_1, Q = s_2$. Therefore, $PQ\pi = s_1$. If $Q \in v(PQ\pi)$, leave P unchanged and update $Q = s_2s_3$; if $Q \notin v(PQ\pi)$, add one to $c(n)$, update $P = s_1s_2$ and $Q = s_3$.
2. Assume that $P = s_1s_2\dots s_r, Q = s_{r+1}$. If $Q \in v(PQ\pi)$, leave P unchanged and update $Q = s_{r+1}s_{r+2}$, and again assess whether $Q \in v(PQ\pi)$. Repeat until $Q \notin v(PQ\pi)$ and update $c(n) = c(n) + 1$. Q is now equal to $s_{r+1}s_{r+2}\dots s_{r+i}$ and P should be updated to $P = s_1s_2\dots s_r s_{r+1}s_{r+2}\dots s_{r+i}$. Q is then set equal to s_{r+i+1}
3. Repeat step 2 until Q reaches the last substring of S .

(Lempel and Ziv, 1976) demonstrate that for sequence of length n consisting of d symbols (for a binary sequence $d = 2$)

$$c(n) < \frac{n}{(1 - 2^{\frac{(1+\log_d \log_d(dn))}{\log_d(n)}}) \log_d n} \quad (4.8)$$

When $n \rightarrow \infty$, $\frac{(1+\log_d \log_d(dn))}{\log_d(n)} \rightarrow 0$; therefore, the normalised algorithmic complexity is calculated as

$$Comp(n) = \frac{c(n)}{n} \log_d(n) \quad (4.9)$$

If algorithmic complexity calculated on the original data is statistically different from algorithmic complexity calculated on the surrogate data, there is evidence of nonlinear deterministic dynamics in the original data.

Nonlinear Prediction Error: The Nonlinear Prediction Error is a measure of the predictability of a time series (Schreiber and Schmitz, 1997). Small (2005) defines the measure for i time steps into the future as

$$e_f(i) = \sum_t \|f(x_t, i) - x_{t+i}\|, \quad (4.10)$$

where $f(x_t, i)$ is the model prediction of the evolution of x_t at time $t + i$. The class of model chosen is important and the nonlinear prediction error is the minimum over all models. However, as testing the prediction error over all classes of model is impractical, Small (2005) suggests using a simple model which states that for an embedded times series, $s_t = (x_t, \dots, x_{t-de+1})$, the best prediction of s_{t+1} will be s_{m+1} , where s_m is the closest neighbour to s_t . In other words, the prediction is to look at the time $t + 1$ position of the time t closest neighbour. The i step nonlinear prediction error is then defined in Small (2005) as

$$E(i) = \sqrt{\frac{1}{N-i} \sum_{t=1}^{N-i} (x_{t+i} - x_{m+i})^2}, \quad (4.11)$$

where the x_{t+i} and x_{m+i} are the first scalar components of s_{t+i} and s_{m+i} respectively and

$$m = \arg \min_{\substack{1 \leq m \leq N-i \\ m \neq t}} \|s_t - s_m\| \quad (4.12)$$

If the nonlinear prediction error calculated on the original data is statistically different from the nonlinear prediction error calculated on the surrogate data, there is evidence of nonlinear predictability in the original data.

4.3 Data

We obtain 10 year daily sovereign bond yield data for Greece, Ireland and Portugal from Thomson Datastream. We apply tests for nonlinear determinism to bond yield returns series of length 500 trading days, 1,000 trading days and 2,500 trading days preceding the onset of the sovereign debt crisis. The shorter samples allow for testing of short term nonlinear determinism in the lead-up to the crisis, but the results may be less robust. The longer samples are more robust, but results may be less sensitive to short run nonlinear deterministic behaviour. Small and Tse (2003) use surrogate data analysis to test for non-linear determinism in financial time series, finding the results to be less robust for sample sizes less than 800 observations. This finding is due to the critical dependence of algorithms, such as the correlation dimension and the nonlinear prediction error, on the quantity of data available for testing.

The final observation for each of our samples is the 30th of April 2010. This cut-off date was chosen due to the fact that the Troika bailout of Greece occurred in May 2010. This event lead to significant increases in bond yields in Portugal and Ireland, culminating in bailouts in Ireland in December 2010 and in Portugal in May 2011. Thus, the end of April 2010 is a logical cut-off point.

We also obtain daily stock index data for the Dow Jones Industrial Average, FTSE 100 and NIKKEI 225, and apply the tests for nonlinearity to the daily log-returns. Again we use sample sizes of 500, 1,000 and 2,500 with a cut-off date of the 28th of September 2007. This date broadly corresponds to the end of a growth period in each of the stock markets, which was almost uninterrupted since Q1 2003, and to the onset of the global financial crisis. Moloney and Raghavendra (2011) apply tests for nonlinearity to the S&P 500, the FTSE 100 and the NIKKEI 225 using sample sizes of 14,821, 6338 and 6,127 observations respectively, with a cut-off date of the end of November 2008. While the larger sample is more robust, we hope our approach may be more sensitive to short-term nonlinear dynamical behaviour which may have emerged closer to the global financial crisis.

Table 4.1 presents descriptive statistics for each of the bond and stock market samples. The data are characterised by volatility clustering, fat tails and skewness as

evidenced by the ARCH LM test p-value and the Jarque-Bera test statistic. This is in line with the stylised facts of financial time series discussed in Chapter 1.

TABLE 4.1: Stock and bond time series: sample and descriptive statistics

Sample and Descriptive Statistics									
	Dow Jones			FTSE 100			NIKKEI 225		
Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Std. dev.	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Skewness	-0.49	-0.27	-0.14	-0.42	-0.42	-0.19	-0.35	-0.49	-0.17
Kurtosis	5.45	4.27	6.7	5.62	6.12	5.58	4.47	4.72	4.79
JB Stat	144.91	79.88	1435.2	157.36	436.38	710.74	54.71	162.5	347.91
JB p-val	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SC LM p-val	0.00	0.23	0.16	0.31	0.08	0.00	0.2	0.47	0.08
ARCH LM p-val	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Start Date	05/10/2005	13/10/2003	29/10/1997	05/10/2005	13/10/2003	29/10/1997	27/07/2005	13/05/2003	27/09/1996
End Date	28/09/2007	28/09/2007	28/09/2007	28/09/2007	28/09/2007	28/09/2007	28/09/2007	28/09/2007	28/09/2007
No. Obs.	500	1,000	2,500	500	1,000	2,500	500	1,000	2,500
	Greece 10 yr			Ireland 10 yr			Portugal 10 yr		
Mean	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Std. dev.	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Skewness	0.46	0.58	0.68	0.61	0.48	0.61	0.5	0.46	0.57
Kurtosis	6.8	9.6	10.5	5.14	5.47	6.29	6.96	8.13	7.77
JB Stat	319.28	1869.87	6055.36	126.03	293.12	1283.78	347.66	1132.32	2512.04
JB p-val	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
SC LM p-val	0.00	0.00	0.00	0.00	0.00	0.01	0.00	0.00	0.01
ARCH LM p-val	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00
Start Date	07/05/2008	12/05/2006	31/05/2000	07/05/2008	12/05/2006	31/05/2000	07/05/2008	12/05/2006	31/05/2000
End Date	30/04/2010	30/04/2010	30/04/2010	30/04/2010	30/04/2010	30/04/2010	30/04/2010	30/04/2010	30/04/2010
No. Obs.	500	1,000	2,500	500	1,000	2,500	500	1,000	2,500

4.4 Results

4.4.1 BDS Test

We apply the GARCH(1,1) model, as outlined in Chapter 1, to each of sample sizes for our sovereign bond and stock market log-returns series. Serial Correlation LM and ARCH LM tests are then used to test for remaining linear serial correlation and conditional heteroskedasticity. The GARCH(1,1) specification was sufficient to remove all conditional heteroskedasticity, as evidenced by insignificant nR^2 statistics in the ARCH-LM test. An AR1 term was added to remove remaining linear serial dependence for the 500 observation sample of the Dow Jones Industrial Average, the Greek 10 year bond and the Irish 10 year bond. An AR1 term was also added for the 1,000 observation sample of the Greek 10 year bond.

The standardised residuals of each model are obtained and the log-squared transformation applied, as recommended in Caporale et al. (2005). The BDS statistic is then applied to the 18 time series (six stock and bond market series over the three sample sizes), with embedding dimensions varying from 2 to 7 and scaling regions 0.5σ , 1σ , 1.5σ and 2σ . Brock et al. (1996) state that the standardised BDS test statistic follows a standard normal distribution if $n/de > 200$. We note that $n/de < 200$ for sample size $n = 500$ with $de = 3, \dots, 7$ and sample size $n = 1,000$ with $de = 6, 7$. Therefore, the BDS statistic is estimated and compared to the critical values estimated by Kanzler (1999) for small samples. We estimate the BDS test statistic on the log-squared GARCHs residuals using the tseries package in R. The results of the BDS test using the Kanzler (1999) critical values were compared to the BDS results using bootstrapped p-values in Eviews. In all cases the significance levels were equivalent or more conservative using the Kanzler (1999) critical values. A two tail test is applied, where for a 5% significance level the 2.5% and 97.5% critical values are used. The results of the BDS test on the GARCH (1,1) residuals are presented in Table 4.2, Table 4.3 and Table 4.4 for $n=500$, $n=1,000$ and $n=2,500$ respectively. A sensitivity test was carried out by applying the BDS test to the log-squared standardised residuals of an EGARCH model applied to the time series under analysis. The results were not qualitatively different from those presented for the GARCH model. The results for

the EGARCH model are presented in the appendices to this chapter. We also compare all significant results to those obtained from a GARCH(1,2) model, finding that the results are consistent with those presented in Tables 4.2– 4.4. Similarly to the results presented in Moloney and Raghavendra (2011), in all cases the null hypothesis of IID residuals is accepted for the Dow Jones Industrial Average, FTSE 100 and NIKKEI 225. This indicates that the GARCH (1,1) model sufficiently captures nonlinear dependence in the stock market indices under examination in the lead-up to the global financial crisis.

TABLE 4.2: BDS Statistic for Log Squared GARCH Residuals, N=500. *, ** and *** indicate significant at the 10%, 5% and 1% respectively using a two-tailed test.

m	Dow	FTSE	NIKKEI	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	-0.68	-1.34	-0.18	1.8	-1.7	0.7
3	-0.87	-0.68	-0.39	2.04	-2.38*	1.01
4	-0.53	-0.02	-0.58	2.18	-2.50*	1.31
5	-1.31	0.51	0.25	1.55	-2.22	0.89
6	-1.71	-0.01	0.18	0.71	-1.97	0.16
7	-2.91	-0.17	0.5	0.56	-2.65	-0.45
$\epsilon/\sigma = 1$						
2	-1.21	-0.51	-0.36	1.88*	-1.63	-0.46
3	-0.67	-0.33	-0.77	2.11*	-2.28**	-0.3
4	-0.25	0.33	-0.69	2.35**	-2.36**	-0.28
5	-0.58	0.85	-0.5	2.08*	-2.29**	-0.37
6	-0.69	0.81	-0.57	1.8	-2.39**	-0.6
7	-0.7	0.64	-0.58	1.65	-2.08**	-0.93
$\epsilon/\sigma = 1.5$						
2	-0.64	-0.24	-0.75	1.95*	-0.59	-0.49
3	-0.43	-0.18	-0.88	2.26**	-1.2	-0.4
4	-0.02	0.05	-0.69	2.33**	-1.47	-0.65
5	-0.19	0.42	-0.56	2.02*	-1.48	-0.82
6	-0.22	0.32	-0.63	1.75*	-1.45	-1.13
7	-0.33	0.14	-0.61	1.61	-1.18	-1.38
$\epsilon/\sigma = 2$						
2	0	-0.06	-0.7	1.83*	0.02	-0.08
3	0	0.06	-0.76	1.76*	-0.44	-0.14
4	0.22	0.3	-0.47	1.59	-0.97	-0.47
5	0.15	0.49	-0.3	1.18	-1.15	-0.76
6	0.34	0.4	-0.42	1.04	-1.08	-1.12
7	0.36	0.34	-0.51	0.96	-0.84	-1.39

The results for the sovereign bond markets are somewhat more interesting. For the Greek 10 year bond, using a sample size of 500, the null hypothesis of IID residuals is rejected at a 10% significance for $de = 2, 3$ and 5, and at a 5% significance for $de = 4$ with $\epsilon = 1\sigma$. However, bearing in mind that the BDS statistic should have power to reject the null hypothesis independent of the choice of de , the negative result for $de = 6$ and 7 does not allow for a definitive conclusion of non-IID residuals. For $\epsilon = 1.5\sigma$

4.4. Results

the results are again inconclusive due to the failure to reject the null of IID residuals for $de = 7$. For $n = 1,000$ the results are much more conclusive for the Greek sovereign bond data. Using ϵ greater than or equal to 1σ , the null hypothesis of IID residuals is rejected for all embedding dimensions with at least 10% significance.

TABLE 4.3: BDS Statistic for Log Squared GARCH Residuals, N=1000.
*, ** and *** indicate significant at the 10%, 5% and 1% respectively using a two-tailed test.

m	Dow	FTSE	NIKKEI	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	-0.42	-1.05	0.31	1.89*	-1.15	0.5
3	-1.32	-0.91	-0.29	1.78	-1.67	0.31
4	-0.81	-0.2	0.11	2.62**	-1.2	0.39
5	-0.7	-0.14	0.51	2.36	-0.87	0.4
6	-1.12	-0.43	0.28	2.12	-0.93	0.13
7	-1.62	-0.38	0.03	1.99	-0.7	-0.09
$\epsilon/\sigma = 1$						
2	-0.8	-0.24	0.21	1.85*	-0.83	0.54
3	-0.74	-0.18	-0.27	2.05*	-1.66	0.36
4	-0.36	0.5	0	2.73**	-1.47	0.43
5	-0.47	0.77	0.3	2.82**	-1.08	0.6
6	-0.57	0.73	0.2	2.92**	-0.78	0.71
7	-0.72	0.63	0.25	2.91**	-0.58	0.7
$\epsilon/\sigma = 1.5$						
2	0.07	0.26	0.16	1.89*	-0.74	1.08
3	0.3	0.62	-0.21	2.26**	-1.29	1.05
4	0.68	1.4	-0.11	2.73**	-1.15	1.04
5	0.62	1.77*	-0.02	2.79**	-0.85	1.03
6	0.52	1.73*	-0.15	2.92***	-0.54	1.03
7	0.22	1.61	-0.13	2.96***	-0.36	1.04
$\epsilon/\sigma = 2$						
2	0.79	0.7	-0.23	2.03*	-0.14	2.15**
3	1.02	0.94	-0.4	2.34**	-0.59	2.16**
4	1.24	1.64	-0.23	2.48**	-0.58	1.98*
5	1.11	1.93*	-0.11	2.39**	-0.43	1.73*
6	1.08	1.80*	-0.26	2.43**	-0.27	1.54
7	0.95	1.79*	-0.28	2.53**	-0.16	1.51

The results for Ireland are negative for $n = 1,000$ and $n = 2,500$. However for $n = 500$ and $\epsilon = 1\sigma$ the null hypothesis is rejected at the 5% level for all but $de = 2$. These mixed results do not provide conclusive evidence of remaining dependence structure in the GARCH residuals of the Irish sovereign bond. Finally, the results for Portugal are negative for $n = 500$ and $n = 1,000$. However, for $n = 2,500$ the null of IID residuals is rejected across all scaling regions and embedding dimensions, providing strong evidence of remaining dependence structure in the GARCH residuals.

The results presented above indicate that for the Greek and Portuguese government bond samples, the GARCH(1,1) model cannot capture the entire dependence

TABLE 4.4: BDS Statistic for Log Squared GARCH Residuals, N=2500.
 *, ** and *** indicate significant at the 10%, 5% and 1% respectively
 using a two-tailed test.

m	Dow	FTSE	NIKKEI	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	0.3	-0.06	-0.17	0.12	-0.31	2.97***
3	-0.67	-0.5	-0.39	0.56	-1.03	3.02***
4	-0.82	-0.3	-0.02	1.05	-1.21	2.89**
5	-1.09	0.08	0.55	1.43	-0.97	3.24**
6	-1.07	0.21	0.39	1.74	-0.55	3.68***
7	-1.26	0.29	0.66	1.72	-0.52	3.81**
$\epsilon/\sigma = 1$						
2	-0.08	-0.06	0.15	0.43	0.12	4.31***
3	-0.59	-0.5	-0.09	0.67	-0.8	4.12***
4	-0.45	-0.3	0.34	0.91	-0.99	3.85***
5	-0.36	0.08	0.94	1.18	-0.63	3.96***
6	-0.34	0.21	1.02	1.54	-0.42	3.98***
7	-0.41	0.29	1.25	1.64	-0.36	3.84***
$\epsilon/\sigma = 1.5$						
2	0.72	0.05	0.06	1.04	0.48	6.79***
3	0.37	0.12	-0.16	0.98	-0.44	6.42***
4	0.56	0.51	0.07	1	-0.68	5.93***
5	0.67	0.93	0.49	1.05	-0.46	5.71***
6	0.73	1	0.54	1.4	-0.36	5.43***
7	0.61	1.18	0.72	1.49	-0.37	5.14***
$\epsilon/\sigma = 2$						
2	1.3	-0.61	0.04	1.49	0.7	9.28***
3	1.04	-0.08	0	1.08	0.06	9.07***
4	1.13	0.47	0.13	0.93	-0.02	8.60***
5	1.21	0.82	0.39	0.85	0.14	8.17***
6	1.3	0.88	0.47	1.06	0.13	7.60***
7	1.21	1.14	0.63	1.08	0.12	7.10***

structure of the data in the lead-up to the sovereign debt crisis. The mixed results for the Irish government bond sample with $n = 500$ requires further investigation. In the following section we will further investigate the results presented above using surrogate data methods.

4.4.2 Surrogate Data Analysis

A key weakness of the BDS test is that it cannot tell us anything about the type of dependence structure in the standardised GARCH residuals. It can only tell if the data is IID or not. Small and Tse (2003) state that acceptable alternate hypotheses include linear noise, nonlinearity, chaos, non-stationarity or persistent processes. Moreover, where the null hypothesis is rejected when applied to standardised residuals, the rejection of the null may be due to misspecification of the model or the wrong choice model from a particular family of models. We attempt to offset the risk

4.4. Results

of misspecification, or poor model choice, by applying both GARCH and EGARCH models and by testing the residuals using the serial correlation LM test and ARCH LM test. However, surrogate data analysis can provide us with more information regarding the type of dynamical system which is generating the results.

In this section we present the results of two null hypotheses using surrogate data analysis. Firstly, using the Kugiumtzis (2002a) STAP surrogates, that the financial time series represent a static transformation of linearly filtered noise¹¹. Secondly we use the PPS algorithm of Small, Yu, and Harrison (2001), with the null hypothesis that the data may be adequately modelled by local linear models, autoregressive models or state dependent noise processes, such as GARCH processes. Moreover, our choice of invariant statistics provides more information on the type of dynamical system which is generating the data, if the null hypothesis is rejected.




FIGURE 4.1: Dow Jones Industrial Average log-returns for $N = 2500$ up until the 28/09/2007 (top panel). Simulated GARCH data generated using GARCH(1,1) model fitted to the Dow Jones Industrial Average log-returns (bottom panel)

¹¹The IAAFT framework of Schreiber and Schmitz (1996) is also consistent with this hypothesis and the results of surrogate data analysis using this framework are presented in the appendices to this chapter. The results are not qualitatively different from the STAP results.

In order to demonstrate the power of the STAP and PPS data to test against the above null hypotheses, we fit a GARCH(1,1) model to the Dow Jones Industrial Average 2,500 observation sample described in the Section 4.3. We use the parameter values for the model to generate a random time series with similar GARCH properties to the real data. The original data and simulated time series are presented in Figure 4.1. We calculate Pearson's cumulative autocorrelation (5 lags), cumulative mutual information (5 lags), algorithmic complexity and nonlinear prediction error¹² for the GARCH time series and also for surrogate data sets generated using the STAP and PPS methods. 100 surrogate time series were generated for each of the surrogate data algorithms. The results are presented in Figures 4.2 and 4.3, with p-values generated using t-tests.

The Pearson cumulative autocorrelation function can be seen as a test for the correct specification of the surrogate data sets. If the Pearson cumulative autocorrelation calculated on the original data is statistically different from the statistics calculated on the surrogate data sets, this means there is linear dependence in the data which is not preserved in the surrogates (Kugiumtzis and Tsimpiris, 2010). In both Figure 4.2 and Figure 4.3 the autocorrelation statistic falls within the distribution; therefore, the surrogate data sets are correctly generated.

The results of the simulation exercise using the STAP surrogates presented in Figure 4.2 demonstrate that we can reject the null hypothesis of a static transformation of linearly filtered noise for the simulated data, using cumulative mutual information and the nonlinear prediction error. This is expected because data generated by the GARCH family of models do not fall within this hypothesis. Interestingly, the algorithmic complexity statistic does not detect any differences between the GARCH data and the surrogate time series. If algorithmic complexity were used in isolation, this would indicate that no nonlinear deterministic dynamics were present in the data. These results imply that, if we wish to further investigate the results of the BDS tests using STAP generated surrogates, we must apply the test to the standardised GARCH residuals rather than to our original financial time series.

¹²An embedding dimension of 7 and time lag of 1 were used for both the calculation of the nonlinear prediction error and the generation of PPS surrogates.

4.4. Results




FIGURE 4.2: Results of STAP surrogate data analysis applied to simulated GARCH(1,1) process.

PPS data can capture non-stationarity and drifts, locally linear nonlinearities and non-stationarity in the variance (Small and Tse, 2003). As expected, the PPS surrogate results in Figure 4.3 fail to reject the null hypothesis that the simulated GARCH data may be adequately modelled by local linear models, autoregressive models or state dependent noise processes. Therefore, the PPS surrogate data sets have preserved the conditional heteroskedasticity in the simulated GARCH series. This implies that we can apply PPS surrogate data methods directly to our original log-returns series in order to test for nonlinearities in the data, beyond those which can be explained using GARCH modelling techniques.

The results of the surrogate data analysis of the stock and bond market samples are presented in Tables 4.5 and 4.6. As discussed above, the STAP surrogate testing was applied to the standardised GARCH residuals, whereas the PPS analysis was applied directly to the log-returns of the original data. 1,000 surrogate data series are generated using each of the surrogate algorithms, and invariant statistics are calculated on the original data (standardised GARCH residuals or raw returns

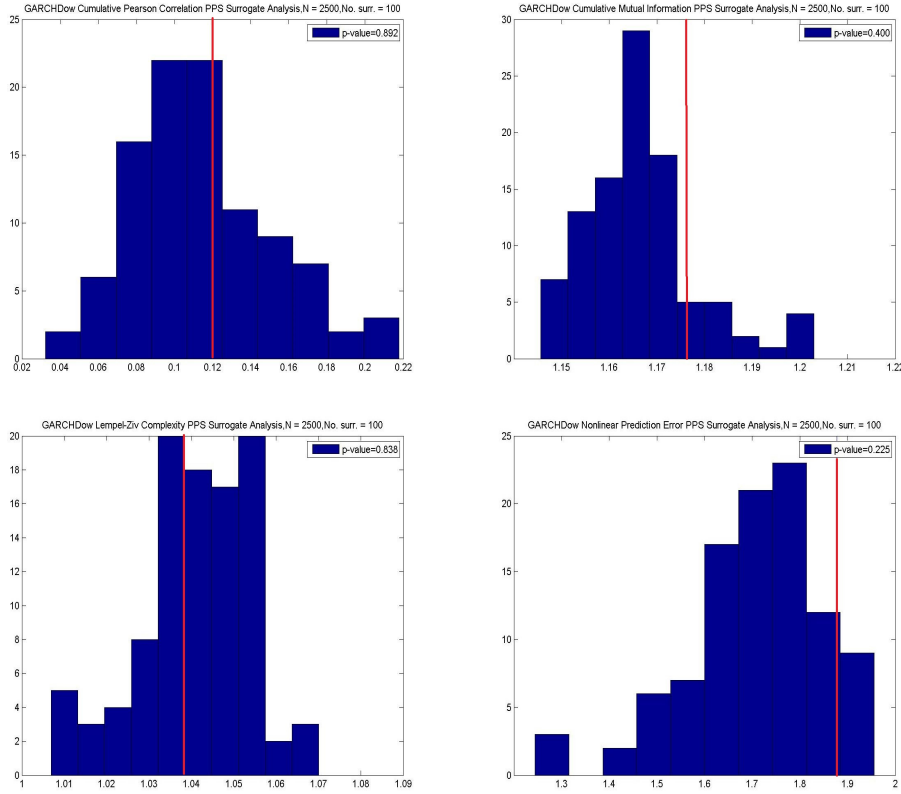


FIGURE 4.3: Results of PPS surrogate data analysis applied to simulated GARCH(1,1) process

depending on the algorithm). False Nearest Neighbours and the first zero of Average Mutual Information were used to determine the embedding dimension and time lag for generation of the PPS data and the calculation of the nonlinear prediction error. An embedding dimension of 7 and a time lag of 1 were used in all cases. A probability value is obtained by calculating a test statistic and applying a Student's T test.

$$TestStatistic = \frac{\mu - \mu_s}{\sigma_s}, \quad (4.13)$$

where μ is the invariant statistic calculated on the original time series, μ_s is the mean of the invariant calculated on the surrogate data sets, and σ_s is the standard deviation of the invariant calculated on the surrogate data series. This method for generating p-values assumes a Gaussian distribution for the invariant statistic, which may not hold up in reality. Kugiumtzis (2002b) suggest this approach is valid, provided the invariant statistics calculated on the surrogate data sets are approximately normally distributed. For simplicity of presentation we present the p-values in Tables 4.5 and 4.6; however, we have also examined the rank ordering of the invariant statistics to

confirm significance in all cases.

Turning our attention to Table 4.5, looking at the equity indices we find evidence of nonlinearity in the data. For example, the insignificant Pearson's cumulative autocorrelation statistic indicates that the surrogates are preserving the linear dependence structure in all cases. The results of the STAP surrogate data analysis for the stock market differ from those generated using the BDS test. For the Dow Jones Industrial Average we find a significant algorithmic complexity statistic and a significant cumulative mutual information statistic for $n = 500$. This indicates that nonlinear deterministic dynamics and nonlinear serial dependence are detectable in the standardised GARCH residuals of the returns series. Moreover, as we have seen from the simulation exercise above, the algorithmic complexity variable does not detect GARCH effects. Therefore, the remaining structure in the data is likely not conditional heteroskedasticity caused by a mis-specified GARCH model. For the Dow Jones $n = 2,500$ series, the FTSE $n = 2,500$ series, the NIKKEI $n = 500$ and the NIKKEI $n = 2,500$, the cumulative mutual information statistic is the only significant invariant statistic. A significant cumulative mutual information statistic, coupled with an insignificant Pearson's cumulative autocorrelation statistic, indicates that there is some remaining serial dependence in the data and that the dependence is nonlinear.

The result for the sovereign bond series in the lower half of Table 4.5 also show evidence of nonlinearity and are somewhat similar to the results of the BDS test. The nonlinear prediction error is significant for the Irish 10 year bond with $n = 500$, as is the BDS statistic. The nonlinear prediction error uses a local model to determine the level of nonlinear predictability in the data. The results for the Irish 10 year bond indicates that the predictive power of the data is worse than the surrogates, as the prediction error is in the upper tail of the surrogate distribution. This may seem counterintuitive; however, Small and Tse (2003) find similar results in financial time series. They state that if the predictive power of the data is greater than that of the surrogates, it would indicate that a local linear model could be used for profitable trading. Therefore, the lower predictive power indicates that local linear models cannot adequately capture the dynamics of the underlying data generating process.

For the Portuguese 10 year bond $n = 2,500$ series, we find a highly significant cumulative mutual information statistic, implying nonlinear serial dependence in the standardised GARCH residuals. This is consistent with the results of the BDS test. Interestingly, none of the statistics are significant for the Greek 10 year series. This means that the significant results of the BDS test for $n = 1,000$ may be due to some remaining linear structure not captured in the GARCH modelling framework, despite insignificant serial correlation LM tests. The result for the Greek series, coupled with significant results for some of the stock market indices, may indicate a difference in the power BDS test and the STAP surrogate test.

Finally, we turn our attention to the results of the pseudo-periodic surrogates in Table 4.6. The results are very different from those presented for the BDS test and the STAP surrogate data analysis. For the stock indices, the nonlinear prediction error is significant for the NIKKEI $n = 500$ series, and the algorithmic complexity statistic is significant for the NIKKEI $n = 2,500$ time series. For the bond markets, only the nonlinear prediction error is significant for the Ireland 10 year $n = 2,500$ series. The variation in the results may be due to the fact that the previous tests were conducted on model residuals, whereas the PPS surrogate data analysis was conducted directly on the returns series. As discussed in Section 2, this may be due to spurious determinism introduced due to the pre-whitening of the data. Small and Tse (2003) offer an alternate explanation: that the local linearity of the PPS series mean that data generated by mildly nonlinear systems may also lead to insignificant results. Kantz and Schreiber (2004) explain that a local linear model fit to a particular point in phase space represents a linear approximation of the global nonlinear dynamics at that point. The collection of all local linear models constitutes a nonlinear dynamic. Therefore, additional nonlinear dependence, above that explained by the GARCH modelling framework, may not be picked up by the PPS surrogate data framework if the strength of the nonlinear signal is weak compared to the large scale dynamics.

TABLE 4.5: STAP Surrogate p-values for the nonlinear prediction error (NLPE), algorithmic complexity (Comp), Pearson's cumulative mutual information (Pear) and cumulative mutual information (Mut). 1000 surrogates are used in testing. 10% significance is highlighted in yellow, 5% in orange and 1% in red.

N	Dow Jones				FTSE				NIKKEI			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.21	0.04	0.45	0.02	0.12	0.47	0.19	0.64	0.18	0.49	0.12	0.02
1000	0.16	0.97	0.45	0.44	0.22	0.88	0.11	0.96	0.38	0.64	0.25	0.80
2500	0.41	0.72	0.19	0.09	0.64	0.12	0.51	0.04	0.64	0.36	0.15	0.02
N	Greece				Ireland				Portugal			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.18	0.97	0.16	0.51	0.06	0.47	0.65	0.64	0.88	0.19	0.47	0.83
1000	0.75	0.65	0.31	0.79	0.18	0.86	0.13	0.48	0.87	0.32	0.12	0.31
2500	0.98	0.68	0.17	0.99	0.55	0.36	0.25	0.48	0.59	0.98	0.23	0.00

TABLE 4.6: Pseudo-Periodic Surrogate p-values for the nonlinear prediction error (NLPE), algorithmic complexity (Comp), Pearson's cumulative mutual information (Pear) and cumulative mutual information (Mut). 1000 surrogates are used in testing. 10% significance is highlighted in yellow, 5% in orange and 1% in red.

N	Dow Jones				FTSE				NIKKEI			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.30	0.81	0.96	0.83	0.33	0.70	0.51	0.52	0.05	0.65	0.48	0.91
1000	0.12	0.79	0.96	0.58	0.34	0.61	0.51	0.98	0.27	0.48	0.13	0.95
2500	0.24	0.92	0.91	0.42	0.39	0.48	0.41	0.30	0.23	0.07	0.64	0.82
N	Greece				Ireland				Portugal			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.60	0.33	0.73	0.66	0.64	0.81	0.88	0.52	0.84	0.84	0.93	0.87
1000	0.54	0.20	0.82	0.36	0.31	0.76	0.96	0.77	0.93	0.73	0.84	0.79
2500	0.25	0.16	0.87	0.54	0.05	0.68	0.98	0.21	0.33	0.68	0.97	0.28

4.5 Conclusion

A key finding of this chapter is the difficulty providing robust results when testing for nonlinearity in financial time series. This is due to the finite, noisy nature of financial data. That said, we provide some mixed evidence for (weak) nonlinearity, particularly in sovereign bond yield data in the lead-up to the Eurozone sovereign debt crisis. The BDS test and the STAP surrogate data test provide consistent evidence of nonlinear structure in the standardised GARCH residuals of the Irish 10 year bond returns for $n = 500$, and the Portuguese 10 year bond returns for $n = 2500$. The lack of significant results for the PPS surrogate data tests applied to the same samples may be due to the fact that the null hypothesis of the PPS surrogates includes mildly nonlinear data generating processes. Thus, if the nonlinear signal in the data is weak relative to level of noise, or if the data generating process is high dimensional deterministic, the PPS tests may not have sensitivity to detect the nonlinearity. The significant BDS results for the Greek 10 year bond returns is contradicted by both the STAP and PPS surrogate data analysis. This indicates that the BDS results may be due to remaining linear dependence in the standardised GARCH residuals.

The results for the stock indices are even more inconclusive. The BDS test finds no evidence of nonlinear dependence in stock indices. This is in line with the results presented in Moloney and Raghavendra (2011). However, the STAP surrogate testing indicates evidence of nonlinear serial dependence and nonlinear determinism in the standardised GARCH residuals of the Dow Jones index for $n = 500$, and nonlinear serial dependence in the residuals of the Dow Jones ($n = 2, 500$), FTSE ($n = 2, 500$) and NIKKEI ($n = 2, 500$). Similarly to the results for the bond samples, the PPS surrogate data tests applied to the stock market results are much less significant, with evidence of nonlinearity found only in the NIKKEI $n = 500$ and $n = 2, 500$ time series. We can only hypothesise that the mixed results are due to a weak signal to noise ratio, or high dimensional deterministic dynamics. The results indicate that the signal to noise ratio may be weaker for stock market indices than for sovereign bond markets, as evidenced by insignificant BDS test results for the indices.

The implication of these results are two-fold: firstly, there is large scope for further

research into nonlinearity in financial time series in the lead-up to financial crises, using more sophisticated filtering techniques to improve the signal to noise ratio without introducing spurious determinism, or destroying the signal. For example, Antoniou and Vorlow (2005) propose that wavelet methods may reduce the level noise and preserve delicate nonlinear deterministic dynamics in the data, which may be masked by the amplification of noise following log-differencing. Secondly, the fact that some mixed evidence exists for nonlinearity in stock and bond market data in the lead-up to crises opens up the possibility that such time series may be generated by nonlinear dynamical systems. It is well documented in the literature that both stochastic and deterministic nonlinear dynamical systems experience bifurcations. As a result, we believe that further research is warranted into the possibility that financial crises can be explained as bifurcations in nonlinear dynamical systems.

A key question in this regard is: if nonlinearity in financial time series is weak, is it ok to use linear models as a simple proxy for the deterministic structure? There are two potential perspectives. Firstly, linear models may be appropriate if we are looking to fit a model and test for a relationship between dependent and independent variables. However, a second perspective is if linear models are fit to the data, and the residuals are assumed to be IID, then the use of VaR or other stochastic risk indicators may provide inaccurate results, particularly during times of crisis. The second perspective is particularly important if the nonlinearity that is present in the financial series pushes the market towards a critical threshold, where instabilities in the system build-up over time and lead to an abrupt regime shift. The remainder of this thesis examines this perspective further by analysing sovereign bond and stock markets for evidence of a build-up of instabilities, and regime shifts, during financial crises.

Chapter 5

Critical Transitions in Financial Markets

5.1 Introduction

The financial history of the last century unequivocally points to the stylized fact that financial markets undergo large and abrupt transitions following a prolonged phase of stability, or bullish growth. As discussed in Chapter 1, the most common approach to financial market analysis is the linear stochastic framework, which is essentially an equilibrium approach and examines the issue of financial crises by looking at the deviations from the steady state (i.e. “normal times”) following a stochastic shock. The alternative is to take a complex systems approach, examine the possibility of multiple steady states in financial markets, and investigate market crashes using the concepts of bifurcations, or regime shifts, in nonlinear dynamical systems. Using this framework we can investigate the use of an emergent property of complex systems approaching regime shifts, namely critical slowing down, to develop early-warning signals of financial crises¹.

When a regime shift occurs in a dynamical system, the system switches abruptly between alternate steady states (equilibria) and displays qualitatively different behaviour to the original state. Once the regime shift has occurred, the system can linger in the alternate state for a long period of time. For instance, consider the sovereign bond yield to maturity for Greece, Ireland and Portugal, the so called

¹Refer to Chapter 1 for a discussion of emergent phenomena in complex systems

“peripheral countries”, before and during the Eurozone debt crisis (see Figure 5.1). It is clear that the bond yields go from a period of low stable yields to a period of high volatile yields. The stable yield regime, which begins soon after the introduction of the Euro in 2001, ends in late 2009 – early 2010 and is followed by sharp increases in sovereign yields, with increased volatility across all the three countries. After the transition the markets remain in the high yield state for a prolonged period of time.

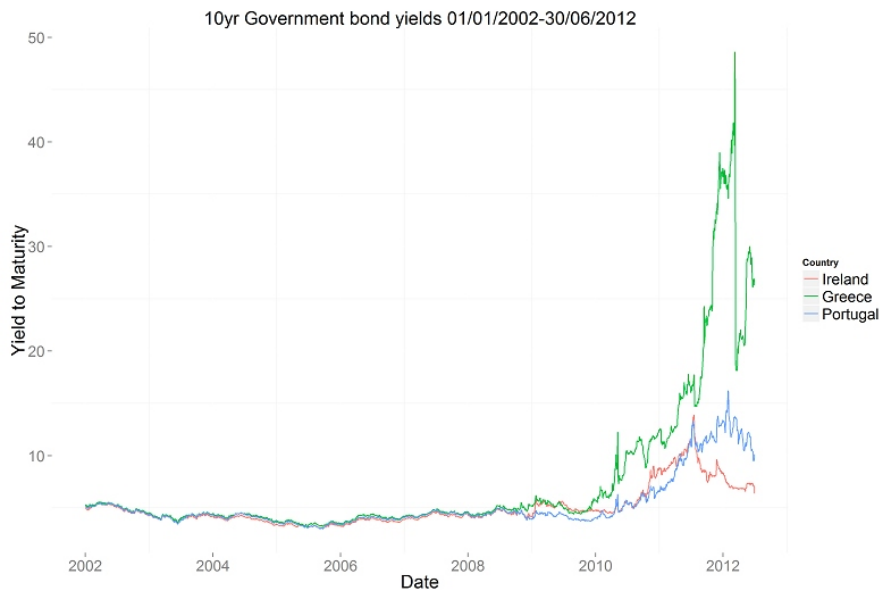


FIGURE 5.1: 10yr Sovereign Benchmark Bond Yield for Greece, Ireland and Portugal 01/01/2002-28/05/2012

Calvo (1988), in a seminal paper, proposed that multiple equilibria may exist in sovereign debt markets in which public debt is auctioned without an effort to manage market expectations or interest rates. The author postulates that the role of interest rates, and thus the burden of debt, is central to expectations of investors and the probability that sovereign debt will be (partially) repudiated. De Grauwe (2011) discussed how the Eurozone sovereign debt crisis may provide an example of multiple equilibria in financial markets. The author outlines that in a monetary union, such as the Eurozone, countries lose control over their own money supply. Therefore, in the event of a national sovereign debt crisis the central bank cannot step in to provide liquidity. Debt dynamics may become self-fulfilling, causing the economy to switch from a “good” equilibrium (solvency) to a “bad” equilibrium (default) when faced with a sudden stop of capital. This happens if liquidity suddenly dries up for the country’s sovereign bond market, causing an increase in

the interest rate at which the country can access debt markets. When this happens, what was originally a liquidity crisis can transform into a solvency crisis. In the absence of the central bank stepping in to provide liquidity, this can cause default or a bailout such as those seen for Greece, Ireland and Portugal during the Eurozone crisis. In this situation the investors' fears regarding the solvency of a country become self-fulfilling.

Models of self-fulfilling expectations gained popularity in the literature on speculative attacks on a currency within a fixed exchange rate system. Obstfeld (1996) develop a model which allows for the possibility of multiple equilibrium devaluation probabilities when faced with a speculative attack. In the Obstfeld (1996) model, the level of economic fundamentals determines the range of possible devaluation probabilities, or equilibria, in the system. Depending on fundamentals, and in turn level of reserves the government is willing to commit to defend the currency, there are multiple possible outcomes. The intermediate state of fundamentals allows for two possible equilibria, where a currency collapse is possible given a speculative attack, but not certain. This is similar to the "good" and "bad" equilibria discussed in De Grauwe (2011) for sovereign debt markets.

While the Obstfeld (1996) model is not directly framed in terms of bifurcations, as discussed in Chapter 3, bifurcations are the mechanism by which equilibria are created and destroyed in dynamical systems. Jeanne (1997) presents a second-generation currency crisis model, describing how currency crises can be a result of either self-fulfilling speculative attacks or bad economic fundamentals. Jeanne (1997) explicitly uses the framework of bifurcations in economic systems. The author describes a discontinuous expansion of the set of equilibrium devaluation probabilities occurring due to changes in economic fundamentals, relating the dynamical behaviour to the theory of bifurcations in nonlinear systems. In the model, market expectations become self-fulfilling when fundamentals are within a specific range, and when the expected net benefit of the fixed peg is negative.

A number of recent contributions to the literature on sovereign debt crises have developed models which follow the spirit of the second-generation currency crisis models. In such models, the government minimises a loss function where they will

continue debt repayments if the net benefit of doing so is positive. Gros (2012) for example, develops a theoretical model where multiple equilibria in sovereign bond markets are caused by changes in investors' required risk premia, which are a function of a country's perceived probability of default, the level of debt in the economy and the cost of raising taxes. The model specifies parameter values for which both single and multiple equilibria are possible. McHale (2012) outlines a model which incorporates conditional support from a lender of last resort, for example the European Central Bank in the case of Eurozone states, in which shifts in non-interest rate fundamentals can cause the probability of sovereign default to shift into a region where multiple possible equilibria can occur. In the model, multiple equilibria occur due to feedback loops between the expected probability of default and the interest rate, i.e. investors' expectations become self-fulfilling.

A key challenge in the analysis of economic, financial or other dynamical systems is the early detection of regime shifts between alternate equilibria. As discussed above, a number of theoretical models of multiple equilibria in financial markets have been developed. However, there is little empirical evidence that such models can effectively predict when the market will undergo a regime shift. Recent progress has been made in this regard in other fields of study. A large body of literature exists which provides direct empirical evidence that ecosystems and climate systems undergo catastrophic regime shifts, where the system abruptly transits from one steady state to another, referred to as "bi-stable" systems. It is documented that these bi-stable systems exhibit sudden structural change as the underlying parameters of the system evolve and reach a certain threshold. The gradual evolution of these systems prior to the transition masks the underlying changes that are happening, and makes it difficult to anticipate the impending catastrophic transition. In recent years, much effort has been expended in the attempt to find consistent, reliable indicators of regime shifts (Scheffer et al., 2001; Rietkerk et al., 2004; Lenton et al., 2008; Brock and Carpenter, 2006; Dakos et al., 2008; Scheffer et al., 2009; Hirota et al., 2011). These indicators, the so-called "early warning signals", are based on the statistical properties of the underlying dynamics that drive the evolution of the system.

The primary aim of this chapter is to examine sovereign bond yields of Greece,

Ireland and Portugal in the lead-up to the Eurozone sovereign debt crisis and ask whether such early warning signals were present in the case of bond market. In particular, we test for a phenomenon called critical slowing down (CSD) which has been shown to precede regime shifts in many real systems (Wissel, 1984; Dakos et al., 2008; Carpenter et al., 2011; Lenton et al., 2012). We ask the question for two main reasons, which are in turn the objectives of the chapter. Firstly, were early warning signals present before the crisis? Secondly, how reliable are these indicators and can they be used to provide consistent warnings of impending crisis in sovereign bond markets, or financial markets in general? Furthermore, our research aims to inform debate regarding the explicit use of bifurcation theory in the development of multiple equilibria models of sovereign bond markets.

The remainder of the chapter is structured as follows: In Section 5.2 we discuss theoretical framework that underlie the property of CSD in dynamical systems. In Section 5.3 we describe the empirical methodology used in the chapter. In Section 5.4 we examine the data used for empirical testing in more detail. In Section 5.5 we present the results of the empirical testing. Finally, in Section 5.6 we provide some concluding remarks.

5.2 Theoretical Framework: Bifurcation Theory

The theory of bifurcations has provided a taxonomy of various types of bifurcations. Chapter 3 of this thesis outlines a number of primary classes of bifurcations which occur in nonlinear dynamical systems. A particular class of bifurcation that can cause abrupt regime shifts in dynamical systems is called the fold catastrophe. In the literature, these are often referred to as critical transitions or catastrophic bifurcation because the system undergoes abrupt qualitative change as some underlying parameter passing through a critical threshold (Scheffer et al., 2009). In the following, we will briefly explain the fold catastrophe and illustrate it with a few examples.

The fold catastrophe can be described in simple terms. Consider a simple one-dimensional system $\dot{x} = h + rx - x^3$. Given $r > 0$, when $h = 0$ the system has two stable fixed points at $x = \pm\sqrt{r}$ and an unstable fixed point at the origin ($x = 0$).

The unstable fixed point acts as a divider that separates the basins of attraction² of the two fixed points. The trajectories starting from initial conditions to the left of the unstable fixed point at the origin, i.e. $x_{t=0} < 0$, go to the stable fixed point at $x = -\sqrt{r}$, while the trajectories starting from initial conditions to the right of the unstable fixed point at the origin i.e., $x_{t=0} > 0$ go to the stable fixed point at $x = +\sqrt{r}$. However, by varying the parameter h the curve in Figure 5.2 shifts, changing the location, stability and existence of fixed points. In particular, when h reaches a critical value, $|h_{crit}| = \frac{2r}{3}\sqrt{\frac{r}{3}}$, the unstable fixed point and the right-hand side stable fixed point converge and form a semi-stable fixed point. For $|h| < h_{crit}$ there is only one fixed point in the system. Essentially, a saddle node bifurcation occurs as the stable and unstable fixed point converge. As discussed in Chapter 3, the saddle node bifurcation is the primary mechanism by which equilibria are created and destroyed in a nonlinear dynamical system.




FIGURE 5.2: $\dot{x} = h + rx - x^3$ plotted for $r > 0$ and $h = 0$, $|h| = h_{crit}$ and $|h| > h_{crit}$. The black dots indicate a stable fixed point, the white dots an unstable fixed point and the half-black half-white, a semi stable fixed point.

The fold catastrophe can be understood by examining the stability of the equilibria in the system using basin of attraction diagrams. Imagine the basin of attraction as a well: if the walls are steep and we place a ball anywhere other than its deepest point, the ball will quickly roll down to the deepest point (or global minimum) of the well.

²The basin of attraction of a fixed point is the set of initial conditions that asymptotically lead to that point.

However, if the walls are relatively flat, the ball will take relatively more time to reach the minimum of the well. This is analogous to physical systems where we know that the potential energy of the well is instrumental in bringing the ball back to the minimum point in the well. The steeper the well, the higher the potential energy and the more likely that the ball will roll back at an exponential speed to the minimum. In the above example, the shape of the basin of attraction of the fixed points evolve as the h parameter is varied. When the stable and unstable fixed points converge, the slope of the line representing the rate of change of the system asymptotically approaches 0. This can be seen as a flattening of the basin of attraction in Figure 5.3. Once the critical threshold is passed, the right-hand stable and the unstable fixed point collapse and we are left with a single basin of attraction representing a system with a single equilibrium.

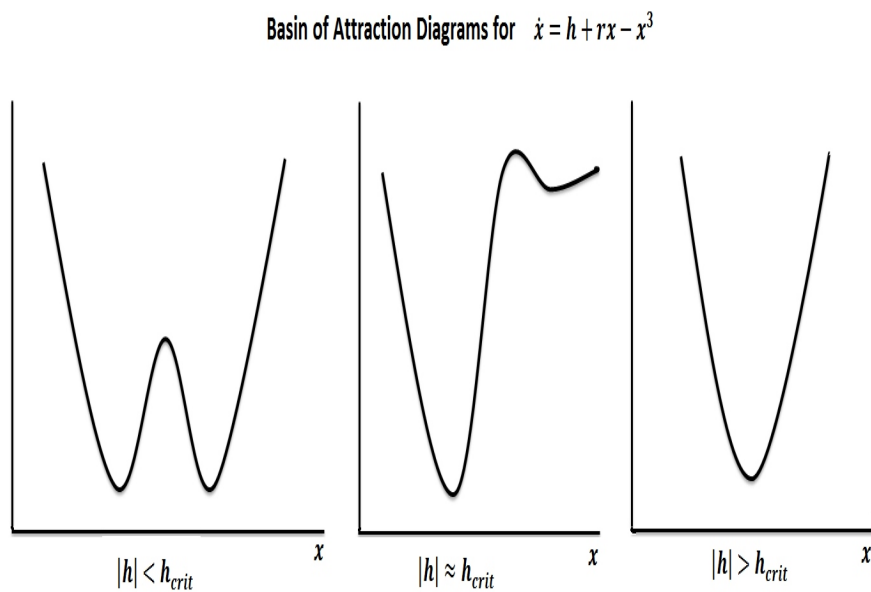


FIGURE 5.3: Basin of attraction diagrams for $\dot{x} = h + rx - x^3$ plotted for $r > 0$ and $h = 0$, $|h| = h_{crit}$ and $|h| > h_{crit}$. A local minimum indicates a stable fixed point and a local maximum an unstable fixed point.

In Figure 5.4, we have shown this dynamical behaviour using the bifurcation diagram. The y-axis represents the state of the system (x), and the values of the parameter (h) on the x-axis. The solid line represents the location of the stable fixed points, whereas the broken line represents unstable fixed points. If the system rests at fixed point A , after a small perturbation it will return to its stable fixed point. At this point there is only one stable equilibrium in the system. If the parameter increases

and enters the range of values between the vertical lines at B and C there is the possibility for multiple equilibria in the system. Now if the system is at point B , a small stochastic shock or a small increase in the parameter will cause the system to jump from the upper branch of the diagram to the lower. This large change in the state of the system is called a catastrophe and is due to the existence of multiple stable equilibria in the system.




FIGURE 5.4: Bifurcation diagram for fold catastrophe model. The graph represents the number and location of fixed points in the system with a solid line representing a stable fixed point and broken line representing an unstable fixed point. Between the arrow lines at points B and C there exists the possibility for multiple equilibria in the system.

It is difficult to predict when a fold catastrophe bifurcation will occur, as the transition is usually preceded by smooth or gradual change in the state of the system. A defining feature of such bifurcations is that, after the transition has occurred, reducing the parameter below point B is not enough to reverse the transition. It must be reduced past point C in order for the system to return to its original state. This is called hysteresis and could explain why financial markets appear to linger in the recessionary state following a financial market crash. These catastrophic bifurcations are often referred to as critical transitions in the literature, as the crash occurs once the control parameter passes the critical or tipping point. Henceforth in the chapter we shall refer to them simply as saddle node bifurcations or critical transitions.

5.2.1 Critical Slowing Down

The flattening of the basin of attraction as the critical threshold is approached is known in the literature as CSD. Wissel (1984) found that this phenomenon is a universal property of systems approaching a critical transition. CSD refers to how long it takes a system to return to equilibrium following a perturbation. For example, in our case we can think of the equilibrium price as the price at which the bond markets clear, i.e. demand equals supply at that price. Hence, we would say that the bond market is in equilibrium if, following a perturbation, the market returns back to its equilibrium level of price. However, a crucial question is, after the perturbation, how long will the system take to return to equilibrium? The return time very much depends on the shape of the basin of attraction, or the curvature of the basin of attraction.

CSD has been detected prior to regime shifts in real world dynamical systems which are believed to have passed through a critical transition (Wissel, 1984; Dakos et al., 2008; Carpenter et al., 2011; Lenton et al., 2012). CSD has also been detected in data generated from common ecological and climate models, where the system is made pass through a critical transition by varying a control parameter (Kleinen, Held, and Petschel-Held, 2003; Brock and Carpenter, 2006; Guttal and Jayaprakash, 2008; Dakos et al., 2012b). Recently the framework has been applied to study stock markets and real estate market crashes. Diks, Hommes, and Wang (2015) find evidence of CSD prior to the 1987 stock market crash, the dot.com crash and the Asian financial crisis. However, Guttal et al. (2016) find no evidence of CSD prior to stock market crises in major stock markets over the past century. Difference in the results may be due to differences in methodologies used to detrend the data. Guttal et al. (2016) apply the analysis to detrended logged returns, while Diks, Hommes, and Wang, 2015 analyse the log of the price series detrended using Gaussian kernel smoothing. Tan and Cheong (2014), also detect CSD in the U.S. housing market in the lead up to the sub-prime crisis, with weaker signals detected for the 1997–1998 Asian financial crisis and the 2000–2001 technology bubble.

Research has emerged which demonstrates that CSD can occur prior to other types of

transitions, which may not lead to abrupt change or hysteresis (Chisholm and Filotas, 2009; Kuehn, 2011; Kéfi et al., 2013; Boettiger and Hastings, 2012). This may lead to false positive results if indicators of CSD are used as early warnings of regime shifts. Moreover, other types of dynamical behaviour that generates rapid regime shifts will not be preceded by CSD, resulting in the possibility of false negative results.

Boettiger, Ross, and Hastings (2013) provide a classification system for different types of transitions. There are three broad, overlapping categories: 1) rapid regime shifts, 2) generated by bifurcations, and 3) preceded by CSD (See Figure 5.5 below). Saddle node bifurcations alone fall into the sub-category that undergo rapid regime shifts, are generated by bifurcations and are preceded by CSD. Using this classification system we use the convention that given a rapid regime shift has occurred, the presence of slowing prior to the shift is an indication that the system has undergone a critical transition. This is due to the fact that the other classes of transitions, prior to which CSD can be detected, do not generate abrupt regime shifts and hysteresis. Therefore, given an abrupt transition has occurred, and the system lingers in an alternate state for a prolonged period of time, it is reasonable to explore the possibility that the regime shift is generated by the system undergoing a critical transition.

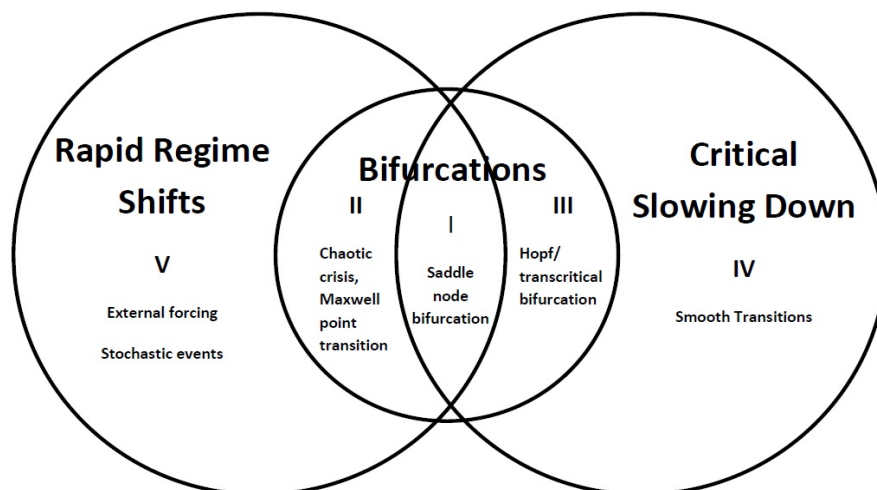


FIGURE 5.5: Classification system for different types of transitions. Adapted from Boettiger, Ross, and Hastings (2013). There are 3 overlapping categories: Rapid Regime Shifts, Bifurcations and Critical Slowing Down. The classification system shows that saddle node bifurcations are caused by bifurcations, resulting in rapid regime shifts and are preceded by critical slowing down. Examples of types of dynamical behaviour are provided for each of the 5 sub-categories.

5.3 Empirical Methodology

The recent ecological literature provides us with methods and tools to detect the phenomenon of CSD from the statistical properties of a system. We examine the statistical properties of the 10 year bond yield to maturity data for Greece, Ireland and Portugal prior to the abrupt rise in bond yields experienced in 2009-2010. In the critical transition literature, it is well documented that the phenomenon of CSD is related to the increase in the correlation of the system. This is due to the fact that, as the critical transition is approached, the rate at which the system recovers from shocks reduces steadily. This reduction in the recovery rate making successive observations of the state of the system closer in value to one another.

We follow the methodology of Dakos et al. (2012b), who provide an overview of a range of the available methods used to identify impending critical transitions. This is the main methodological guide for pre-processing the data, testing of the data using rolling windows, sensitivity analysis and significance testing. Dakos and Lahti (2013) provide a toolbox of code for the R statistical package, called the Early Warnings Toolbox. According to the authors, non-stationary time series cannot be used to test for impending critical transitions as this might result in spurious results. Therefore, we test for unit roots using the Augmented Dickey Fuller test. In all three cases we found that a unit root was present in the yield series, and we calculate the log-returns of the series. This removes the unit root and is the equivalent of the continuously compounded rate of return from one time period to the next.

In order to ensure there are no remaining trends in our log-return time series, we apply Gaussian kernel smoothing to further filter out any trends that may cause spurious results. Gaussian kernel smoothing means applying a moving average smoother, which uses a Gaussian weighting system to average observations (Shumway and Stoffer, 2006). The average series is calculated as follows,

$$\bar{f}_t = \sum_{i=1}^n w_t(i)x_i, \quad (5.1)$$

where \bar{f}_t is the Gaussian smoother at time t , x_i is the value of x at time i , $w_t(i)$ is the weighting given to observation i at time t , and is calculated as follows:

$$w_t(i) = K\left(\frac{t-i}{\lambda}\right) / \sum_{i=1}^n K\left(\frac{t-i}{\lambda}\right), \quad (5.2)$$

where

$$K(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \quad (5.3)$$

This is the Naradaya-Watson estimator, where K is the kernel function and λ is the bandwidth which determines the smoothness of the estimator. The smoothed time series is obtained by subtracting \bar{f}_t from x_t for all values of t . This should eliminate any remaining trends in the data. The challenge is to choose a value for the bandwidth which will remove trends but not filter out the dynamics of interest. As suggested in Dakos et al. (2012b), in our baseline analysis we use a bandwidth value of 10% of our data sets. A sensitivity analysis shows the results are robust to choice of bandwidth.

One of the most common indicators of CSD is autocorrelation (Dakos et al., 2008; Dakos et al., 2012b). Slowing down indicates that the recovery rate in a system decreases. This means that state variable over successive time periods will become more and more alike as the critical transition is approached. This will be picked up as an increase in autocorrelation. The recovery rate in the system should go smoothly to 0 as a critical transition is approached (Scheffer et al., 2009); therefore, we would expect the autocorrelation to approach 1 in a continuous manner. We use the first lag of the autocorrelation function (ACF1) over a rolling window to probe for increasing autocorrelation in a time series. This is given by the following formula

$$\rho_x = \frac{E[(X_t - \bar{X})(X_{t-1} - \bar{X})]}{\sigma_X^2}, \quad (5.4)$$

where \bar{X} is the mean value of X and σ_X^2 is the variance of X .

Another possible method to capture the change in the correlation structure prior to the critical transition is to fit a simple AR model of order 1 (AR1) to the data using a least squares methodology. The coefficient α is equivalent to ρ_X (Dakos et al., 2012b).

$$X_t = \alpha X_{t-1} + \epsilon_t, \quad (5.5)$$

where ϵ_t is an IID stochastic shock.

Brock and Carpenter (2006) and Carpenter and Brock (2006) propose rising variance as a leading indicator of impending critical transitions. Carpenter and Brock (2006) outline an inverse proportionality between the variance of the state a system and the rate at which it returns to a steady state following a small shock. This effect is referred to as squealing in the literature and has been found to be present prior to critical transitions in ecological and climate systems (Brock and Carpenter, 2006; Brock and Carpenter, 2010; Dakos et al., 2012b). Skewness measures the asymmetry of a distribution. As a critical transition is approached, shocks to the state variable of a system can push it towards the boundary between two alternative states. As a result, the critical transitions literature predicts that in certain cases skewness should increase prior to a critical transition (Dakos et al., 2012b). Similarly, as the rate of recovery decreases, one would expect more observations in the tails of the distribution, causing kurtosis to increase (Dakos et al., 2012b)³.

Therefore, we follow the results outlined in the literature and examine changes in the AR1 coefficient, ACF1 (Dakos et al., 2008), standard deviation (SD) (Brock and Carpenter, 2006), skewness (Guttal and Jayaprakash, 2008) and kurtosis (Biggs, Carpenter, and Brock, 2009). We can calculate the recovery rate of the system for any perturbation from the equilibrium point using the AR1 coefficient, which is simply the reciprocal of the AR1 coefficient. So, we expect the recovery rate to tend to zero as the system approaches the critical threshold. The existing literature states that these indicators should increase in a continuous manner as a critical transition is approached.

³The equations for variance, skewness and kurtosis are provided in Equations 1.6–1.8 of Chapter 1

It is necessary to quantify the trend in each of the indicators outlined above. We use Kendall's Tau correlation coefficient to test the strength of the trend in each of the indicators. To understand Kendall's Tau we must consider the pair (t, Y_t) , where t is the time period and Y_t is the value of a variable at time t . Let's say the number of time periods we have is 10 ($n = 10$). Starting at $t = 1$, a concordant pair is any pair where t is greater than 1 and Y_t is also greater than Y_1 . A discordant pair occurs if $t > 1$ but $Y_t < Y_1$. This procedure is carried out for all values of t up to $t = 10$ and the numbers of concordant and discordant pairs are totalled. Kendall's Tau is calculated as follows

$$KendallTau = \frac{(No.Concordant) - (No.Discordant)}{\left(\frac{1}{2n(n-1)}\right)} \quad (5.6)$$

If we have a continuously increasing indicator, i.e. as we increase t then Y_t also increases for all values of t , then Kendall's Tau will approach 1. We use surrogate data testing to determine the significance of the trends calculated on the indicators. The surrogates are created to have the same ARMA structure as the detrended yield to maturity data. This is done by selecting the ARMA(p,q) model with the lowest Akaike Information Criterion. 1000 different surrogate time series are created for each significance test by adding random shocks to the ARMA model. The indicator we wish to test for significance is calculated on each of the surrogates and the trend in the indicator is evaluated using Kendall's Tau. We then compare the trend statistics calculated on the original time series with the distribution of those calculated on the surrogates (Dakos et al., 2012b).

5.4 Data

We obtain 10/11 year sovereign bond yield to maturity daily data from Thomson Datastream for Greece, Ireland and Portugal (See Table 5.1). The sample bonds all have at least three years to maturity when the transition occurs. We follow the literature and choose our critical threshold based on a visual inspection of the data. We find that the Greek sovereign bond yields begin to rise at the end of October 2009. This corresponds broadly to announcements by the new Papandreou government

5.5. Results

Data and Descriptive Statistics			
	Greece	Ireland	Portugal
Date of issue	14/01/2003	28/01/2002	24/10/2003
Maturity date	20/05/2013	18/04/2013	16/06/2014
Sample	01/01/2007-30/10/2009	01/01/2007-31/03/2010	04/01/2007-31/03/2010
no. observations	740	848	848
Descriptive statistics of the bond yield sample (%)			
Mean	4.21	3.80	3.82
Median	4.29	3.94	3.98
SD	0.59	0.60	0.61
Skewness	-0.79	-0.63	-0.41
Kurtosis	3.33	2.54	2.05
Minimum	2.74	2.30	2.62
Maximum	5.63	4.93	4.97

TABLE 5.1: Data and descriptive statistics

that deficits were on track to more than double the previous government's forecasts. For both Ireland and Portugal the bond yields began to rise in Q2 2010. This timeline corresponds to the agreement for a Greek rescue package, causing uncertainty about the sustainability of the fiscal position of the peripheral Eurozone countries.

We analyse the trend in the indicators of CSD from the 01/01/2007–20/10/2009 for Greece and from the 01/01/2007–31/03/2010 for Ireland and Portugal. Choosing a starting date of 2007, with our baseline 50% rolling window size, means that we will be able to detect changes in the statistical properties of the data from Q2 2008 for Greece and from Q3 2008 for Ireland and Portugal, in all cases approximately 6 quarters from the critical threshold.

5.5 Results

Figures 5.6–5.8 display the results of the testing for CSD in the bond yield data. The first pane displays the detrended log-returns data. In the following panes we have the rolling window AR1, ACF1, SD, skewness and kurtosis. The results are for the baseline case of 50% rolling window size and 10% bandwidth for the Gaussian kernel smoother. The Kendall-tau value is also provided for each indicator. The Kendall-tau values, alongside their p-values, are summarised in Table 5.2. The p-values are generated using surrogate significance testing, as described above. Results are presented for the baseline specification and also based on a 10% bandwidth with 6

Significance of Indicators of Critical Slowing Down								
	6 month		1 year		50% Sample		75% Sample	
Greece	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
ACF1	0.626**	0.014	0.825***	0.004	0.943***	0.001	0.853**	0.011
AR1	0.625***	0.006	0.825***	0.002	0.943***	0.002	0.853**	0.024
SD	0.68***	0.006	0.901***	0.002	0.855***	0.007	0.977***	0.001
Skew	0.254	0.189	0.35	0.192	0.237	0.39	-0.177	0.564
Kurt	0.45*	0.074	0.361	0.278	0.249	0.389	-0.828	0.99
Ireland	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
ACF1	0.439**	0.046	0.682**	0.027	0.844***	0.008	-0.01	0.56
AR1	0.439*	0.052	0.683**	0.024	0.844***	0.001	-0.004	0.5
SD	0.465*	0.052	0.504	0.112	0.76**	0.038	0.973***	0.001
Skew	0.302	0.136	0.6**	0.05	0.54	0.157	-0.091	0.578
Kurt	0.344*	0.094	0.273	0.284	-0.244	0.646	-0.906	0.999
Portugal	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
ACF1	0.451**	0.038	0.666**	0.034	0.796**	0.026	0.188	0.436
AR1	0.449*	0.056	0.666**	0.026	0.795**	0.032	0.12	0.446
SD	0.629***	0.006	0.759***	0.01	0.855***	0.004	0.978***	0.002
Skew	0.516***	0.008	0.772***	0.01	0.661	0.114	0.218	0.366
Kurt	0.501**	0.028	0.289	0.298	-0.081	0.549	-0.229	0.644

TABLE 5.2: Kendall's tau coefficients and p-values for trend in indicators of critical transitions. Results provided for rolling windows of 6 month, 1 year, 50% of sample, 75% of sample. P-values obtained by calculating Kendall's tau coefficients from 1000 surrogate data sets with same ARMA structure as bond yield data. *=10%, **=5% and ***=1% significance. Baseline scenario of 50% rolling window size highlighted in bold font.

month rolling window, 1 year rolling window and rolling window equal to 75% of sample size.

Upon initial examination of Figures 5.6–5.8, it is clear that there is a very strong positive trend in the ACF1 coefficient, the AR1 coefficient and SD across our three sample countries. Apart from the clear visual trend, this is evidenced by high, positive Kendall's tau coefficients. Referring to the results of the significance testing for the baseline analysis (50% sample column) in Table 5.2, we can see that for Greece the trends in AR1, ACF1 and SD are all significant at a 1% level. For Ireland the trends in AR1 and ACF1 are significant at a 1% level, while the trend in SD is significant at a 5% level. For Portugal the trends in AR1 and ACF1 are significant at a 5% level, with 1% significance for the trend in SD. Trends for skewness and kurtosis are insignificant across all sample countries.

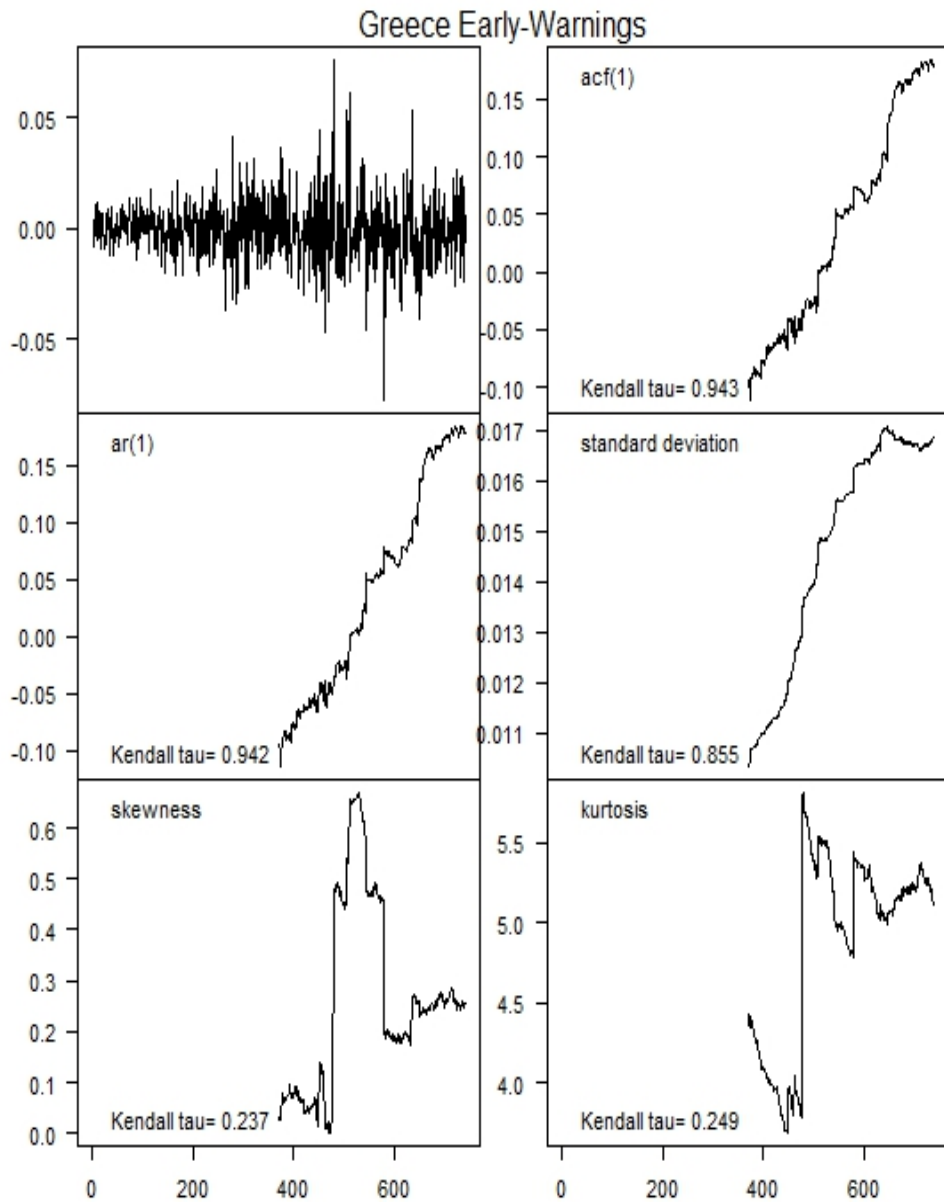


FIGURE 5.6: Indicators of critical transitions Greece 01/01/2007-30/10/2009. On the x-axis is the observation number and on the y-axis is the indicator value




FIGURE 5.7: Indicators of critical transitions Ireland 01/01/2007-31/03/2010. On the x-axis is the observation number and on the y-axis is the indicator value.

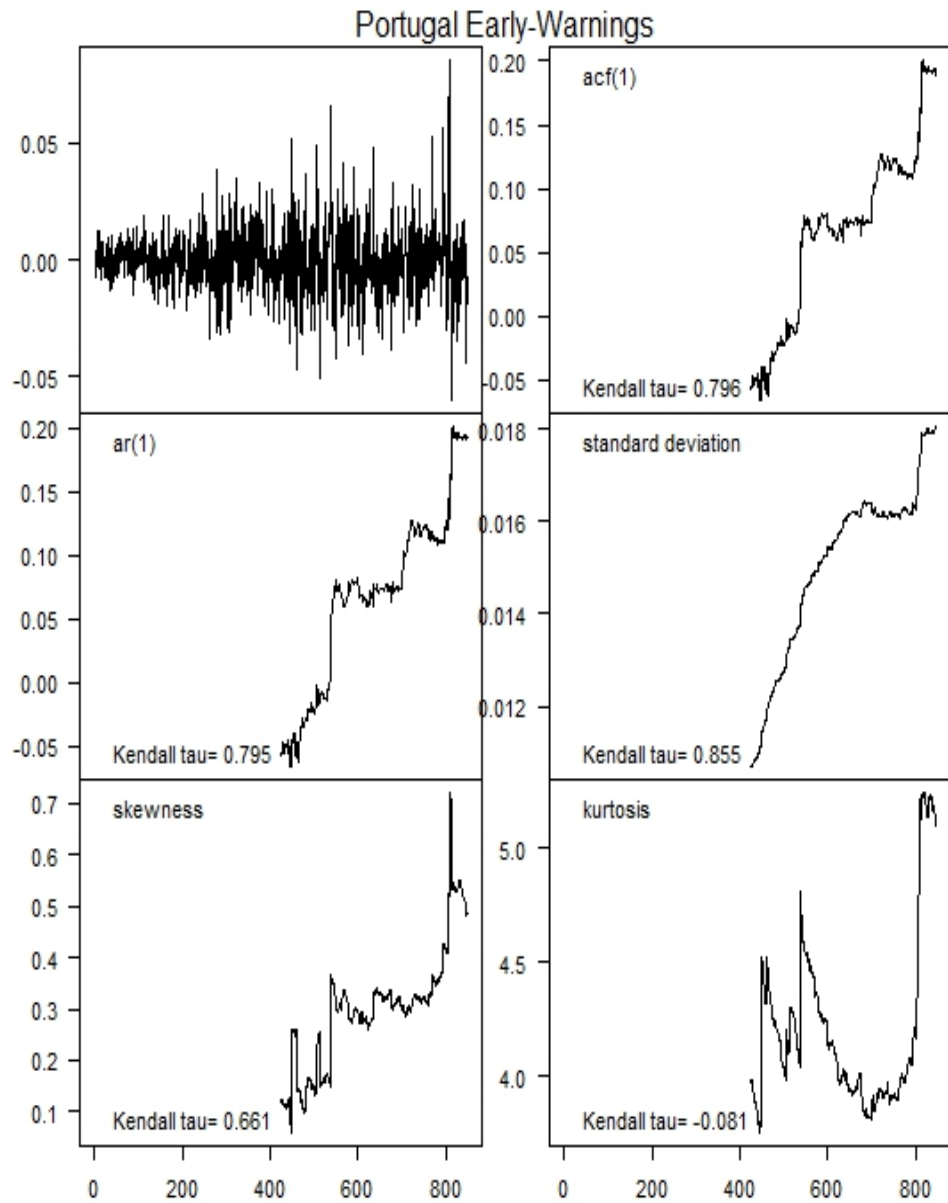


FIGURE 5.8: Indicators of critical transitions Portugal 01/01/2007-31/03/2010. On the x-axis is the observation number and on the y-axis is the indicator value

5.5.1 Sensitivity to Parameter Choice

Quax, Kandhai, and Sloot (2013) propose an alternate measure of detecting critical transitions in financial markets, using an information theoretic approach. Their new measure, information dissipation length, appears to predict the Lehmann Brothers crash 118 trading days in advance. When investigating autocorrelation and variance as indicators of CSD, the authors find that these indicators detect impending regime shifts in interest rate swap data prior to the Lehmann collapse only for certain maturities, and only for certain values of the rolling window used to calculate the indicator. These results point to a problem of choice of testing parameter values used to evaluate early warning signals in financial data. Therefore, we perform a sensitivity analysis to check the robustness of our results by varying the rolling window size and kernel smoother bandwidth.

Table 5.2 provides the results of a test for significance of the Kendall's tau coefficients calculated on the indicators for a 6 month rolling window, a 1 year rolling window, and a rolling window length equal to 75% of the sample size. The indicators for Greece are robust given any of the choices of rolling window size with ACF1, AR1 and SD displaying trends which are significant at the 5% level in all cases. For Ireland, the significance of these indicators decrease for the 6 month and 1 year rolling window lengths. However, with the exception of SD for a 1 year rolling window, all are still significant at the 10% level or higher. The trends become insignificant when a 75% rolling window is chosen. Given a 1 year rolling window, the trend in skewness becomes significant at a 5% level. Using a 6 month window, the trend in kurtosis is significant at a 10% level. Finally, examining the results of the Portuguese data, it is clear that the indicators are robust for rolling window sizes shorter than the baseline scenario of 50% of sample size. For the 6 month rolling window results, both skewness and kurtosis become significant. Similar to the Irish results, the significant trends disappear for the 75% rolling window size, with the notable exception of SD which remains significant at the 1% level. Figure 5.9 provides further results for sensitivity testing of ACF1 and SD indicators, carried out by varying the rolling window size and kernel smoother bandwidth. The results are similar to those presented above, with some sensitivity to choice of rolling window.

5.5. Results

The indicators are not sensitive to choice of bandwidth.

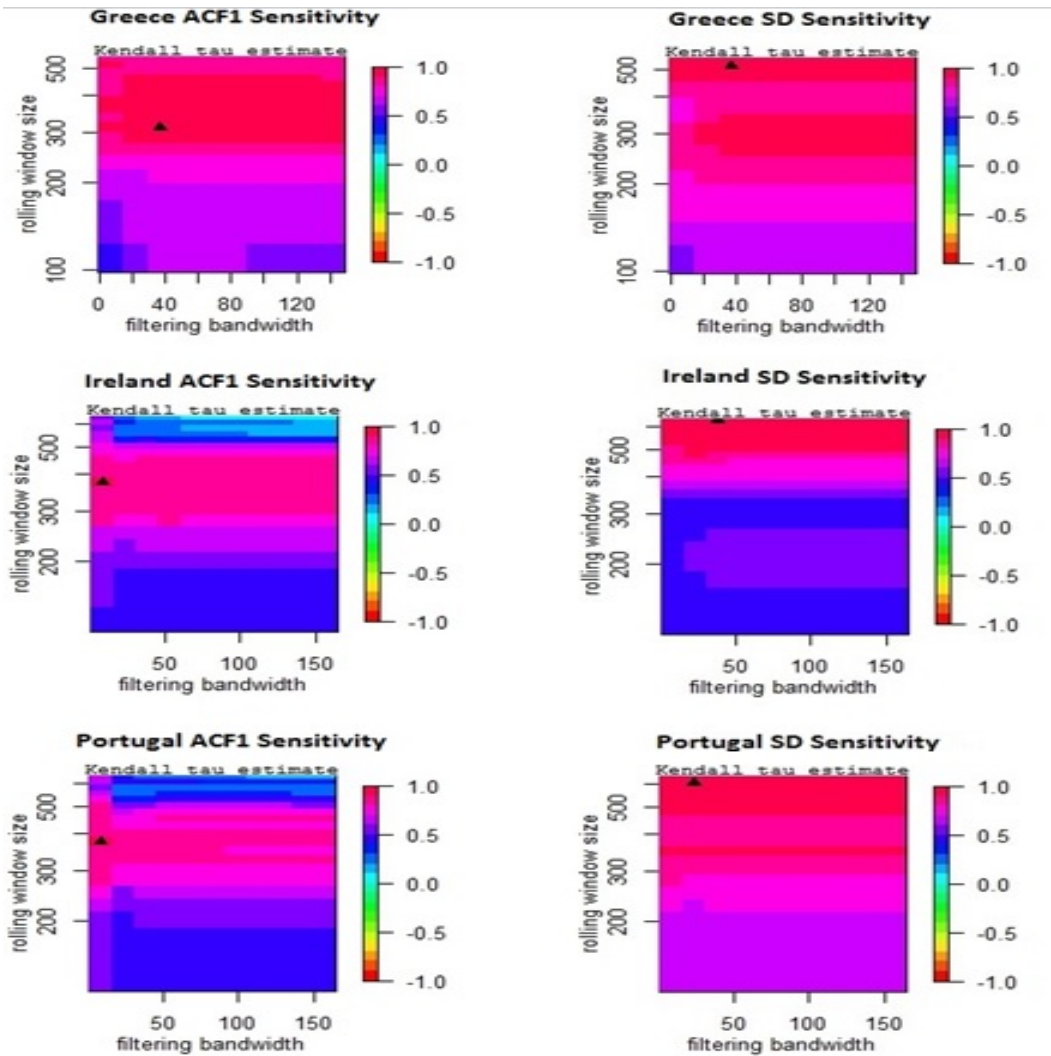


FIGURE 5.9: Sensitivity analysis carried out by varying rolling band width from 1% to 20% of sample size in increments of 15 observations and by varying the rolling window size from 25% to 75% of sample size in increments of 25 observations. The black triangle indicates the combination of rolling window size and bandwidth that gives the highest Kendall's tau coefficient

5.5.2 Comparative Analysis

As a further check on our results, we compare our findings with Kendall's tau statistics for ACF1 coefficients estimated in rolling windows for a number of Eurozone and international 10–11 year sovereign bonds. We use the ACF1 indicator because it has been found to be the most robust early warning signal (Dakos et al., 2012a). The chosen bond yield to maturity data are plotted in Figure 5.10. Up until late 2008 the Eurozone bond yields closely follow the same path. The yields diverge

after the onset of the global financial crisis in 2007–2008, with investors charging higher risk premia for Ireland and Greece in particular, but then follow a similar downwards trend up until the 31/03/2010. The notable exception is the Greek yield to maturity which begins to increase in late 2009. The black vertical line in Figure 5.10 marks the end of the period of analysis (30/10/2009) of the Greek bond yield data.

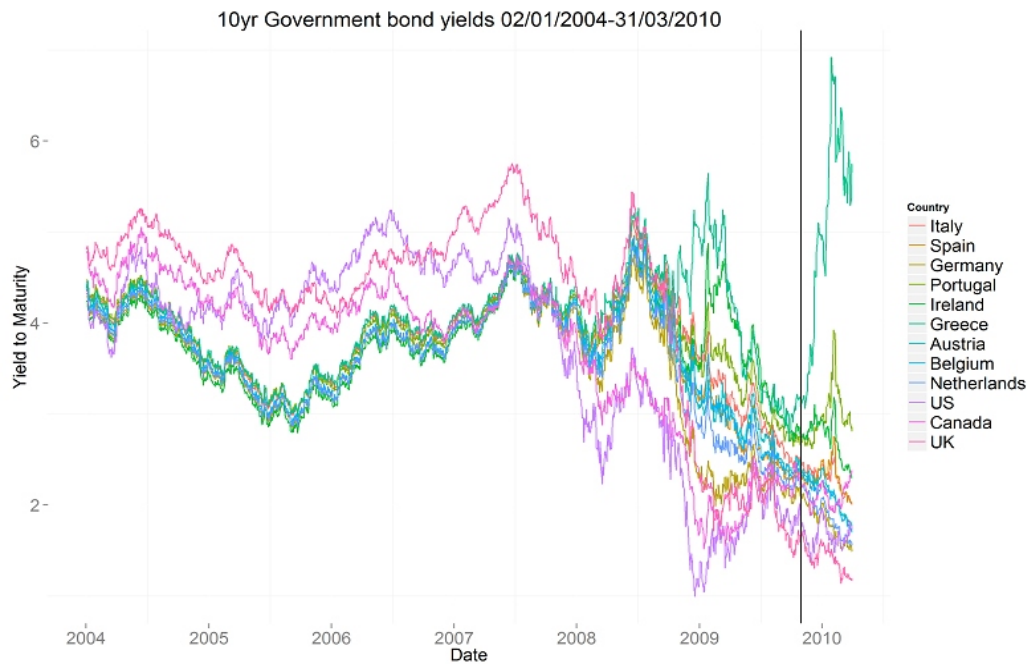


FIGURE 5.10: Sovereign 10-11yr bond yield data for Austria, Belgium, Canada, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, UK and US form 01/01/2004-31/10/2010.

Table 5.3 displays the results of significance tests completed for 6 Eurozone and 3 international bond yield time series. Using the same methodology and sample size described for Greece, Ireland and Portugal, we find that all the Eurozone countries have significant Kendall’s tau coefficient when calculated using a 50% rolling window size. The coefficients of Germany and Italy are significant at a 1% level, Spain and Belgium at 5% and Austria and Netherlands at 10%. With the exception of Germany, which displays significant trends across all choices of window size, these results are not robust to changes in rolling window size. The Spanish results are also significant at a 10% level using a 1 year rolling window. All other results become insignificant when the window size is varied. These results are interesting, as none of these countries entered a bailout programme or experienced the large abrupt increase in yields that occurred in the Greek, Irish or Portuguese markets. Therefore, we

5.5. Results

Significance of Trend in Autocorrelation Over Varying Rolling Window Sizes						
	Austria		Germany		Spain	
Window size	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
6 month	0.133	0.306	0.484**	0.038	0.239	0.198
1 yr	0.446	0.0178	0.788***	0.002	0.588*	0.078
50% sample	0.716*	0.062	0.858**	0.008	0.794**	0.024
75% Sample	-0.488	0.7774	0.792**	0.032	-0.007	0.5
	Belgium		Italy		UK	
Window size	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
6 month	0.145	0.31	0.224	0.248	-0.045	0.578
1 yr	0.452	0.16	0.503	0.114	0.0400	0.472
50% sample	0.737**	0.04	0.811***	0.001	-0.180	0.614
75% Sample	-0.488	0.772	-0.040	0.496	-0.523	0.772
	Canada		Netherlands		US	
Window size	Kendall's Tau	p-value	Kendall's Tau	p-value	Kendall's Tau	p-value
6 month	0.362*	0.094	0.219	0.232	-0.133	0.694
1 yr	0.572*	0.072	0.409	0.148	-0.043	0.5554
50% sample	0.409	0.23	0.666*	0.0727	0.186	0.428
75% Sample	-0.430	0.766	0.131	0.428	-0.035	0.514

TABLE 5.3: Significance tests on sovereign 10-11yr bond yield data for Austria, Belgium, Canada, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, UK and US. All bonds were issued in 2002-2003 and matured in 2013. Significance test carried out using surrogate data tests.

would not expect the presence of early warning signals in these markets. The UK and the US results show no significant trends in autocorrelation from 2007–2010 for any rolling window size. Canada is significant at a 10% level using 6 month and 1 year rolling window but this significance disappears for larger window sizes.

In order to gain a better understanding of the significant trends found in the above analysis, we plot the ACF1 coefficient for the 3 bailout countries and the 9 countries in our comparative sample (Figure 5.11). A 400 day rolling window is used to calculate the ACF1 from 02/01/2004–31/03/2010. This window choice is approximately equal to the 50% window used in the statistical testing above and, therefore, will help us understand the significant trends found across the Eurozone countries. The significance of the ACF1 coefficients is estimated at a 5% level using Bartlett (1946) standard errors. Periods during which ACF1 is significant are shaded grey. The ACF1 coefficients of Greece, Ireland and Portugal trend upwards from 2008 and breach the 5% significant level in early 2009. The coefficients continue to trend upwards until the end of the sample period (31/03/2010).

Out of the 9 countries in the comparative sample, only Spain and Germany display significant ACF1 coefficients at the end of the sample period. These significant ACF1

coefficients are coupled with significant Kendall's tau coefficients, as seen in 5.3. One potential explanation is that spillover effects from the Greek, Irish and Portuguese markets were beginning to be felt in the Spanish and German markets, with well documented fragilities in the Spanish market pushing it closer to a critical transition and flight to quality effects being detected in the German market. We find that the significant Kendall's tau coefficients obtained for Austria, Belgium, Canada, Italy and the Netherlands in Table 5.3 cannot be seen as positive early warning signals due to insignificant ACF1 coefficients. However, the results presented above provide us with extra confidence in the robustness of our results with respect to the detection of CSD for Greece, Ireland and Portugal.

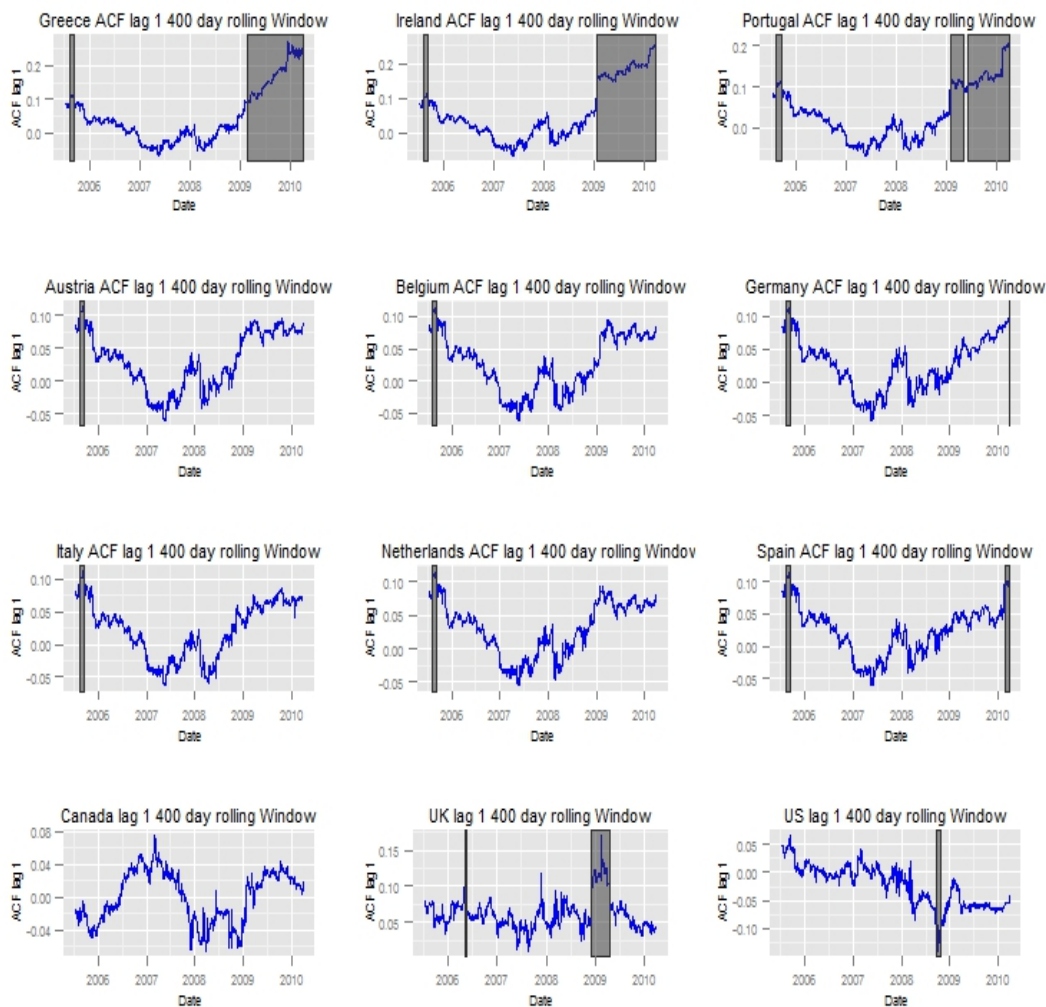


FIGURE 5.11: Rolling window ACF1 coefficient using 400 day window from 01/01/2004-31/03/2010 for Austria, Belgium, Canada, Germany, Greece, Ireland, Italy, Netherlands, Portugal, Spain, UK and US. Grey regions indicate periods where ACF1 coefficient is significantly different from 0 with 5% confidence intervals using Bartlett (1946) standard errors.

Positive Early Warning Signals				
	Start date	Stop date	Succeeded by transition	False Positive
Dax 30	19/05/1993	10/02/1994	N	Y
	23/12/1999	07/03/2000	Y	N
	02/01/2012	27/04/2012	N	Y
Dow Jones	21/01/1995	08/03/1996	N	Y
Eurostoxx 50	03/05/1994	21/02/1995	N	Y
	11/05/1999	26/10/1999	N	Y
	24/01/2012	17/04/2012	N	Y
FTSE 100	05/05/1998	03/12/1998	Y	Y
NIKKEI 225	N/A	N/A		
Toronto SE	23/10/1997	27/10/1997	N	Y
	11/05/1999	26/10/1999	N	Y
	24/01/2012	17/04/2012	N	Y

TABLE 5.4: Incidence and duration of positive early warning signals in daily stock price returns of Dax 30, Dow Jones, Eurostoxx 50, FTSE 100, NIKKEI 225 and Toronto SE.

5.5.3 False Positive Early Warning Signals

We further investigate rising autocorrelation as a leading indicator of critical transitions in financial markets, examining 6 international stock market indices over the period 01/01/1990 to 12/02/2014. A 400 day rolling window is used to calculate ACF1 on detrended stock index log-returns, returning a positive early warning signal when both the Kendall's tau statistic and ACF1 coefficient are significant at a 5% level. The results are plotted in Figures 5.12–5.17. The shaded section in the plot indicates that, during this time period, both the Kendall's tau statistic and ACF1 coefficient are significant at a 5% level. Thus, they meet our criteria for the occurrence of a positive early warning signal.

There are 10 occasions where positive signals are recorded across the 6 stock markets. When a number of thin shaded regions are clustered closely together this is due to a significant Kendall's tau coupled with an ACF1 that is flickering between significant and insignificant. We count this as a single (weak) positive signal. Table 5.4 summarises the occurrence of the positive signals over the sample period.

There are 2 occasions where a significant early warning is followed by a financial crisis. The first is for the DAX 30 stock exchange (Figure 5.12) in the period leading up to the bursting of the tech bubble. This crisis resulted in a bull market in the DAX which caused it to lose over 50% of its value between 2001 and 2003. During this

period the ACF1 coefficient becomes significant from the 23/12/1999 until the 07/03/2000, leading to a positive early warning signal for over 3 months before the crash. This financial crisis would meet the criteria for a critical transition due to the large abrupt drop in stock prices and the market remaining in a bull phase for a prolonged period of time.

The second early warning signal which precedes a crash is for the FTSE (Figure 5.14) in 1998. The trend in ACF1 became significant on the 05/05/1998 and remained significant as the FTSE lost 20% of its value from August to October 1998 during the Russian crisis. However, by November 1998 the FTSE had returned to its previous peak. This stock market crash does not appear to meet the criteria set out by Boettiger, Ross, and Hastings (2013) for a catastrophic bifurcation. While there is a large abrupt change in the stock index, and evidence of CSD, there is no evidence of hysteresis. None of the other positive signals are followed by regime shifts that meet the criteria for a critical transition. This leads us to conclude that using rising autocorrelation as a leading indicator of impending critical transitions for stock indices provides false positive signals.

The results of this section are broadly similar to those outlined in Guttal et al. (2016), who find either no trends or weak trends in the ACF1 prior to stock market crashes over the past century. The authors however, found evidence of strong trends in both variance and power spectrum prior to a number of stock market crashes. The authors suggest that stock market crashes may be better explained as a stochastic transitions. In such a transition, a regime shift occurs in a bi-stable system while the system is still far from a critical point, driven by increasing trends in stochastic perturbations. The results presented in this chapter imply that different modelling approaches are required to capture the dynamics of regime shifts in bond and stock markets.




FIGURE 5.12: Dax 30 index positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014.




FIGURE 5.13: Dow Jones index positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014.




FIGURE 5.14: FTSE 100 index positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014.




FIGURE 5.15: Eurostoxx 50 positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014

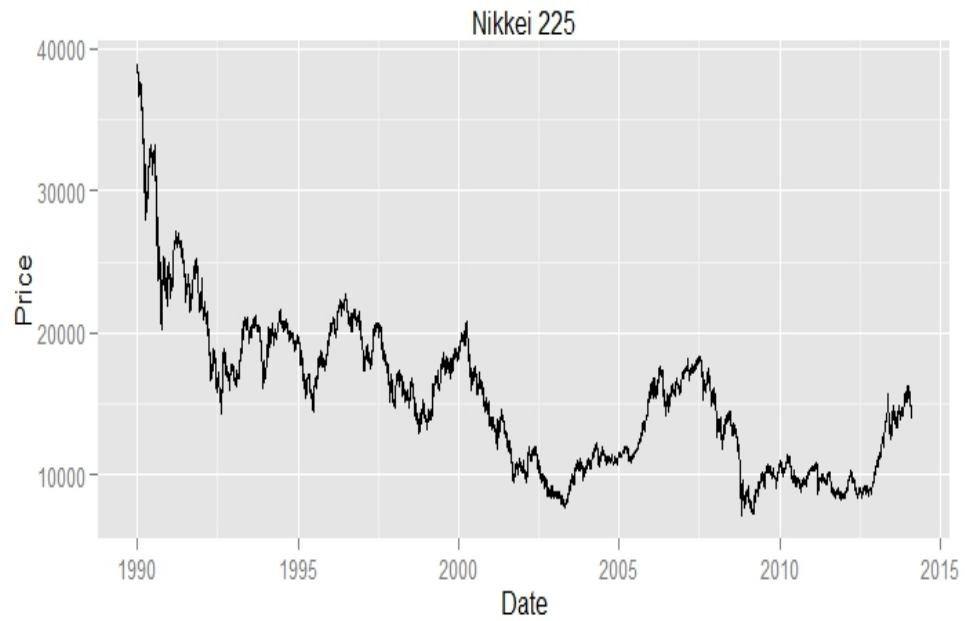


FIGURE 5.16: NIKKEI 225 index positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014.




FIGURE 5.17: Toronto Stock Exchange index positive trends in autocorrelation calculated on log-returns of international stock indices. Shaded area indicates period over which both ACF1 and Kendall's are significant at a 5% level. ACF1 calculated using a 400 day window and trend evaluated using Kendall's tau over 400 days from 01/01/1990-12/02/2014.

5.6 Conclusion

CSD has been used in recent ecological, climatology and financial literature as a potential early warning signal for impending regime shifts. The assumption is that abrupt change is generated endogenously by some underlying parameter reaching a critical threshold, causing the system to switch to an alternate equilibrium. Our results provide evidence that the sovereign debt markets of Greece, Ireland and Portugal experienced critical transitions during the Eurozone sovereign debt crisis. Statistically significant positive trends occur in indicators of CSD calculated from detrended bond yield returns. These trends begin up to 18 months prior to the abrupt increase in bond yields experienced by Greece in Q4 2009, and by Ireland and Portugal in Q2 2010. For instance, the Kendall tau coefficient for ACF1, AR1 and SD indicate strong positive trends that are robust to changes in rolling window size and Gaussian kernel smoother bandwidth. Further examination of the results reveal that in all cases ACF1 begins rising in 2008, becomes statistically significant in early 2009 and continues to rise up until the end of testing period. This indicates a positive early warning signal in the three markets. In all cases the positive signal was followed by abrupt change in the state of the market, with the market remaining in the high yield state for a long period of time.

We examine trends in ACF1 for 9 sovereign bond markets which did not experience abrupt regime shifts during the Eurozone sovereign debt crisis. The results reveal positive Kendall's tau statistics for 6 Eurozone countries and Canada. Out of these, only in the cases of Germany and Spain does the ACF1 coefficient become significant at a 5% level, indicating positive early warning signal of very short duration. This appears to be a false positive signal as neither market underwent a regime shift. However, the result may be due to spillover effects from the sovereign debt crisis. Further analysis of the use of ACF1 as a leading indicator of critical transitions in stock markets reveals a propensity for false positive signals, with significant ACF1 coefficients trending upwards for prolonged periods where no regime shift occurs. In fact, only in two cases is a positive signal followed by a market crash, and only in the case of the Dax 30 in March 2000 does the crash meet the criteria for a saddle node bifurcation as outlined in Boettiger, Ross, and Hastings (2013).

Our results have a number of implications: firstly, we have shown that CSD may provide a useful indicator of impending regime shifts in sovereign bond markets, and should be further investigated as a tool for regulators and financial market participants in the detection and monitoring of sovereign bond market crises.

Secondly, we find a very poor detection rate and prevalence for false positive results when estimating CSD as an early warning signal in stock indices. Chapter 6 of this thesis proposes a number of alternate, network theoretic, early warning signals for stock markets. However, in order to better understand the dynamics of regime shifts in stock markets, further research is required into the potential modelling of stock market crashes using the framework of stochastic transitions in bi-stable systems as proposed in Guttal et al. (2016). Our results suggest that the theory of bifurcations, and in particular catastrophe theory, may be used in the development of theoretical multiple equilibria models of sovereign debt crises. We propose future research into the development of such models. In particular, we suggest that future research in this area draw upon models and results from the ecological and other sciences, which have successfully simulated the dynamics of regime shifts in complex systems. Using this approach, one could model the behaviour of sovereign debt markets approaching a critical transition, capturing phenomena such as CSD.

Finally, we note that in both Chapter 4 and Chapter 5, our results are either weak or insignificant for stock market indices. One potential factor is that fluctuations in stock index returns are generally calculated as a weighted average of fluctuations in the underlying component stocks. This additional level of aggregation may destroy interesting dynamics in the underlying component time series. Therefore, in Chapter 6 we examine system dynamics at the level of component stocks of global indices, and examine co-movements of log-returns using minimum spanning tree analysis of correlation networks.

Chapter 6

Network Topology and Systemic Risk

6.1 Introduction

The question of how a crisis in a small sub-section of the U.S. financial system, the subprime mortgage market, set off a chain of events that led to a global financial crisis is still fresh in the minds of policymakers, financial regulators and academics. Ex post analyses of the crisis postulates that trends of global deregulation of markets and institutions, the growth of the shadow banking industry, and the ongoing phenomenon of financial innovation led to increasing fragility of the financial system as a whole (Bisias et al., 2012). This has prompted the realisation that the subprime crisis was not the driving force, but rather the trigger for the global financial crisis (Karmin and Pounder DeMarco, 2010). Substantial effort has since been expended on finding measures to detect and monitor the build-up of instabilities in the financial system, which can lead to a shock or stress event in financial markets becoming systemic¹. Figure 6.1 demonstrates the scale of disruption to global stock markets during the global financial crisis.

In Chapter 2 of this thesis we discussed that, in order to provide robust early warning signals of financial crises, we must move beyond traditional risk models which assume that financial data can be modelled as stochastic processes, depending solely on past observations of itself and other market data. These models fail to take into

¹For a review of systemic risk research and measures see Bisias et al. (2012).




FIGURE 6.1: Price Index (01/01/2003=100) for four major stock market indices - The S&P 500, FTSE 100, NIKKEI 225 and SBF 120

account the complexity of economic and financial systems, where the interaction of many market participants can lead to behaviour which can very different from the norm, and exceedingly difficult to predict. One method for the detection of instabilities in complex systems, which can capture changes in the interactions of market participants, is analysis of changes in network topology. Schweitzer et al. (2009), for example, argue that changes in network topology can push a system closer to a critical point. As discussed in Chapter 5 of this thesis, when a system reaches a critical point even a small shock can cause large, discontinuous change in the state of the system. However, our investigation of nonlinear dependence and CSD in stock market indices returned largely insignificant results. Therefore, in this chapter we apply our analysis to the individual component stocks of major global stock indices, investigating changing correlation network topology as an early warning signal for financial crises.

A body of recent literature has examined the correlation network of financial markets for evidence of changes in network topology in the lead-up to the global financial crisis, finding evidence of changing network topology, phase transitions, and the increasing influence of the financial sector in financial networks in the lead-up to the global financial crisis. In this chapter, we draw upon a number of previous results in the literature to assess potential early warning signals of systemic risk. For example,

Onnela et al. (2003a) use minimum spanning trees (MST) estimated for U.S. stock markets and find a topological shrinking of the MST during financial crises. Kaya (2015) and Wiliński et al. (2013) propose decreasing spread of the MST as a leading indicator of systemic risk, using evidence from global asset indices and stock market returns from the Frankfurt Stock Exchange respectively. Kennett et al. (2010) and Musmeci, Aste, and Di Matteo (2015) present results for the U.S. stock market related to the influence of the financial sector in the lead-up to the crisis, as well as results related to sectoral dynamics in general.

The primary objectives of this chapter, and our contributions to existing literature, are as follows: firstly, we perform a static MST analysis to determine whether a detectable phase transition occurred in the correlation network structure of the S&P 500 during the global financial crisis. We follow the methodology of Wiliński et al. (2013), who fit a power law distribution to the degree distribution of MSTs extracted from the Frankfurt Stock Exchange, finding the emergence of a super-connected hub in the lead-up to the crisis². We also assess the measures of MST centrality, spread and sectoral clustering, estimated in overlapping, rolling windows from the component stocks of the S&P 500, FTSE 100, NIKKEI 225 and SBF 120. The purpose of the analysis is to determine whether existing results, such as those found in Onnela et al. (2003a), Wiliński et al. (2013), Kaya (2015) and Musmeci, Aste, and Di Matteo (2015), related to changing correlation network topology can be used to obtain reliable indicators of systemic risk in stock markets.

The results of the static analysis on the S&P 500 reveals a clear MST power law degree distribution in the pre-crisis period. Following the onset of the crisis, a power law scaling region is detectable; however, the extrema of the degree distribution are not power law distributed. Finally, in the post-crisis period we find a power law scaling region with one highly connected node. The results show that there is significant variation in network topology over the period. This is confirmed by dynamic analysis of the mean occupation layer, a measure of the spread of the network, which shows significant variability but no clear phase transition prior to the crisis. We find that the most suitable measure of systemic risk for all four markets is the normalised tree

²Wiliński et al. (2013) also refer to the topology of the network containing the super-connected hub as a superstar structure.

length, which measures the overall interconnectedness of the network. Our results indicate that the normalised tree length provides the most reliable leading indicator for systemic risk, with a strong decrease across all four markets in the lead-up to the crisis. This decrease signals a shrinking of the MST and an increased likelihood for a shock to propagate through the system. For the FTSE 100, NIKKEI 225 and SBF 120, the decreasing trend begins in early to mid-2006, providing a clear early warning signal long before the onset of the crisis. Other measures, such as the average degree of the financial sector and the modularity of Industry Classification Benchmark (ICB) industry groupings within the MST, appear to provide early warning signals for the S&P 500 and the FTSE 100; however, further research is required into these measures to determine whether they are suitable leading indicators of systemic risk.

The remainder of the chapter is structured as follows: Section 6.2 introduces the concept of systemic risk and provides an overview of relevant literature. Section 6.3 presents the data and details the MST methodology. Section 6.4 presents the main results of the analysis, and Section 6.5 provides some concluding remarks.

6.2 Literature Review

6.2.1 Systemic Risk

Systemic Risk has been defined by three major policy institutions, the IMF, BIS, and FSB (2009), as the risk of

“the disruption to the flow of financial services that is (i) caused by an impairment of all parts of the financial system and (ii) has the potential to have serious negative consequences for the real economy.”

This definition says little about the mechanisms by which instabilities build-up in the financial system over time, leaving it susceptible to a systemic event. However, there is general consensus amongst policy makers and researchers that systemic risk has both a cross-sectional and a time dimension.³

³See for example Caruana (2010).

In the cross-sectional dimension risks are related to both common exposures, and to the complex network of transactions and balance sheet exposures between financial institutions. In the time dimension there is a pro-cyclicality to systemic risk (Caruana, 2010). During prolonged boom periods financial institutions lend more freely and increase leverage, investors become less risk averse, and often increased financial innovation and deregulation occurs (Bisias et al., 2012). Fragilities build-up within the system and a relatively minor stress event can turn into a systemic crisis, caused by positive feedback loops. These positive feedback loops are partially related to large numbers of market participants following similar trading or risk management strategies, with resultant large swings in asset prices and lack of liquidity (Danielsson, 2002). Zigrand (2014) provides a useful explanation of systemic risk as capturing the exogenous risks that may prevent the system from functioning properly, as well as the endogenous risk that is generated by the system itself. He states that endogenous risk is related to positive feedback loops and cascades.

The build-up of systemic risk through endogenous mechanisms related to positive feedback loops is not a new concept in finance. Minsky's financial instability hypothesis, (Minsky, 1977), views financial instability as a fundamental part of a functioning financial system. He proposes that when the financial system is robust, and the economy is a state of growth, optimism can lead to positive feedback loops between credit and asset prices. Capital gains and financial success from rising asset prices incentivises riskier lending thereby increasing the financial fragility of the system. A *Minsky moment* occurs when cash flows from assets become unable to meet debt repayments, forcing some investors to sell off assets to cover these repayments. Asset prices begin to fall and confidence dries up in the markets, causing a liquidity crisis as lenders are no longer willing to refinance debt. This leads to a vicious downwards asset price spiral and credit crunch with implications for the real economy.

6.2.2 The Correlation Network Approach

In this chapter we analyse a number of potential systemic risk indicators estimated on equity market correlation networks. Measures based on equity market returns are

suitable for detecting systemic risk as they are forward looking. Moreover, equity return correlations are used as an input for default correlation models, as stocks can be seen as a call option on the assets of a firm (Merton, 1973). Therefore, correlation of equity returns reflects the correlation between the value of firms' assets (Patro, Qi, and Sun, 2013). By definition, systemic risk measures should detect potential spillovers to the real economy. Therefore, equity market based measures are a logical choice due to strong links to the real economy through wealth effects (Modigliani, 1971) and financial accelerator effects (Bernanke, Gertler, and Gilchrist, 1994).

Previous studies have used correlation between equity returns as a measure of systemic risk. De Nicolo and Kwast (2002) examine equity returns' correlations between large and complex banking organisations from 1988 to 1999, finding a trend of increasing correlation over the period. They relate this to an increase in systemic risk potential, as increased correlation means exogenous shocks can better propagate through the system. Patro, Qi, and Sun (2013) use a similar methodology by examining the correlations between the equity returns of bank holding companies and investment banks from 1988 to 2008. Their study reveals that mean correlations increased significantly since 1996, indicating higher systemic risk potential. They also note a large spike in mean correlations occurred in the second half of 2007. While analysis of mean correlation coefficients can provide useful information regarding the interconnectedness of financial markets and systemic risk potential, MST analysis allows us to delve deeper into the network structure of markets and detect dynamical behaviour at a firm and sectoral level which may not be apparent at an aggregate, market level.

MSTs were first applied in the analysis of financial markets by Mantegna (1999). Mantegna showed that a static analysis of the stock market using MST could identify hierarchal structures in the market, with economically meaningful clustering of stocks. In more recent times, dynamic MST analysis has been used to investigate the changing nature of financial asset interdependencies over time (Onnela et al., 2003b; Coelho et al., 2007). MST analysis is particularly suited to the analysis of systemic risk as it filters the correlation matrix of asset returns, extracting the most important information, and identifies the shortest and most probable path for the transmission

of price shocks throughout the system (Lautier and Raynaud, 2013). Moreover, Pozzi, Di Matteo, and Aste (2013) contend that MSTs are effective at distinguishing the highly connected, important and influential vertices at the centre of a network.

A number of recent papers use MST analysis to examine changes in network topology during the global financial crisis. Kaya (2015), for example, use a mutual information based MST to assess the synchronisation of different asset classes. They examine the average eccentricity of minimum spanning trees estimated from networks of stock, bond and commodity asset classes, and networks based on regional/sectoral equity indices. Eccentricity is a measure of the centrality and spread of the nodes in the network, and the authors find that asset sell-offs are more prevalent following a decrease in network eccentricity. Clear indications of decreasing eccentricity are present prior to the crisis period.

Wiliński et al. (2013) find evidence of two separate phase transitions in the Frankfurt Stock Exchange during the global financial crisis. The first phase transition is observed as the crisis unfolds, with the network degree distribution switching from a scale free distribution to a superstar structure, where one dominant hub appears at the centre of the network. The second transition occurs after the crisis period. In this period, the network is characterised by a scale free degree distribution, with a number of smaller star-like hubs. The authors state that the emergence of the superstar structure is a precursor to the 2008 crash. They claim that the mean occupation layer, a measure of network centrality and spread which is similar in nature to network eccentricity, is more sensitive than the normalised tree length in the detection of topological phase transitions.

The MST is not the only or even the most efficient method for filtering a correlation network (Kaya, 2015). Kennett et al. (2010) analyse correlations between stock prices in the NYSE over the period 2001 to 2003 using Partial Correlation Threshold Networks (PCTNs) and Partial Correlation Planar maximally filtered Planar Graphs (PCPGs). Partial correlation networks correct for indirect effects when calculating the correlation between two stocks, and can simplify the description of the system. Both static and dynamic analysis of the partial correlation networks reveal a dominant financial sector in terms of centrality in the network, particularly the investment

services sub-sector. The authors find that the significant influence of the financial sector is consistent over the entire sample period.

Musmeci, Aste, and Di Matteo (2015) analyse the correlation structure of the S&P 500 using Planar Maximally Filtered Graphs (PMFG) and their associated clustering structure, the Directed Bubble Hierarchical Tree (DBHT). The PMFG algorithm is less restrictive than the MST methodology, filtering out less information. The MST algorithm starts with a random node and adds the most highly connected nodes such that no loops can occur. The PMFG algorithm adds nodes in a similar manner but must only ensure the planarity of the resulting network. Thus, the MST network is a subgraph of the PMFG network. The authors find that the similarity between DBHT clusters and the Industry Classification Benchmark (ICB), while initially high, drops off from the crisis onwards. The authors conclude that ICB may be becoming a less informative benchmark to diversify risk, and that persistence in the DBHT clusters could be a potential indicator of systemic risk. The authors also find that the cluster containing financial sector stocks becomes larger and contains a larger number of industries from 2007. The authors relate the increase in size of the financial sector cluster to the role of the financial sector in the propagating the subprime crisis to other sectors of the economy.

Pozzi et al. (2008) assess the stability of edges in PMFG and MST methodologies, and their ability to reproduce the correlation dynamics of the top 300 stocks of the NYSE over the period 2001–2003. The authors find that the PMFG performs slightly better in terms of persistence of edges, and incorporates a larger number of economically meaningful connections. However, there is a cost in terms of visualisation, as the PMFG network has triple the number of edges as that of the MST. The authors also find that both algorithms perform well in terms of identifying sectoral clusters.

6.3 Data

We analyse the correlation network topology of the components of four major stock market indices over the period 2003–2010 inclusive. Our sample includes the components of the S&P 500, FTSE 100, NIKKEI 225 and the SBF 120. The sample

period allows us to detect evidence of fragility building up in the financial system in the period before the subprime mortgage crisis, as well as to analyse developments in the pre-crisis, crisis and post-crisis periods. We obtain a list of tickers, market capitalisation, and ICB industry and sub-sector information for the each of the stock indices from Thomson Reuters Datastream as of 31/12/2010. Stocks which don't have data for the entire period are excluded. Data is not available for stocks which have delisted between the end of the sample and the date the data was downloaded. Hence, there is a survivorship bias in the sample which is an unavoidable caveat of the research.

Data is available for the full sample period for 399 of the S&P 500 stocks, 88 of the FTSE 100 stocks, 205 of the NIKKEI 225 stocks and 97 of the SBF 120 stocks. Table 6.1 provides a breakdown of the data by ICB industry classification. The largest two industries by number of stocks in the FTSE 100 as at 31/12/2010 are financials and consumer services, with 23.8% and 17% respectively. Industrials and consumer goods comprise the largest portion of the NIKKEI 225, with 31.2% and 20% of stocks in the sample respectively. The largest two industries in the SBF 120 sample are industrials and consumer services with 18.5% and 16.4% of stocks respectively. Finally, the largest two industries in the S&P 500 are financials and industrials with 17.7% and 16.2% respectively.

TABLE 6.1: Breakdown of FTSE 100, NIKKEI 225, SBF 120 and S&P 500 samples by % of stocks and % of Market Capitalisation in ICB industry group

Industry	FTSE 100		NIKKEI 225		SBF 120		S&P 500	
	% Market Cap	% Stocks	% Market Cap	% Stocks	% Market Cap	% Stocks	% Market Cap	% Stocks
Basic Materials	7.8	9.0	6.4	12.6	4.8	6.1	3.3	5.0
Consumer Goods	21.5	10.2	26.0	20.0	25.6	14.4	10.1	11.7
Consumer Services	8.5	17.0	9.4	11.2	8.6	16.4	12.3	14.5
Financials	22.6	23.8	13.8	8.7	18.0	15.4	16.1	17.7
Health Care	10.6	4.5	6.0	4.8	8.9	6.1	10.3	8.7
Industrials	4.9	14.7	19.6	31.2	17.2	18.5	12.7	16.2
Oil & Gas	11.6	6.8	0.1	0.4	7.9	5.1	12.4	7.7
Technology	2.2	4.5	5.4	6.3	5.2	12.3	16.3	10.5
Telecommunications	4.6	2.2	11.5	1.9	2.3	1.0	2.8	1.0
Utilities	5.2	6.8	1.5	2.4	0.9	4.1	3.2	6.5

6.4 Methodology

6.4.1 Minimum Spanning Trees

In this chapter we assess the MST methodology as a candidate for detecting and monitoring systemic risk in financial markets. As discussed in Section 5.3, other methods such as PFMG are marginally more effective at detecting economically meaningful connections. However, we choose the MST methodology due to its computational efficiency and robustness for large numbers of assets (Kaya, 2015). In order to extract the MST from the fully connected correlation matrix we must calculate a measure of distance between two stocks based on the correlation coefficient. As discussed above, we first calculate the logarithmic return of the equity prices in our sample.

A number of papers propose that financial returns be detrended prior to filtering of the correlation network. Wiliński et al. (2013), for example, state that if correlation with the overall market is not removed, this can lead to spurious correlations between stocks related to bubble dynamics. Borghesi, Marsili, and Micciché (2007) find that subtraction of a global component, for example the average market return, from stock return series improves the ability of MSTs to detect industry clustering, particularly for intraday returns. However, although there is an improvement in detection of industry clustering across all sample sizes, the authors find that for samples where the number of stocks is less than 500, the degree distribution of the MST is not statistically different from that calculated from random walks. The results of the study show that the effect of noise is significantly reduced for the non-detrended stock returns. Moreover, the MST for non-detrended returns with a daily frequency performs quite well both in terms of detecting economically meaningful clusters and in terms of the persistence of intra-industry connections over time.

We choose not to remove the market (global) factor from our log-returns sample for two primary reasons: firstly, the analysis is being performed on four samples of size 88 (FTSE 100), 205 (NIKKEI 225), 97 (SBF 120) and 399 (S&P 500). Therefore, as shown in Borghesi, Marsili, and Micciché (2007), for the smaller sample sizes the effect of

noise may over power the topological dynamics if the global component is removed. Secondly, as the focus of this chapter is on systemic risk analysis, the removal of dynamics related to overall market movements may destroy interesting trends in the network metrics related to bubble dynamics. Furthermore, Borghesi, Marsili, and Miccichè (2007) show only a marginal improvement in persistence and clustering by removal of the global factor when daily data is used, as in this chapter.

Next the Pearson's correlation coefficient is calculated between the logarithmic returns of stock X_i and stock X_j as

$$\rho_{ij} = \frac{\langle r_i r_j \rangle - \langle r_i \rangle \langle r_j \rangle}{\sqrt{(\langle r_i^2 \rangle - \langle r_i \rangle^2) - (\langle r_j^2 \rangle - \langle r_j \rangle^2)}}, \quad (6.1)$$

where $\langle \dots \rangle$ denotes the expectation or time operator and r_i is the log returns series for stock i .

Given some distance d_{ij} between stocks $i = 1, \dots, N$ and $j = 1, \dots, N$ the minimum spanning tree is a network of nodes which are all connected by at least one edge. The edge weights are given by the distance between the vertices, and the sum of the edges is the minimum possible to connect all the vertices without any loops (Sandoval, 2012). Mantegna and Stanley (1999) outline a distance metric between two stocks

$$d_{ij} = \sqrt{2(1 - \rho_{ij})} \quad (6.2)$$

The distance measure proposed by Mantegna and Stanley (2000) treats negative correlations as large distances between stocks. However, negative correlations play a large role in the dispersion of price shocks throughout the system. Therefore, recent applications of MST analysis to correlation networks of stock returns use $1 - \text{abs}(\rho_{ij})$ as a distance metric. This metric treats positive and negative correlations as equally important. However, as the returns of the individual stock pairs in our sample are predominantly positively correlated, the aggregate trends detected from both distance metrics are virtually identical. We apply the $1 - \text{abs}(\rho_{ij})$ distance measure in this chapter.




FIGURE 6.2: Relationship of correlation coefficient to the distance measure over the range -1 to 1

The Prim (1957) algorithm is applied to extract the MST from the network of $\frac{(N-1)N}{2}$ distance measures. Prim's algorithm provides an iterative, computationally efficient method for building a MST. This method is chosen over the Kruskal (1956) algorithm as it is more computationally efficient for fully connected graphs (Onnela, 2002). G denotes the fully connected graph with N vertices and edge weight between vertex i and vertex j , d_{ij} . S denotes the minimum spanning tree, which is a sub-graph of G . We begin building the tree by selecting an arbitrary vertex, V_1 , which can be called the root of the MST. We add V_1 to S and then follow a 2 step iterative process.

1. Find the edge with the shortest distance between two vertices such that one vertex is in S and the other is in G . Add this edge and vertex to S .
2. If all vertices are contained in S then S is the minimum spanning tree of G . If not, repeat step 1.

The results of the algorithm applied to a network containing loops is displayed in Figure 6.3.

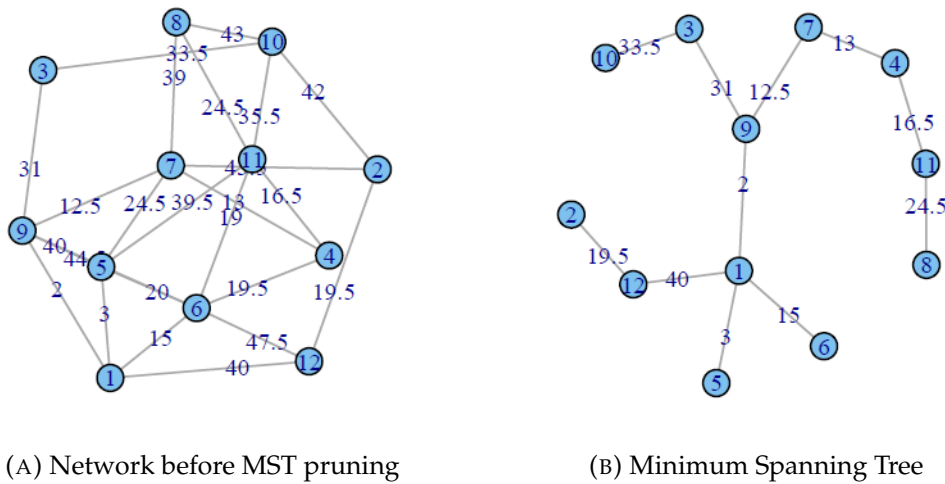


FIGURE 6.3: Extraction of the minimum spanning tree

6.4.2 Network Metrics

Onnela et al. (2003b) found that, due to increased correlation during financial crises, minimum spanning trees shrink topologically. As the tree shrinks the shortest path by which a shock can propagate through the system decreases. This shrinking increases the risk of a systemic event. Therefore, we analyse the evolution of the minimum spanning tree by calculating the normalised tree. For a MST with N vertices and $N - 1$ edges the normalised tree length L is defined by Onnela et al. (2003b) as

$$L = \frac{1}{N - 1} \sum_{d_{ij} \in S} d_{ij} \quad (6.3)$$

A decrease in L over time indicates that the MST is shrinking and, on average, the equity returns are becoming more interconnected. An increase in correlation leads to a decrease in the distance d_{ij} between stock i and stock j . The normalised tree length sums the distances (or edge weights) in the MST for all i and j , normalising by the total number of edges in the tree. Another statistic outlined in Onnela et al. (2003b), which quantifies the spread of the network, is the mean occupation layer. This statistic calculates the average number of steps between the nodes in the network and the central node. Using a similar method, asset eccentricity, Kaya (2015) found evidence of a decrease in the spread of MSTs calculated on global financial indices

prior to crises. The mean occupation layer is defined as

$$Occ(V_t) = \frac{1}{N} \sum_{V_i \in S} Lev(V_{i,t}), \quad (6.4)$$

where S is the MST, Occ is mean occupation layer, $Lev(V_{i,t})$ is the level of node V_i with respect to the central node and measures the shortest path from node $V_{i,t}$ to the central node. The central node is defined as the node with the highest degree, weighted by its edge weights.

One of the main advantages of the MST is that it can detect clustering of highly correlated stocks. In particular, for equity markets the MST analysis can detect sectoral clustering using industry classification benchmarks (Onnela, Kaski, and Kertész, 2004; Coelho et al., 2007; Pozzi, Di Matteo, and Aste, 2013). By plotting statistics which quantify the level of sectoral clustering in rolling windows, changes in network structure can be detected. In order to quantify the level of sectoral clustering we calculate the modularity of the ICB industry and sub-sector groupings over time. Modularity quantifies both the similarity of a particular community of nodes in a network and dissimilarity with other communities which are present in the network. Modularity (Q) is calculated as

$$Q = \frac{1}{2m} \sum_{i,j} (A_{ij} - \frac{Deg_i * Deg_j}{2m}) \delta(c_i, c_j), \quad (6.5)$$

where m is the total number of edges in the network, A_{ij} is the element of adjacency matrix A for node i and node j , Deg_n is the degree of node n in the network and $\delta(c_i, c_j)$ takes a value of 1 if both node i and node j belong to the same community and 0 otherwise. An increase (decrease) in modularity in the MST indicates increased (decreased) partitioning along industry groupings.

An increasingly dominant position in the centre of the tree can indicate that a sector is exerting large influence on the market, and may be a key propagator of price shocks throughout the system. We calculate the average financial sectoral degree to detect the relative influence of the financial sector over the sample period. The average

degree for a sector (\overline{Deg}_z) is a measure of the centrality of a particular sector (or industry) with respect to the rest of the network and can be described by the equation

$$\overline{Deg}_z = \frac{1}{N_{i,z}} \sum_{i=1}^{N_z} Deg_{i,z}, \quad (6.6)$$

where $N_{i,z}$ is the number of stocks in industry z and $Deg_{i,z}$ is the degree of stock i in industry z . An increase in \overline{Deg}_z means that sector z is taking a more central role in the network.

As there are a fixed number of vertices and edges in a minimum spanning tree for a fully connected correlation matrix, an increase in the average degree for one sector necessitates a decrease for another sector. Therefore, the average sectoral degree is a relative measure of centrality in the network.

6.4.3 Time Parameter

MSTs are obtained from correlation matrices and the normalised tree length, mean occupation layer, average sectoral degree and modularity are calculated in rolling windows in order to visualise the evolution of stock return co-movements over the period 2003–2010. The sample is divided into rolling windows of width T . Both T and the step size, δT , act as smoothing parameters, with a larger value of either parameters giving a smoother graph for the mean of the correlation matrix plotted over time. However, once δT is relatively small the smoothing effects of this parameter are negligible (Onnela et al., 2003a). Therefore, we choose a step size of 1 trading day to retain the maximum possible information.

The choice of T takes a little more effort. Too large a window width and changes in the network metrics may be over-smoothed, destroying patterns in network topology which could indicate an increase in systemic risk. Too small a window width could lead to an unstable MST, results dominated by measurement noise, or over-sensitivity of results to price shocks.

Previous studies have used a range of time parameters in order to study the evolution of correlation based financial networks. Onnela et al. (2003b) analyse daily

stock returns from the NYSE using a 1000-day rolling window. However, using a *Single Step Survival Ratio*, they find that the minimum spanning tree is relatively stable for rolling windows as low as 250 trading days. Coelho et al. (2007) calculate the normalised tree length for the FTSE 100 using a 500-day window. The authors recognise the trade-off between stability of the network measure and over-smoothing. Lautier and Raynaud (2013) examine the correlation dynamics of the energy derivatives market using a 480-day rolling window, Lee et al. (2013) examine CDS markets with a 300-day rolling window, and Jang, Lee, and Chang (2011) use a 250-day window for currency markets.

In order to be consistent with our analysis in Chapter 5, we choose a rolling window width of 400 trading days for calculation of the network metrics. We follow the methodology set out in Onnela et al. (2003b) and perform edge survival analyses to assess the stability of the MST over time using parameter $T = 400$. The *Single Step Survival Ratio* is defined in Onnela et al. (2003b) as the fraction of edges found in common in two consecutive trees. It can be calculated between times t and $t - \delta t$ as

$$\sigma(t) = \frac{1}{N-1} |E(t) \cap E(t-1)|, \quad (6.7)$$

where $E(t)$ is the edges in the MST at time t and $|\dots|$ is the number of elements in the set.

Although we have chosen a time step of $\delta t = 1$ for our analysis, we assess the single step survival ratio using $\delta t = 21$, or approximately one trading month to ensure stability of our results. This means that the single step survival ratio results presented in Figure 6.4 provide the fraction of edges in the MST which are still present in consecutive months. The results reveal that for a 400-day window width the MST is stable, with a mean survival ratio of above 0.75 for the S&P 500, the FTSE 100, the NIKKEI 225 and the SBF 120. This is a close approximation to the results presented for the NYSE in Onnela et al. (2003a) where they use a 1000-day window. Therefore, we conclude that the 400-day window provides suitably stable results, and any further increases in window width may provide only marginal increases in stability with a cost of decreased sensitivity.

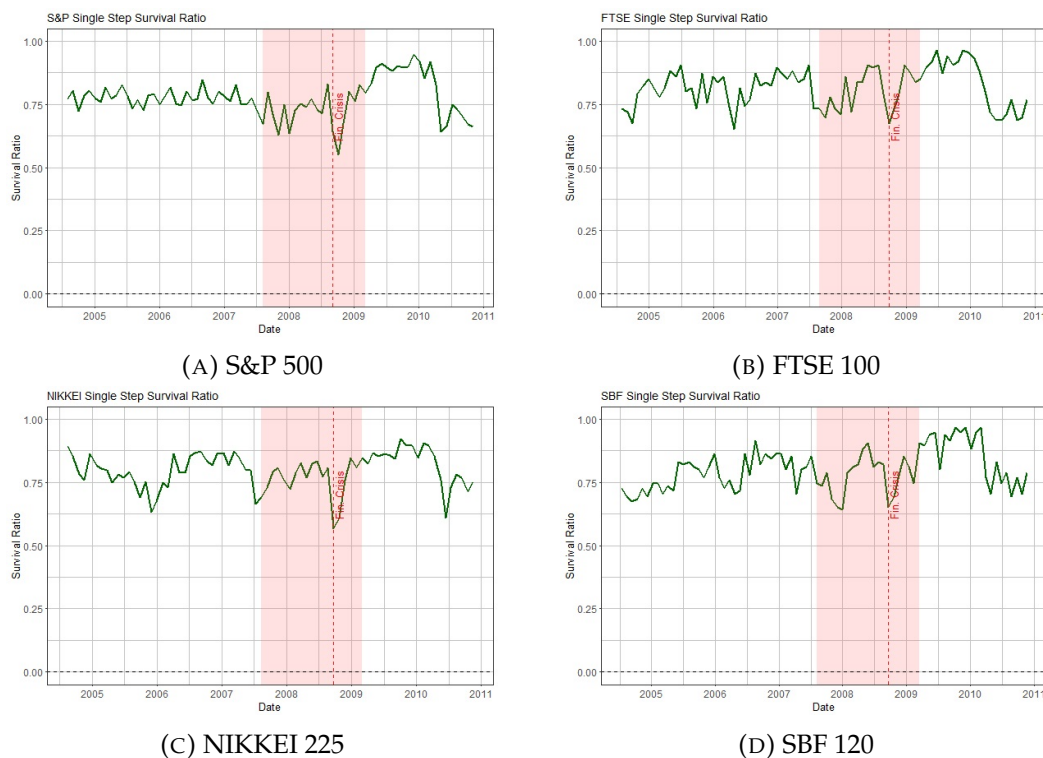


FIGURE 6.4: Single step survival ratio estimated using $T = 400$ and $\delta T = 21$ for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 from 2003–2010. Shaded area indicates the crisis period, from August 2007 – March 2009. The red dashed line indicates the Lehmann Brothers crash of the 14th of September 2008

6.5 Results

6.5.1 S&P 500 Static Analysis

In this section we present the results of a static minimum spanning tree analysis conducted on the S&P 500. As discussed above, Wiliński et al. (2013) detected the emergence of a superstar structure in MSTs extracted from the correlation matrix of the Frankfurt Stock Exchange during the global financial crisis. We test this result for the S&P 500 by plotting and assessing the degree distribution of MSTs extracted in three discrete time periods. The static MSTs are extracted using 400-day windows for periods up to end December 2005 (pre-crisis period), the end of August 2008 (crisis period), and the end of December 2010 (post-crisis period)⁴.

Visual analysis of the MSTs in Figure 6.5 panes (a), (c) and (e) reveals strong sectoral clustering and the presence of highly connected hubs. Financial sector stocks and

⁴We have tested multiple variations of these three periods, finding no significant difference in results. The results of these variations are not presented as they are captured in the mean occupation layer results presented in Section 5.2.

6.5. Results

TABLE 6.2: Percentage of time periods where central node of the S&P 500 MST has ICB industry classification equal to Financials. Estimated from overlapping rolling windows with $T = 400$ and $\delta T = 1$

Financial Sector Centrality			
Date Range	Aug 04–Dec 06	Jan 07–Sept 08	Sept 08–Dec 10
No. Days in Range	825	429	578
% Center	29.2%	67.4%	2.0%

industrial stocks tend to be closer to the centre of the network in the pre-crisis and crisis periods, with oil and gas, utilities and technology sectors residing on the peripheries of the network. Sectoral clustering in the consumer goods sector appears to be lower than in other sectors. Interestingly, in the post-crisis period the financial sector stocks have moved to the outer branches of the network. These initial results indicate that in the lead-up to the global financial crisis, stocks in the financial sector are strong propagators of shocks but once the crisis hits we detect a relative decoupling from the financial sector.

We also find further evidence of the changing role of the financial stocks over the period, 2003–2007, by analysing the central vertex in each time period. The central vertex in a MST is the vertex with the highest number of connections, weighted by the edge weights of those connections. In Table 6.2 we present the results for three time periods. From August 2004 to December 2006 we find that central vertex comes from the financial sector 29.2% of time periods. This increases to 67.4% of time periods from January 2007 to September 2008. Finally, in the period from the end of September 2008 to December 2010, the central vertex comes from the financial sector only 2% of the time. As discussed above, due to the fixed number of nodes and edges in the MST, the migration of the financial sector to the outer branches of the MST indicates a relative, rather than an absolute decoupling of the sector. The increase in influence of the financial sector during the crisis is in accordance with the results of Musmeci, Aste, and Di Matteo (2015) for the S&P 500. Furthermore, our results are similar to those of Aste, Shaw, and Di Matteo (2010), who found an overall decreasing trend in the relative degree of the financial sector of the U.S stock market from 1996–2010, with a spike in centrality over the period 2007–2008 followed by a sharp decrease in the post-crisis period.

In order to detect topological changes in the degree distribution of MSTs extracted from the S&P 500 over the sample period, we fit power law distributions to the degree distribution of the pre-crisis, crisis and post-crisis MSTs using the `powerLaw` package in R. As discussed in Chapter 2, power law network degree distributions (scale-free networks) are ubiquitous features of complex systems due to their inherent stability to random deletion of a node. In a power law distribution, the tail of the distribution $P(X = x)$ decays proportionally to $x^{-\alpha}$. The approach for fitting the power law distribution is described in Gillespie (2015) as follows:

The power law probability mass function is given by

$$P(X = x) = \frac{x^{-\alpha}}{\sum_{n=0}^{\infty} (n + x_{min})^{-\alpha}}, \quad (6.8)$$

where x_{min} identifies where the tail or scaling region of the distribution begins, $\alpha > 1$ and $x_{min} > 0$ and $\sum_{n=0}^{\infty} (n + x_{min})^{-\alpha}$ is the normalisation constant.

The slope of the power law distribution α is estimated using maximum likelihood estimation approximation.

$$\hat{\alpha} \approx 1 + N \left[\sum_{i=1}^N \ln \frac{x_i}{x_{min} - 0.5}^{-1} \right] \quad (6.9)$$


x_{min} is estimated using Kolmogorov-Smirnov approach as the value which maximises the distance between the data and the fitted model CDFs. The probability that the observed degree distribution follows a power law distribution is estimated using a bootstrapping approach by simulating M power law distributions using the estimated x_{min} and α , calculating the Kolmogorov-Smirnov value on the original and simulated distributions, and generating a p -value. When p is large, we fail to reject the null hypothesis that the data follows a power law distribution. For $p \approx 0$ we reject the null. In this chapter we estimate p -values using $M = 100$.

Panels (b), (d) and (f) of Figure 6.5 display the power law distributions fitted to our pre-crisis, crisis and post-crisis MSTs respectively. In the pre-crisis period the degree distribution has a clear power law tail with $\alpha = 3.26$. This finding is similar to that

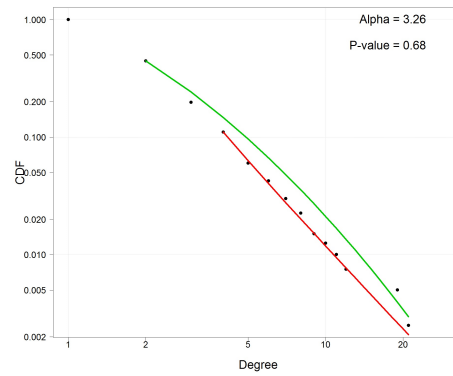
outlined in Wiliński et al. (2013) who find a power law distribution with $\alpha = 3.0 \pm 0.21$ for the pre-crisis period in the Frankfurt Stock exchange. The authors state that this power law exponent is a good approximation to the Barabási-Albert scale-free complex network. If we examine the network estimated for the crisis period we find that, while there still appears to be a power law scaling region, the extrema of the distribution are much smaller than would be expected from a power law distribution. Moreover, the exponent for the region that does scale according to a power law is significantly lower than in the pre-crisis period, with a value of 2.71.

This result is significantly different from that outlined in Wiliński et al. (2013), who split the crisis period into two distinct time periods and find the emergence of a super-connected node (superstar structure), followed by a power law distribution decorated by large hubs. The authors state that the emergence of the superstar hub may be an indication that a stock market crash is imminent within a few months. However, our results show that while the emergence of such a topological structure may be an indication of a stock market crash, it is not a necessary condition and, in fact, a reduction in the size of the hubs in the S&P 500 degree distribution occurs prior to the crash. Finally, in the post-crisis period, we again find a power law scaling region with an exponent of 3.44. However, in this time period a large hub with a degree of approximately 30 has appeared.


The results of this section demonstrate significant variation in the topological structure of the MST in the pre-crisis, crisis and post-crisis periods. However, while the topological structure of the MST changes during the crisis, we find no evidence of a topological phase transition, such as the emergence of a superstar structure as found in Wiliński et al. (2013). We further investigate this result in Section 6.2 for a number of major stock market indices by estimating measures of centrality and spread, such as the normalised tree length and mean occupation layer.




(A) Pre-Crisis MST




(B) Pre-Crisis Degree Distribution




(C) Crisis MST



(D) Crisis Degree Distribution



(E) Post-Crisis MST



(F) Post-Crisis Degree Distribution

FIGURE 6.5: Static MST analysis: Pre-Crisis, Crisis and Post-Crisis periods estimated using 400-day samples up the end of December 2005, August 2008, December 2010 respectively

6.5.2 Dynamic Network Metrics

The normalised tree length metrics for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 are presented in panels (a), (b), (c) and (d) of Figure 6.6 respectively. The graphs present evidence of a significant decrease in the normalised tree length for the FTSE 100, NIKKEI 225 and SBF 120 beginning in 2006. This result is confirmed by the Kendall's tau correlation coefficients in Figure 6.7, which shows that by the start of 2007 the normalised tree length has been displaying a strong decreasing trend for approximately 250 trading days in the three markets. The decreasing trend in the S&P 500 begins later, in mid-2007, just before the onset of the crisis period. A second discontinuous decrease occurs in all markets, corresponding with the stock market crash in September 2008.

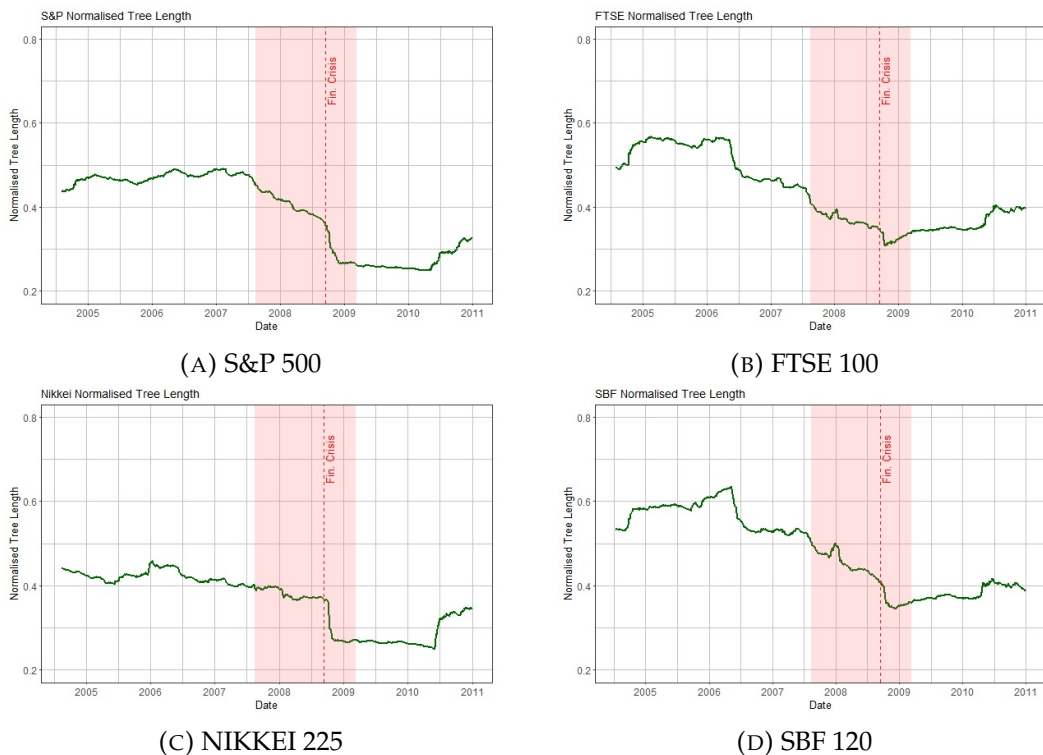


FIGURE 6.6: Normalised Tree Length estimated using $T = 400$ and $\delta T = 1$ for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 from 2003–2010. Shaded area indicates the crisis period, from August 2007 – March 2009. The red dashed line indicates the Lehman Brothers crash of the 14th of September 2008

These trends indicate an increasing co-movement between the stocks in each market, which begins before the onset of the crisis. This result may indicate that common factors related to a stock market bubble are driving dynamics in the markets in the lead-up to the crisis period. As discussed above, the decrease in the normalised tree

length increases the probability that a shock will propagate through the system and, can be seen as an increase in systemic risk. It is interesting that the trend in the S&P 500 begins later than in other markets, given the common wisdom that the sub-prime crisis was the spark that ignited the global financial crisis, and the U.S. stock market, particularly the financial sector, was the epicentre of the crisis.

The results indicate that the normalised tree length is a potential early warning signal for stock market crashes. The indicator provides a stable yet sensitive measure of interconnectedness in the network and, in three out of the four markets studied, displays strong trends long before the onset of the crisis period. Furthermore, the measure is also a useful measure for ongoing crisis monitoring. For example, the S&P 500 entered a bear phase in October 2007. Following this period, central banks of major economies offered unprecedented levels of direct funding to banks, due to non-existent financial sector liquidity, and brokered deals such as the March 2008 acquisition of Bear Stearns by JP Morgan Chase. While major adverse developments in the financial sector continued, the Federal Reserve backed acquisition of Bear Stearns and the opening of new Federal Reserve lending facilities restored some calm to the markets (Mishkin, 2010). The normalised tree length results indicate; however, that markets remained highly vulnerable to adverse shocks propagating through the system, despite the actions of central banks. This vulnerability became apparent in September 2008 with the decision to allow Lehmann Brothers to fail, when the shock cascaded through global stock markets (see Figure 6.1).

6.5. Results

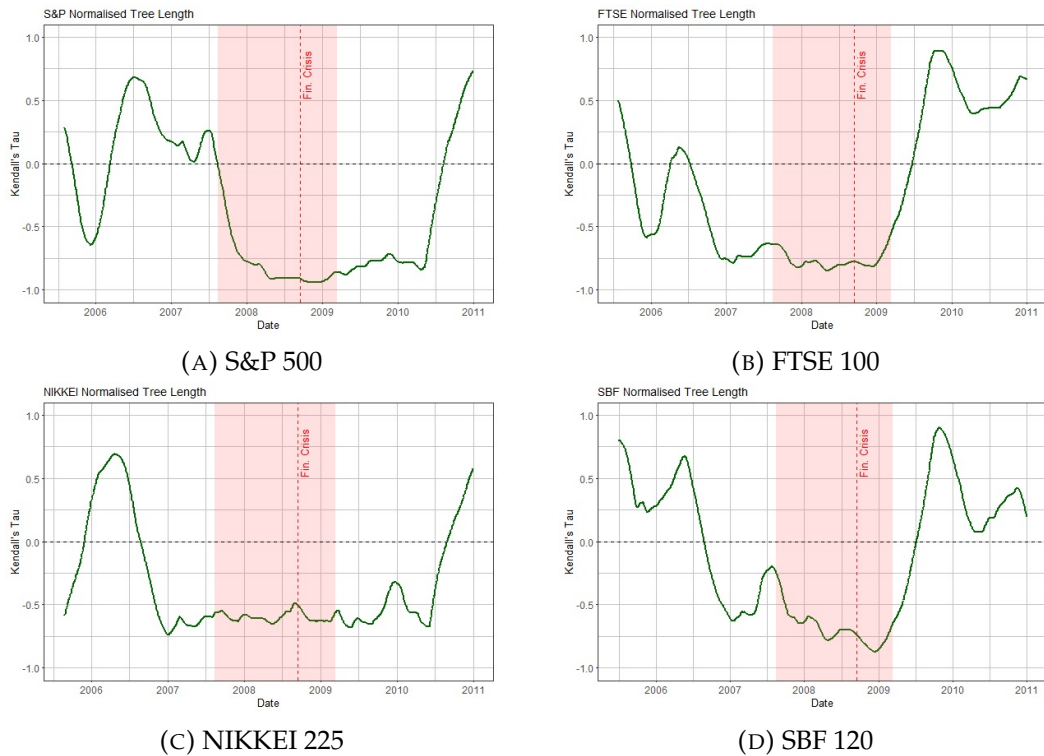


FIGURE 6.7: Kendall's tau coefficient estimated for the normalised tree length with window width of 250 trading days for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120

The mean occupation layer measures the average level of the tree, or the average number of steps to the central vertex. Kaya (2015) and Wiliński et al. (2013) found a prevalence for decreasing trends in the spread of MSTs calculated on global financial indices prior to crises. Onnela et al. (2003b) state that an increase in the mean occupation layer indicates an increase in diversification potential in the market. Figure 6.8 presents the results of the mean occupation layer, with Figure 6.9 providing the Kendall's tau coefficients for the metric.

The results indicate a decrease in the mean occupation layer for a number of the markets either directly prior to, or during, the global financial crisis. Panel (a) of Figure 6.9 for the S&P 500 shows a Kendall's tau coefficient of almost -0.5 for the 250 days up to the onset of the crisis in August 2007. The trend in the FTSE 100 (panel (b)) is even stronger with a Kendall's tau of approximately -0.7. A second negative trend occurs in the FTSE 100 in the 250 days leading up to the Lehmann Brothers crash, with a Kendall's tau of approximately -0.6. For the Nikkei 225 (panel (c)) a negative trend occurs in the 250 days leading up to the end of Q1 2008, continuing until the crash. There is no evidence of a strong decrease in mean occupation layer

for the SBF 120 over the crisis period (panel (d)).

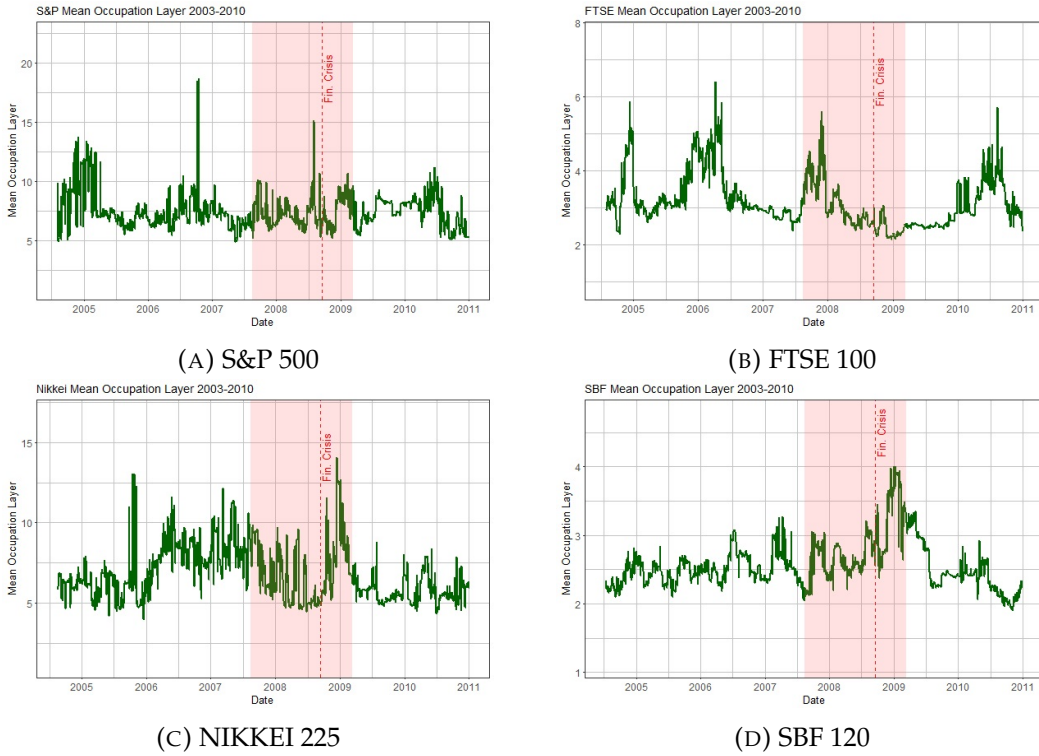


FIGURE 6.8: Mean Occupation Layer estimated using $T = 400$ and $\delta T = 1$ for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 from 2003–2010. Shaded area indicates the crisis period, from August 2007 – March 2009. The red dashed line indicates the Lehmann Brothers crash of the 14th of September 2008

While there is some evidence of decreasing trends in the mean occupation layer in the periods up to and including the crisis, with the exception of the FTSE 100, stronger negative trends occur in the mean occupation layer in non-crisis periods than in crisis periods. Moreover, visual analysis of Figure 6.8 reveals that the Mean Occupation Layer is quite volatile over the period of analysis, with large peaks and troughs before, during and after the crisis period. Thus, we find that the mean occupation layer may be a less informative indicator of systemic risk than the normalised tree length.

6.5. Results

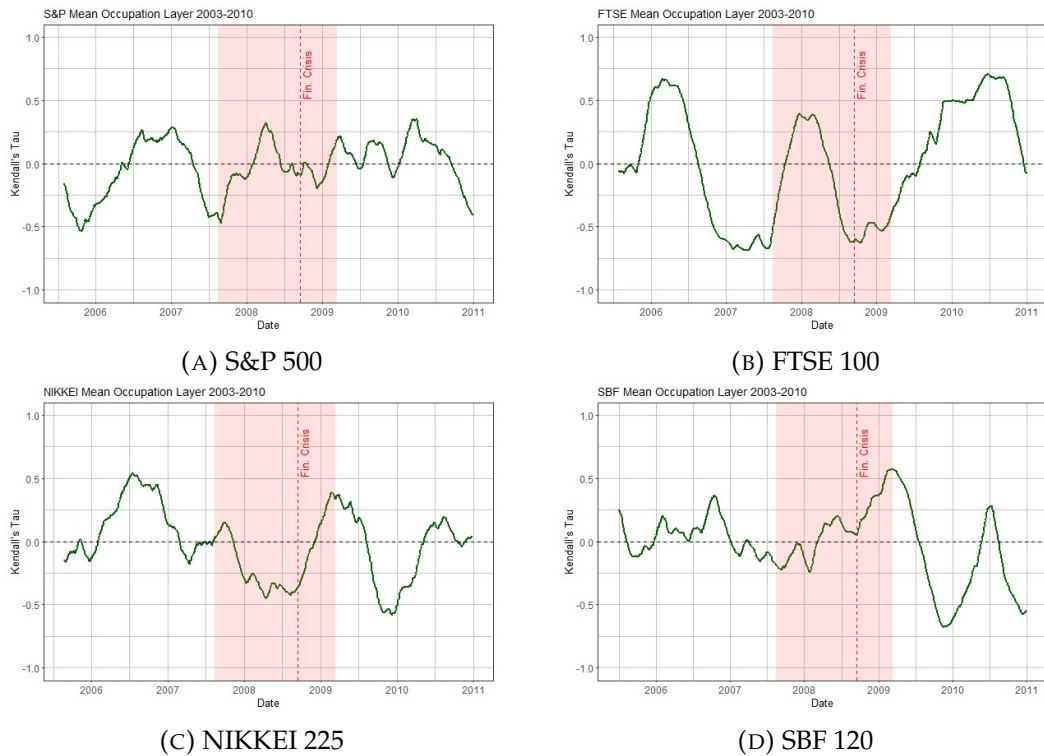


FIGURE 6.9: Kendall's tau coefficient estimated for the mean occupation layer with window width of 250 trading days for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120

In Figure 6.10, we examine the modularity of the ICB industry and sub-sector groupings. As outlined in Section 6.4, the modularity of a network community is a measure of similarity of the nodes within that community and dissimilarity with other communities within the network. Musmeci, Aste, and Di Matteo (2015) found that for the S&P 500, the informativeness of the ICB industry and sub-sector groupings decreased in the lead-up to the financial crisis, recovering somewhat after the crisis. We further investigate this result for our four markets of interest to determine if this is a general trend in financial markets in the lead-up to the financial crisis, and thus, a potential indicator of systemic risk. Our hypothesis is that a decrease in the degree to which stocks from the same sector cluster together may indicate bubble dynamics, where the market factors, such as risk aversion or credit availability, become more important than sector specific factors.

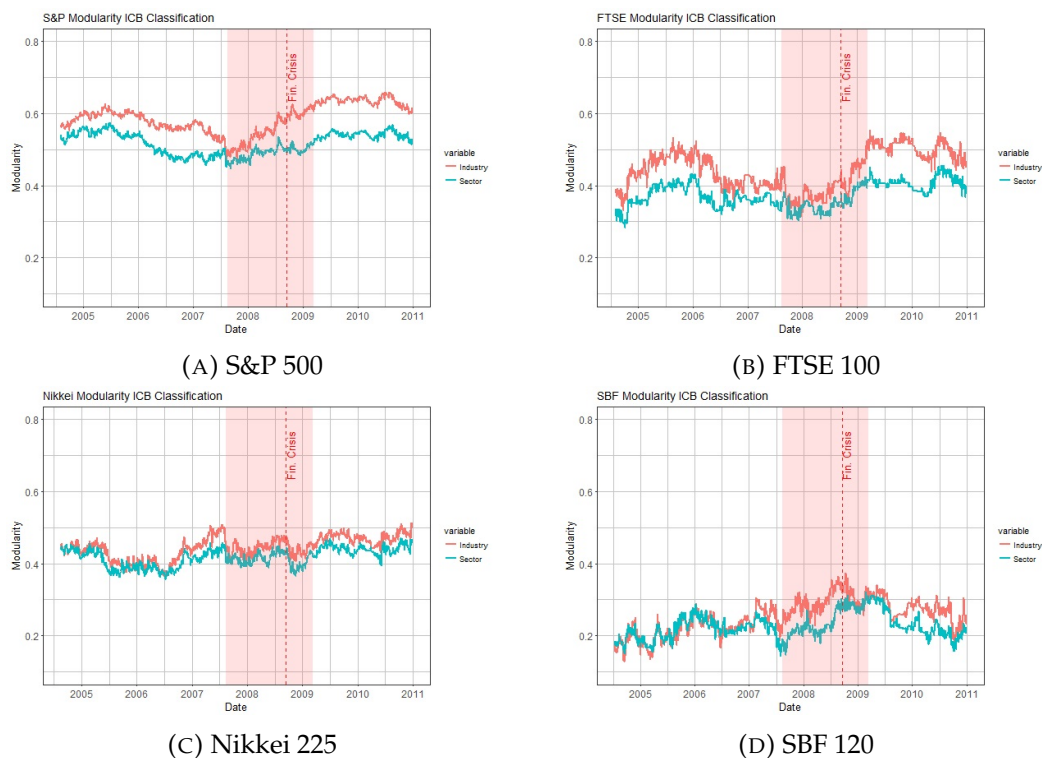


FIGURE 6.10: ICB Industry and sub-sector modularity estimated using $T = 400$ and $\delta T = 1$ for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 from 2003–2010. Shaded area indicates the crisis period, from August 2007 – March 2009. The red dashed line indicates the Lehman Brothers crash of the 14th of September 2008

The results presented in Figure 6.10 for the S&P 500 broadly correspond to those presented in Musmeci, Aste, and Di Matteo (2015) for both industry and sub-sector groupings. Examining the Kendall’s tau coefficient in Figure 6.11 for the industry level grouping, we can see a strong decreasing trend in the period up to Q4 2006, followed by a short levelling out period, and another strong decreasing trend up to the end of 2007. From the onset of the crisis there is a strong recovery in modularity, particularly at industry level. The results for the FTSE 100 display similar trends to those for the S&P 500, indicating similar underlying dynamics in the market. No such decrease is present for the NIKKEI 225 or the SBF 120, with an overall increase in industry modularity in the lead-up to the crisis. All four markets experience a sharp, short lived decrease in modularity with the onset of the crisis in August 2007.

6.5. Results

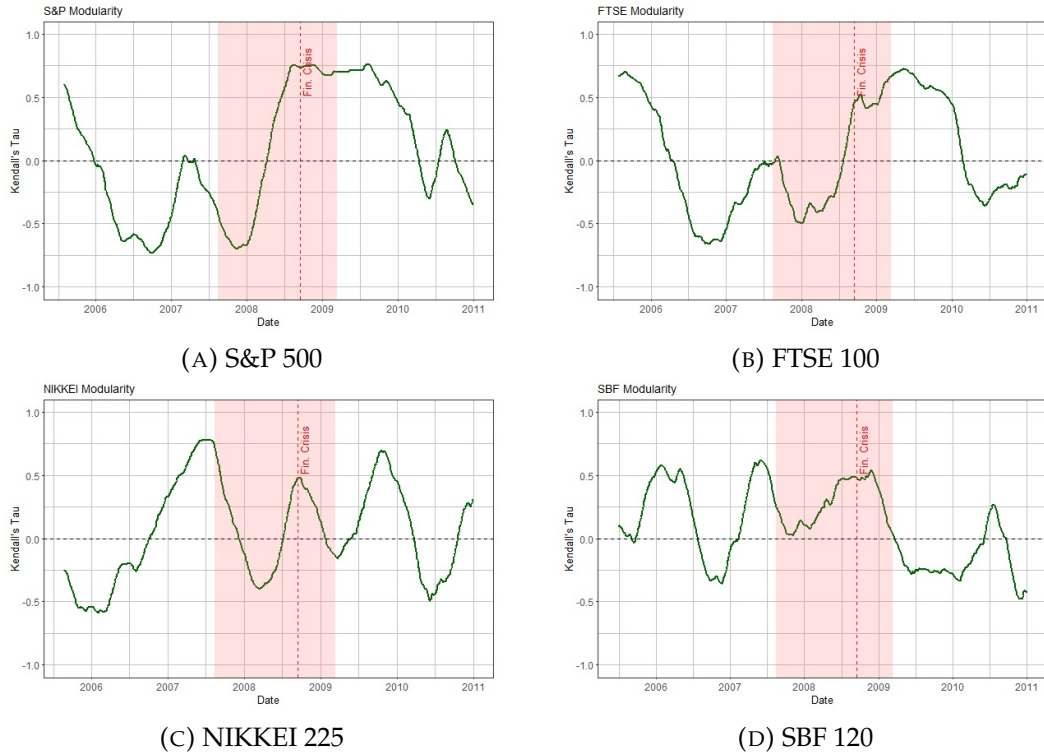


FIGURE 6.11: Kendall's tau coefficient estimated for the ICB industry and sub-sector modularity with window width of 250 trading days for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120

We also detect a divergence in trends between the S&P 500 and the FTSE 100 on one hand, and the NIKKEI 225 and SBF 120 on the other, in terms of the relative influence of the financial sector in the market. The average financial sector degree displayed in Figure 6.12 indicates the relative influence of the financial sector in the MST, with Kendall's tau coefficients displayed in Figure 6.13. After an initial decrease in financial sector centrality for the S&P 500 between 2004 and late 2005, the average degree of the financial sector trends upwards and reaches a peak in 2007, corresponding to the onset of the financial crisis (panel (a)). A similar trend is detectable for the FTSE 100 in panel (b). This trend may indicate stress from the financial sector increasingly being transmitted to the rest of the market. We detect no corresponding increase in financial sector centrality for MSTs estimated for the NIKKEI 225 (panel (c)) or SBF 120 (panel (d)). Interestingly, all four stock markets experience a migration of the financial sector to the outer branches of the MST in the period following the Lehmann Brothers crash. This reflects the global nature of the impairment of the financial sector in this period, which may have recovered slower relative to other sectors, due the requirement for significant deleveraging and

disposal of toxic assets.

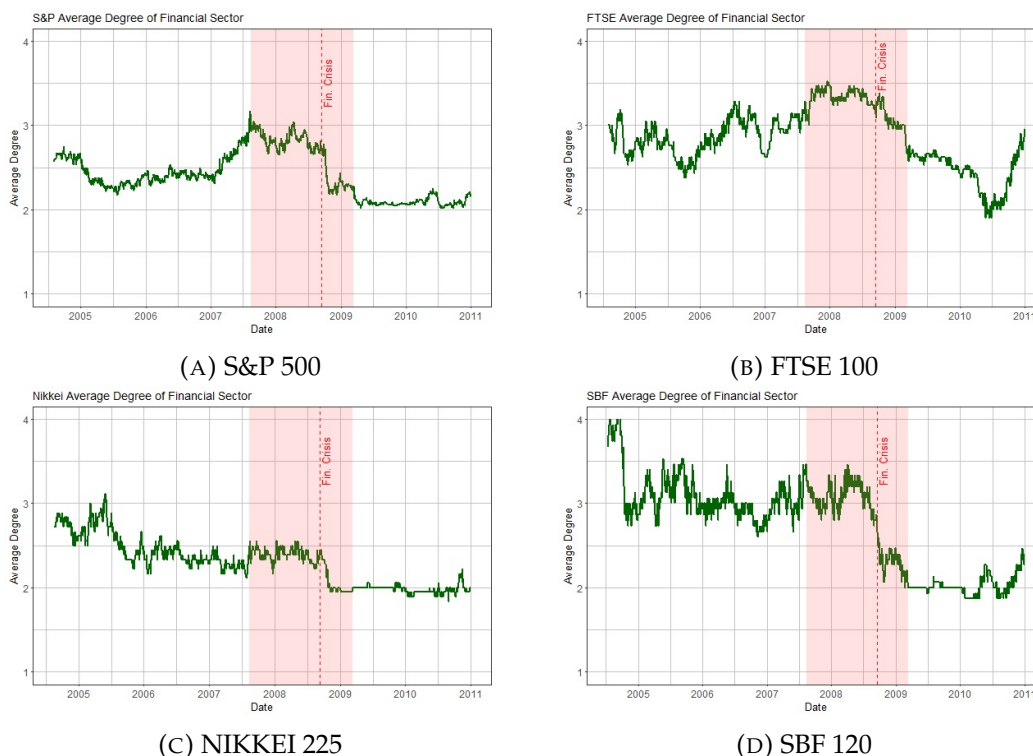


FIGURE 6.12: Average Financial Sector Degree estimated using $T = 400$ and $\delta T = 1$ for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120 from 2003–2010. Shaded area indicates the crisis period, from August 2007 – March 2009. The red dashed line indicates the Lehmann Brothers crash of the 14th of September 2008

The negative trends in modularity in both the S&P 500 and the FTSE 100 broadly correspond to increases in financial sector centrality in those markets. This is evidenced by the fact that peaks in the Kendall's tau coefficients for the average financial sector degree occur in approximately the same time periods as the troughs in modularity in both markets. After the onset of the crisis in August 2007, modularity begins to increase and financial sector centrality begins to decrease. There is a -0.88 Pearson's correlation over the entire sample period between modularity and average financial sector degree for the S&P 500 and -0.67 for the FTSE 100. This is contrasted with a correlation of 0.3 for the SBF 120 and 0.13 for the NIKKEI 225. While this does not provide direct evidence of a negative causal relationship between the network measures, it provides an interesting avenue for future research.

6.5. Results

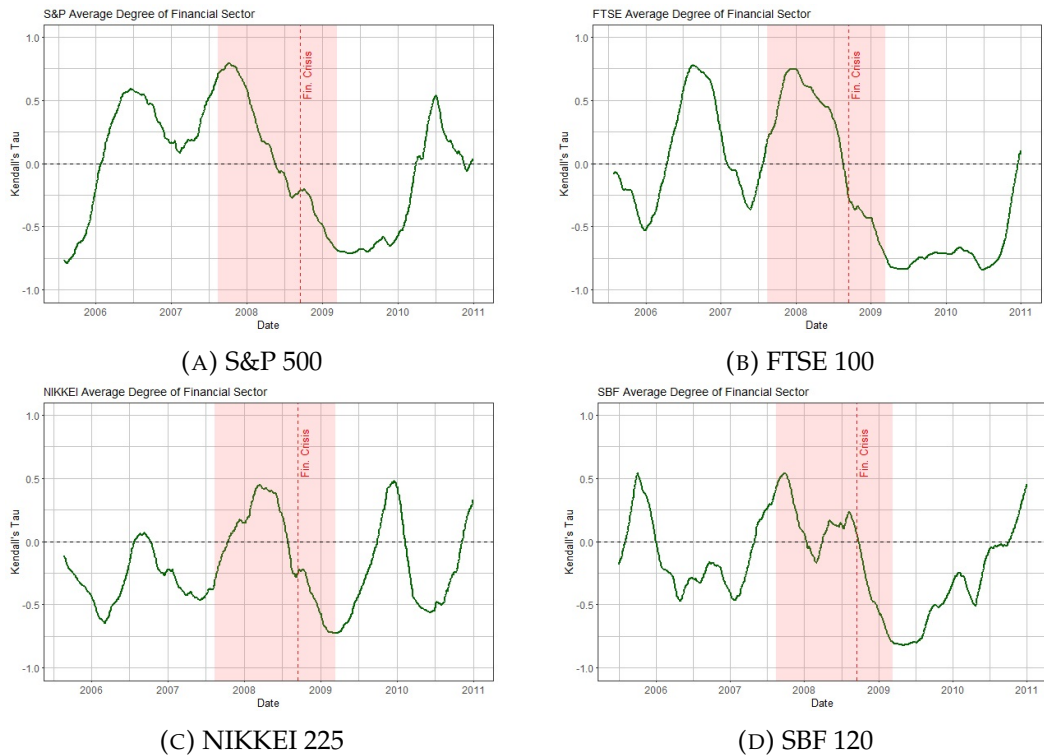


FIGURE 6.13: Kendall's tau coefficient estimated for the average financial sector degree with window width of 250 trading days for the S&P 500, FTSE 100, NIKKEI 225 and SBF 120

One potential hypothesis is that the increase in financial sector influence, coupled with decreasing modularity, is related to financial factors, such as credit availability, becoming more important than sector specific factors. As a narrative, the U.S. mortgage market peaked in mid-2006, followed by a steep decline, causing fears regarding U.S. banks exposures to sub-prime mortgages. This led to a reduction in interbank lending and general tighter credit conditions. The results presented for financial sector centrality indicates that tightening credit conditions may have been felt first in the U.S. and U.K. stock markets, with shocks being transmitted to non-financial sectors. If this is the case, then the average financial sector degree and modularity of industry groupings may provide useful early warning signals for determining when financial sector impairment becomes systemic, with spillover effects to the real economy. The fact that Financials comprise the largest component of the FTSE 100 and S&P 500 samples but only the third and fifth largest in the SBF 120 and NIKKEI 225 may contribute to the divergence in results discussed above. However, further research is required into the role of the financial sector in pre-crisis dynamics in all four stock markets to determine why trends are detectable in only

two out of the four markets in the lead-up to the crisis.

6.6 Conclusions

Detecting systemic risk and instabilities in financial markets requires a multifaceted approach due to the complex nature of financial markets. In this chapter we assessed dynamic minimum spanning tree analysis of stock markets as a potential method for the detection of systemic risk. Due to wealth effects, financial accelerator effects and their use in credit risk estimation, equities make a natural candidate for detecting developments in financial markets with potential spillover effects to the real economy. Moreover, minimum spanning trees, which extract the shortest and most probable path through which a shock can pass through a system, are capable of detecting sectoral clustering in stock markets.

Our analysis reveals a large degree of variability in the topological structure of MSTs estimated from the correlation networks of the S&P 500, FTSE 100, NIKKEI 225 and SBF 120, as evidenced by a large number of peaks and troughs in measures of network spread. However, we find strong evidence that a measure of overall interconnectedness of the MST, the normalised tree length, provides a clear early warning signal for the global financial crisis in all four markets. Furthermore, we find some evidence of early warning signals from the average financial sector degree, and ICB industry grouping modularity for the S&P 500 and FTSE 100. These signals are not present for the NIKKEI 225 and SBF 120; however, the absence of such signals may be related to the role of the financial sector in pre-crisis market dynamics. We propose further research into this result to assess the suitability of these measures as leading indicators of the transmission of shocks from the financial sector to non-financial sectors.

The results related to the increasing interconnectedness of the network in the pre-crisis period are particularly interesting when framed in the light of previous literature on phase transitions in complex networks. As discussed in Chapter 2, May (1972) and Haldane (2009) suggest that the level of interconnectedness, or strength of connections, in a complex network can reach a critical level, causing a phase

transition in the network. Our results highlight a decreasing trend in the normalised tree length for all four market under analysis in the build-up to the crisis. This could potentially indicate that a critical level of interconnectedness in the financial system was reached in 2007–2008, causing a phase transition in financial markets that materialised as the global financial crisis.

Overall our results indicate that dynamic minimum spanning tree analysis provides an intuitive method for assessing financial fragility and systemic risk, and would form a useful addition to the toolbox of financial regulators and supervisors alike. Used in isolation it cannot provide direct evidence of the presence of asset bubbles; however, it can provide a measure of the dominance of sectors within the market. Furthermore, MST analysis can detect increasing interconnectedness, and the increased transmission of shocks between sectors, all of which may contribute to a financial system which is more vulnerable to a systemic event. Finally, our results indicate that methods which take into account the non-stationary nature of dependence structures in financial markets can add significant value to the development of early warning signals of financial crises.

Chapter 7

Conclusions

7.1 Overview of Thesis and Contributions

The global financial crisis of 2007–2009 highlighted the need for the development of new approaches in the measurement of risk in financial markets. Many common risk models relied on the linear stochastic modelling framework, which assumes that financial data depends only on past observations of itself and other market data (Danielsson, 2002). Such models utilised assumptions, such as IID normally distributed modelling error terms, stationarity of empirical distributions and co-dependence structures, and that when taken on aggregate the actions of market participants are random and cannot influence the market. Chapter 1 of this thesis outlined a number of stylised facts for financial data, indicating that these assumptions do not always hold, and may lead to an underestimation of risk.

Empirical research has found fat tailed distributions with power law tails, long-memory dependence structures, and non-stationarity in variance and covariance across many asset classes and regions (Cont, 2001). Combining such features with the fact that financial markets tend to undergo abrupt price changes, remaining in a recessionary state for a prolonged period, points to the need for an alternate framework for risk measurement and early detection of financial crises.

Chapter 2 outlines an alternate framework for the analysis of risk in financial markets. This approach is to treat financial markets as complex systems, i.e. systems in which the interaction of the elements of the system at a micro level leads to very different behaviour at higher levels of aggregation. This approach is very different

from the linear stochastic framework in that it allows for endogenous dynamical behaviour (Arthur, 1999). Market participants, in general, are faced with uncertainty regarding the outcomes of their actions, due to limited information and limited cognitive ability (Simon, 1957). This leads to the use of heuristics and imitation of other market participants (Tversky and Kahneman, 1975; Shiller, 2002). As a result, market participants are constantly adapting and reacting to the outcomes of their own actions and the actions of other market participants. This constant cycle of reacting and adaption can lead to endogenously driven fluctuations, and have been related to bubble dynamics in financial markets.

When endogenous dynamics, such as positive feedback loops between asset price increases and investor behaviour, push the system towards a critical point, changes in underlying parameters can lead to structural change in system and regime shifts between alternate equilibria. In Chapter 2, we outlined previous literature that found evidence of emergent phenomena in financial markets, such as log-periodic oscillations, critical slowing down, power law distributions and scale free networks. These phenomena, which are common across many asset classes and time periods, provide evidence and motivation for the study of financial markets as complex dynamical systems and financial crises as critical phenomena.

It is clear from the discussion in Chapter 2 that financial markets are highly complex systems, likely displaying path dependence, making it difficult to extrapolate from past behaviour. However, the emergence of behaviours or trends which are ubiquitous in complex systems motivates the objectives of this thesis. The central research question of this thesis is: Can we exploit certain universal features of systems approaching a regime shift to detect the build-up of instabilities in financial markets and generate consistent, reliable early warning indicators of financial crises? This central research question is in turn broken down into three further research objectives, which are addressed in Chapter 4 – Chapter 6 of this thesis as follows:

1. To determine if nonlinear dependence structure are present in financial time series in the lead-up to financial crises.
2. To investigate if early warning signals of regime shifts were present prior to the

Eurozone sovereign debt crisis and, if so, whether or not these signals can be used as robust indicators of financial crises in general.

3. To examine the usefulness of measures of time-varying correlation network topology of equity returns as a measure of systemic risk in financial markets.

In answer to the central research question, we have found a number of promising early warning indicators for financial crises in stock and sovereign bond markets. For example, Chapter 5 of this thesis finds evidence of critical slowing down in sovereign bond markets in the lead up-to the Eurozone sovereign debt crisis, and Chapter 6 finds that a topological shrinking of MSTs estimated from the correlation network of global stock markets prior to the global financial crisis. We discuss the contributions of each of the empirical chapters (Chapter 4 – Chapter 6) below.

The objective of Chapter 4 is to determine if we can detect nonlinearity in financial time series in the periods directly preceding financial crises. The motivation is due to the fact that regime shifts are triggered by strong nonlinear responses in a system pushing it towards a critical threshold (Dakos et al., 2012b). Moreover, bifurcations cannot occur in purely linear systems. Our primary contribution is to investigate the presence of nonlinear dependence structures in stock index and sovereign bond time series in the periods leading up to the global financial crisis and the sovereign debt crisis respectively. Our results find mixed evidence for weak nonlinear dependence in a number of the financial time series under investigation. The evidence is strongest for sovereign bond markets, with stock markets yielding largely inconclusive or negative results.

The BDS test (Brock, Dechert, and Scheinkman, 1987), applied to log-squared GARCH residuals of the stock and bond time series, displayed evidence of nonlinearity in the Greek, Irish and Portuguese bond data for different sample sizes. No evidence of nonlinear dependence was found in the Dow Jones, FTSE 100 or NIKKEI 225. The results for the Irish and Portuguese bond markets appear to be confirmed by STAP surrogate data tests (Kugiumtzis, 2002a). However, insignificant results for the Greek market indicate that the significant result in the BDS test may be due to remaining linear dependence following the application of the GARCH model.

The STAP methodology may have more power than the BDS test to detect nonlinearity in stock market indices, with some significant nonlinear serial dependence detected for the Dow Jones, the FTSE 100 and the NIKKEI 225. When combined with the results of the BDS test, this indicates that the signal to noise ratio for stock market indices may be weaker than for sovereign bond yields. Finally, the application of PPS surrogate data tests (Small, Yu, and Harrison, 2001) directly to the log-returns finds no significant results for bond series and only weak evidence of nonlinear dependence in the NIKKEI 225. The PPS method is used due to the fact that it can capture volatility clustering effects and does not have to be applied to GARCH residuals. However, it may not have the power to detect weak nonlinearity in noisy time series.

A second key finding of this chapter is the difficulty providing robust results when testing for nonlinearity in financial time series. This is due to the finite, noisy nature of financial data. Where the nonlinear signal in the data is weak relative to level of noise, or if the data generating process is high dimensional deterministic, tests for nonlinear dependence may not have sensitivity to detect the signal in the data. While the results of Chapter 4 are mixed, we do find some evidence of weak nonlinearity in the data, particularly for sovereign debt markets. Therefore, in Chapter 5 we investigate if nonlinear responses in the system pushes the market towards a critical threshold, where instabilities in the system build-up over time and lead to an abrupt regime shift.

Recently, the Eurozone debt crisis sparked conversation surrounding the possibility of multiple equilibria in sovereign bond markets, particularly in a monetary union where the sovereigns do not have the guarantee of a lender of last resort (De Grauwe, 2011). Comparisons have been made between the sovereign debt crisis and the self-fulfilling expectations framework set out in second-generation currency models. In such models, multiple equilibria are possible in the market for a particular range of economic fundamentals. In effect, a bifurcation occurs when economic fundamentals enter a certain range (Jeanne, 1997). Motivated by the recent literature, and our findings in Chapter 4 regarding evidence of nonlinear dependence in sovereign bond markets, we test for evidence of CSD in the sovereign debt markets

of Greece, Ireland and Portugal.

The primary finding of Chapter 5, and our contribution to the literature, is to present evidence of the presence of CSD in the lead-up to the debt crisis for Greek, Irish and Portuguese sovereign bond yield returns. CSD has been used as a leading indicator of critical transitions in many real world dynamical systems (Scheffer et al., 2009). In all three markets we detect evidence of strong positive, statistically significant trends in ACF1, AR1 and SD. The significant trends, coupled with the fact that all three markets experienced large increases in yields, and linger in the high yield state for a prolonged period of time, indicates that CSD may provide a useful indicator of impending transitions in sovereign debt markets.

However, we find that the use of ACF1 as a leading indicator of critical transitions in stock markets reveals a propensity for false positive signals. This result confirms recent findings in Guttal et al. (2016). This secondary finding is interesting when coupled with the findings of Chapter 4. The mixed results of the tests for nonlinearity, coupled with the finding that CSD is not a useful indicator for crises in stock market indices, points to the need for a different approach. Price changes in stock indices are generally a weighted average of the changes in price of the stocks in the underlying market. It is possible that the aggregation method washes out interesting dynamics in the underlying stock returns. Therefore, in Chapter 6 we go to a lower level of aggregation, investigating changes in correlation network topology as a potential early warning indicator of stock market crashes.

In Chapter 6 we investigate dynamic minimum spanning tree analysis of stock markets as a potential method for the detection of systemic risk. Our primary contribution to the literature is to apply a range of different measures of changing network topology to the components of four stock market indices, the S&P 500, the FTSE 100, the NIKKEI 225 and the SBF 120. The motivation for the empirical methodology stems from results of network analysis in other complex systems, which find that the interconnectedness of a network can reach a critical level of interconnectedness and cause a phase transition in the system. For example, May (1972) and Haldane (2009) suggest that the level of interconnectedness, or strength of connections, in a complex network can reach a critical level, causing a phase

transition in the network. Furthermore, we assess whether a number of recent results related to topological change of correlation networks are generalisable across multiple markets, thereby providing us with consistent, reliable early warning signals of financial crises.

The primary finding of Chapter 6 is a strong decreasing trend in the normalised tree length for all four markets under analysis in the build-up to the crisis. This provides some support for the theory that a critical level of interconnectedness was reached in the financial system, causing a phase transition in financial markets. Furthermore, we find some evidence of continuous change in the average financial sector degree and ICB industry grouping modularity for the S&P 500 and FTSE 100 in the lead-up to the financial crisis. While these signals were not present in the SBF 120 and the NIKKEI 225, we suggest that this may be due to the less significant role of the financial sector in these markets. If one were to take the dominance of the financial sector as a proxy for credit availability, or credit driven growth, an increase in the centrality of the financial sector, coupled with a decrease in sectoral modularity, may be an indication that systematic factors are becoming more important in asset price dynamics.

Another key finding of this chapter is that the S&P 500 displays power law degree distributions in the pre-crisis period, similar to the results found in Wiliński et al. (2013) for the Frankfurt Stock Exchange. However, we do not detect the same sharp decrease in the spread of the network or the emergence of a super-connected hub that is highlighted for the German market. This indicates that the dynamics of the crisis may be different in the two markets.

When taken on aggregate, the results highlighted in the empirical chapters provide a number of key contributions to the literature. Firstly, we highlight the challenge of detecting nonlinearity in financial markets, and the strong dependence on the methods and parameters used. More importantly, the results indicate that when conducting risk analysis of financial markets, the level of aggregation at which we apply our methodologies can play a key role in the success of our results. The results of the Chapter 4 and Chapter 5 indicate nonlinearity and early warning signals of critical transitions in sovereign bond markets in the lead-up to the Eurozone sovereign debt crisis. However, the results are either insignificant or mixed for stock

market indices. Reducing the level of aggregation to the individual components of the stock market indices, we find interesting trends and potential early warning indicators for the financial crisis. Finally, the results indicate that moving beyond the linear stochastic framework and applying methods which can capture evolving, complex interactions, and resultant emergent patterns, adds significant value in the development of early warning signals of financial crises.

7.2 Limitations

While we have attempted to ensure the robustness of results and suitability of methodologies utilised in this thesis, there are a number of limitations which must be discussed. In Chapter 4, as we have already discussed, the methodologies for detecting nonlinearity in financial time series are limited due to the finite, noisy nature of such series. Many methods which are prevalent in nonlinear time series analysis struggle to distinguish between stochastic and deterministic dynamics when the signal to noise ratio is low. We have attempted to overcome this limitation by applying a range of different methods, for example by applying BDS testing, as well as STAP and PPS surrogate data analysis. However, each of the methods has its own limitations. The BDS and STAP tests are both applied to model residuals, which has the potential to either bleach interesting dynamics or add spurious dependence structure to the data. The PPS test was applied directly to the log-returns; however, it may not have the power to distinguish between time series with a weakly nonlinear deterministic component and one which is purely stochastic.

In Chapter 5, we assess the possibility of a fold catastrophe (or saddle node bifurcation) occurring in financial markets in the lead-up to the Eurozone sovereign debt crisis. We apply this methodology due to the qualitative similarities between these bifurcations and financial crises. The CSD methodology indicates dynamics are present which support the hypothesis that sovereign bond markets were approaching a critical transition in the periods leading up to the crisis. However, as discussed in Boettiger, Ross, and Hastings (2013), CSD can occur prior to other types of transitions, including continuous transitions. In order to get around this limitation

we have used the classification system outlined in Boettiger, Ross, and Hastings (2013) to assess whether a saddle node bifurcation has occurred. This involves the use of judgement and, as such, has the potential for errors.

Finally, in Chapter 6 we assess changes in correlation network topology as a potential indicator of systemic risk. In this chapter we use Pearson's linear correlation coefficient to measure the co-movement of log-returns. Linear correlation coefficients may not detect the presence on nonlinear dependence between the stocks. A methodology such as mutual information, which captures both linear and nonlinear dependencies may also be useful in this regard. We use linear correlation due to the large body of literature which support its use in analysis of correlation network topology. However, we recognise that ignoring nonlinear dependencies is a potential limitation of this work. Finally, due to the unavailability of data for prices of stocks which have been delisted from the stock indices under examination, an unavoidable survivorship bias exists for the results of Chapter 6.

7.3 Future Research

This thesis has provided a number of important avenues for future research. Chapter 4 highlights the difficulty in detecting nonlinearity in financial time series due to their finite, noisy nature. A number of methods are available, such as wavelet methods, for reducing the noise levels in time series, while preserving any deterministic components. We propose future research to assess the suitability of these methods for financial time series. Significant research is required in this area as many noise reduction methodologies are known to introduce spurious determinism. Where a suitable noise reduction methodology is found, we propose that further research be conducted into the presence of nonlinearity in filtered financial time series in the build-up to crisis periods. This may provide stronger evidence for, or against, the presence of nonlinearity in financial markets in the lead-up to crisis periods.

Chapter 5 of this thesis provides evidence that Eurozone sovereign debt crisis may be explained as a system passing through a bifurcation point, where some underlying parameter (such as the level of economic fundamentals) reaches a critical threshold

and causes the system to switch between alternate equilibria. Second generation currency models have recently been applied to explain sovereign debt crises. These models allow for a range of values for economic fundamentals which lead to multiple equilibria in bond markets. When fundamentals pass a critical value, the system collapses to a single, high yield equilibrium state. However, these models do not explicitly apply the framework of bifurcations to explain price dynamics as the transition is approached. We propose future research into the development of modelling approaches which directly apply bifurcation theory to model the dynamics sovereign debt crises. Such an approach would leverage the large body of literature in the ecological and other sciences that have successfully applied bifurcation theory to model regime shifts in complex systems.

In Chapter 6 we highlight an interesting result where the S&P 500 and the FTSE 100 display an apparent relationship between sectoral modularity and financials sector centrality. We hypothesise that the increase in financial sector influence, coupled with decreasing modularity, is related to systematic factors, such as credit availability, becoming more important than sector specific factors. The same relationship is not present in the Nikkei 225 or the SBF 120, where the financial sector is a less sizeable component of the overall index. We propose further research into the role of the financial sector in pre-crisis stock market dynamics, to determine if the apparent relationship between financial sector centrality and modularity can be used as an early warning signal for credit bubbles in markets with dominant financial sectors.

Appendix A

Appendix Chapter 4

TABLE A.1: BDS Statistic for Log Squared EGARCH Residuals, N=500

m	Dow	FTSE	Nikkei	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	-1.21	-1.06	0.28	1.72	-2.15	0.28
3	-1.39	-1.02	-0.36	1.54	-2.43**	0.62
4	-1.33	-0.4	-0.69	2.05	-2.96**	1.19
5	-1.88	-0.09	0.43	1.45	-2.36	1.23
6	-2.09	-0.5	0.69	1.04	-2.26	0.1
7	-2.21	-0.57	-0.54	0.76	-2.03	0.06
$\epsilon/\sigma = 1$						
2	-1.38	-0.71	-0.35	1.95*	-1.44	-0.1
3	-1.27	-0.53	-1.01	2.28**	-2.02**	0.07
4	-1.03	0.12	-1.02	2.45**	-1.97*	-0.02
5	-1.52	0.57	-0.7	2.18*	-1.92*	-0.1
6	-1.72	0.61	-0.69	1.90*	-1.97*	-0.29
7	-1.762*	0.39	-0.37	1.63	-1.69	-0.56
$\epsilon/\sigma = 1.5$						
2	-0.35	-0.87	-0.99	2.01*	-0.92	-0.67
3	-0.55	-0.16	-1.51	2.49**	-1.57	-0.48
4	-0.41	0	-1.56	2.54**	-1.80*	-0.68
5	-0.73	-0.03	-1.26	2.24**	-1.71*	-0.87
6	-0.86	-0.22	-1.04	1.99*	-1.59	-1.09
7	-0.87	-0.42	-0.86	1.79*	-1.2	-1.3
$\epsilon/\sigma = 2$						
2	0.44	-0.47	-0.69	1.52	-0.51	-0.47
3	-0.03	0.41	-1.2	1.89*	-0.91	-0.17
4	0.03	0.26	-1.38	1.69	-1.32	-0.38
5	-0.26	-0.06	-1.2	1.24	-1.33	-0.58
6	-0.44	-0.4	-1.01	1.03	-1.21	-0.9
7	-0.43	-0.59	-0.86	0.88	-0.91	-1.12

TABLE A.2: BDS Statistic for Log Squared EGARCH Residuals, N=1000

m	Dow	FTSE	Nikkei	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	-0.74	-1.22	0.31	2.13*	-1.04	0.17
3	-1.26	-1.04	-0.38	2.02*	-1.71	0.16
4	-1.12	-0.59	-0.14	2.70**	-1.16	0.37
5	-1.4	-0.57	0.21	2.68*	-1.08	0.17
6	-1.44	-0.53	-0.03	2.4	-1.06	-0.04
7	-0.81	-0.45	-0.54	2.17	-1.36	-0.65
$\epsilon/\sigma = 1$						
2	-0.85	-0.83	0.04	1.93*	-0.8	0.4
3	-0.68	-0.41	-0.42	2.13**	-1.61	0.11
4	-0.24	-0.04	-0.1	2.88***	-1.42	0.16
5	-0.22	0.17	0.35	3.00***	-1.02	0.27
6	-0.27	0.13	0.34	3.10***	-0.71	0.37
7	-0.34	-0.01	0.52	3.04**	-0.49	0.32
$\epsilon/\sigma = 1.5$						
2	-0.3	-0.57	-0.04	1.97*	-0.79	0.93
3	0.09	0.3	-0.3	2.39**	-1.28	0.81
4	0.49	0.66	-0.29	2.79***	-1.11	0.83
5	0.43	0.66	-0.07	2.73**	-0.75	0.82
6	0.28	0.41	-0.05	2.74**	-0.38	0.83
7	0.12	0.17	0.06	2.71**	-0.15	0.84
$\epsilon/\sigma = 2$						
2	0.14	-0.47	0.05	1.97*	-0.23	1.80*
3	0.59	0.79	0.01	2.43**	-0.56	1.66
4	0.81	0.95	0.07	2.55**	-0.52	1.49
5	0.68	0.75	0.26	2.40**	-0.3	1.31
6	0.49	0.5	0.23	2.29**	-0.07	1.2
7	0.33	0.32	0.33	2.30**	0.07	1.17

TABLE A.3: BDS Statistic for Log Squared EGARCH Residuals, N=2500

m	Dow	FTSE	Nikkei	Gre	Irl	Port
$\epsilon/\sigma = 0.5$						
2	-0.2	-0.87	-0.03	0.09	-0.28	2.16**
3	-0.91	-1	-0.12	0.57	-0.93	2.74**
4	-0.82	-0.83	0.16	1.04	-1.19	2.69**
5	-1.11	-0.61	0.99	1.41	-0.99	3.15**
6	-1.11	-0.55	0.82	1.58	-0.73	3.56**
7	-0.93	-0.38	0.79	1.56	-0.25	4.27**
$\epsilon/\sigma = 1$						
2	-0.03	-0.39	0.1	0.52	0.18	2.79***
3	-0.37	-0.43	-0.06	0.94	-0.67	3.03***
4	-0.19	-0.31	0.25	1.19	-0.87	2.94***
5	-0.22	0	0.8	1.43	-0.54	3.21***
6	-0.2	0.15	0.92	1.76*	-0.37	3.30***
7	-0.16	0.27	1.21	1.81*	-0.31	3.20***
$\epsilon/\sigma = 1.5$						
2	0.68	0.07	-0.22	1.02	0.59	3.55***
3	0.42	0.39	-0.3	1.11	-0.27	3.81***
4	0.69	0.6	-0.19	1.15	-0.48	3.80***
5	0.74	0.77	0.16	1.21	-0.31	3.93***
6	0.76	0.89	0.28	1.55	-0.25	3.88***
7	0.77	1.09	0.48	1.6	-0.28	3.71***
$\epsilon/\sigma = 2$						
2	1.83*	0.11	-0.67	1.45	1	3.08***
3	1.69	0.56	-0.59	1.27	0.37	3.58***
4	1.93*	0.74	-0.64	1.18	0.29	3.78***
5	1.97*	0.72	-0.41	1.11	0.39	3.94***
6	1.86*	0.8	-0.27	1.26	0.33	3.84***
7	1.76*	0.99	-0.09	1.19	0.28	3.63***

TABLE A.4: IAAFT Surrogate p-values for the nonlinear prediction error (NLPE), algorithmic complexity (Comp), Pearson's cumulative mutual information (Pear) and cumulative mutual information (Mut). 1000 surrogates are used in testing. 10% significance is highlighted in yellow, 5% in orange and 1% in red.

N	Dow Jones				FTSE				NIKKEI			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.16	0.05	0.63	0.01	0.13	0.52	0.56	0.58	0.27	0.47	0.80	0.01
1000	0.16	0.96	0.77	0.42	0.27	0.82	0.72	0.94	0.41	0.64	0.41	0.85
2500	0.43	0.73	0.53	0.10	0.65	0.10	0.66	0.03	0.69	0.34	0.97	0.02
N	Greece				Ireland				Portugal			
	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut	NLPE	Comp	Pear	Mut
500	0.19	0.99	0.76	0.58	0.04	0.52	0.97	0.73	1.00	0.16	0.01	0.68
1000	0.71	0.71	0.77	0.73	0.24	0.83	0.70	0.52	0.96	0.35	0.13	0.36
2500	0.93	0.66	0.61	0.99	0.55	0.37	0.23	0.51	0.59	0.97	0.85	0.00

Bibliography

- Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015). "Systemic risk and stability in financial networks". In: *The American Economic Review* 105.2, pp. 564–608.
- Albert, R. and Barabási, A.-L. (2002). "Statistical mechanics of complex networks". In: *Reviews of modern physics* 74.1, pp. 47–97.
- Anderson, P. W. (1972). "More is different". In: *Science* 177.4047, pp. 393–396.
- Antoniou, A. and Vorlow, C. E. (2005). "Price clustering and discreteness: is there chaos behind the noise?" In: *Physica A: Statistical Mechanics and its Applications* 348.0, pp. 389–403.
- Arthur, W. B. (1999). "Complexity and the economy". In: *science* 284.5411, pp. 107–109.
- Arthur, W. B. (2013). "Complexity economics: a different framework for economic thought". In: *Santa Fe Institute Working Paper Series* 04.12, pp. 1–24.
- Aste, T., Shaw, W, and Di Matteo, T. (2010). "Correlation structure and dynamics in volatile markets". In: *New Journal of Physics* 12.8, pp. 1–21.
- Atkinson, A. B. and Piketty, T. (2007). *Top incomes over the twentieth century: a contrast between continental european and english-speaking countries*. Oxford University Press, Oxford.
- Barberis, N. and Thaler, R. (2003). "A survey of behavioral finance". In: *Handbook of the Economics of Finance* 1.B, pp. 1053–1128.
- Barnett, W., Medio, A., and Serletis, A. (2015). "Nonlinear and Complex Dynamics in Economics". In: *Macroeconomic Dynamics* 18.8, pp. 1749–1779.
- Bartlett, M. (1946). "On the theoretical specification of the sampling properties of autocorrelated time series." In: *Supplement to the Journal of the Royal Statistical Society* 8.1, pp. 27–41.

- Barunik, J. and Kukacka, J. (2015). "Realizing stock market crashes: stochastic cusp catastrophe model of returns under time-varying volatility". In: *Quantitative Finance* 15.6, pp. 959–973.
- BCBS (2004). *International convergence of capital measurement and capital standards: A revised framework*. Bank for International Settlements, Basel. [Online]. URL: <http://www.bis.org/publ/bcbs128.htm>.
- Bernanke, B., Gertler, M., and Gilchrist, S. (1994). "The financial accelerator and the flight to quality". In: *The Review of Economics and Statistics* 78.1, pp. 1–15.
- Biggs, R., Carpenter, S. R., and Brock, W. A. (2009). "Turning back from the brink: Detecting an impending regime shift in time to avert it". In: *Proceedings of the National Academy of Sciences* 106.3, pp. 826–831.
- Bisias, D., Flood, M. D., Lo, A. W., and Valavanis, S. (2012). "A Survey of Systemic Risk Analytics". In: *Annual Review of Financial Economics* 4.1, pp. 255–296.
- Boettiger, C. and Hastings, A. (2012). "Early warning signals and the prosecutor's fallacy". In: *Proceedings of the Royal Society B: Biological Sciences* 279.1748, pp. 4734–4739.
- Boettiger, C., Ross, N., and Hastings, A. (2013). "Early warning signals: the charted and uncharted territories". In: *Theoretical Ecology* 6.3, pp. 255–264.
- Bollerslev, T. (1986). "Generalized autoregressive conditional heteroskedasticity". In: *Journal of econometrics* 31.3, pp. 307–327.
- Borghesi, C., Marsili, M., and Miccichè, S. (2007). "Emergence of time-horizon invariant correlation structure in financial returns by subtraction of the market mode". In: *Physical Review E* 76.2, pp. 1–12.
- Boss, M., Elsinger, H., Summer, M., and Thurner, S. (2004). "Network topology of the interbank market". In: *Quantitative Finance* 4.6, pp. 677–684.
- Brock, W. A. and Carpenter, S. R. (2006). "Variance as a Leading Indicator of Regime Shift in Ecosystem Services". In: *Ecology and Society* 11.2. [online]. URL: <http://www.ecologyandsociety.org/vol11/iss2/art9/>.
- Brock, W. A. and Carpenter, S. R. (2010). "Interacting regime shifts in ecosystems: implication for early warnings". In: *Ecological Monographs* 80.3, pp. 353–367.

- Brock, W. A., Scheinkman, J. A., Dechert, W. D., and LeBaron, B. (1996). "A test for independence based on the correlation dimension". In: *Econometric Reviews* 15.3, pp. 197–235.
- Brock, W. A., Dechert, W. D., and Scheinkman, J. (1987). "A test for independence based on the correlation dimension". In: *SSRI Working Paper, Department of Economics, University of Wisconsin-Madison* 8702.
- Brock, W. A. and Hommes, C. H. (1998). "Heterogeneous beliefs and routes to chaos in a simple asset pricing model". In: *Journal of Economic dynamics and Control* 22.8, pp. 1235–1274.
- Brock, W. A., Hommes, C. H., and Wagener, F. O. (2005). "Evolutionary dynamics in markets with many trader types". In: *Journal of Mathematical Economics* 41.1, pp. 7–42.
- Brock, W. A. and Sayers, C. L. (1988). "Is the business cycle characterized by deterministic chaos?" In: *Journal of Monetary Economics* 22.1, pp. 71–90.
- Brooks, C. (1996). "Testing for non-linearity in daily sterling exchange rates". In: *Applied Financial Economics* 6.4, pp. 307–317.
- Cajueiro, D. O. and Tabak, B. M. (2008). "The role of banks in the Brazilian Interbank Market: Does bank type matter?" In: *Physica A: Statistical Mechanics and its Applications* 387.27, pp. 6825–6836.
- Calvo, G. A. (1988). "Servicing the public debt: The role of expectations". In: *The American Economic Review*, pp. 647–661.
- Campbell, J. Y., Lo, A. W.-C., and MacKinlay, A. C. (1997). *The econometrics of financial markets*. Vol. 2. Princeton University press, New Jersey.
- Caporale, G. M., Ntantmis, C., Theologos, P., and Pittis, N. (2005). "The BDS Test as a Test for the Adequacy of a GARCH(1,1) Specification: A Monte Carlo Study". In: *Journal of Financial Econometrics* 3.2, pp. 282–309.
- Carpenter, S. R. and Brock, W. A. (2006). "Rising variance: a leading indicator of ecological transition". In: *Ecology Letters* 9.3, pp. 311–318.
- Carpenter, S. R., Cole, J. J., Pace, M. L., Batt, R., Brock, W. A., Cline, T., Coloso, J., Hodgson, J. R., Kitchell, J. F., Seekell, D. A., Smith, L., and Weidel, B. (2011). "Early Warnings of Regime Shifts: A Whole-Ecosystem Experiment". In: *Science* 332.6033, pp. 1079–1082.

- Caruana, J. (2010). *Systemic risk: how to deal with it*. Bank for International Settlements. [Online]. URL: <http://www.bis.org/publ/othp08.htm>.
- Chen, Q., Goldstein, I., and Jiang, W. (2010). "Payoff complementarities and financial fragility: Evidence from mutual fund outflows". In: *Journal of Financial Economics* 97.2, pp. 239–262.
- Chiarella, C. and He, X.-Z. (2002). "Heterogeneous beliefs, risk and learning in a simple asset pricing model". In: *Computational Economics* 19.1, pp. 95–132.
- Chisholm, R. A. and Filotas, E. (2009). "Critical slowing down as an indicator of transitions in two-species models". In: *Journal of Theoretical Biology* 257.1, pp. 142–149.
- Coelho, R, Hutzler, S, Repetowicz, P, and Richmond, P (2007). "Sector analysis for a FTSE portfolio of stocks". In: *Physica A: Statistical Mechanics and its Applications* 373, pp. 615–626.
- Colander, D., Goldberg, M., Haas, A., Juselius, K., Kirman, A., Lux, T., and Sloth, B. (2009). "The financial crisis and the systemic failure of the economics profession". In: *Critical Review* 21.2-3, pp. 249–267.
- Cont, R. (2001). "Empirical properties of asset returns: stylized facts and statistical issues". In: *Quantitative Finance* 1, pp. 223–236.
- Cover, T. M. and Thomas, J. A. (1991). *Elements of information theory*. John Wiley & Sons, New Jersey.
- Creedy, J. and Martin, V. (1994). *Chaos and non-linear models in economics: theory and applications*. Edward Elgar Publishing, Brookfield.
- Cvitanovic, P., Artuso, R., Mainieri, R., Tanner, G., and Vattay, G. (2005). *Chaos: Classical and quantum*. ChaosBook.org (Niels Bohr Institute), Copenhagen.
- Dakos, V., Nes, E. H. v., D'Odorico, P., and Scheffer, M. (2012a). "Robustness of variance and autocorrelation as indicators of critical slowing down". In: *Ecology* 93.2, pp. 264–271.
- Dakos, V. and Lahti, L. (2013). *R Early Warning Signals Toolbox*.
- Dakos, V., Scheffer, M., Nes, E. H. van, Brovkin, V., Petoukhov, V., and Held, H. (2008). "Slowing down as an early warning signal for abrupt climate change". In: *Proceedings of the National Academy of Sciences* 105.38, pp. 14308–14312.

- Dakos, V., Carpenter, S. R., Brock, W. A., Ellison, A. M., Guttal, V., Ives, A. R., Kéfi, S., Livina, V., Seekell, D. A., Nes, E. H. van, and Scheffer, M. (2012b). "Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data". In: *PLoS ONE* 7.7, pp. 1–20.
- Danielsson, J. (2002). "The emperor has no clothes: Limits to risk modelling". In: *Journal of Banking and Finance* 26.7, pp. 1273–1296.
- Das, A. and Das, P. (2007). "Chaotic analysis of the foreign exchange rates". In: *Applied mathematics and computation* 185.1, pp. 388–396.
- Day, R. H. et al. (1994). *Complex Economic Dynamics-Vol. 1: An Introduction to Dynamical Systems and Market Mechanisms*. The MIT Press, Cambridge.
- De Grauwe, P. (2011). "Managing a fragile Eurozone". In: *CESifo Forum*. Vol. 12. 2, pp. 1–40.
- De Masi, G. (2009). "Empirical Analysis of the Architecture of the Interbank Market and Credit Market Using Network Theory". In: *Networks, Topology and Dynamics*. Springer, Berlin, pp. 241–256.
- De Nicolo, G. and Kwast, M. L. (2002). "Systemic risk and financial consolidation: Are they related?" In: *Journal of Banking and Finance* 26.5, pp. 861–880.
- Diamond, D. W. and Dybvig, P. H. (1983). "Bank Runs, Deposit Insurance, and Liquidity". In: *The Journal of Political Economy* 91.3, pp. 401–419.
- Diebold, F. X., Schuermann, T., and Stroughair, J. D. (2000). "Pitfalls and opportunities in the use of extreme value theory in risk management". In: *The Journal of Risk Finance* 1.2, pp. 30–35.
- Diks, C., Hommes, C., and Wang, J. (2015). "Critical Slowing Down as Early Warning Signals for Financial Crises". In: *Center for Nonlinear Dynamics in Economics and Finance Working Paper Series* 15.4, pp. 1–31.
- Ding, Z., Granger, C. W., and Engle, R. F. (1993). "A long memory property of stock market returns and a new model". In: *Journal of empirical finance* 1.1, pp. 83–106.
- Dorogovtsev, S. N., Goltsev, A. V., and Mendes, J. F. (2008). "Critical phenomena in complex networks". In: *Reviews of Modern Physics* 80.4, pp. 1275–1335.
- Durlauf, S. N. (2005). "Complexity and empirical economics". In: *The Economic Journal* 115.504, pp. 225–243.

- Easley, D. and Kleinberg, J. (2010). *Networks, crowds, and markets: Reasoning about a highly connected world*. Cambridge University Press, Cambridge.
- Eidelson, R. J. (1997). "Complex adaptive systems in the behavioral and social sciences." In: *Review of General Psychology* 1.1, p. 42.
- Engle, R. F. (1982). "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation". In: *Econometrica: Journal of the Econometric Society*, pp. 987–1007.
- Farmer, J. D., Gallegati, M., Hommes, C., Kirman, A., Ormerod, P., Cincotti, S., Sánchez, A., and Helbing, D. (2012). "A complex systems approach to constructing better models for managing financial markets and the economy". In: *The European Physical Journal Special Topics* 214.1, pp. 295–324.
- Fontana, M. (2010). "Can neoclassical economics handle complexity? The fallacy of the oil spot dynamic". In: *Journal of Economic Behavior & Organization* 76.3, pp. 584–596.
- Frankel, J. A. and Froot, K. A. (1990). "Chartists, fundamentalists, and trading in the foreign exchange market". In: *The American Economic Review* 80.2, pp. 181–185.
- Fraser, A. and Swinney, H. (1986). "Independent coordinates for strange attractors from mutual information". In: *Physical Review A* 33.2, pp. 1134–1140.
- Fujiwara, Y., Aoyama, H., Ikeda, Y., Iyetomi, H., and Souma, W. (2009). "Structure and temporal change of the credit network between banks and large firms in Japan". In: *Economics: The Open-Access, Open-Assessment E-Journal* 3.7, pp. 1–18.
- Gabaix, X. (2009). "Power Laws in Economics and Finance". In: *Annual Review of Economics* 1.1, pp. 255–294.
- Gabaix, X. and Landier, A. (2008). "Why has CEO pay increased so much?" In: *Quarterly Journal of Economics* 123.1, pp. 49–100.
- Gabaix, X., Gopikrishnan, P., Plerou, V., and Stanley, E. (2007). "A unified econophysics explanation for the power-law exponents of stock market activity". In: *Physica A: Statistical Mechanics and its Applications* 382.1, pp. 81–88.
- Gai, P. and Kapadia, S. (2010). "Contagion in financial networks". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 466.2120, pp. 2401–2423.

- Gaunersdorfer, A. (2000). "Endogenous fluctuations in a simple asset pricing model with heterogeneous agents". In: *Journal of Economic Dynamics and Control* 24.5, pp. 799–831.
- Gillespie, C. S. (2015). "Fitting Heavy Tailed Distributions: The powerLaw Package". In: *Journal of Statistical Software* 64.02, pp. 1–16.
- Gopikrishnan, P., Plerou, V., Amaral, L. A. N., Meyer, M., and Stanley, H. E. (1999). "Scaling of the distribution of fluctuations of financial market indices". In: *Physical Review E* 60.5, pp. 5305–5316.
- Gopikrishnan, P., Plerou, V., Gabaix, X., and Stanley, H. E. (2000). "Statistical properties of share volume traded in financial markets". In: *Physical Review E* 62.4, pp. 4493–4496.
- Grassberger, P. and Procaccia, I. (1983). "Characterization of Strange Attractors". In: *Physical Review Letters* 50.5, pp. 346–349.
- Gros, D. (2012). "A simple model of multiple equilibria and default". In: *CEPS Working Document* 2012.366, pp. 1–21.
- Guttal, V. and Jayaprakash, C. (2008). "Changing skewness: an early warning signal of regime shifts in ecosystems". In: *Ecology Letters* 11.5, pp. 450–460.
- Guttal, V., Raghavendra, S., Goel, N., and Hoarau, Q. (2016). "Lack of Critical Slowing Down Suggests that Financial Meltdowns Are Not Critical Transitions, yet Rising Variability Could Signal Systemic Risk". In: *PloS one* 11.1, pp. 1–20.
- Haldane, A. (2009). "Why banks failed the stress test". In: *BIS Review* 18.
- Hanias, M. P., Curtis, G, and Thallasinos, J. (2007). "Non-linear dynamics and chaos: The case of the price indicator at the Athens Stock Exchange," in: *International Research Journal of Finance and Economics* 11, pp. 154–63.
- Hirota, M., Holmgren, M., Van Nes, E. H., and Scheffer, M. (2011). "Global Resilience of Tropical Forest and Savanna to Critical Transitions". In: *Science* 334.6053, pp. 232–235.
- Hirshleifer, D. and Hong Teoh, S. (2003). "Herd behaviour and cascading in capital markets: A review and synthesis". In: *European Financial Management* 9.1, pp. 25–66.
- Holt, R. P., Rosser Jr, J. B., and Colander, D. (2011). "The complexity era in economics". In: *Review of Political Economy* 23.3, pp. 357–369.

- Hommes, C. H. (2001). "Financial markets as nonlinear adaptive evolutionary systems". In: *Quantitative Finance* 1, pp. 149–167.
- Hsieh, D. A. (1991). "Chaos and Nonlinear Dynamics: Application to Financial Markets". In: *The Journal of Finance* 46.5, pp. 1839–1877.
- Huang, Y, Saleur, H, and Sornette, D (2000). "Reexamination of log periodicity observed in the seismic precursors of the 1989 Loma Prieta earthquake". In: *Journal of Geophysical Research: Solid Earth* 105.B12, pp. 28111–28123.
- IMF, BIS, and FSB (2009). "Guidance to assess the systemic importance of financial institutions, markets and instruments: initial considerations". In: *Report to G20 finance ministers and governors*. [Online]. URL: <http://www.bis.org/publ/othp07.htm>.
- Iori, G., De Masi, G., Precup, O. V., Gabbi, G., and Caldarelli, G. (2008). "A network analysis of the Italian overnight money market". In: *Journal of Economic Dynamics and Control* 32.1, pp. 259–278.
- Jang, W., Lee, J., and Chang, W. (2011). "Currency crises and the evolution of foreign exchange market: Evidence from minimum spanning tree". In: *Physica A: Statistical Mechanics and its Applications* 390.4, pp. 707–718.
- Jeanne, O. (1997). "Are currency crises self-fulfilling?: A test". In: *Journal of International Economics* 43.3, pp. 263–286.
- Jiang, Z.-Q., Zhou, W.-X., Sornette, D., Woodard, R., Bastiaensen, K., and Cauwels, P. (2010). "Bubble diagnosis and prediction of the 2005–2007 and 2008–2009 Chinese stock market bubbles". In: *Journal of economic behavior & organization* 74.3, pp. 149–162.
- Johansen, A, Saleur, H, and Sornette, D (2000). "New evidence of earthquake precursory phenomena in the 17 January 1995 Kobe earthquake, Japan". In: *The European Physical Journal B-Condensed Matter and Complex Systems* 15.3, pp. 551–555.
- Johansen, A, Sornette, D, and Ledoit, O (1999). "Predicting financial crashes using discrete scale invariance". In: *Journal of Risk* 1.4, pp. 5–32.
- Johansen, A., Ledoit, O., and Sornette, D. (2000). "Crashes as critical points". In: *International Journal of Theoretical and Applied Finance* 3.02, pp. 219–255.
- Johansen, A. and Sornette, D. (2000). "Critical ruptures". In: *The European Physical Journal B-Condensed Matter and Complex Systems* 18.1, pp. 163–181.

- Jondeau, E., Poon, S.-H., and Rockinger, M. (2007). *Financial modeling under non-Gaussian distributions*. Springer, London.
- Kantz, H. and Schreiber, T. (2004). *Nonlinear time series analysis*. 2nd ed. Cambridge University Press, Cambridge.
- Kanzler, L. (1999). "Very Fast and Correctly Sized Estimation of the BDS Statistic". In: *SSRN Electronic Journal*. [Online]. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=151669.
- Karmin, S. B. and Pounder DeMarco, L. (2010). *How did a Domestic Housing Slump turn into a Global Financial Crisis*. International Finance Discussion Papers. Board of Governors of the Federal Reserve System.
- Kaya, H. (2015). "Eccentricity in asset management". In: *The Journal of Network Theory in Finance* 1.1, pp. 1–32.
- Kennel, M. B., Brown, R., and Abarbanel, H. D. (1992). "Determining embedding dimension for phase-space reconstruction using a geometrical construction". In: *Physical review A* 45.6, pp. 3403–3411.
- Kennett, D. Y., Tumminello, M., Madi, A., Gur-Gershgoren, G., Mantegna, R., and Ben-Jacob, E. (2010). "The Dominating Clasp of the Financial Sector Revealed by Parital Correlation Analysis of the Stock Market". In: *Plos One* 12.5, pp. 1–14.
- Kenny, O., Killeen, N., and Moloney, K. (2016). "Network analysis using EMIR credit default swap data". In: *IFC Bulletin Chapters* 41, pp. 1–20.
- Keynes, J. M. (1936). *The general theory of interest, employment and money*. Macmillan, London.
- Kéfi, S., Dakos, V., Scheffer, M., Van Nes, E. H., and Rietkerk, M. (2013). "Early warning signals also precede non-catastrophic transitions". In: *Oikos* 122.5, pp. 641–648.
- Kleinen, T., Held, H., and Petschel-Held, G. (2003). "The potential role of spectral properties in detecting thresholds in the Earth system: application to the thermohaline circulation". In: *Ocean Dynamics* 53.2, pp. 53–63.
- Kruskal, J. B. (1956). "On the shortest spanning subtree of a graph and the traveling salesman problem". In: *Proceedings of the American Mathematical society* 7.1, pp. 48–50.

- Kuehn, C. (2011). "A mathematical framework for critical transitions: Bifurcations, fast–slow systems and stochastic dynamics". In: *Physica D: Nonlinear Phenomena* 240.12, pp. 1020–1035.
- Kugiumtzis, D (2002b). "Surrogate data test on time series". In: *Modelling and Forecasting Financial Data*. Springer, pp. 267–282.
- Kugiumtzis, D (2002a). "Statically transformed autoregressive process and surrogate data test for nonlinearity". In: *Physical Review E* 66.2, p. 025201.
- Kugiumtzis, D. (2008). "Evaluation of Surrogate and Bootstrap Tests for Nonlinearity in Time Series". In: *Studies in Nonlinear Dynamics & Econometrics* 12.1, pp. 1–24.
- Kugiumtzis, D. and Tsimpiris, A. (2010). "Measures of Analysis of Time Series (MATS): A MATLAB Toolkit for Computation of Multiple Measures on Time Series Data Bases". In: *Journal of Statistical Software* 33.5, pp. 1–30.
- Kyrtsou, C. and Serletis, A. (2006). "Univariate tests for nonlinear structure". In: *Journal of Macroeconomics* 28.1, pp. 154–168.
- Lautier, D. and Raynaud, F. (2013). "Systemic Risk and Complex Systems: A Graph-Theory Analysis". In: *Econophysics of Systemic Risk and Network Dynamics*. Springer, pp. 19–37.
- LeBaron, B. (1994). "Chaos and nonlinear forecastability in economics and finance". In: *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 348.1688, pp. 397–404.
- Lee, S., Kim, M. J., Lee, S. Y., Kim, S. Y., and Ban, J. H. (2013). "The effect of the subprime crisis on the credit risk in global scale". In: *Physica A: Statistical Mechanics and its Applications* 392.9, pp. 2060–2071.
- Lempel, A. and Ziv, J. (1976). "On the complexity of finite sequences". In: *IEEE Transactions on information theory* 22.1, pp. 75–81.
- Lenton, T. M., Livina, V. N., Dakos, V., Nes, E. H. van, and Scheffer, M. (2012). "Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 370.1962, pp. 1185–1204.
- Lenton, T., Held, H., Kriegler, E., Hall, J., Lucht, W., Rahmstorf, S., and Schellnhuber, H. (2008). "Tipping elements in the Earth's climate system". In: *Proceedings of the National Academy of Sciences* 105.6, pp. 1786–1793.

- Lima, P. J. F de (1996). "Nuisance parameter free properties of correlation integral based statistics". In: *Econometric Reviews* 15.3, pp. 237–259.
- Liu, Y., Gopikrishnan, P., Stanley, H. E., et al. (1999). "Statistical properties of the volatility of price fluctuations". In: *Physical Review E* 60.2, pp. 1390–1400.
- Lucas, R. E. (1976). "Econometric policy evaluation: A critique". In: *Carnegie-Rochester conference series on public policy*. Vol. 1. Elsevier, pp. 19–46.
- Lux, T. and Alfarano, S. (2016). "Financial power laws: Empirical evidence, models, and mechanisms". In: *Chaos, Solitons & Fractals* 88, pp. 3–18.
- Malkiel, B. G. and Fama, E. F. (1970). "Efficient capital markets: A review of theory and empirical work". In: *The Journal of Finance* 25.2, pp. 383–417.
- Mandelbrot, B. B. (1963). "The Variation of Certain Speculative Prices". In: *The Journal of Business* 36.4, pp. 394–419.
- Mantegna, R. N. (1999). "Hierarchical structure in financial markets". In: *The European Physical Journal B-Condensed Matter and Complex Systems* 11.1, pp. 193–197.
- Mantegna, R. N. and Stanley, H. E. (1999). *Introduction to econophysics: correlations and complexity in finance*. Cambridge University Press, Cambridge.
- May, R. M. (1972). "Will a large complex system be stable?" In: *Nature* 238, pp. 413–414.
- May, R. M., Levin, S. A., and Sugihara, G. (2008). "Complex systems: Ecology for bankers". In: *Nature* 451.7181, pp. 893–895.
- McHale, J. (2012). *A Model of Debt Sustainability and Creditworthiness in the Eurozone*. Presentation at the Economic Policy Conference of the Dublin Economics Workshop, Galway, 12th-14th October 2012.
- Merton, R. C. (1973). "Theory of rational option pricing". In: *The Bell Journal of economics and management science* 4.1, pp. 141–183.
- Miller, J. H. and Page, S. E. (2009). *Complex adaptive systems: an introduction to computational models of social life: an introduction to computational models of social life*. Princeton University Press, New Jersey.
- Minsky, H. P. (1977). "The financial instability hypothesis: an interpretation of Keynes and an alternative to "standard" theory". In: *Challenge* 20.1, pp. 20–27.
- Mishkin, F. S. (2010). "Over the cliff: From the subprime to the global financial crisis". In: *Journal of Economic Perspectives* 25.1, pp. 49–70.

- Mitchell, M. (2009). *Complexity: A guided tour*. Oxford University Press, Oxford.
- Modigliani, F. (1971). "Monetary policy and consumption". In: *Consumer spending and monetary policy: the linkages*. Boston Federal Reserve, Boston, pp. 9–84.
- Moloney, K. and Raghavendra, S. (2011). "Testing for Nonlinear Dependence in the Credit Default Swap Market". In: *Economics Research International* 708704, pp. 1–12.
- Musmeci, N., Aste, T., and Di Matteo, T. (2015). "Risk diversification: a study of persistence with a filtered correlation-network approach". In: *Journal of Network Theory in Finance* 1.1, pp. 77–98.
- Obstfeld, M. (1996). "Models of currency crises with self-fulfilling features". In: *European economic review* 40.3, pp. 1037–1047.
- Onnela, J.-P., Kaski, K., and Kertész, J. (2004). "Clustering and information in correlation based financial networks". In: *The European Physical Journal B-Condensed Matter and Complex Systems* 38.2, pp. 353–362.
- Onnela, J.-P., Chakraborti, A., Kaski, K., Kertesz, J., and Kanto, A. (2003b). "Dynamics of market correlations: Taxonomy and portfolio analysis". In: *Physical Review E* 68.5, p. 056110.
- Onnela, J.-P., Chakraborti, A., Kaski, K., and Kertesz, J. (2003a). "Dynamic asset trees and Black Monday". In: *Physica A: Statistical Mechanics and its Applications* 324.1, pp. 247–252.
- Onnela, J.-P. (2002). *Taxonomy of financial assets*. A Thesis for the Master of Science in Engineering, Helsinki University of Technology.
- Papana, A. (2009). "Nonlinear Statistical Analysis of Biological Time Series". PhD thesis. Aristotle University of Thessaloniki.
- Patro, D. K., Qi, M., and Sun, X. (2013). "A simple indicator of systemic risk". In: *Journal of Financial Stability* 9.1, pp. 105–116.
- Peltonen, T. A., Scheicher, M., and Vuillemeij, G. (2014). "The network structure of the CDS market and its determinants". In: *Journal of Financial Stability* 13, pp. 118–133.
- Plerou, V., Gopikrishnan, P., and Stanley, H. E. (2005). "Quantifying fluctuations in market liquidity: Analysis of the bid-ask spread". In: *Physical Review E* 71.4, p. 046131.

- Plerou, V. and Stanley, H. E. (2007). "Tests of scaling and universality of the distributions of trade size and share volume: Evidence from three distinct markets". In: *Physical Review E* 76.4, p. 046109.
- Pozzi, F, Di Matteo, T, and Aste, T (2013). "Spread of risk across financial markets: better to invest in the peripheries". In: *Scientific reports* 3.1665, pp. 1–7.
- Pozzi, F., Aste, T., Rotundo, G., and Di Matteo, T. (2008). "Dynamical correlations in financial systems". In: *Proc. SPIE*. Vol. 6802. 68021E.
- Prechter Jr, R. R. (2001). "Unconscious herding behavior as the psychological basis of financial market trends and patterns". In: *The Journal of Psychology and Financial Markets* 2.3, pp. 120–125.
- Prim, R. C. (1957). "Shortest connection networks and some generalizations". In: *Bell system technical journal* 36.6, pp. 1389–1401.
- Quax, R., Kandhai, D., and Sloot, P. M. A. (2013). "Information dissipation as an early-warning signal for the Lehman Brothers collapse in financial time series". In: *Scientific Reports* 3, pp. 1–7.
- Reinhart, C. M. and Rogoff, K. (2009). *This time is different: eight centuries of financial folly*. Princeton University Press, New Jersey.
- Rietkerk, M., Dekker, S. C., Ruiten, P. C. de, and Koppel, J. van de (2004). "Self-Organized Patchiness and Catastrophic Shifts in Ecosystems". In: *Science* 305.5692, pp. 1926–1929.
- Rosser, J. B. (2000). *From Catastrophe to Chaos: A General Theory of Economic Discontinuities: Mathematics, Microeconomics and Finance*. Vol. 1. Springer, Netherlands.
- Rössler, O. E. (1976). "An equation for continuous chaos". In: *Physics Letters A* 57.5, pp. 397–398.
- Ruelle, D. (1989). "The Claude Bernard Lecture, 1989. Deterministic Chaos: The Science and the Fiction". In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 427.1873, pp. 241–248.
- Sandoval, L. (2012). "Pruning a minimum spanning tree". In: *Physica A: Statistical Mechanics and its Applications* 391.8, pp. 2678–2711.
- Sauer, T., Yorke, J. A., and Casdagli, M. (1991). "Embedology". In: *Journal of statistical Physics* 65.3, pp. 579–616.

- Scheffer, M., Carpenter, S., Foley, J. A., Folke, C., and Walker, B. (2001). "Catastrophic shifts in ecosystems". In: *Nature* 413.6856, pp. 591–596.
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., Nes, E. H. van, Rietkerk, M., and Sugihara, G. (2009). "Early-warning signals for critical transitions". In: *Nature* 461.7260, pp. 53–59.
- Scheinkman, J. A. and LeBaron, B. (1989). "Nonlinear dynamics and stock returns". In: *Journal of Business* 62.3, pp. 311–337.
- Schreiber, T. (1998). "Constrained randomization of time series data". In: *Physical Review Letters* 80.10, pp. 2105–2108.
- Schreiber, T. and Schmitz, A. (1996). "Improved Surrogate Data for Nonlinearity Tests". In: *Physical Review Letters* 77.4, pp. 635–638.
- Schreiber, T. and Schmitz, A. (1997). "Discrimination power of measures for nonlinearity in a time series". In: *Physical Review E* 55.5, pp. 5443–5447.
- Schreiber, T. and Schmitz, A. (2000). "Surrogate time series". In: *Physica D: Nonlinear Phenomena* 142.3–4, pp. 346–382.
- Schweitzer, F., Fagiolo, G., Sornette, D., Vega-Redondo, F., Vespignani, A., and White, D. R. (2009). "Economic networks: The new challenges". In: *science* 325.5939, pp. 422–425.
- Shannon, C. E. (1948). "A mathematical theory of communication". In: *Bell Systems Technical Journal* 27.4, pp. 623–666.
- Shiller Robert, J (2002). "Bubbles, human judgment, and expert opinion". In: *Financial Analysts Journal* 58.3, pp. 18–26.
- Shiller, R. J. (1995). "Conversation, information, and herd behavior". In: *The American Economic Review* 85.2, pp. 181–185.
- Shumway, R. and Stoffer, D. (2006). *Time Series Analysis And Its Applications: With R Examples*. Springer, Switzerland.
- Simon, H. A. (1957). *Models of man; social and rational*. Wiley, New Jersey.
- Simon, H. A. (1962). "The Architecture of Complexity". In: *Proceedings of the American Philosophical Society* 106.6, pp. 467–482.
- Simon, H. A. (2000). "Bounded rationality in social science: Today and tomorrow". In: *Mind & Society* 1.1, pp. 25–39.

- Small, M. (2005). *Applied Nonlinear Time Series Analysis Applications. in Physics, Physiology, and Finance*. Vol. 52. World Scientific, Singapore.
- Small, M. and Tse, C. K. (2003). "Determinism in Financial Time Series". In: *Studies in Nonlinear Dynamics & Econometrics* 7.3, p. 5.
- Small, M., Yu, D., and Harrison, R. (2001). "Surrogate Test for Pseudoperiodic Time Series Data". In: *Physical Review Letters* 87.18, p. 4.
- Solomon, S. and Golo, N. (2014). "Microeconomic structure determines macroeconomic dynamics: Aoki defeats the representative agent". In: *Journal of Economic Interaction and Coordination*, pp. 1–26.
- Soofi, A. S. and Galka, A. (2003). "Measuring the complexity of currency markets by fractal dimension analysis". In: *International Journal of Theoretical and Applied Finance* 6.06, pp. 553–563.
- Soofi, A. S., Galka, A., Li, Z., Zhang, Y., and Hui, X. (2014). "Applications of Methods and Algorithms of Nonlinear Dynamics in Economics and Finance". In: *Complexity in Economics: Cutting Edge Research*. Springer, Switzerland, pp. 1–30.
- Sornette, D. (2006). *Critical phenomena in natural sciences: chaos, fractals, selforganization and disorder: concepts and tools*. Springer, Heidelberg.
- Sornette, D. (2014). "Physics and financial economics (1776–2014): puzzles, Ising and agent-based models". In: *Reports on Progress in Physics* 77.6, p. 062001.
- Sornette, D., Johansen, A., and Bouchaud, J.-P. (1996). "Stock market crashes, precursors and replicas". In: *Journal de Physique I* 6.1, pp. 167–175.
- Sornette, D., Woodard, R., and Zhou, W.-X. (2009). "The 2006–2008 oil bubble: Evidence of speculation, and prediction". In: *Physica A: Statistical Mechanics and its Applications* 388.8, pp. 1571–1576.
- Sornette, D., Woodard, R., Yan, W., and Zhou, W.-X. (2013). "Clarifications to questions and criticisms on the Johansen–Ledoit–Sornette financial bubble model". In: *Physica A: Statistical Mechanics and its Applications* 392.19, pp. 4417–4428.
- Sprott, J. C. (2003). *Chaos and time-series analysis*. Oxford University Press, Oxford.
- Stanley, H., Amaral, L., Buldyrev, S. V., Gopikrishnan, P., Plerou, V., and Salinger, M. (2002). "Self-organized complexity in economics and finance". In: *Proceedings of the National Academy of Sciences* 99.1, pp. 2561–2565.

- Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering*. Perseus: Westview Press, Cambridge.
- Sussmann, H. J. and Zahler, R. S. (1978). "A critique of applied catastrophe theory in the behavioral sciences". In: *Behavioral science* 23.4, pp. 383–389.
- Takens, F. et al. (1981). "Detecting strange attractors in turbulence". In: *Lecture notes in mathematics* 898.1, pp. 366–381.
- Taleb, N. N. (2007). *The black swan: The impact of the highly improbable fragility*. Penguin Random House, New York.
- Tan, J. P. L. and Cheong, S. S. A. (2014). "Critical slowing down associated with regime shifts in the US housing market". In: *The European Physical Journal B* 87.2, pp. 1–10.
- Taylor, J. B. (2009). "ECONOMIC POLICY AND THE FINANCIAL CRISIS: AN EMPIRICAL ANALYSIS OF WHAT WENT WRONG". In: *Critical Review* 21.2-3, pp. 341–364.
- Taylor, M. P. and Allen, H. (1992). "The use of technical analysis in the foreign exchange market". In: *Journal of international Money and Finance* 11.3, pp. 304–314.
- Theiler, J., Eubank, S., Longtin, A., Galdrikian, B., and Doynne Farmer, J. (1992). "Testing for nonlinearity in time series: the method of surrogate data". In: *Physica D: Nonlinear Phenomena* 58.1–4, pp. 77–94.
- Thom, R. (1972). *Structural Stability and Morphogenesis*. Benjamin, New York.
- Tversky, A. and Kahneman, D. (1975). "Judgment under uncertainty: Heuristics and biases". In: *Utility, probability, and human decision making*. Springer, pp. 141–162.
- Wiliński, M, Sienkiewicz, A, Gubiec, T., Kutner, R, and Struzik, Z. (2013). "Structural and topological phase transitions on the German Stock Exchange". In: *Physica A: Statistical Mechanics and its Applications* 392.23, pp. 5963–5973.
- Wissel, C. (1984). "A Universal Law of the Characteristic Return Time near Thresholds". In: *Oecologia* 65.1, pp. 101–107.
- Wooldridge, J. (2012). *Introductory Econometrics: A Modern Approach*. Cengage, Boston.
- Zeeman, E. C. (1974). "On the unstable behaviour of stock exchanges". In: *Journal of Mathematical Economics* 1.1, pp. 39–49.
- Zeeman, E. C. (1977). *Catastrophe theory: Selected papers, 1972–1977*. Addison-Wesley, MA.

- Zhang, Y., Soofi, A. S., and Wang, S. (2011). "Testing for nonlinearity of exchange rates: an information-theoretic approach". In: *Journal of Economic Studies* 38.6, pp. 637–657.
- Zhao, Y., Small, M., Coward, D., Howell, E., Zhao, C., Ju, L., and Blair, D. (2006). "Identifying deterministic signals in simulated gravitational wave data: algorithmic complexity and the surrogate data method". In: *Classical and Quantum Gravity* 23.5, pp. 1801–1814.
- Zhou, W.-X. and Sornette, D. (2003). "2000–2003 real estate bubble in the UK but not in the USA". In: *Physica A: Statistical Mechanics and its Applications* 329.1, pp. 249–263.
- Zhou, W.-X. and Sornette, D. (2006). "Is there a real-estate bubble in the US?" In: *Physica A: Statistical Mechanics and its Applications* 361.1, pp. 297–308.
- Zigrand, J.-P. (2014). "Systems and Systemic Risk in Finance and Economics". In: *Systemic Risk Centre Special Paper* 3.1, pp. 1–77.