

# **Influence in economic and political systems: A network scientific approach**

*Stefano Gurciullo*

A dissertation submitted in partial fulfillment  
of the requirements for the degree of  
**Doctor of Philosophy**  
of  
**University College London.**

School of Public Policy & Financial Computing Centre  
University College London

December 21, 2016



I, Stefano Gurciullo, confirm that the work presented in this thesis is my own. Where information has been derived from other sources, I confirm that this has been indicated in the work.



# Abstract

Complex social systems thrive by exchanging information and resources. By means of the exchange, some actors in the system are able to at least partially determine the behaviour of another actor, thereby influencing it. Both the information exchange process and the degree of actors influence are latent, unobserved phenomena in many instances of real-world systems. This thesis presents a framework that intends to unearth the two hidden properties. It does so by introducing a Network Inference and Influence Framework (NIIF), which makes use of graph-based methods to derive a latent network in a social system, and measure the influence of its elements. The framework is applied on three case studies where the latency problem translates into research questions with importance for public policy making. The first case study uses NIIF to estimate the latent network of interdependency across financial institutions, and measures the extent to which a bank may negatively influence the system after an economic distress. In the second case study, a network of information diffusion is extracted from House of Commons parliamentary debates, testing the relation between the resulted metric of influence and speakers positions in government. The last case study builds a network of semantic and ideological affinity across UN General Assembly members, showing how graph-based methods can detect global political change. The thesis concludes with a discussion of potential future usages of the framework, as well as ameliorations.



# Acknowledgements

Every kid dreams of doing marvellous things in the future. Among my goals, there was a PhD. It meant, to the younger me, pushing the boundaries of knowledge, contributing to the intellectual and technological advancement of society. Many years later, the feeling has surprisingly remained intact. What I learnt through the journey is that leaps forward in knowledge can only happen with the support of others. And I have many to thank. Slava Mikhaylov has truly been the best supervisor I could wish for. He pointed me to the opportunity that financed my doctorate, and guided me constantly throughout my research efforts, always sharing his feedback, opinions, and above all his excitement. I thank him for having introduced me to the beautiful idea that text can be used as data, something which has already started to disrupt many disciplines and practices. I am grateful to my co-supervisor Robert Smith, someone with whom I could share my passion for complexity science, evolution and uncertainty. When my mind was wandering too far from my PhD topic, he was there helping me setting back to the right track. I had the privilege to work alongside great minds. Stefano Battiston and Marco D'Errico at the University of Zurich, and Guido Caldarelli at IMT Lucca gave me the unique possibility to help in the fight against financial crises. I thank them for the data and teamwork that led to the development of the first case study of this thesis. With Alexander Herzog, Peter John and Slava Mikhaylov I ventured myself in the study of public policy through natural language processing, which inspired the work behind Chapter 5 and 6 of this thesis. I could have not done any of this work without financial independence. I am deeply grateful to the UCL Financial Computing Doctoral Training Centre and the UK Engineering and Physical Sciences Research Council for their financial support, and to Philip Treleaven, director of the Centre, who makes sure that many young bright minds can pursue research. The UCL School of Public Policy has been the ideal haven to develop my research. I thank the institution for the financial support, and the faculty and colleagues there, who provided me with precious feedback and a great time during the construction of this work. In Summer 2015 I was privileged to spent time at the Santa Fe Institute, the legendary research centre that hosts some of the greatest minds doing ground-breaking work in the study of

complex systems. I am deeply grateful to its faculty and the fellow PhD candidates with whom I shared the experience. What I learnt there has inspired much of my work on transfer entropy, and infused in my brain many more ideas to be explored beyond this thesis. I acknowledge the use of the UCL Legion High Performance Computing Facility (Legion@UCL), and the associated support services, thanks to which I was able to perform the technical tasks in Chapter 5. I am also deeply grateful to the open source community that strives to provide the best libraries for data analysis in Python, most notably scikit-learn, pandas and networkX. We take it for granted, but without their work my research, nor the one of many others, would have been severely impeded.

I thank my parents, Graziella and Nuccio, for always supporting me, whatever my goal is. They made me the person I am now. The same goes for my brother Antonio, always happy to listen to my progresses and share the same intellectual enthusiasm. Anna has relentlessly been by my side, despite everything. Of her strength and love I am debtor.

I thank all the friends and people whose name I am not mentioning, but that in one way or another contributed to the ideas and applications exposed here. All mistakes and errors are solely my responsibility, and not in any way attributable to the persons acknowledged.

Thank you all.



# Contents

<b>1</b>	<b>Introduction</b>	<b>15</b>
<b>2</b>	<b>Background</b>	<b>21</b>
2.1	Complex systems: Definition and properties . . . . .	21
2.2	Markets and politics as complex systems . . . . .	27
2.3	Complex graphs: Definition and applications . . . . .	30
<b>3</b>	<b>Methodology: The network inference and influence framework</b>	<b>37</b>
3.1	Network inference methods . . . . .	37
3.2	Network influence metrics . . . . .	53
<b>4</b>	<b>First case study: Financial distress contagion</b>	<b>63</b>
4.1	Case study background . . . . .	64
4.2	Data and network reconstruction . . . . .	70
4.3	Influence measurement specifications . . . . .	75
4.4	Experiment settings and results . . . . .	77
4.5	Practical implementations . . . . .	82
4.6	Case study discussion . . . . .	86
<b>5</b>	<b>Second case study: Influence in parliamentary debates</b>	<b>89</b>
5.1	Case study background . . . . .	90
5.2	Data and network reconstruction . . . . .	95
5.3	Analysis settings and results . . . . .	99
5.4	Practical implementations . . . . .	111
5.5	Case study discussion . . . . .	114
<b>6</b>	<b>Third case study: Global political change in the United Nations</b>	<b>117</b>
6.1	Case study background . . . . .	118
6.2	Data and network reconstruction . . . . .	122
6.3	Analysis settings and results . . . . .	127

6.4	Practical implementations . . . . .	134
6.5	Case study discussion . . . . .	139
<b>7</b>	<b>Concluding remarks</b>	<b>143</b>
	<b>Bibliography</b>	<b>147</b>

# List of Figures

2.1	Properties of a complex system. . . . .	22
2.2	Examples of undirected and directed graphs. . . . .	31
3.1	Steps to infer the topology of a graph through the fitness model. . .	40
3.2	Steps to infer a latent information flow network using transfer entropy.	48
3.3	Steps to infer the topology of a graph through discrete state neural learning. . . . .	51
3.4	Steps to obtain a model-based influence metric. . . . .	58
4.1	Frequency distributions of equity and total asset size, years 2008 and 2013. . . . .	71
4.2	Frequency distributions of total interbank assets and liabilities, years 2008 and 2013. . . . .	72
4.3	Interbank network generated through the fitness model. . . . .	74
4.4	Degree frequency of a generated interbank network. . . . .	74
4.5	Evolution of the DebtRank influence metric and default contagion for five selected banks. . . . .	78
4.6	Frequency distributions of the influence evaluated by the default contagion algorithm for node 1. . . . .	80
4.7	Frequency distributions of the influence evaluated by DebtRank for node 1. . . . .	81
4.8	Capital loss distribution and 5% VaR, CVaR for HSBC. . . . .	83
4.9	Capital loss distribution and 5% VaR, CVaR for Santander. . . . .	83
4.10	Capital loss distribution and 5% VaR, CVaR for Intesa San Paolo. .	84
4.11	Top ten most systemically influential banks, years 2008 and 2013. .	85
5.1	Network of House of Commons speakers, Human Rights Act - 2001–2005. . . . .	104
5.2	Network of House of Commons speakers, Nuclear power - 2005–2010. . . . .	105

5.3	Network of House of Commons speakers, Youth unemployment - 2010–2014. . . . .	106
5.4	Degree distributions for six networks of parliamentary speakers. . .	108
5.5	Conditional probability distribution of eigenvector centralities. . . .	110
5.6	Conditional probability distribution of weighted eigenvector centralities. . . . .	111
5.7	Most influential actors, Al-Qaeda - 2001-2005. . . . .	112
5.8	Most influential actors, Human Rights Act - 2005-2010. . . . .	113
5.9	Most influential actors, Syrian humanitarian crisis - 2010-2014. . . .	113
6.1	Number of nations and mean token frequency per document, 1970-2014. . . . .	123
6.2	UNGA General Debate inferred network, 1970. . . . .	125
6.3	UNGA General Debate inferred network, 2010. . . . .	126
6.4	Wald test statistics for each year. . . . .	130
6.5	Density index of the UNGA dynamic networks, 1970-2014. Base year = 1970. . . . .	132
6.6	Eigenvector centrality indices for USA and USSR/Russian Federation. Base year = 1970. . . . .	132
6.7	Topic-related semantic indices for education, health, nuclear weapons and Islamic terrorism. Base year = 1995. . . . .	137

# List of Tables

3.1	NIIF network inference methods and their input features. . . . .	38
3.2	Pseudocode to find parameter $z$ for the fitness model. . . . .	41
3.3	Pseudocode to generate a network ensemble. . . . .	42
3.4	Toy adjacency matrix for iterative proportional fitting. . . . .	43
3.5	Toy adjacency matrix after one round of iterative proportional fitting. . . . .	44
3.6	Pseudocode to generate a transfer entropy matrix. . . . .	49
3.7	Pseudocode for the simulation generating the model-based influence metric. . . . .	59
4.1	Change in GDP and average stock market values in selected countries (Rose and Spiegel, 2012). . . . .	64
5.1	UK House of Commons Debates - Summary statistics. . . . .	95
5.2	Selected policy themes for each governmental session. . . . .	96
5.3	Quadgrams for selected policy themes, and respective relative frequencies. . . . .	98
5.4	Summary statistics for the inferred networks of speakers. . . . .	103
5.5	Kolmogorov–Smirnov test for exponentiality results. . . . .	107
5.6	OLS models testing MP position as a predictor of influence. . . . .	109
6.1	Structural break test statistics and p-values. . . . .	129
6.2	Policy themes considered and respective key words . . . . .	135



## Chapter 1

# Introduction

On Monday, September 15 2008, Lehman Brothers, then the fourth largest investment bank in the United States, announced its intention to file for bankruptcy (Lehman Brothers, 2008). Amid the panic of other financial and economic actors, the market for Credit Default Swaps (CDS) – a derivative owned by Lehman in large quantities – froze. What followed was the most severe financial crisis of modern times, which saw the failure of 371 commercial banks in the US alone (FDIC, 2011). ). It remains an open-ended question as to whether it is possible to properly quantify the economic costs of the financial crisis. Atkinson et al. (2013) suggest a loss between \$6 trillion and \$14 trillion in the United States; a similar fate hit the majority of OECD countries (Claessens et al., 2010).

On Thursday, 23 June 2016, the United Kingdom was politically shaken by the results of the Brexit referendum. With 52% of registered votes in favour of leaving the European Union (The Guardian, 2016a), the leaders of the Conservative, Labour and UK Independence Party have already resigned or face the likelihood of being replaced in the near future (BBC, 2016a; The Guardian, 2016b; Anushka et al., 2016). The political decision has already had short-term negative effects on the stock and foreign exchange markets, and Dhingra et al. (2016) forecast a much graver long term scenario. According to economists, with the loss of access to the EU single market and annexed privileges, even under an optimistic scenario the UK is bound to lose between 6.3% and 9.5% of its 2015 GDP. The occurrence of Brexit came very unexpectedly in the eyes of many political actors, its designers included (MacShane, 2015). It is still an unsolved question as to how a political idea, deemed to be unrealistic until a few years ago, became a concrete policy action with immense effects on British and European society.

On Monday, 11 March 1985, Mikhail Gorbachev was appointed General Secretary of the Soviet Unions Community Party. Later that year he initiated with Ronald Reagan an unprecedented shift of policy away from nuclear weaponry

(Gorbachev, 2007). Such efforts were heavily dismissed by senior figures from both sides as infeasible due to the possibility of cheating (McNamara, 1986). Yet, the subsequent years witnessed a historic breaking point with the collapse of the Soviet Union, the nuclear weapon race, and a major change in the ideas that governed inter-state politics (Lebow, 1994). Human rights and development issues increasingly took prominence among nations, culminating in global efforts such as the Millennium Development Goals (MDGs) (Cingranelli and Richards, 1999; Hulme, 2007).

The three, large-scale societal events described above have several features in common. They all involve a multitude of agents - be they companies, people, or nations - that interact with each other in unpredictable ways, leading to novel, unprecedented or unexpected events. Because of their very nature, society and the economy are regarded as *complex systems* which can easily bring *emergent* phenomena (Sawyer, 2005), that is, events arising from the cumulative interaction of individual agents who themselves do not possess such properties. A single bank isolated from its counterparties cannot provoke a financial crisis, just as a single political actor cannot move forward with an economy-shaking policy idea and an isolated country cannot generate a paradigm shift in global ideologies. What matters is the set of connections across the elements of a system, not merely the elements themselves.

Societal observers have both a scientific and ethical duty to untangle its complex web of interactions. It is scientific as it poses one of the most difficult challenges yet to be solved. The outcomes of people and institutions sharing information is much more nonlinear than those found in physical systems. It is ethical, since a better understanding of complex social systems can confer upon humankind the ability to better deal with current seemingly-intractable problems (West, 2013).

This doctoral thesis is a modest attempt to push forward the understanding of complex social systems, with a major focus on the practical policy implications that may arise. It tackles the challenge by first making one epistemological assumption: that a complex system of social actors can be modelled as a *network* where the nodes are its elements, and the edges incorporate information about the nature of the relation between any pair of elements. Having done so, the thesis focuses on providing solutions to two main problems:

**Problem 1:** To uncover the latent structure of a social complex network.

**Problem 2:** To measure the extent to which one or more elements are able to



influence an emergent outcome in the social complex network.

Problem 1 arises from the common inability to know how social actors might be interacting, and is mainly due to lack of data about such relations. Not even regulatory financial bodies knew the degree of credit exposures which banks had towards each other during the 2008 collapse (Commission and Commission, 2011). In a similar vein, the degree to which members of parliament share information about policy issues is not often apparent, as informal political exchanges, party dynamics and institutional rules make such relations fuzzy and difficult to detect (Reh et al., 2013).

Problem 2 lies at the core of the study of emergent phenomena, and is of great practical use in policy-making. While emergent events definitively do not arise from a single system element, they may be brought about by or be dependent upon certain elements that feature a special position within the complex network (Scheffer, 2009). Banks that retain a great amount of assets may be more influential in spreading financial shocks. Parliamentary speakers who are better connected across the political community may have greater opportunities to affect the direction of government than those who do not.

The solutions to problems 1 and 2 require a multi-methodological framework. The reason is twofold. First, complex social systems may be governed by different patterns and laws, and while a unified theory of their behaviour and a general method with which to investigate it are not impractical (Norberg and Cumming, 2013), it is beyond the scope of this work. Second, and most importantly, the very kind and context of the data collected about a social complex system can allow a particular solution to the two problems. Sticking to one and only one method while the empirical world is so diversified would impede the practical use of this research.

This work, therefore, advances a general framework to infer the latent structure of a complex social network and measure the influence of its elements with respect to an emergent event. The framework conjugates a typology of methods to address the two problems, each based on the nature of the data at hand and the type of event that is of interest.

The framework – which will be referred to as Network Inference and Influence Framework (NIIF) – is partitioned into two parts and summarised as follows. The first part relates to the inference of hidden relationships across social actors, and comprises a model-based approach known as the fitness inferential model, and two data-driven approaches that are respectively based on entropic analysis and discrete-state neural learning. The second part of NIIF focuses on how to derive a

metric of influence of network nodes. These comprise a contagion-based metric, and a topologically-based metric. Further explanation regarding the conditions in which the framework should be used is included in the following subsection, along with an exposition of the plan of this thesis.

Including the introduction, this doctoral thesis is arranged into seven Chapters. The first three provide the necessary theoretical and methodological background regarding the application of complex systems and complex scientific networks in the study and modelling of society, as well as an overview of NIIF. Chapters 4-6 feature three case studies in which the framework is applied. Chapter 7 offers a conclusion. A detailed overview is provided in the paragraphs below.

Chapter 2 covers the scientific context for this research and is structured into two main parts. The first provides a general introduction to complex systems science, focusing on its applications within socio-economic systems. The concept of the complex system is defined and the Chapter explains how social systems fall in this category. Emergence is also exemplified through examples coming from the socio-economic realm. The second part of the chapter conveys a review of graphical models, with a focus on the definition of a complex network and related topological analysis. It also reviews how scholarly efforts have applied complex network analysis to the problems stated in the first part of this introduction.

Chapter 3 gives an overview of the methodology behind the framework. It first contributes by detailing the network inference module of NIIF. The module consists of three potential techniques to infer the latent structure of a social complex network. The first, the inferential fitness method, is a model-based method that reconstructs a network through knowledge or assumed knowledge solely about its latent density. The method is to be used when data about the social system of interest is not enough to recreate the latent network. The second and third methods are entirely data-driven, and infer the network of a system respectively by harnessing transfer entropy and neural discrete state embedding . The influence module is then introduced. It comprises a topologically-based set of metrics based on measures of centrality, and a contagion-based set wherein the measure is extracted by simulating an event spreading throughout the social network. Among the five modules, the fitness model is an original methodological contribution by the author, expanding

previous work done in network reconstruction. The remaining four modules, while not representing novel methodological contributions, are harnessed for three original case study applications.

Chapters 4, 5 and 6 illustrate three different applications of NIIF on real datasets and problems. Each of the chapters contains a background section, providing context to the specific domain of the case study, and a brief literature review on how the case study has been tackled by previous research. The first relates to the problem of identifying what financial institutions are more influential in spreading an adverse event across a financial system. Data about 182 EU-regulated banks is obtained and used to reconstruct probabilistically their latent interbank network. The contagion-based influence metric is then used to simulate economic shocks and single out which set of banks may be more systemically crucial in affecting the overall health of the system. This work has been performed as part of the EU-funded collaboration project SIMPOL<sup>1</sup>, led by Prof. Stefano Battiston at the University of Zurich. The author is responsible for the network inference and analysis presented in this Chapter, while the data on which it is based has been collected by collaborators. The practical applications have been developed by the author, then better elaborated with other colleagues in order to be presented as a standalone paper (Battiston et al., 2016).

The second case study tackles the problem of measuring the influence of political actors in relation to determined policy topics. It focuses on the UK House of Commons, and reconstructs a network of how information is likely to spread through speakers in the House of Commons based on debates examined during the last three governments, from 2001 to 2014. For each governmental mandate, six policy topics are identified, and the latent network for each is inferred through transfer entropy. The application of the framework in this case study suggests novel insights into the dynamics of political leadership and policy dynamics, by identifying which political actors feature a more central position in the network.

The third case study confers insights about the spread and importance of policy ideas in inter-state politics. It utilises General Debate speeches by members of the United Nations General Assembly between 1970 and 2014 to derive a network of ideological and semantic proximity among countries. In this context, neural discrete state embedding is performed for the task. As in the previous case study, this Chapter builds upon topological measures of influence to deliver metrics that are important to both the country and its policies, showing that it is possible to track

---

<sup>1</sup><http://simpolproject.eu/>

the salience of political ideas across several decades.

In line with the practical focus of this PhD thesis, each case study provides an attentive outline of how the framework outputs can be used to build instruments valuable to policy-making, transparency and accountability.

The concluding chapter (7) ends the presentation of the research with a synopsis of its findings and contributions, and a discussion on further research. The discussion first concentrates on the main limitations of the framework, then elaborates further on future applications. Rather than in an Appendix, code, additional figures and analyses can be found in the author's GitHub page<sup>2</sup>.

---

<sup>2</sup><https://github.com/etstefano/PhD>

## Chapter 2

# Background

This chapter presents a review of the key theoretical and methodological concepts that lie at the core of the research, essential to the comprehension of its contributions. It first reviews the notion of a complex system, providing a qualitative definition and examining some of the scholarly effort that identifies social systems as complex. It then proceeds on to the introduction of graphical models for the study of such systems, with a specific focus on the concepts of the graph and its components as understood in the field of complex network analysis.

### 2.1 Complex systems: Definition and properties


Defining a complex system has been a challenge that has eschewed a widely-accepted consensus for decades (Gallagher et al., 1999). Different proposals coming from the empirical experience of different disciplines have been advanced, often promoting one property of complex systems over another.

A set of definitions focuses on the intuition that complex systems arise where structure is observed along with variations (Goldenfeld and Kadanoff, 1999). Structure is seen as the occurrence of some form of order, of regular patterns and behaviour within the system (McKelvey, 2001). This property is definitively a compelling one, as many scholars regard complex systems to be found at the edge between complete order and chaos (Hayles, 1991). Yet, it seems not to be a sufficient condition in which a complex system can exist. A chess game possesses both structure in terms of the roles played by each piece, and the rules of the game and variations in relation to patterns played. Yet, a chess game system is not regarded as complex, since the output of the game does not lead to novelty within the system.

Another category of definitions focuses on interaction across multiple actors, which leads to adaptation and new systemic patterns (Arthur, 1999; Rind, 1999). This line of thought comes from the observation of economic and ecological phe-

nomena, where evolutionary-like dynamics appear to take place. The definition of a complex system as an adaptive one is at least partly correct: adaptation implies novelty and therefore emergence. Yet, it is incomplete, since it assumes away other key properties that still characterise what scholars may define as a complex system in non-evolutionary domains (Whitesides and Ismagilov, 1999), such as hierarchy and nonlinearity.

Accepting that no single property or small group of such is likely to be sufficient, drawing on Ladyman et al. (2013) this section lays down a properties-based definition of complex systems which implies the necessary conditions in order for them to exist. Not all of the properties must be found in a complex system; they can certainly be present to a lesser or greater extent. Yet, complexity science scholars appear to agree that a complex system should be characterised by at least the majority. Figure 2.1 illustrates the properties defining a complex system, which are discussed in the following paragraphs<sup>1</sup>.



**Figure 2.1:** Properties of a complex system.

**Nonlinearity and feedback.** Nonlinearity is considered to be a baseline property of a complex system (Sengupta, 2006). A system is considered linear if one can add any two solutions to the equations that describe it and obtain another, and

<sup>1</sup>It must be noted that the definition is not a formal but rather a qualitative one. The reason is that efforts to provide a pure quantitative method of the degree of system complexity have yet to come to a consensus and empirical validity. For a review of the attempts to provide metrics of complexity, most of them based on information theoretical insights, see Lloyd (2001) and Vovk et al. (2015).

multiply any solution by any factor and obtain another (Farina and Rinaldi, 2011). Nonlinearity implies that this principle – known as superposition – does not apply.

The property entails a breakdown of causation, or a difficulty in tracing causal processes much higher than in linear systems. One event does not clearly lead to another in a complex system (Barringer et al., 2013). Also, the magnitudes of initial phenomena and the systemic outputs originating from them can differ exceedingly. This peculiarity is popularly known as the butterfly effect (Lorenz, 2000), and can be routinely observed in several physical and social systems. One instance is found in Fiedler and Bukovsky (2011), who demonstrate how the installation of a wind farm in the Midwest of the United States had an enormous impact on the amount of precipitation for an entire season, and therefore on the long-term weather forecast. In finance, a single fake tweet stating an explosion at the White House in April 2013 caused the New York Stock Exchange to plunge by 143 base points (Kelly, 2013). While the piece of information was online for just a couple of minutes and came from a notoriously fraudulent account, automatic trading strategies picked up the news and, reinforced by the momentum, caused one of the most curious shocks in the history of trading. Hong and Sun (2000) make the case of how very small – perhaps insignificant in the eyes of many – events and agents have severely shaped long-term China-U.S. relations. The authors describe the role of players and journalists revolving around the world of table tennis, who paved the way for a cultural and political exchange that formed non-hostile relations between the two nations.

The nonlinear dynamics associated with the so-called butterfly effect is known as feedback. The examples above relate to positive feedback whereby a relatively small perturbation in the system can lead to massive changes. Negative feedback can also occur and engenders the opposite type of dynamics: relatively huge perturbation leading to small changes within the system (Zeigler et al., 2013). Examples include the classic stabilising mechanism related to relative  $CO_2$  and  $O_2$  presence in the atmosphere and temperatures (Walker et al., 1981), financial regulatory bodies manipulating market expectations to diminish price volatility (Heemeijer et al., 2009), the rules of debate in parliaments, which prevent the escalation and disruption of the bill-passing system even after disruptive political events (Rhee, 2000).

Nonlinearity and feedback are not sufficient conditions for a system to be complex. Systems may well be nonlinear, yet be affected by pure chaotic motion, therefore not incorporating any kind of order (Devaney, 1992). To be complex, a system ought to feature some form of structure. This is discussed in the next subsection.

**Self-organisation and hierarchy.** Along with nonlinearity, self-organisation is another ground property of complex systems. It is referred to as a process by which some kind of order or coordination is born out of the interactions between the elements of a system (Nicolis et al., 1977; Kauffman, 1993a).

Certainly, the most significant form of self-organisation found in nature is that which led to life itself. Through some endogenous triggering mechanisms still partially beyond the comprehension of the scientific community, particles arranged themselves to form a closed system whose internal dynamics dictated the perpetuation of its existence over time by means of consuming external resources (Kauffman, 1993b). Through evolutionary pressures, those particle arrangements that made the organism more likely to survive were selected. The result is enormously complex life systems such as the human body, containing and preserving an order constituted of specialised cells and elaborate interactions with external signals (Ingber, 1998).

Just as feedback is a special brand of nonlinearity, hierarchy can be seen as a special form of self-organisation<sup>2</sup>. Hierarchy is defined as a form of order by which certain elements of a system find themselves with more formal control or leverage over its whole dynamics than other elements (Flood and Carson, 2013). The constitution of hierarchy is very typical of social systems and appears to be born out of two potential phenomena. The first is through positive feedback about a system's agent status, the so-called "rich-get-richer" effect (Jiang and Huang, 2012). The second is by selection: hierarchy can sometimes be deemed to be a more efficient way of flowing information within a system (Stiglitz, 1975).

Financial systems provide an interesting example of self-organisation that takes a hierarchical form. Among the global network of banks, a small number are selected with the label of Globally Systemic Important Banks, or G-SIBs (Moshirian, 2012). These are banks that have accumulated so much capital and exposure that regulatory bodies are forced to pay a different sort of attention to their actions for the sake of the entire system's stability. The list of G-SIBs includes institutions such as HSBC, Barclays, Unicredit and Merrill Lynch (Weistroffer et al., 2011). In political systems self-organisation has been more evident than in other social and physical systems, since a defining pillar of decision-making is itself some form of social order (Fukuyama, 2011). In all functioning democracies, political

---

<sup>2</sup>Some scholars indulge by contrasting self-organisation and hierarchy, where the latter is seen as the child of external design (Valverde and Solé, 2007). This research does not espouse this view, as it is deemed not to be consistent with the very fact that hierarchical order is found in animal systems (Theraulaz et al., 1995).



decision-making is organised into parties. In the United Kingdom, for instance, parties in the House of Commons organise into party leaders and backbenchers, with the aim of making bill proposals and debates more efficient (Norton, 2015). In the United Nations, important changes in policy ideas and priorities appear to be communicated by top level country representatives during the initial General Debate (Lindemann, 2014).

Especially in social systems, the dynamics of self-organisation does not occur randomly. On the contrary, it is informed by its members' past experience (Solé and Bascompte, 2006). That is the reason behind the importance of the path-dependency property to define a complex system.

**Path-dependency.** Path-dependency is the restriction of future action by previous action. Stated more clearly, the property entails that the dynamics characterising a system at a given time shapes the range of potential future dynamics towards one specific trajectory rather than others (Ottosson and Magnusson, 1997). First introduced in organisational theory (Dopfer, 1991; Sydow et al., 2005), path-dependency is heavily explored in physical and biological science as well.

A beautiful and remarkable example of path-dependency is found in the chirality of the biomolecules that characterise life as we know it on Earth. A chiral biomolecule is an asymmetric one, and can be left-handed or right-handed. Life makes use of left-handed molecules (Quack, 2012), the reason for this appears to be purely down to statistical chance. Simply, in primordial ponds filled with organic material, by random fluctuation a small imbalance between left- and right-handed molecules occurred, leading to organic matter having a preference for the former (Gibney, 2014), and conditioning the entire spectrum of complex life forms hereafter.

In complex social systems, path-dependency has been studied extensively. The very dynamics of institutional formation are intuitively informed and restricted by past experience and history. Morgan and Kubo (2005) provide an in-depth case study on the institutional changes of Japanese capital markets over the years. During the 1990s, Japan began promoting a sort of divergent capitalism (Gerlach, 1992; Carr, 2005), with different sets of rules and norms than found in the West. The authors provide an explanation of how past economic changes in Japan have shaped the current financial market. Bevan and Robinson (2005) turn to path-dependence to explain the sub-optimality of healthcare reforms pushed by the British government over the years. According to the scholars, the policy path reforming the National Health Service (NHS) is inefficient in terms of the control of total costs,

the equitable distribution of hospital services, and efficiency in delivery. None of the reforms applied between Thatcher's and Blair's years have substantially improved the situation as they were all incremental and based on previous, biased ones. In a similar vein, path-dependency has heavily influenced the terms of the UN Global Compact, a voluntary initiative based on CEO commitments to implement universal sustainability principles (Rasche et al., 2013). UN business relation norms have been informed by the experience of the international organisation with US companies during the Cold War, and by the increasing momentum of human rights and sustainability values over the last fifteen years.

**Emergence.** As previously mentioned in the introductory chapter, emergence engenders the materialisation of a qualitatively novel phenomenon out of the interactions of system agents where such a phenomenon cannot be replicated by the single agents themselves (Yates, 2012). The property of wetness of an object cannot be reproduced by the single water particles or the single object. It is the interaction between the multitude of water particles and the object that allows it to emerge. Similarly, it is the intricate interaction of the actors in a complex system that bring about emergent phenomena.

The questions tackled in the case studies of this doctoral thesis concern emergent phenomena, and can be used as examples. When a market shock causes sudden adverse changes in security prices, financial institutions that are heavily exposed may become insolvent or not be able to repay their debts in full (Moussa, 2011). This causes other banks to become distressed, thus fuelling a process of contagion. Major financial crises are emergent events arising from the interaction of their assets and liabilities. In both government and foreign relations, policy themes and ideas can be thought of as emergent properties of the system of people and institutions dealing with them. Policy ideas and themes are informational constructs that spread and retract during the informational exchange that occurs within the political system (Speel, 1997). Be they topics discussed in the UK House of Commons or the UN General Debate, they are emergent properties of their respective systems, since they solely exist and spread through the interaction of the political actors in whose cognitive content they exist. A policy topic that lies in the mind of a single actor would be non-existent with respect to the system, as it is not spread.

This section has defined a complex system as a set of elements interacting with one another, exhibiting nonlinear behaviour, a form of order in terms of self-

organisation, path-dependent evolution and emergent properties. The attentive reader would justifiably reason about the evident overlapping of the four key properties identified. Emergent phenomena may well arise from nonlinear behaviour, as they pose a rupture from the previous state of the system; and self-organisation is an emergent property of complex systems, being born out of the elements' interactions. Yet, a somewhat definite distinction of the concept is required: conceptual clarity provides the right cognitive framework to analyse an aspect of a complex system distinctively.

Throughout Section 2.1, examples from systems close to those inspected from Chapter 4 have been advanced, with the aim of demonstrating how social systems - financial and political ones, particularly - can be seen from the lenses of complexity science. The following section briefly reviews part of the key literature that has most contributed to the study of complex economic and political systems.

## **2.2 Markets and politics as complex systems**

The literature around markets as complex systems is a vast and established one. It is not the aim of this section to review it all<sup>3</sup>, rather, they will focus on a number of key contributions that particularly characterised financial systems, the focus of the first case study of the thesis.

Seminal work has been undertaken in the design of artificial stock markets, showing that if built with rules of the dynamics typical of complex systems, they will feature the same kind of signals observed in reality. The first attempts were made in the 1980s, by Cohen et al. (1983), who looked at the impact on various market structures of random-behaving agents, and by Kim and Markowitz (1989), who examined the interaction of specific trading strategies. A breakthrough in the field was achieved by Brian Arthur and John Holland who finalised the Santa Fe Institute artificial stock market in the early 1990s. The first version of the market is presented in Palmer et al. (1994). What the authors create is a simple world populated by agents that could trade an asset at a price which is set by the market-maker to all traders. Agents find a matching rule for current market conditions and come to the market in order to buy or sell one share of the asset. Backed by a utility function, each agent would over time develop a different trading strategy, thus recreating an ecology of traders as often observed in real markets (Conrad and Kaul, 1998). The work of Arthur and Holland paved the way to the current study of financial systems

---

<sup>3</sup>Readers are invited to consult Arthur (2014) for a comprehensive review.

using agent-based modelling (ABM), i.e. the bottom-up simulation of economic agents and assets who co-evolve to generate the kind of complex dynamics seen in real life (Chen et al., 2012). The state of the art of ABMs applied to finance explores the contagion of economic shocks and the impact of policy decisions. Gangel et al. (2013), for instance, use ABM to explore the impact of real estate foreclosures in the US for the wider financial market. They simulated agents and assets resembling the actual property market in the United States, and performed experiments that began with different types of adverse conditions. Fagiolo and Roventini (2012) propose the use of ABM to substitute the standard macroeconomic policy tool currently used to regulate financial markets and dynamic stochastic general equilibrium models (DSGEs). The authors provide a compelling justification of how their approach is better than DSGEs because they relax the equilibrium hypothesis contained in the latter.

Another influential strand of literature studies financial systems with the same theoretical concepts and methods used to inspect complex physical systems. Unsurprisingly, this approach is known as econophysics (Stanley et al., 1999). Econophysics has made substantial contributions in explaining the emergence of many phenomena in finance, such as the occurrence of bubbles and crashes. Mandelbrot (1997), Muzy et al. (2006) and Focardi and Fabozzi (2013) are just some of the many scholars who have provided extensive empirical evidence towards the existence of fat tails, that is, a universal skew in the distribution of historical asset prices suggesting that extreme adverse shocks are more likely than what is assumed in mainstream economic studies. This skew, ignored by neoclassical economics, becomes apparent when financial markets are observed under the lens of complexity: positive feedback behaviour causes sudden endogenous crashes, generally due to a relatively small initial shock (Wray and Bishop, 2016). Some econophysics scholars have also devoted themselves to uncovering an emergent property out of different financial bubbles. Of fascinating beauty is the work by Sornette and collaborators. Sornette (2009) advances the idea that a latent distribution of hyper-shocks is hidden within the wider distribution of financial shocks in the market. In other words, the very fat tails in market prices may be due to an emergent property that periodically leads to extreme events. This financial crash model was applied to variously explain and predict – to a certain extent – the Nasdaq crash in 2000 (Johansen and Sornette, 2000), the Latin American financial bubble (Johansen and Sornette, 2001), and the 2006-2008 oil bubble (Sornette et al., 2009), to cite a few. Bianchetti et al. (2016) have even used it to predict the financial crash occurring after the Brexit vote.

In the study of political systems, the use of complex systems approaches has not yet matured to the level seen in finance. The reason is quite straightforward: the amount and diversity of empirical data points within political science had not reached levels seen in other sciences until very recently (Grimmer, 2015). However, the deficit has not impeded the elaboration of theories and respective applications that emphasise some of political system properties that characterise them as complex.

The first is the punctuated equilibrium hypothesis. First introduced in Baumgartner and Jones (1993), this is a theoretical framework explaining policy processes as characterised by periods of stable, incremental development, then ruptured by sudden policy changes that lead to a novel status quo. Needless to say, punctuated equilibria are nonlinear dynamic processes. In their ambitious work, the scholars build the hypothesis upon case studies on US agendas about civilian nuclear power, urban affairs, smoking, and automobile safety. What is observed is always the same pattern: small changes to the policies related to any such topic suddenly being swept by abrupt change. The reason behind such nonlinearities is not external; in fact, it is endogenous. According to the scholars, the gradual building of the salience of a certain policy topic - in the form of more discussion by the media, greater awareness by citizens and political actors - causes a precipitous phase shift in the system. One of the ways punctuation manifests itself is through changes in budget allocation. With this insight, Citi (2013) provides a compelling analysis of the European Union policy process. Budget data ranging from 1984 to 2011 allowed the scholar to observe structural changes in resource allocation that cannot be explicated by an incremental agenda process, but rather by a mix of inertia and sudden bursts. On a similar line, John et al. (2014) use content analysis of policy textual data and budgetary information to discover the agenda setting dynamics at the UK House of Commons, suggesting the occurrence of punctuation.

In the study of foreign relations, complex systems theory has gained the status of a new paradigm (Kavalski, 2007). The perception of pervasive randomness in international life has challenged the dominant frameworks for the study of world politics, turning some commentators towards complexity theory. Works such as Geyer (2003), Chesters (2004) and Flockhart (2006) have built the ontological and epistemological foundations of this approach. They all envision political actors in the international realm as aggregate cognitive agents that share, shape and diffuse social norms. An example is the monograph by Clemens Jr (2001) on the peaceful transition that Baltic states experienced from being under Soviet rule to a successful trio of democracies. The author suggests that at the core of their success lies a complex political system that has managed to form and diffuse norms about peace, while at the same time being punctuated with policies to make social capital flourish

within their borders.

There are several approaches to depict and extract insights from a complex system. The one espoused by this research relies on complex networks. The next section dives into the definition of the method.

## 2.3 Complex graphs: Definition and applications

Graph approaches are used to summarise and analyse relations across multiple actors within a system in a comprehensible, tractable way. In this section, graphs and their properties are defined, along with a review of how they have been applied to understand real world social systems.

**Definition of a graph.** A graph (or network) is a collection of entities, generally known as nodes, and a collection of relationships between nodes, known as edges or links (Wasserman and Faust, 1994). Analytically, a network can be simply represented as a set of sets containing all nodes and edges,

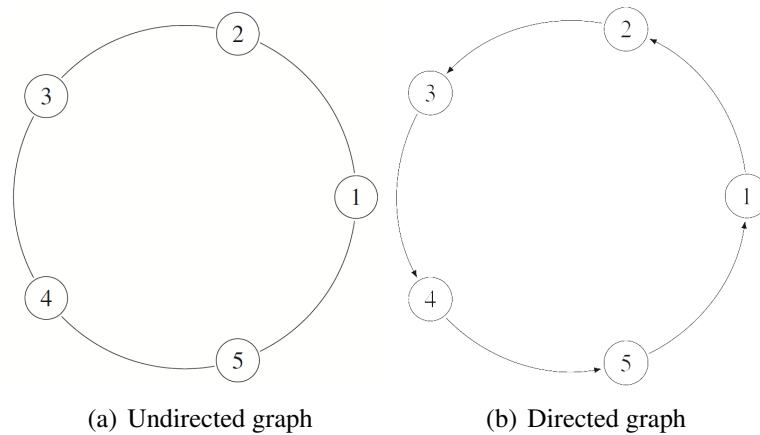
$$\begin{aligned} G &= \{N, E\}, \\ N &= \{n_1, n_2, \dots, n_g\}, \\ E &= \{e_1, e_2, \dots, e_k\}, \end{aligned} \tag{2.1}$$

where  $n_g$  is a single node in the network, and  $e_k$  is an edge between two nodes. An alternative is to represent a network using an adjacency square matrix, with dimensions equal to the total number of nodes in the network. At the intersections between the rows and columns of the adjacency matrix edge, values are reported.

Each node can store attributes, qualitative or quantitative properties that provide specific information about the nature of such nodes (Brandes and Wagner, 2004). In a financial network, the node may store the identifier of a bank and the level of capital it possesses at a given point in time. In a political network, nodes can feature the identifier and length of speech of a political actor over a set of debates.

Edges store information about the nature of the relationship between any two elements of a system. They can be undirected or directed. If directed, the kind of relationship represented may be asymmetric. Figure 2.2(a) provides a toy example wherein each node has an indirect relationship with another, while figure 2.2(b) shows a directed toy graph. Furthermore, edges can also be weighted with a value

explicating the intensity of the relationship (Palla et al., 2007). A graph is defined as complex when the distribution of its edges is not purely random or if it does not follow a lattice pattern, that is, when it forms regular tiling when embedded in Euclidean space (Caldarelli and Catanzaro, 2012).



**Figure 2.2:** Examples of undirected and directed graphs.

**Applications of graphs in complex social systems.** Graphs have traditionally had a strong record of applications within the social sciences. The reason is intuitive: the most crucial feature of social systems is the relations they engender. In the following paragraphs, a brief review is given on the use of networks in the realms of finance, and national and international politics.

Finance implicates the flow of capital from an investing agent to a recipient, and single agent behaviour shaped by the actions of other agents. Generally speaking, the literature revolves around three main areas: to discover and explore the structure of financial systems, to detect underlying correlations across financial securities, and to explore their systemic stability under distressing conditions.

Contributions to the first line of research have indicated that financial systems tend to have a core-periphery structure, that is, a topology where a minority of agents are heavily connected both to each other and to those that are less connected (Csermely et al., 2013). In their work, Minoiu and Reyes (2013) analyse the global banking network across 184 countries through the use of data on cross-border banking flows from 1978 to 2010. Among their findings, the scholars suggest that the core-periphery structure of the system (dominated by those institutions found in financial epicentres such as New York, London, and Hong Kong) is pro-cyclical, as the network density – the number of links in the network – increases during

periods of global economic growth. They also suggest that network density accelerates before imminent financial crises. A core-periphery structure is also found in national financial markets. Fricke and Lux (2015a) demonstrate it for the Italian financial interbank system e-MID, where banks headquartered in the country can obtain short-term credit from one another. What is most revealing in their work is that larger banks tend to borrow from smaller ones at the periphery, thus making the latter more vulnerable to any shocks occurring at the centre. The market for equities appears to be affected by the same kind of structure. Glattfelder (2013) and Vitali et al. (2011) create a network of corporate control on equity data about 43,060 transnational institutions based on equity ownership data. In this system, the core is formed by 1,347 companies that control three quarters of the remainder of the network. This interesting insight seems to validate the thesis for the existence of a highly unequal control of capital in the hands of a small economic elite (Korten, 2015).

In order to make better investment decisions in an uncertain environment, it is important to obtain a glimpse of how securities' prices fluctuate conditional upon other securities in the market. This problem led several scholars and practitioners to construct graphs of stock market products where the edges are correlations in a given frame. This strand of work is usually categorised according to its geographical scope. Preis et al. (2012) provide an example of an analysis taking into account securities mirroring the performance of global stock markets. The team analyses 72 years of daily closing prices of the 30 stocks forming the Dow Jones Industrial Average (DJIA), elaborating correlation coefficients for each pair of securities. Their work finds that portfolio diversification efforts may become vulnerable in the face of distress, as stocks previously assumed to be statistically independent feature nonlinear latent relations. Nobi et al. (2014) provide another contribution by examining how the correlation and network structure of global stock market indices and local Korean indices have changed between the years 2000 and 2012. In line with the previously-cited work, they find that during adverse times the topology of the graph alters since financial crises make indices much more correlated, nullifying or trivialising the effect of diversification. Other approaches have focused on reconstructing the relations of only national stock markets. Vizgunov et al. (2014), for instance, attempt to use correlations of stock returns for the Russian market for the years ranging between 2007 and 2011. Their main finding informs investment managers that stocks with more coupled returns appear to perform better than the rest. Jallo et al. (2013) present a more elaborate analysis with the American and Swedish stock markets as their case studies. They construct a depiction of the financial systems based on stock returns, stock returns with vertices weighted with a



liquidity measure, and correlations of volume fluctuations. Among their results, the authors observe statistically significant changes in network densities as a function of important economic phenomena leading to sudden price changes.

The flow of a systemic shock in a financial system can be studied as information flowing through a graph, as the case study in chapter 4 demonstrates. The intuition has been picked up by several scholars, especially after the 2008 financial crisis. Studies in this field can be taxonomised according to the type of relationship they use to build networks. Some works build a model of the financial system and its vulnerability under shock by taking into account portfolio commonalities (Allen et al., 2012). If two banks invest in the same class of asset, and one of the banks is in risk of default, the distressed bank would offload its assets in a fire sale, causing a downward price shock that would negatively affect the other bank. Based on this mechanism, Caccioli et al. (2014) build a bipartite graph model of financial systems, whereby banks are linked according to their investments. A bipartite graph is one whose nodes can belong to two categories (Bondy and Murty, 1976); in this application the graph involves banks and securities. A further set of applications focuses on the exposures of banks to each other through lending. Financial actors borrow and lend among themselves, causing a chain of interdependence that makes them vulnerable if one or more banks are under stress and cannot repay their debts (Boss et al., 2004b). Landmark studies have investigated the issue in, for instance, the Italian interbank market. Based on data on overnight loans recorded on the e-MID trading platform during the period 1999-2010, Fricke and Lux (2015b) infer the flow of diffusion of a financial shock throughout the system. Georg (2013) contributes with a more general application of interbank networks, as a general model with interactions with a central regulatory bank. In the study, evidence is provided that the central bank stabilizes interbank markets in the short run only. Further instances of this kind of application are provided in the survey by Chinazzi and Fagiolo (2015).

The use of graphs in political science has a strong tradition, despite the fact that it has not been explicitly associated to the wider complexity science domain. Social network analysis has been the focus of political sociologists since the 1970s (Tichy et al., 1979; Snyder and Kick, 1979). The applications of graph methods reviewed here can be categorised into three sets: the modelling of political movement and participation, the analysis of foreign relations, and the mapping of information exchanges in policy systems.

One of the main questions in political systems asks what brings some citizens

to participate in political decisions and what leads them not to engage. Theoretical propositions (Shachar and Nalebuff, 1999; Van Deth, 2001) have advanced that a key factor that causes an individual to join the political discourse is the degree to which their social relations are already active in that respect. This is often known as the network effect (La Due Lake and Huckfeldt, 1998). The hypothesis has been tested by a number of studies. Recent contributions include Valenzuela et al. (2012), who examine the relationship between citizen-to-citizen discussions and online political participation. Using data from national surveys of US citizens, the researchers find that communication among people with strong relationships (e.g. being in the same family, colleagues or friends) increases the chance that the individual will take part in the political process. An extremely similar insight is given by Gil de Zúñiga et al. (2012). The team builds a graph of social relationships across US individuals, and controls for variables such as demographics and traditional media use. Individuals with denser graphs, i.e. with a greater flow of information about politics with others, appear to be more likely to be politically involved. Other seminal work has been published by James Fowler and his research team. Fowler and Christakis (2010) investigate cooperative behaviour for the public good in social networks, by the use of a set of laboratory experiments. They find evidence for the existence of a cascade, by which the cooperative action of a person is positively correlated with the occurrence of cooperative behaviour among individuals in his or her social network. In a later study (Bond et al., 2012), Fowler and colleagues explore the same phenomenon with regards to political mobilisation in online social networks, across over 61 million individuals.

A popular method of empirically studying relations between states is based on the use of dyadic data, that is, variables related to multiple countries rather than just one. In network science, this form of data is simply an edge and specifies a particular kind of relationship among countries. Dyadic data can be about treaties (whether any two countries sign the same treaty), economic flows (net exports between one country and another) or events (whether two countries have gone to war against each other) (Hoff and Ward, 2004). Scholars using dyads essentially build graphs of countries according to one or more types of relationship and attempt to test whether such relations are predictors for their event of interest. A landmark application of dyadic analysis relates to democratic peace theory, the proposition that democratic countries do not go to war against one another (Wiebrecht, 2013). The reason for the absence of conflict among democratic nations would be, according to their supporters, increasing economic interdependencies and the poor likelihood that citizens would decide to risk their own lives because of conflict (Owen, 1994). Cederman (2001) supplies one of the most important empirical contributions to this

theory. The scholar analyses dyadic data about common treaty signatures and mutual conflict, finding that the process involving non-violence between democratic states is not immediate, but a long adaptive one. In the early Twentieth century, democratic nations were much more likely to go to war against one another than in its closing years. ONeal and Russett (2015) provide an updated analysis on the same research question, with a dataset about country relations dating back to 1885. Their results suggest that the progressive benefits of democratic peace were already observable in the Nineteenth century. While the use of dyadic data implies that a graph is studied, often the scholars using them treat them as single variable data, thus falling into methodological problems which invalidate their results. This is especially due to the assumption that dyads are independent events (Erikson et al., 2014). The embrace of a fully graph-theoretic approach would help the community dealing with such obstacles.

In order for a policy to be implemented or gain propensity, information about it must flow across the system of political actors involved. The problem of information flow can be modelled with graphs, and some attempts towards this direction have been performed with regards to policy analysis. Indeed, the second case study of this thesis will contribute to this area of research. A very valuable work is presented by Leifeld and Schneider (2012). The duo proposes that the exchange of information in a policy network also depends on transaction costs, that is, the very cost of delivering a piece of information to another actor. This factor adds to other potential variables such as preference similarity, influence reputation, social trust, and institutional actor roles (Henry et al., 2011). They reconstruct the information flow network in the case of the German toxic chemicals policy domain, finding that political actors choose contacts who minimise transaction costs while maximising outreach. Another elaborated scholarly contribution is authored by Farrell (2016), who is focused on understanding the structure and influence of the policy system around the climate change counter-movement. His research uses network science to uncover the institutional and corporate structure of the counter-movement, and machine-learning text analysis to show its influence in the news media and bureaucratic politics. With data ranging from 1993 to 2013, the author finds that the policy network against climate change prevention policies is heavily tied to elite corporate contributors. Carpenter et al. (2004) use dyadic analysis to uncover information exchange among Washington lobbyists, government agencies and congressional staff. Their results add evidence to the hypothesis that lobbyist disproportionately inform other political agents with similar preferences. Further work on networks and politics is presented and discussed in Lazer (2011).

This section has reviewed some of the application of network analysis in the domains that are of interest to this research. That the relationships across actors in social systems can be explicitly modelled through graphs is an idea welcomed by scholars, and indeed they have delivered promising results particularly in finance and political science. It is worth noting, though, that network analysis has also paved the way towards completely novel research areas that are found in the multi-disciplinary dimension. Very interesting examples include the study of information diffusion across the web (Kim et al., 2016) and in massive social media, such as Twitter and Facebook (Guille et al., 2013). Others have turned onto the observation of norms formation and diffusion through crowdsourced knowledge material such as Wikipedia (Heaberlin and DeDeo, 2016). As the amount of data on human cognitive behaviour grows, it becomes more feasible to model the dynamics of entire societies.

However, until then, a framework is needed to extract as much insight as possible on the relations of social actors with incomplete data, or with a latent network that is not directly observable. This framework is the subject of the next Chapter.

## Chapter 3

# Methodology: The network inference and influence framework

In Chapter 1, it has been stated that the aim of this research is twofold: to uncover the latent structure of relationships across the elements of a complex social system, and determine the extent to which an element can influence the state of any other element. To accomplish the task, this Chapter illustrates a multi-methodological approach, named Network Inference and Influence Framework (NIIF). The task proceeds as follows. Section 3.1 explains the network inference methods of the framework. These are three and comprise a model-based method and two data-driven ones. The model-based method, called the fitness model, is an original methodological contribution of this research. Section 3.2 dwells on the metrics of influence that describe the degree to which actors affect the network. The framework includes two approaches, a topological metric based on the notion of eigenvector centrality, and a model-based metric involving simulated contagion. In each subsection, recommendations about the context and type of data in which each approach should be adopted are provided. The Chapter is then followed in three succeeding Chapters by three case studies applying NIIF.

### 3.1 Network inference methods

NIIF embraces three different network inference methods:

1. The fitness model.
2. The transfer entropy method.
3. Discrete state neural embedding.

Each method has better suitability in different sets of available data; thus, the user should be careful in understanding their own empirical context before opting

for one of the three approaches listed above. Table 3.1 supplies a quick summary of the input features required by the methods. As will become apparent in the following paragraphs, the fitness model is an extremely useful tool when the data about the elements of the social system of interest are at the aggregate level. By that, it is meant that the data points provide information about the agent’s relationship with all other agents, rather than with single ones. The case study of the next Chapter, for instance, makes use of financial data that provides the overall level of exposure that each bank has with all other banks. Ideally, data would be obtained directly about bank exposures to other financial institutions, but that may not be feasible. Other examples may include trading levels between countries. For each country, the data available could consist of the overall net exports value. The fitness model does not require time-varying data, and in its current form it only supports continuous variables.

	Fitness Model	Transfer Entropy	Discrete State Neural Embedding
Aggregate-level data	Yes	No	No
Time series	No	Yes	No
Type of data points	Continuous	Discrete/Continuous	Discrete

**Table 3.1:** NIIF network inference methods and their input features.

The transfer entropy method is designed for individual-level data that are assumed to signal information about the latent network structure of a complex social system. The case study in Chapter 5 exemplifies: in a system of parliamentary speakers, there exists a latent network of information sharing and influence which can manifest itself through the speeches made in the House. Based on this premise, time series data about House of Commons’ speaker activity is used to infer the hidden networks of influence with regard to policy topics. Another potential application is the spread of memes (i.e. popular phrases, images or other media) through the internet, where the latent network is made of websites and social media users. The transfer entropy approach requires time series data, as it infers the strength of a tie between an agent and another based on the former’s behaviour conditional to that of the latter. In this thesis, the transfer entropy approach is applied to discrete variables but can be equally used for continuous data types.

Similar to transfer entropy, discrete state neural word embedding is adopted when the data in hand is at the individual level. As its name suggests though, it is conceived for discrete variables. The third case study of this thesis shows its

application on country-level foreign policy systems, as externalised by UN General Debate speeches. In this instance, the discrete states are words appearing in the speeches. The approach is best suited for data featuring natural language, and where the observer is interested in a non-directional measure of influence, as it will become clearer later. Other applications may go beyond natural language and span over any other context where the variable states can be operationalised as categorical. A potential example is the spread of diseases in social networks where the states assigned to each agent describe their health.

The three approaches are presented in-depth in the following paragraphs. Each consists of an explanation of the nature of the method and its parameters, together with an analytical description and pseudocode of the algorithm implementing it.

**The fitness model.** The fitness model is a graph generation technique that has been explored over the last ten years. While previous work only managed to reconstruct the topology of a network (De Masi et al., 2006), here presented is a novel version, capable of reconstructing both the topology and the weights of the links. The model has three important assumptions. First, as in Musmeci et al. (2013), the model assumes that the probability of any two nodes being connected is proportional to a latent scalar value  $x$ , known as the fitness of the nodes. The relation takes the following form,

$$p_{ij} = \frac{x_i x_j}{1 + x_i x_j}, \quad (3.1)$$

where  $p_{ij}$  is the probability of nodes  $i$  and  $j$  being linked, and  $x_i$  and  $x_j$  are their respective latent fitness values. The form of this equation is inspired by a beautiful pattern found in real-world social complex networks. As seen in Chapter 2, it has been observed that economic, financial, transport and people large scale graphs feature a core-periphery structure (Barabási, 2013). Central nodes are exponentially more inclined to be linked to each other and the rest of the graph than nodes in the periphery. Therefore, in equation 3.1 two nodes with high fitness value would be more likely to be linked in a nonlinear fashion.

Second, because  $x$  is latent, it is assumed the existence of an observable proxy of the fitness value,  $y$ , evaluated from  $x$  through an unknown universal parameter  $z$ ,

as per equation 3.2,

$$\sqrt{zy_i} = x_i. \quad (3.2)$$

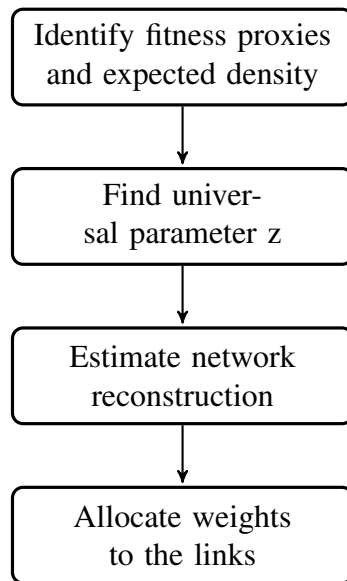
The form of the equation has been chosen after experimenting with a range of other forms. In the experiments, Equation 3.2 demonstrated greater accuracy in inferring networks of known structure. By substituting the latent fitness with the observable proxy and  $z$ , equation 3.1 becomes:

$$p_{ij} = \frac{zy_i y_j}{1 + zy_i y_j}, \quad (3.3)$$

finally, it must be noticed that an interesting property of the analytical specification above is that the sum of the linking probabilities  $p_{ij}$  yields the expected number of links in the network, i.e. the expected density:

$$\sum_i k_i = \sum_i \sum_{j \neq i} p_{ij}, \quad (3.4)$$

where  $k_i$  is the degree of node  $i$ . With equation 3.4 comes the last assumption; the expected density of the graph that has to be inferred must be known or pre-specified.



**Figure 3.1:** Steps to infer the topology of a graph through the fitness model.

Figure 3.1 shows the four steps needed to infer a network using this approach. The first step consists in identifying what observed variables may function as proxies for the nodes' fitness, as well as determining the expected density of the graph. Proxy variables should be proportionally related to the number of ties a node can



have, and requires expert knowledge of the system of interest. In financial networks, for instance, proxy variables can be the amount of total assets of a bank - the larger the bank, the more links characterise it (Caldarelli et al., 2013). In international trade networks, research tends to use GDP values or net exports as proxies (Almog et al., 2015). The density of the graph can be assumed or set by previous research. In finance, scholars have observed that the average density of a network oscillates around 0.1, implying that about 10% of the total possible links generally exists (Montagna and Lux, 2016).

The second step is to evaluate the free parameter  $z$ . This is achieved by numerical simulation. A vector of possible  $z$  values is initialised and substituted in equation 3.4, yielding an expected density value. The  $z$  value that minimises the difference between the expected density and that set by the user is then chosen. Pseudocode for the algorithm to find  $z$  is provided below<sup>1</sup>.

---

```

d ← expected density
N ← number of samples
Z ← z vector
D ← initialised potential density vector
DIFF ← initialised vector of differences between d and solutions

for i in (0,N):
    D[i] ← solution as per equation 3.4 for Z[i]

for i in (0,N):
    DIFF[i] ← |D[i] - d|

return Z[i] for which:
    min(DIFF)

```

---

**Table 3.2:** Pseudocode to find parameter  $z$  for the fitness model.

A careful task in this step is to initialise a vector of  $z$  solutions that assures a satisfactory outcome. The latent  $z$  may vary greatly according to the scale and range of the proxy fitness variables adopted, and may not be within the range originally specified by the user. In order to overcome this obstacle, the proxy fitness variables should be rescaled and normalised, and the vector of  $z$  solutions should feature a wide enough range of numbers. Some trial and error in initialising the vector could

---

<sup>1</sup>Pseudocode follows a Python-like syntax.

be needed at the beginning of the analysis, in order to identify a proper range of values. To provide an example, the initialised  $z$  vector used to infer financial networks in Chapter 4 ranged between 1000 and 100000, a relatively large set of values. At the current state of research, this issue cannot be dealt with automatically. However, because the search for  $z$  is continuous, stochastic optimisation methods avoiding local minima might be well suited to address the task in future work. Examples include Particle Swarm Optimisation (Kennedy, 2011).

Once the free parameter  $z$  is found, a collection of estimated graphs, known as ensemble, is obtained. The task consists of elaborating  $N$  number of networks, where the expected density is as pre-specified, and the probability of any two nodes being connected is determined by equation 3.3. The procedure is explicated by the pseudocode in table 3.3. In a nutshell, given a matrix  $P$  of probabilities for each pair of nodes, a matrix  $R$  of the same size is generated, whose components are random values chosen uniformly from the interval  $[0, 1]$ . All values in the random matrix less than or equal to the respective probability in  $P$  are turned to 1, and 0 otherwise. This represents the estimated adjacency matrix of the graph. Run  $N$  times, an ensemble of probabilistically reconstructed networks is retrieved, over which the distributional character of the connections can be studied.

---

```

P ← probability matrix
N ← size of the ensemble

for n in (0,N):
    R ← random matrix with same size as P
    if R ≤ P: 1 else: 0
    return R

```

---

**Table 3.3:** Pseudocode to generate a network ensemble.

Two things ought to be noticed. First, it should be defined whether the reconstructed network should permit more than a link between two nodes. If not, the adjacency matrix to be estimated can be cut at the diagonal, thus having exactly one potential link per pair of nodes. Second, in the case of a directed network the ensemble estimation does not explicitly take into account the direction of the link. This must be defined by the user once the adjacency matrices are evaluated. The assignation of the link direction can take many forms. It can be random, meaning that there would be a 50% chance that the link would go from  $i$  to  $j$ , and vice versa;

or it can be biased, conditional to some intrinsic nodal quality.

The final stage of this network estimation approach is undertaken if the network features weighted links. This stage is called iterative proportional fitting, and has been inspired by the study of categorical data and contingency tables (Fienberg and Meyer, 2006). The fitting goes as follows. The adjacency matrix representation of a weighted network can be seen as a contingency table. The knowledge of the marginals, i.e. of the total out-strength (sum of the weights of the outward links of a node) and in-strength (sum of the weights of the inward links of a node), allows for an iterative adjustment to the matrix nonzero elements until convergence is reached. Let  $A$  be a toy full adjacency matrix as in table 3.5:

	a	b	c	$A^{out}$
a	1	1	1	11
b	1	1	1	6
c	1	1	1	6
$A^{in}$	10	8	5	

**Table 3.4:** Toy adjacency matrix for iterative proportional fitting.

The full graph has three nodes, with links whose individual weight is unknown. Yet, the sum of their outward and inward link weights,  $A^{out}$  and  $A^{in}$  respectively, is given. The iterative proportional fitting algorithm first adjusts each nonzero element for each row as per expression 3.5. It essentially first divides the nonzero element by the sum of all elements in the row, and then multiplies by the real marginal,

$$e_{ij} \leftarrow \frac{e_{ij}}{\sum_{k \neq i}^N e_{ik}} A_i^{out}, \quad (3.5)$$

secondly, the same adjustment is performed for the inward links, that is, the vertical elements of the matrix:

$$e_{ji} \leftarrow \frac{e_{ji}}{\sum_{k \neq i}^N e_{ki}} A_i^{in}. \quad (3.6)$$

Both steps are repeated until the sum of the calculated weights are equal or arbitrarily close to the real marginals. After a first round of proportional fitting, the toy adjacency matrix would look as follows.

	a	b	c	$\hat{A}^{out}$
a	4.78	3.83	2.39	11
b	2.61	2.09	1.30	6
c	2.61	2.09	1.30	6
$\hat{A}^{in}$	10	8.01	4.99	

**Table 3.5:** Toy adjacency matrix after one round of iterative proportional fitting.

In this trivial exercise, the fitted marginals  $\hat{A}^{out}$  and  $\hat{A}^{in}$  converge easily to the actual ones. In real-world contexts, with great amount of nodes and sparsity, further iterations would be needed.

The fitness model is a useful network estimation tool that allows the inference of the topology of a complex network without any direct or indirect knowledge of the relationship between the agents. However, it comes with a set of assumptions about the latent network density (and link directionality, if the the graph of interest is directed) that are heavily sensitive to human judgement. The tool has been shown to perform extremely well in, for instance, financial systems. Anand (2015) compared the approach presented in this thesis to existing ones, and has shown that it performs extremely well in the estimation of five different real financial networks, with an error rate in link prediction of no more than 10%. It performs much better than competing model-based network estimation approaches in the allocation of link weights. Ultimately, though, its validity depends on the expert knowledge of the user, who must be responsible to set sensible input parameters.

**Transfer entropy estimation.** Contrary to the fitness approach, the transfer entropy method is model-free and has no explicit assumptions about the topology of the latent graph. The method estimates the strength of dependence between any pair of elements solely according to their observed time-varying behaviour (Schreiber, 2000). Before explaining transfer entropy analytically, it is first necessary to define the informational theoretical notion of entropy. As stated by the landmark study by Shannon (1949), entropy is the expected value of information content delivered by a process. The information of a single event  $x$  is postulated as the negative logarithm of the probability of it occurring,<sup>2</sup>

---

<sup>2</sup>Generally, the logarithm at base  $e$ , 2 or 10 are used. This thesis assumes a logarithm at base 10 throughout all Chapters.

$$I(x) = -\log(P(x)). \quad (3.7)$$

The entropy of process  $X$  is therefore the expected value of  $I$  for all  $x$  outcomes observed:

$$H_X = \sum_x P(x)I(x) = -\sum_x P(x)\log(P(x)). \quad (3.8)$$

There is a simple intuition behind the concept of entropy. Events with higher entropy are those which are less likely, therefore less predictable and with a higher informational content. An event with a probability of being observed of 0.1 would yield 1 unit of information, given equation 3.7. On the contrary, one with probability of 0.9 would be characterized by only 0.05 bits of information. A process with a relatively high number of unlikely events will therefore feature a higher entropy than a more ordered one.

Having defined the entropy for one process, it is possible to explore that for two coupled processes. This would take the form of equation 3.9,

$$H_{xy} = -\sum_{xy} P(xy)\log(P(x,y)). \quad (3.9)$$

Assuming that the two processes are independent, the measure would become:

$$H_{xy}^I = -\sum_{xy} P(xy)\log(P(x)P(y)). \quad (3.10)$$

The difference between  $H_{xy}^I$  and  $H_{xy}$  is defined as the mutual information, and yields the additional amount of information gained about process  $x$  or process  $y$  conditional to the other variable (Li, 1990). It is essentially a measure of mutual dependence.

$$\begin{aligned} M_{xy} &= H_{xy}^I - H_{xy}, \\ &= -\sum_{xy} P(xy)\log(P(x)P(y)) + \sum_{xy} P(xy)\log(P(x,y)), \\ &= \sum_{xy} P(xy)[\log(P(x,y)) - \log(P(x)P(y))], \\ &= \sum_{xy} P(xy)\log\left(\frac{P(x,y)}{P(x)P(y)}\right). \end{aligned} \quad (3.11)$$

Two processes with high mutual information allow the observer to know - within a degree of uncertainty - the current state of one process if the other is known, without any prior assumption about the distribution of variables. While being a valuable advance in understanding system dynamics, mutual information suffers from one shortcoming that does not make it suitable for the scope of this research: it is symmetrical.  $M_{xy}$  would be equal to  $M_{yx}$ , with no information about the flow of information from one process to the other. This in turn, impedes any attempt to predict the future state of a variable or identify a causal relation.

Transfer entropy solves the issue by measuring the extra bits of information about the future state of a process based upon knowledge about the current state of another process (Schreiber, 2000). Suppose it is required to know the entropy of variable  $x_{t+1}$  of process  $X$ , given that  $x_t$  and  $y_t$  occurred. This is done as in equation 3.12,

$$H_{x_{t+1}} = - \sum_{x_{t+1}} P(x_{t+1}, x_t, y_t) \log(P(x_{t+1} | x_t, y_t)). \quad (3.12)$$

Assuming that  $x_{t+1}$  is independent of  $y_t$ , the entropy would then become

$$H_{x_{t+1}}^I = - \sum_{x_{t+1}} P(x_{t+1}, x_t, y_t) \log(P(x_{t+1} | x_t)), \quad (3.13)$$

the transfer entropy from  $y$  to  $x$  is then the difference between  $H_{x_{t+1}}^I$  and  $H_{x_{t+1}}$  for all values of  $x$  and  $y$  observed, i.e. the additional amount of bits gained by incorporating knowledge about the past events of  $y$ .

$$\begin{aligned} T_{Y \rightarrow X} &= H_{x_{t+1}}^I - H_{x_{t+1}}, \\ &= - \sum_{x_{t+1}, x_t, y_t} P(x_{t+1}, x_t, y_t) \log(P(x_{t+1} | x_t)) + \sum_{x_{t+1}, x_t, y_t} P(x_{t+1}, x_t, y_t) \log(P(x_{t+1} | x_t, y_t)), \\ &= \sum_{x_{t+1}, x_t, y_t} P(x_{t+1}, x_t, y_t) \log \left( \frac{P(x_{t+1} | x_t, y_t)}{P(x_{t+1} | x_t)} \right). \end{aligned} \quad (3.14)$$

Transfer entropy is, by definition, asymmetrical.  $T_{X \rightarrow Y}$  would then be evaluated as follows

$$T_{X \rightarrow Y} = \sum_{y_{t+1}, x_t, y_t} P(y_{t+1}, x_t, y_t) \log \left( \frac{P(y_{t+1} | x_t, y_t)}{P(y_{t+1} | y_t)} \right). \quad (3.15)$$

It is possible to further simplify the transfer entropy equation. Given the two equalities below

$$P(x_{t+1}|x_t, y_t) = \frac{P(x_{t+1}, x_t, y_t)}{P(x_t, y_t)}, \quad (3.16)$$

$$P(x_{t+1}|x_t) = \frac{P(x_{t+1}, x_t)}{P(x_t)}, \quad (3.17)$$

equation 4.4 then becomes

$$T_{Y \rightarrow X} = \sum_{x_{t+1}, x_t, y_t} P(x_{t+1}, x_t, y_t) \log \left( \frac{P(x_{t+1}, x_t, y_t) \cdot P(x_t)}{P(x_t, y_t) \cdot P(x_{t+1}, x_t)} \right). \quad (3.18)$$

The calculation of transfer entropy is best understood through a quick toy example. Assume the existence of the following time series  $X$  and  $Y$ :

$X$  : 1110110100110011101101,

$Y$  : 1010101011011011011001.

To obtain  $T_{Y \rightarrow X}$ , let all instances of  $P(x_{t+1}, x_t, y_t)$  be evaluated:

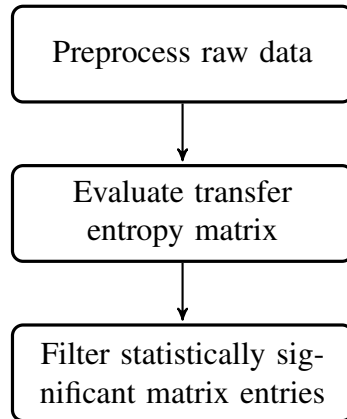
$$\begin{aligned} P(0, 0, 0) &= 0, & P(1, 0, 0) &= 0.143, \\ P(0, 0, 1) &= 0.095, & P(1, 0, 1) &= 0.143, \\ P(0, 1, 0) &= 0.190, & P(1, 1, 0) &= 0.905, \\ P(0, 1, 1) &= 0.095, & P(1, 1, 1) &= 0.238, \end{aligned}$$

$P(x_{t+1}, x_t)$  and  $P(x_t, y_t)$  are then respectively

$$\begin{aligned} P(0, 0) &= 0.095, & P(0, 0) &= 0.136, \\ P(0, 1) &= 0.286, & P(0, 1) &= 0.227, \\ P(1, 0) &= 0.286, & P(1, 0) &= 0.272, \\ P(1, 1) &= 0.333, & P(1, 1) &= 0.364, \end{aligned}$$

with  $P(X = 1) = 0.636$  and  $P(X = 0) = 0.364$ . The analytical combination of all these probabilities as per equation 3.18 yields  $T_{Y \rightarrow X} = 0.044$ . In other words, the information generated by  $Y$  adds 0.044 units of predictability for  $X$ . It is possible to think of  $T_{Y \rightarrow X}$  as the weight of the edge from  $Y$  to  $X$ . The higher is the value of the weight, the more predictable  $X$  is thanks to  $Y$ .

Provided with data about the behaviour of agents of a complex system over time, it is possible to extract the degree to which one element influences the future



**Figure 3.2:** Steps to infer a latent information flow network using transfer entropy.

state of another. The approach comprises of three main steps. The first is to preprocess the raw data, in order to obtain balanced time series for each node of the latent network. The task requires the definition of what constitutes a time step, as well as setting protocols for multiple observations about one node in a single time step. In a corpus of political speeches, for instance, the time step can be a day, a week, or a month, according to the frequency of the data. Multiple documents by the same political actor in one time step can be merged into one. In the context of high-frequency stock price movements, the time step can be more granular, equal to, for instance, one hour of trade. The raw price data can be turned into a dummy variable signifying up and down movements in relation to the price at the previous time step.

Once the preprocessing stage is completed, it is possible to calculate the transfer entropy for each pair of nodes using equation 3.18. The result is a square matrix with size equal to the number of nodes, where the diagonal is filled with zero values. Pseudo code for an algorithm implementing this stage is shown in table 3.6.

A transfer entropy measure can simply be the result of random chance. It is therefore necessary to filter those matrix entries that are not statistically significant. The last stage of the approach deals with this issue through time series randomisation. For each pair of elements, their time series is shuffled  $N$  number of times, and a random distribution of transfer entropies is bootstrapped. A threshold is then set, which is equal to an arbitrary percentile of the distribution. If the estimated transfer entropy value is beyond such threshold, it is assumed to be statistically significant; otherwise, its entry in the transfer entropy matrix becomes zero.



---

```

 $X \leftarrow$  data structure containing the time series
 $T \leftarrow$  initialised square transfer entropy matrix. Size equal to number of elements in  $X$ 

for  $i, j$  in  $(0, \text{length}(X))$ :
    if  $i \neq j$ :
         $T[i, j] \leftarrow$  transfer entropy from  $i$  to  $j$ 

return  $T$ 

```

---

**Table 3.6:** Pseudocode to generate a transfer entropy matrix.

The filtered transfer entropy matrix is the estimated adjacency matrix of the latent network of interest, as signalled by the idiosyncratic behaviour of its elements. The strength of a link between one node and another operationalises the extent to which a node changes its state given the past change of the other, and operationalises the concept of influence. It is not the first time that a transfer entropy approach has been used to detect the hidden web of relations in a system. In the social sciences, Dimpfl and Peter (2013) attempt to measure information flows in the US financial market, with price changes as the manifestation of the latent network. The authors, though, neglect to filter for statistically significant values as proposed in this research. The measure gained much interest in the study of gene expression (Campos and Jacobs-Wagner, 2013), communication across brain cells (Wibrál et al., 2014), and in the quantification of dependencies across geoscientific phenomena such as weather events and ecological changes (Ruddell et al., 2013). It is, indeed, a novel approach to search for coupled behaviour in social systems, and it is hoped that this research would act as a starting point for more elaborated applications in the field.

**Discrete state neural embedding.** Discrete state neural embedding is a machine learning model developed to predict a discrete state given the occurrence of other discrete states within a context window. The model’s goal is best understood under the realm of natural language processing, from which it originates. Essentially, it provides a way to predict the occurrence of a word (the discrete state) given other words around it (Mikolov et al., 2013a). By doing so, it manages to extract the underlying semantics of the text, thus being able to group together words that revolve around the same semantic category.

Discrete state neural embedding transforms each discrete state into a vector representation of real numbers, which locates it into a hyperspace plane. The measurement of the closeness of two vectors operationalises their similarity. In a natural language context, the similarity coincides with the semantic closeness of the two words turned into vectors. In the study of gene activity, instead, because vectors represent sequences of active enzymes, the similarity coincides with coupled gene expressions.

Analytically, the model functions as follows. Let  $s_n$  be a discrete state, and  $k$  the window parameter within which the context for  $s$  is defined. The objective of the model is to maximise the average log probability for  $s_n$ , conditional to its context (Mikolov et al., 2013b),

$$\frac{1}{T} \sum_{n=k}^{n-k} \log P(s_n | s_{n-k}, \dots, s_{n+k}). \quad (3.19)$$

Note the fascinating link with information theory: the log probability expressed above is the negative information yielded by event  $s_n$  conditional to its context. Thus, the objective is to minimise the average information provided by the event, that is, to learn a model whose parameters make  $s_n$  as predictable, less entropic, as possible.

Because the task is to predict a discrete state, the softmax model is used to maximise the objective 3.19. Softmax is an extension of the logistic regression, where the variable can take not only two but arbitrarily many states (Rimer and Martinez, 2004). The conditional probability of state  $s_n$  then turns into

$$P(s_n | s_{n-k}, \dots, s_{n+k}) = \frac{e^{y_{s_i}}}{\sum_i e^{y_i}}, \quad (3.20)$$

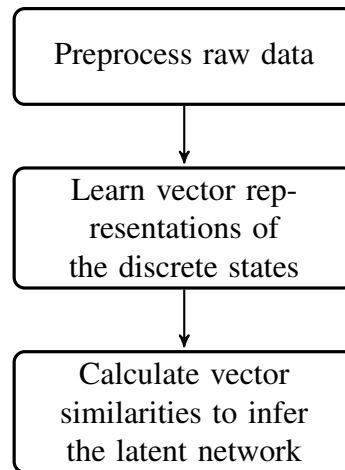
where each  $y_i$  is the normalised log probability for each output state  $i$ , computed as

$$y_i = b + U h(s_{i-k}, \dots, s_{i+k}; W). \quad (3.21)$$

In equation 3.21,  $b$  and  $U$  are the softmax parameters.  $h$  is a concatenation of  $s$  vectors, extracted from  $W$ , which contains all vectors representations of the various discrete states in the system studied.

This machine learning model belongs to the infamous 'black box' family, as it is still not clear - not even to its creators - how exactly it succeeds in turning discrete states into vector representation (Goldberg and Levy, 2014). Here, it is suffice to say that a two-layer neural network is fed the raw dataset containing all

discrete states and their context. In the case of natural language, the data would be a complete corpus, wherein each word stands with its neighbours. Equation 3.20 works as a triggering mechanism, constructing vector representations through its maximisation. The neural network would use backpropagation (Hecht-Nielsen, 1989) to perform the search for parameters that best satisfy the objective. The final output is a matrix in which each column is a vector. It is important to note that the size of the vectors is arbitrarily chosen by the user in this current stage.



**Figure 3.3:** Steps to infer the topology of a graph through discrete state neural learning.

This third approach comprises three steps to infer the latent network of a complex system, as shown in figure 3.3. The first involves the preprocessing of the raw data. In order to use this method, the data should feature only discrete states, such as words. Even entire documents can function as discrete states, as will become apparent with the third case study of this research. The data is assumed to be the manifestation of the hidden states of a system's elements, which in turn signal their relationships. In the third case study, for instance, countries' speeches at the UN General Debate are expressions of their view and ideology with regards to certain foreign policy topics. They in turn provide information about a latent network of ideological similarity across UN members. The exact preprocessing steps are sensitive to the type of data and scope of research, and may require the removal of too frequent and uninformative discrete states (such as stopwords in natural language), or the selection of only pre-specified states (such as nouns and adjectives).

The second step concerns the implementation of the machine learning model described in the paragraphs above, to extract the vector representations of discrete states. Despite the sophistication of the approach, the user is required to tune care-

fully a number of model parameters, particularly the context window size  $k$ , and the size of the vectors. To date, there is no rigorous way to decide the values of these parameters; it all falls on the judgement of the researcher and the inspection of previous related work. Generally, it is regarded as good practice to perform a series of trials with a range of parameter values, and compare the quality of the results. This would involve a loop between this step and the next.

The final stage builds a square matrix with size equal to the number of elements in the system. Recalling that vector representations can be seen as coordinates in hyperspace, calculating the distance between any two vectors, then, is equivalent to evaluating the correlation of two discrete states. In natural language application, the distance would be semantic, i.e. referring to how closely related are two words or documents, given the information gathered through the input corpus. There are several measures of distance that can be used (Wang et al., 2013). This research opts for the cosine distance, in line with the state-of-the-art literature in the field. The reason behind this choice is that the measure neglects absolute frequency difference and instead deals with relative difference (Schakel and Wilson, 2015). This is especially useful when the vector representations have different values because of the frequency size of the original discrete states. If word  $w_1$  is much more frequent than word  $w_2$ , it will be reflected in its coordinates. A Euclidean type of distance would place the two words farther apart, not being able to adjust for the relative difference in frequency. Analytically, the cosine distance between two discrete states,  $a$  and  $b$ , is the dot product of their vectors, divided by the product of the vector norms:

$$C(a, b) = \frac{\vec{a} \cdot \vec{b}}{\|\vec{a}\| \|\vec{b}\|}. \quad (3.22)$$

Akin to the process shown in table 3.6, a square cosine similarity matrix the same size as the number of elements is produced. Each entry contains the cosine distance between a pair of elements, and the diagonal is filled with zero values. The matrix is equivalent to the adjacency matrix of a fully connected network, with weighted, undirected edges that represent the degree of similarity between two nodes.

It is important to highlight two characteristics of the resulting matrix. First, it is symmetrical, implying the absence of directionality in the information shared in the inferred network. Discrete state neural embedding is unable to provide insights onto the information flow like the previous two approaches, and the network should

be interpreted accordingly. In the context of the third case study, for instance, the links between nations and policy concepts imply ideological and semantic closeness. Second, the issue of filtering the adjacency matrix remains just as relevant as in the transfer entropy approach. Ideally, it is possible to bootstrap probability distributions for each single entry of the matrix by shuffling the order of appearance of the discrete states  $N$  number of times, and recalculate the cosine similarity of the resulting vector representations at each iteration. The matrix would then be filtered for those values that fall beyond a pre-specified percentile threshold of the respective bootstrapped distribution. This solution has, however, proved to be impractical for the scope of this research, as it required more computing power than available, making it impossible to implement. Future work may well experiment with this approach. Therefore, more basic filtering has been opted for, such as setting a cosine similarity threshold value, above which an edge is defined as statistically significant. Ultimately, the filtering approach would depend on the domain of application and expert knowledge.

The validity of this machine learning model is emphasised by the increasing interest that it is receiving from industry and scholars alike. In economics, for instance, it has been used to develop a novel method of technology forecasting (Nikitinsky et al., 2015) More specifically, it has been applied on documents about government industrial contracts, identifying technology areas which appear to be of greater interest to public institutions, therefore more likely to experience more funding and faster development. Finance scholars have used neural word embeddings on documents from the Multiple Listing Service (MLS), the most important real estate database in the United States, in order to perform property ranking and assess their investability (Shahbazi et al., 2016). They do so by measuring the similarity of documents related to any one property to a pre-determined set of words that defines good estate conditions. With this thesis, the hope is to introduce discrete state neural embedding to scholars interested in the extraction of hidden relationships in social systems, especially in the case when such relationships become apparent through text.

## **3.2 Network influence metrics**

The relationships between actors in social complex systems are often latent because of the intrinsic quality of a relationship, or because collecting data about it is infeasible for technical or ethical reasons. The problem of inferring the structure of the hidden social network is addressed by the first module of NIIF which comprises three different network extraction approaches. Once the web of relationships is re-

constructed, a set of metrics are required to measure the degree of influence exerted by a node onto the whole graph or single node.

Influence is a loose concept, and must be defined more clearly. This research adopts a soft definition: influence refers to the ability of one element or a group of elements to affect the state of others in a system. Influence, at its very core, concerns the exchange of information or resources. If the flow of information or resources from actor *A* to actor *B* leads to a change in behaviour in *B*, then *A* has influence over *B*. Note that influence, as defined here, overlaps very closely with the notion of systemic importance in finance (Drehmann and Tarashev, 2013), and also the concept of power in the study of politics (Shively, 2011). This is no coincidence: both phenomena involve the change of behaviour of some elements in a system due to information or resources received from other elements.

Even with a hidden network partially uncovered, influence remains a latent event. Because of the very properties of complex systems – nonlinearity in primis – distinguishing its occurrence is a hard task. Take, for instance, the failure of a commercial bank several months after a financial shock hits the system. The default must have been caused by the distress of other banks that propagated the adverse conditions throughout the system. Untangling a single or group of institutions that caused the commercial bank to go down is a daunting task, because the very notion of linear causality may not hold in this context. What can be done, though, is to analyse the known network of relationships to formulate a hypothesis of what financial institution had more leverage in propagating the distress. In politics, asserting that an agent has been led to favour a specific policy issue because of the actions of another agent is often an empirically-unverifiable statement. However, by analysing the topology of the network arising from the patterns of political agents, it is possible to identify those elements that appear to be more central to the flow of information, thereby more likely to be influential.

This section introduces two measures that capture the latent influence arising from a complex network. The first measure is topological and is based on the assumption that the very structure of the inferred networks provides information about each actor's degree of influence. The second measure is model-based and incorporates not only topology, but also a simulation of how the resources or information diffuse through the network given the occurrence of an event.

**Eigenvector centrality.** Eigenvector centrality is a measure of topological

importance. It is used in this research as a measure of influence, because it is assumed that the network inferred with the first step of NIIF embodies the structure of information flow throughout the system. Early pioneers of this metric in graph theory have been Seeley (1949) and Leontief (1951), and its versatility has led it to be the ground for some of the algorithms that created the current digital age, such as Google's PageRank (Page et al., 1999). Eigenvector centrality defines the topological importance of a node according to the centrality of its neighbours (Bonacich, 2007). It is a recursive form of centrality, with the intuition that the most important nodes within a network should be those that are more connected to other important nodes.

Before presenting the analytical form of eigenvector centrality, it is first necessary to introduce the concepts of eigenvalue and eigenvector of a matrix. Let  $M$  be a square matrix:

$$M = \begin{pmatrix} 1 & 3 \\ 5 & 3 \end{pmatrix}.$$

$\lambda$  is an eigenvalue of  $M$  if it yields a solution to the equation

$$M \cdot \vec{v} = \lambda \vec{v}, \quad (3.23)$$

where  $\vec{v}$  is a non-zero  $n \times 1$  vector with real numbers, known as the eigenvector. In this example, solving equation 3.23 yields two eigenvector values,

$$\lambda_1 = 6, \quad \lambda_2 = -2,$$

together with their respective eigenvectors

$$\vec{v}_1 = \begin{pmatrix} 3 \\ 5 \end{pmatrix}, \quad \vec{v}_2 = \begin{pmatrix} -1 \\ 1 \end{pmatrix}.$$

The eigenvector centrality of a graph are contained within the eigenvector associated to the largest eigenvalue of its adjacency matrix. The reason is that the eigenvector incorporates the recursive process of adding up the centrality of a node's neighbour. The mathematical proof for this property is not within the scope of this thesis, but a clear analytical explanation is found in Spizzirri (2011).

Let  $A$  be the adjacency matrix of a directed network with five nodes,

$$A = \begin{pmatrix} 0 & 0 & 1 & 1 & 0 \\ 1 & 0 & 1 & 1 & 0 \\ 1 & 1 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \end{pmatrix},$$

the matrix possesses four different eigenvalues:

$$\lambda_1^A = 2.414, \quad \lambda_2^A = 0, \quad \lambda_3^A = -0.414, \quad \lambda_4^A = -1.$$

Being the largest,  $\lambda_1^A$  is chosen, thus resulting with the following eigenvector centralities:

$$\vec{v}_1^A = \begin{pmatrix} 2.414 \\ 3.414 \\ 2.414 \\ 3.414 \\ 1 \end{pmatrix}.$$

With its recursive formulation, this topological metric is able to capture the structurally most important agents within a graph. As a node with a high eigenvector centrality is well-connected to other nodes with many ties, it is more likely that its information or resources pass through the rest of the network. The agent is in fact in a better position to influence the state of the system. A political actor whose policy ideas are echoed by other prominent political members with several relationships is more likely to influence decision-making than someone whose message is listened by many poorly-connected actors. A small financial institution in distress, which owes significant amounts of money to large banks with ties all over the world's financial system, is more likely to cause a severe systemic disruption than a larger bank whose debt is owned only by banks of the same nationality.

Despite its potential, eigenvector centrality might not be enough to measure the manifestation of influence in complex social networks. This is especially true



when the inferred network has not been constructed in a data-driven mode, but instead is based on model assumptions. Another condition can be the opportunity to investigate influence by simulating the actual flow of information. In such cases, a model-based metric is necessary. This is presented in the following paragraph.

**Model-based influence.** The model-based influence metric constructs a model of how a piece of information flows through the network and, given an event simulation, it measures the extent to which a node or group of nodes contribute to the overall behavioural change within the system. The method developed here is born within research efforts on financial systems, yet, its elasticity allows it to be applicable to several other domains if expert knowledge is at hand.

In contrast with the topological measure of influence, which simply requires some matrix algebra on the inferred graph, the model-based metric necessitates four main steps. The first and most important involves the definition of a flow function, i.e. a function describing how resources or information flow through the network. The function's parameters include the inferred link between the nodes and should also feature some intrinsic quantities related to the individual nodes. A novel matrix that is the same size as the inferred adjacency matrix is then evaluated using the flow function for each pair of nodes that registers a relationship. The matrix is called flow matrix, and is referred to as  $F$ . Essentially, the flow matrix is the inferred network in which the edges have been transformed to define how information propagates from one node to another,

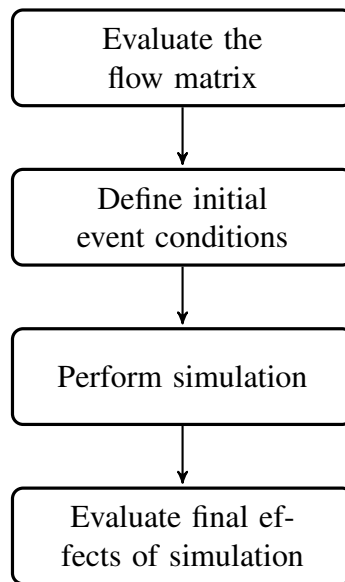
$$F_{i,j} = f(A_{ij}, I_i, I_j). \quad (3.24)$$

Equation 3.24 shows the general form of the flow function, where  $A_{ij}$  is the inferred adjacency matrix, and  $I_i, I_j$  are vectors of quantities characterising the single nodes. Because the flow function affects the state of the receiving node, it can be stated that  $I_j$  is a function of  $F_{i,j}$ , as well as of  $\varepsilon$ , which for the sake of analytical simplicity is assumed to be a vector containing the dependencies upon other nodes and exogenous factors,

$$I_j = f(F_{i,j}, \varepsilon). \quad (3.25)$$

In Chapter 4 the first case study, related to financial shock propagation through an interbank network, provides an example of how the flow function can be specified.

In that context, the function describes how the distress of a bank propagates through another due to the inability to repay in full the amount borrowed from the latter. In a system of parliamentary speakers, the flow function may engender the probability for a speaker to mention a policy idea put forward by another speaker, conditional upon the intrinsic properties of both, such as party membership, constituency represented, seniority etc.



**Figure 3.4:** Steps to obtain a model-based influence metric.

The second stage of the approach is to define the initial event that would pass through the flow matrix. The event is a change in value of an intrinsic property of a node in vector  $I$ , which would then affect other nodes by means of the relationship described in the flow function. In the context of a financial system, the initial event can be the sudden decrease of capital of a pre-selected bank  $a$ . With less capital, the bank does not return the money borrowed from bank  $b$ . The second bank, thus, experiences a decrease in assets and a related decrease in capital. Therefore,  $I_b$  is updated with the novel capital value. The change would reverberate to the next bank with which  $b$  has a borrowing relationship, call it  $c$ , by means of the computation of  $F_{b,c}$ . The contagion of the event would go on in this fashion through the network.

In a graph of political actors inferred through speeches, the initial event can be the mentioning of a policy proposal by speaker  $s$ . This would be reflected in its property vector  $I_s$ . By computing  $F_{s,q}$ , the connected speaker  $q$  has a given probability of echoing the message delivered by  $s$ , and if it does so, it will influence the likelihood that other speakers will repeat the policy proposal.

The initial event can be a single event or a random one drawn from a probability distribution. In finance, for example, market shocks can be drawn from an empirical

distribution (fitted with historical data), or from a theoretical one that satisfies the researcher's requirements. In politics, the mention of one policy idea rather than another can be drawn from a multinomial distribution. In this way, it is possible to evaluate the measure of influence with several scenarios and experiments in hand.

Analytically, an initial event starting from node  $i$  can be denoted as the change in value of  $I_i$  to the event term  $\psi_i$  at the zeroth iteration of a simulation,

$$I_i \leftarrow \psi_i.$$

Once the initial conditions are set, the simulations can take place. These involve the iterative updating of the nodes' properties in  $I$  due to changes resulting from the evaluation of the flow function. Table 3.7 shows how the simulation algorithm works. After setting up the initial conditions, the neighbours of the starting nodes are identified, and the impact onto them is evaluated through the flow function values embedded in the flow matrix. At the next iteration, the neighbours that experience the flow would spread to their own neighbours.

---

$F \leftarrow$ flow matrix
$I \leftarrow$ vector with nodes' properties
$i \leftarrow$ index of starting node or nodes
$I[i] \leftarrow \psi_i$ #assigning starting event
<b>while</b> halting conditions are not met:
find $i$ 's graph neighbours
update $I$ for neighbours according to the flow matrix
set neighbours as new starting points for the next iteration
<b>return</b> $I$

---

**Table 3.7:** Pseudocode for the simulation generating the model-based influence metric.

The cycle ends when a pre-specified halting condition is met. Depending on the context, this can relate to a maximum number of iterations, a convergence set of values for  $I$ , or be linked to a dynamic rule. In the first case study, for instance, each financial institution in the network is given a state variable, labelling it as to whether they have experienced, are experiencing or have not experienced the distress. The iteration stops when no new, undistressed financial institutions are reachable through the topology of the graph.

In the final stage, the impact impressed by the initial condition upon the whole system is evaluated. The metric of influence is therefore the extent to which the properties of the nodes have changed because of a node or nodes that began the dynamics. Contrary to eigenvector centrality, this metric can have a very practical interpretation in the application domain. Let a bank be in distress, causing other banks to suffer from the reverberating shock. Once the contagion has taken its course, the final, systemic amount of damage caused by the initial shock is the amount of capital lost during the process. Analytically, let  $C_0$  be the total sum of capital in the financial system, equivalent to the sum of the elements in vector  $I_0$ , which contains all capital values for each financial institution. Let, also,  $C_t$  be the total sum of capital at the end of the simulation. Given an initial shock, departed from node  $i$ , its influence in diffusing the market shock throughout the system is simply equal to the relative amount of capital lost, as per equation 3.26,

$$H_{it} = \frac{C_0 - C_t - \psi_i}{C_0}. \quad (3.26)$$

In a parliamentary system, the metric can be the relative number of speakers who mention the policy given that a pre-specified actor introduces it. In international trade, it can be the relative amount of GDP losses due to a decrease in imports from a pre-determined country.

The formulation of this metric for the financial system has already resulted in valuable real-world applications for central banks which can use it to monitor the health of the institutions they are regulating (Battiston et al., 2016). Yet, its elasticity allows it to be applicable to any system for which a network representation is possible, and wherein a function of how information or resources flow can be formulated.

This Chapter has introduced a framework that addresses the two problems first exposed in the introductory paragraphs: to uncover the latent structure of relations in a complex system, and to quantitatively measure the degree to which one or more agents influence the rest of the system. The framework, denominated Network Inference and Influence Framework (NIIF), includes three different network inferring methods and two approaches to measure the influence of actors. Each features specific properties that allow them to be applied to different empirical contexts, making NIIF flexible enough for several real-world social systems. The next three Chap-

ters will substantiate this contention by providing three different comprehensive applications. These applications will also allow a better understanding of all the techniques theoretically explored in the previous paragraphs.



## Chapter 4

# First case study: Financial distress contagion

The world of finance perhaps best exemplifies the two problems tackled by this thesis. The globalised interconnections of financial instruments and institutions are enormously difficult to track, let alone control. Most of the relevant data is, after all, confidential for market or regulatory purposes (Haldane, 2013). Financial systems complexity tends to obfuscate how information and resources flow through the various agents, especially when it comes to adverse events. Indeed, only with hindsight has the scholarly and policy-making community acknowledged the link between the first subprime mortgage security defaults and the subsequent financial crisis (Colander et al., 2009). In this Chapter, the NIIF framework is applied to a European-based system of financial actors. More specifically, the network of interbank dependencies is reconstructed using the fitness model, and it is shown how the model-based influence metric – presented in the last section of Chapter 3 – is capable of providing a valuable way to measure the impact of a shock spreading across the network. The framework is also used to identify those institutions that have a particularly high impact in diffusing distress.

The Chapter is structured as follows. A background section provides the rationale behind the choice of this case study, and reviews part of the relevant literature that tackled the same problem from other perspectives. This is followed by a section describing the data collected. Based on that, a fitness model is used to estimate an ensemble of interbank networks. The third section specifies the model-based influence metric used. It explains the exact form of the flow function, its output and interpretation. The ability of the model to detect distress flows is then tested against a competing measure, the default contagion model. It is shown that the former is much more sensitive to adverse events, being capable of detecting their impact at earlier time steps. The fifth section provides instances of how this particular

declination of the NIIF framework can be used for policy purposes, by detecting which financial institutions are more influential in spreading shocks, and for better risk assessment measuring. A discussion of the overall findings of the case study follows.

## 4.1 Case study background

Financial crises have deep, long-lasting consequences for the well-being of societies. The challenge of developing ways to prevent them is not only intellectual, but foremost ethical. Their destabilising effects, in terms of lives negatively impacted, are probably not directly quantifiable, but recent research efforts can provide a general idea. Reviews such as Reinhart and Rogoff (2009, 2013) supply stylised facts about the effect of financial turmoil on GDP growth. The two scholars record a persistent decline in real GDP for several countries that were affected by a financial crisis between 1977 and 2009, with an average annual decline of 9.3% lasting for 1.9 years. Rose and Spiegel (2012) validate the negative pattern by analysing the annual changes in GDP and stock market values in selected non-US countries, believed to have been principally affected by the 2008 financial crisis.

Country	% GDP change, 2008	% Stock market value change, 2009
<b>Iceland</b>	-4.7	-90
<b>Estonia</b>	-2.8	-63
<b>Latvia</b>	-4.6	-55.1
<b>Ireland</b>	-2.8	-7.8
<b>New Zealand</b>	-0.9	-37.4
<b>Italy</b>	-0.6	-49.5
<b>Denmark</b>	-0.9	-48.6
<b>Japan</b>	-0.5	-42.1

**Table 4.1:** Change in GDP and average stock market values in selected countries (Rose and Spiegel, 2012).

Historical evidence regarding a link between financial stress and unemployment is shown in an OCSE study by Choudhry et al. (2010). The authors investigate the change in youth unemployment in South European countries, namely Italy, Greece, Spain, Portugal, conditional to the occurrence of negative financial events. After controlling for other youth unemployment determinants, such as GDP growth, gross capital formation, inflation and openness to trade, financial crisis appears to



increase the variable on average between 1.81% and 1.29%.

Needless to say, the research and policy making community has devoted a significant amount of time researching financial crises and tackling the problem of understanding how distress propagates through a financial system. In this section, three strands of relevant literature are reviewed. Each, just as this research, attempts to yield policy instruments that can be used by the macroprudential<sup>1</sup> policymaker.

**Indicator-based approaches.** Financial policymakers generally opt for frameworks that can be easily and robustly implemented, which require only public or readily-accessible data, and are computable by non-scientific software (Adolfson et al., 2007). Because of such constraints, model-based approaches are not often widely popular in central banks and other regulatory bodies dealing with systemic risk. A number of institutions confronted the issue by developing indicator-based approaches. The most relevant attempt is that presented by the Basel Committee on Banking Supervision (BCBS), which identifies systemically important financial institutions (Basel Committee, 2011). The question as to how an adverse event affects the whole financial system is then answered by stating the extent to which the institution hit by the initial shock is systemically important. The more it is, the more severe will be the threat to the whole system. This line of literature has little explicit groundings on economic and financial theory; it is instead developed through several rounds of consultations with policy experts and empirical trials and evaluations (Financial Stability Board, 2009)<sup>2</sup>.

The Basel Committee's indicator conglomerates five distinct sets of categories, each arithmetically contributing to its value equally. They are, namely, cross-jurisdictional activity, size, interconnectedness, substitutability and complexity. Cross-jurisdictional activity indicates the importance of bank activities outside of their home territories (Weistroffer et al., 2011). A bank is deemed to be more systemically important if its activities are distributed across the globe. Size is another factor that, intuitively, might affect bank importance. When banks possess

---

<sup>1</sup>Macroprudential policy making includes regulatory mechanisms focused onto ensuring the overall stability of the financial system (Clement, 2010). This is juxtaposed with microprudential policy, aimed at monitoring the compliance of single banks with current regulatory requirements (Osinski et al., 2013).

<sup>2</sup>Admittedly, the judgement of policy experts that participated in the construction of the indicator may well be informed by theory, but a clear theoretical ground is neither declared nor discernible in the technical documents (Basel Committee, 2013).

high levels of total exposure, they are more likely to generate widespread distress in financial and economic systems. Interconnectedness relates to how a financial institution interacts with the rest of the system. According to the Basel Committee, the more an institution is connected to others, the higher is its systemic impact in the case of failure or distress. Substitutability refers to how easily the services of a financial institution can be replaced with those of others. It is likely that the failure of a less substitutable bank would affect the health of the system more heavily. A negative correlation, thus, between substitutability and systemic importance exists. Finally, complexity regards the intricacy and ramifications of the bank's organisational structure and operations. Higher complexity would mean that the resolution of the default in a bank incurs higher costs and negative externalities.

The Basel Committee is not the sole body advancing indicator-based approaches to systemic risk. In fact, a number of studies followed its example to derive alternative indicators. All of these, nonetheless, tend to maintain the same categories of analysis, challenging only their operationalisation by proposing alternative sources of data or different calculation methods. A recent instance is Masciantonio (2015), which adapts the Basel Committee's methodology such that only publicly available data is needed to construct the index. A further example is Brämer et al. (2012), which, in a similar way, proposes an indicator-based measurement of systemic banks and their impact that adopts only already-existing indices for each category. The approach has also been used to identify systemically important non-banks financial institutions such as insurance companies (Acharya, 2009).

Indicator-based measurements feature at least two critical shortcomings, respectively methodological and conceptual in nature. These, arguably, prevent this stream of literature to properly reconstruct distress propagation and influence in distressed financial systems. The first weakness is the underlying assumption that every category characterising the systemic importance of a financial institution is equally relevant. Even presuming that each is actually relevant, there is absolutely no a priori reason to justify equal weights. This only apparently naive methodological decision is most probably dictated by the current uncertainty that still pervades the causal processes behind systemic risk and failures (Schwarcz, 2008). Some studies such as Stern and Feldman (2004), Zhou (2009), Hart and Zingales (2011) point towards the significance of the size and complexity of financial institutions', whereas others such as Drehmann and Tarashev (2013), Dungey et al. (2012) point to their interconnectedness and international presence. To the knowledge of the author, currently no study or sensitivity test has been specifically designed to shed light on what indicator categories contribute the most to the systemic stability of the financial system, and in what conditions. Nonetheless, this does not justify the

choice of equally weighting the factors. Such methodology can lead to biases in the estimation of the impact of shocks to the system, thereby causing erroneous policy decisions and detrimental consequences for the economy. The second weakness relates to the conceptual implicit assumption made by indicator-based approaches: knowing the systemic importance of single financial institutions is enough to understand how a shock affects the whole financial system. That is dangerously wrong. The failure of Northern Rock, a very systemically important bank, in late 2007 did not allow policy-makers to assess the subsequent systemic consequences. Even assuming that indicator-based approaches correctly identify bank relevance, there is not enough information to comprehend how the shock propagates through other institutions, as it is not known how they are linked to one another (Minoiu and Reyes, 2013). That is why a network-based approach could be more suitable for the problem.

**Dynamic stochastic general equilibrium models.** The second line of study investigates financial distress propagation through the use of dynamic stochastic general equilibrium (DSGE) models. These are able to simulate the main actors within an economy – households, firms, governments – and their interactions by the computation of differential equations and utility functions which describe their behaviour, according to neoclassical or neo-Keynesian economic theories (Sbordone et al., 2010). Because DSGE models allegedly possess the faculty to capture in a stylised way the dynamics of an economy in its entirety, several central banks have developed their own variants. Examples are the Bank of England (Harrison et al., 2005), the Bank of Canada (Moran, 2001) and the Central Bank of Chile (Chumacero and Hebbel, 2004). Usually, DSGE models do not include financial institutions in their set of actors. If they do, they are assumed to be perfectly identical and frictionless, merely acting as intermediaries in the flow of money among all other economic agents (Tovar, 2009). This is certainly questionable, given the not-so-irrelevant frequency of economic shocks sparked within the financial realm, such as the 1980-90s Japanese crisis, the 1980s Latin American crisis, the 1997 Asian crises, the dot-com bubble and the recent 2008 sub-prime crisis (Kindleberger, 2000).

Motivated by the gap in the literature, some scholars embarked onto the construction of novel DSGE models which would account for more sophisticated financial markets, thereby understanding how financial systems are impacted by shocks. It is likely that the most representative instance of this line of research is by De Walque et al. (2010). The trio devised a DSGE model with a heteroge-

neous banking sector, that is, with banks whose portfolios and profit-maximising behaviours are not identical. They also introduce endogenous default probabilities for both firms and banks, and allow for the simulation of banking regulations and liquidity injections in the interbank market. Such novelties permit the observation of the impact of external or internal shocks in the financial sectors, and would even enhance policy making through the experimentation of anti-systemic risk practices in the model. An analogous scholarly effort has been produced by Gerali et al. (2010). The paper introduces a DSGE model with financial frictions and imperfectly competitive financial markets. The assumption of banking existing within a market regime of perfect competition is relaxed, allowing for the existence of monopolistic structures. Theoretically, such a feature is instrumental in the simulation of realistic shock dynamics, as the size and market concentration of the financial institution may well affect the health of the overall system.

A DSGE approach to systemic risk models has severe limitations in assessing the systemic effects of shocks in finance. Two are explored in this paragraph. DSGE models equations are deeply rooted in neoclassical and, more recently, neo-Keynesian theories (Wickens, 2011). This means that any bias affecting such theories is also reflected in the computational models. The most critical assumption is that of general equilibrium: it is granted that there exists a steady state for the economy, towards which it tends to converge over time (Starr, 2011). The assumption is regarded as unrealistic. As seen in Chapter 2, closer empirical inspection suggests that economic and social processes are dynamic and in a state of evolvable change and multiple equilibria (Beinhocker, 2006). The practical implication of the general equilibrium assumption is an underestimation of the propagation of shocks within the system. Simulated shocks cause the system to diverge from its steady state, whereas the equations dictating the behaviour of financial actors will act in such a way as to allow the system to return it (Colander et al., 2009). This negative feedback is likely to reduce or nullify the positive feedback of the shock, acting as a bias. The second weakness relates to the over-simplicity featured in DSGE models. The one in De Walque et al. (2010) counts only two firms and two banks. Also Gerali et al. (2010) simulates only a very small number of financial institutions. Without a realistic number of financial actors, it is not possible to obtain informed accounts on how shocks affect the financial system. The main reason being that with a larger number of market participants, the shock dynamics may substantially change qualitatively. It is indeed, an emergent phenomenon.

**Network analysis.** The shortcomings of the two approaches explored above are mainly due to the disregard of the financial system's complexity. Both lines of contribution have been used in policy-making because of the simplicity of their elaborations, and the compatibility with already-existing policy instruments at the cost of compromising the ability to understand the dynamics of financial crises. The costs far outweigh the benefits and an alternative, more empirically-valid and elaborate, approach is needed. Complex network analysis, arguably, is the best candidate (Haldane and May, 2011).

Some early studies showed the importance of interbank relations with regards to systemic risk. Allen and Gale (2000) provided the theoretical foundations of how shocks can spread through a network of financial institutions, whereas Freixas and Parigi (1998) and Eisenberg and Noe (2001) devised the very first financial shock contagion models. The latter study inspired the empirical application of network analysis in financial systems. Examples include (Upper and Worms, 2004) for the German financial market, and (Boss et al., 2004a) for the Austrian one. The lack of appropriate data slowed down the advancement of this line of research. Information about interbank linkages, such as short and long-term lending and portfolio commonalities, is often lacking (Upper, 2007). In the last few years, though, scholars have produced seminal work with the potential to properly answer the research question, thanks to new data gathered by central banks and the appearance of novel statistical methods to treat partial information (Servedio et al., 2004).

Two research papers are of key interest. The first is Battiston et al. (2012), which develops a new measure of systemic risk known as DebtRank. DebtRank is a recursive algorithm that is able to provide estimates of the total systemic loss caused by the failure or distress of one or more banks. The second research study was produced by Roukny et al. (2013), who explore the stability of different network topologies against default cascades, that is, a chain of financial institutions failing due to a market shock. The authors simulate default cascade dynamics in banking networks under several market conditions, such as market illiquidity and sudden capital ratios changes. The topologies of the simulated networks are scale-free or random. Their results show that no architecture can always be preferred over another, yet their methodology provides a valuable starting point to assess how market shocks propagate among financial networks.

In addition, these works are not immune from complications, giving rise to some of the gaps that this research intends to fill. The most critical still being the poor availability of interbank data. Battiston et al. (2012) rely on a chunky dataset, where only partial data is available and most probably does not represent the totality of the interbank financial network. Other related works, such as Huang et al. (2013)

and Caccioli et al. (2014), are solely based on theoretical models, with no empirical information, thus undermining the validity of their insights. This case study extends the literature by providing an empirical application of complex network analysis in financial distress contagion.

## 4.2 Data and network reconstruction

This case study makes use of data referring to the top 182 financial institutions by total asset size as in the year 2008, that are headquartered in the European Union, and publicly listed on a stock market. The data is collected from BankScope<sup>3</sup>, a private data provider service by the Bureau van Dijk company. Two years are taken into consideration in the case study analysis: 2008 and 2013. For both time frames and for each bank, the following information is collected.

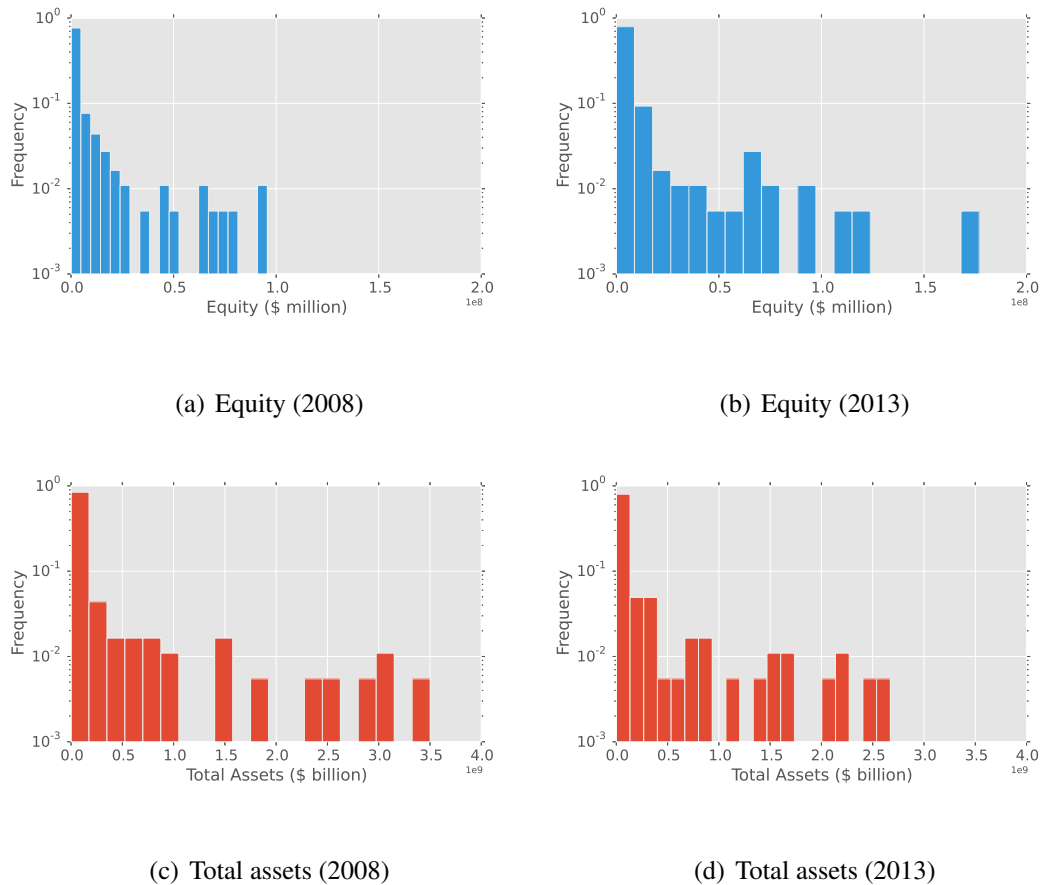
1. *Equity*. Equity is the capital received by a bank for a share of its ownership. It is regarded as a key proxy for bank stability. In financial policy making, banks are deemed to be in default if the amount of asset losses are greater than their total equity (Avellaneda and Zhu, 2001). While in reality the conditions leading to the closure of a bank may be more intricate, this stylised rule is the standard definition of a default both in academia and policy, and is adopted by this study too.
2. *Total asset size*. Total asset size refers to the sum of the value of all assets owned by the financial institution in a given year. The asset portfolio of a bank is generally diversified across several classes, such as sovereign securities, company shares, loans, and derivative instruments (Van Horne et al., 1975). A sudden decrease in value in any of the assets contained in a portfolio can cause distress within the bank, if large enough.
3. *Total interbank assets*. Total interbank assets are a subcategory of total assets, and refer to the sum of the value of loans given to other financial institutions. Interbank lending is not uncommon and is regularly used to accumulate cash against liquidity risk, i.e. the likelihood of experiencing a severely high number of clients demanding cash back than previously expected (Afonso et al., 2013). When individual banks lend to another bank, the former is exposed to the latter, with the risk of being distressed if the latter is unable to return the amount loaned plus interest within the agreed timeframe. Interbank lending is regarded as the most rapid way by which a financial shock propagates

---

<sup>3</sup><https://bankscope.bvdinfo.com>

through a system of banks (Rochet and Tirole, 1996).

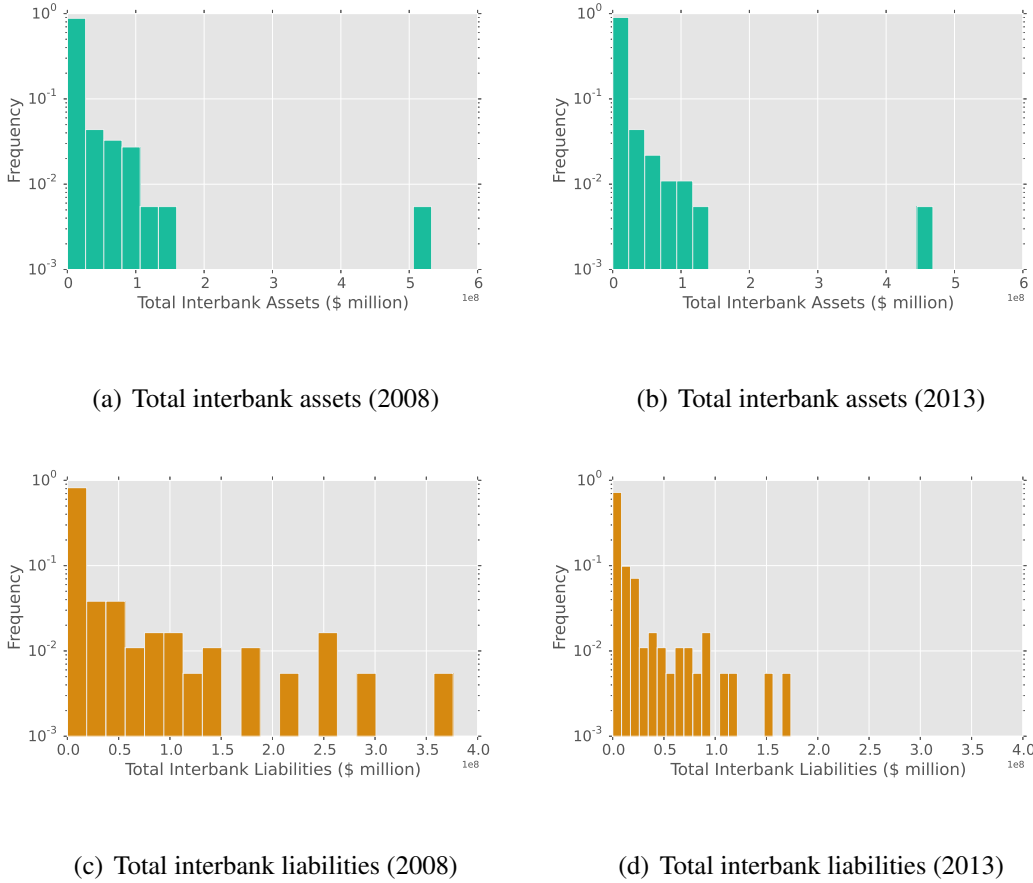
4. *Total interbank liabilities.* Total interbank liabilities are the sum of the value of the loans owed by one bank to others. It is simply the other side of interbank lending.



**Figure 4.1:** Frequency distributions of equity and total asset size, years 2008 and 2013.

Figures 4.1 and 4.2 provide a snapshot as to the nature of the distribution of the four key variables involved in this study. In both years, the distribution of all quantities is heavily skewed to the right. This implies that only a minority of banks accounts for a significant proportion of the total assets in the system. The same goes for interbank assets and liabilities. Also, it is possible to notice a regime shift between 2008 and 2013. The distribution of equity becomes more skewed to the right, implying that more banks adopt higher capital values to shelter themselves against shocks. Also, lower levels of interbank assets and liabilities are recorded, meaning the overall amount of lending has diminished. In other words, the banks

in these financial systems were more connected and exposed to one another at the onset of the crisis than towards its end.



**Figure 4.2:** Frequency distributions of total interbank assets and liabilities, years 2008 and 2013.

When it comes to interbank lending, there exists an actual network of dependencies from one bank to another. This is hidden, however, as data at such a level of granularity is not available, not even for regulatory bodies. Yet, given the data in hand, it is possible to reconstruct probabilistically the hidden graph. In this case study, the fitness model is chosen. In a network where nodes are the financial institutions and the edges represent a bank lending to one another, total interbank assets and liabilities function as the sums of out- and in-strengths (i.e. the weights) for each node. The problems turn, therefore, into the evaluation of a contingency table in which the marginals are known.

It is necessary, however, to elaborate a hypothesis about the density of the latent network. Incidentally, extensive research has been performed on the network structure of interbank networks, which tend to feature extremely similar topologies



and densities, regardless of the country and specific context. On average, interbank networks appear to have a 5% density (van Lelyveld et al., 2014), implying that only 5% of all possible connections occur. In this case study, therefore, the fitness model would have a set sum of links equal to 1,665 links.

The second input of the fitness model is the proxy fitness  $y$ . Because both the values of interbank assets and liabilities inform the propensity of a bank having connections, they are conjunctively used. More specifically, their normalised sum is used as the fitness proxy, hereby named  $IB$ . The fitness equation 3.3 thus specified becomes


$$p_{ij} = \frac{zIB_iIB_j}{1 + zIB_iIB_j}, \quad (4.1)$$

where  $z$  is found by numerically solving


$$1665 = \sum_i \sum_{j \neq i} p_{ij}. \quad (4.2)$$

Once  $z$  and the probability for each pair of banks are evaluated, they are used to generate probabilistically an ensemble of 100 networks for each year. The ensemble has a mean density of 1,665.5 nodes, with a standard deviation of 28.8. Due to lack of information about what could determine the direction of a link, this is determined with a 0.5 chance of originating from one bank rather than another. The weights are distributed across the links according to the iterative proportional fitting algorithm described in section 3.1. The correspondent marginals are, obviously, the total interbank assets and liabilities.

Figure 4.3 shows the topology of one of the ensemble networks. It is important to note that it features, predictably, a core-periphery structure. Banks with relatively higher rates of lending and borrowing end up being more connected than the rest, forming the highly dense core in the graph. The intuition is confirmed by the very skewed degree frequency, shown in figure 4.4. The ensembles are the starting point for the validation experiments, the settings of which are described in the next section.



**Figure 4.3:** Interbank network generated through the fitness model.



**Figure 4.4:** Degree frequency of a generated interbank network.

### 4.3 Influence measurement specifications

In the domain of this case study, influence is specifically defined by the systemic impact that a financial institution has on the whole system, should it experience distress. The type of information that flows through the network relates to the inability to repay interbank loans, which in turn hurts the asset portfolios of the exposed banks. How distress propagates through financial networks is modelled through a flow function, which is defined in the paragraphs below. The framework adopted is then tested against a competing measure of influence, also based on graph theory, known as default contagion.

Recall from section 3.2 that a model-based influence metric requires a flow function which characterises how exposure occurs from one node to another. The flow function is then used to transform the adjacency matrix into a flow matrix, essentially a network with the same topology as the original but where the nodes already embody how information flows. In interbank networks, the extent to which, say, bank  $a$  is exposed to bank  $b$  depends upon the loan amount from  $a$  to  $b$ , relative to the equity level registered by  $b$ . The reason is intuitive: the loss incurred by  $b$  if  $a$  does not repay all or part of its debts is covered by  $b$ 's capital. Should the total amount not repaid be greater than the equity, bank  $b$  defaults. This can analytically be translated by a flow function specified as below,

$$f(l_{ba}, c_b) = \min \left[ 1, \frac{l_{ba}}{c_b} \right], \quad (4.3)$$

where  $l_{ba}$  is the estimated amount of interbank loan from  $b$  to  $a$ , and  $c_b$  is the capital or equity level of  $b$ . Equation 4.3 defines a ceiling value equal to 1, implying a total default when the amount of interbank loans is equal or greater than the capital level.

With the flow function thus specified, the flow matrix  $F$  can be elaborated for all the adjacency matrices in the ensemble. The next step is to simulate a distress event departing from one or more nodes, and allow it diffuse through a flow matrix. Let two state variables be assigned to each node in the network: a continuous variable  $h_i \in [0, 1]$ , and a discrete variable  $s_i \in [U, D, I]$ . The first determines the relative quantity of the impact that flows through the matrix, while the second one determines whether a node is respectively undistressed (i.e. not yet hit by a shock), distressed or inactive. At time 1, one or more nodes for which influence shall be measured are selected, belonging to set  $S_f$ . They are subject to a shock  $\psi_n$  and defined as distressed, while the rest is assigned a vector of 0 values  $\psi_0$ .

$$\begin{aligned}
h_i(1) &= \psi_n, & \forall i \in S_f, \\
h_i(1) &= \psi_0, & \forall i \notin S_f, \\
s_i(1) &= D, & \forall i \in S_f, \\
s_i(1) &= U, & \forall i \in S_f.
\end{aligned} \tag{4.4}$$

For each node, then, the shock received at time  $t$  is determined by the shock it had at time  $t - 1$ , plus the distress flowing through the neighbours that have been hit,

$$h_i(t) = \min \left[ 1, h_i(t-1) + \sum_{j|s_j(t-1)=D} F_{ij} h_j(t-1) \right]. \tag{4.5}$$

At each time step, the state variable  $s$  is also updated, such that the distressed nodes become inactive, not being able to propagate any more shocks into the flow matrix:

$$s_i(t) = \begin{cases} D, & \text{if } h_i(t) > 0, \quad s_i(t-1) \neq I, \\ I, & \text{if } s_i(t-1) = D, \\ s_i(t-1), & \text{otherwise.} \end{cases} \tag{4.6}$$

The simulation ends when all reachable nodes become inactive. The final influence of nodes in set  $S_f$  is sum of the amount of capital lost, relative to the total capital present in the system before the contagion. It is simply a specification of equation 3.26. The closer its value is to 1, the greater the initial nodes' impact in the system. In this very case, it would mean that the entire capital value of the financial system was burnt,

$$H_{S_f}(T) = \frac{\sum_i h_i(T) c_i - \sum_i h_i(1) c_i}{\sum_i c_i}. \tag{4.7}$$

This very specification of the model-based influence metric is known as DebtRank 2.0, and has been developed by the author in collaboration with Professor Stefano Battiston and his team in Battiston et al. (2016).

The validity of DebtRank is here tested against another conventional metric, that of default contagion (Moussa, 2011). This is the standard, graph-model based measure of financial impact currently used in policy making. Essentially, it is possible to describe default contagion in the exact analytical manner shown above. Both methods adopt the use of network topology, and simulate contagion. The crucial difference lies in how the shock propagates. In the default contagion case, the distress reverberates through the system if and only if the nodes default. Analytically, the only difference would then be in equation 4.6:

$$s_i(t) = \begin{cases} D, & \text{if } h_i(t) \geq 1, \quad s_i(t-1) \neq I, \\ I, & \text{if } s_i(t-1) = D, \\ s_i(t-1), & \text{otherwise.} \end{cases} \quad (4.8)$$

## 4.4 Experiment settings and results

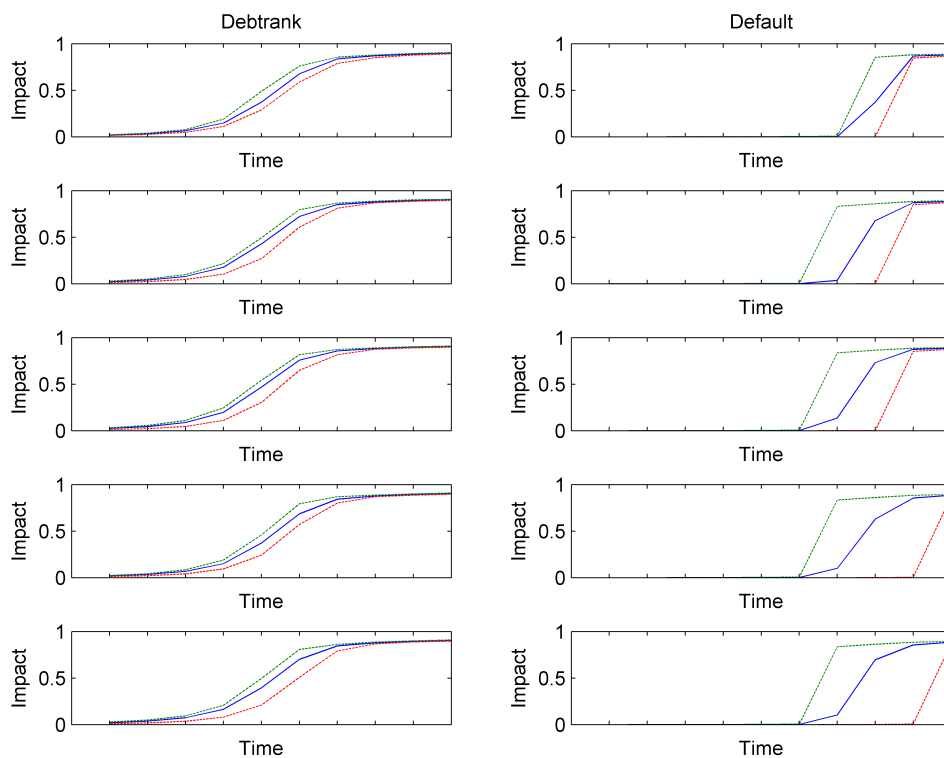
This section introduces a synthetic experiment that tests the performance of the model-influence metric in detecting the impact of the five selected banks against the more conventional default contagion algorithm. The first part of the section introduces the experiment settings. The second part shows and interprets the results.

**Experiment settings.** Into this experiment is built an artificial time series of network snapshots based on the 2013 ensemble, with gradually worse systemic conditions which will render each of the five banks extremely systematically influential. In other words, by the end of the time series, the failure of each bank would cause almost the entire financial network to collapse. Operationally, this scenario can be simulated in two ways: by gradually increasing the amount of interbank lending, or gradually decreasing the market capitalisation of the banks. In both cases, the relative dependency from one bank to another increases, thereby causing greater distress in the case of a shock. In this experiment, it was decided to opt for the latter method. Ten time steps are constructed. At each step, the capital of all banks becomes 30% of the value in the previous step. Analytically:

$$c_i(t) = 0.3 \cdot c_i(t-1). \quad (4.9)$$

Apart from the change in capitalisation, all conditions remain the same. The networks at each time step feature same topologies and balance sheet attributes. The experiment then proceeds as follows. For each time step, each starting node  $i$  is assumed to fail, which is analytically translated into  $h_i(t) = \psi_n = 1$ . The proportion of the system's losses caused by the failure is evaluated using both model-based influence and the default contagion algorithm. As the scenarios are simulated for all 100 graphs in the ensemble, the losses shown are the arithmetic averages across them, with maximum and minimum variables being respectively the upper and lower uncertainty bounds.

**Experiment results.** Figure 4.5 shows values of DebtRank (left columns) and default contagion (right columns) for the three ensembles. The horizontal axes are time steps, while the vertical ones represent the proportion of total market capitalisation in the system which has been burnt as a consequence of bank failure. It is evident that DebtRank can capture the increasing systemic importance of the collapsed banks in a more continuous and sensible way than the default contagion algorithm. In this experiment, therefore, DebtRank is a more sensitive measure of the financial institution's influence in propagating a financially-adverse event.



**Figure 4.5:** Evolution of the DebtRank influence metric and default contagion for five selected banks.

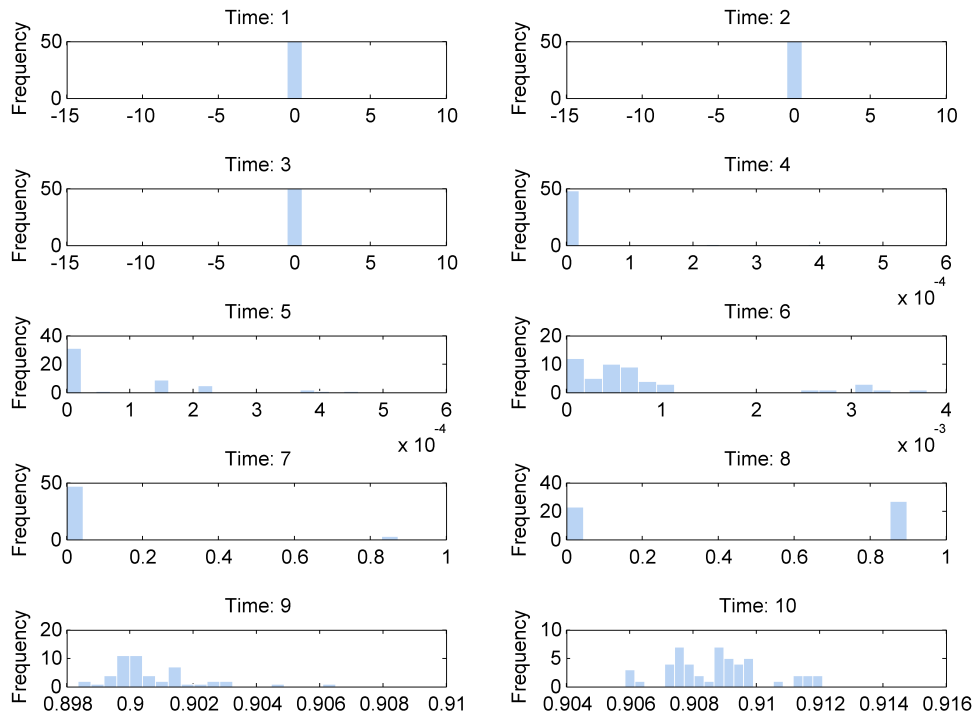
From a deeper scrutiny of the results, the following observations can be made:

- (i) Both metrics converge at the final time steps when the financial institution is already capable of allowing almost all of the interbank network to collapse. This shows the value of the default contagion method in capturing the systemic importance of a bank and, therefore, the fragility of the system with

respect to it. Nevertheless, it does this too late - when the losses are already at great magnitude.

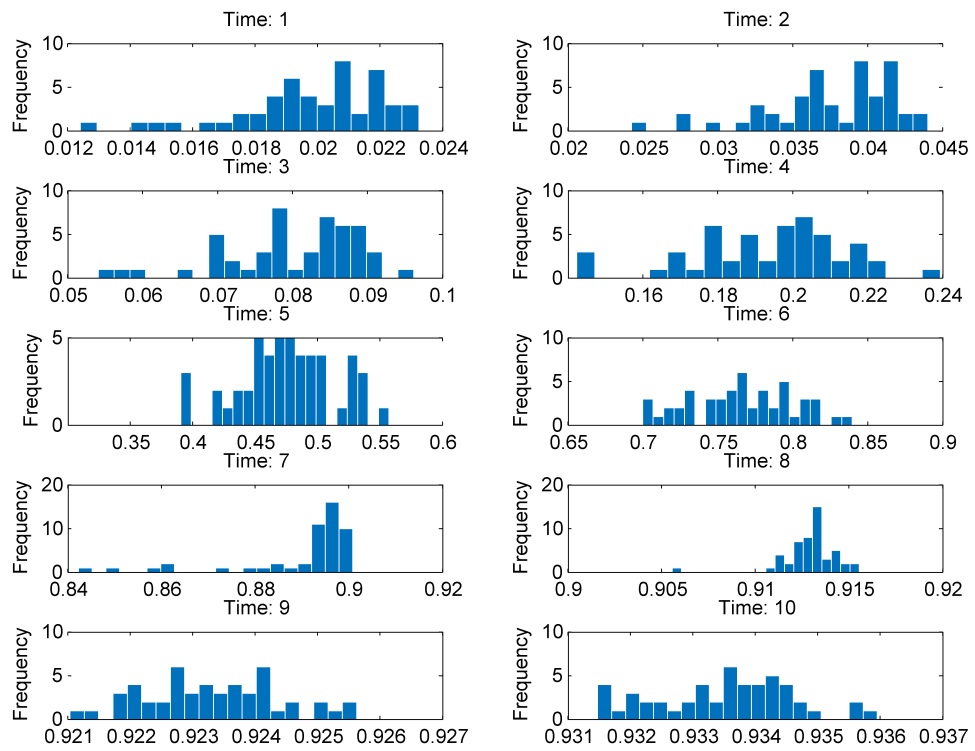
- (ii) When the default contagion algorithm still registers zero systemic influence, this research's model-based metric features increasingly steep values, thereby acting as an early warning signal of the later total collapse. It is possible to note the increasing nonzero values of the impact at time step 4 for all nodes, for instance.
  
- (iii) The uncertainty intervals of the default cascade algorithm are extremely large, while those of DebtRank are narrower, with a convergence towards the final time steps. This suggests a greater ability of this metric to capture the systemic importance of banks despite the lack of knowledge of its exact links configuration, a property that is extremely valuable in a context where the lack of fine-grained data is persistent.

Point (iii) becomes more apparent by analysing how the qualitative properties of the impact distribution estimated by the algorithms change over time. As a demonstration, figures 4.6 and 4.7 show the variation of this quantity for node 1 in the third ensemble. The default contagion algorithm features zero extremely minute values until time step 6, after which it begins to exhibit a quasi-binary distribution. Only at the final two time steps does it generate a more continuous range of values. The same does not hold for DebtRank which at every time step exhibits a continuous distribution of variables, making the construction of uncertainty intervals a more meaningful exercise.



**Figure 4.6:** Frequency distributions of the influence evaluated by the default contagion algorithm for node 1.





**Figure 4.7:** Frequency distributions of the influence evaluated by DebtRank for node 1.

## 4.5 Practical implementations

The last section performed a validity experiment that demonstrated the greater sensitivity of the application of the NIIF-based influence metric against the conventional default contagion algorithm. Under stressful conditions, the influence metric proposed by this thesis is able to capture the ability of a node to propagate the shock in a more granular way. Over the course of this research, this particular implementation of the NIIF framework has been presented in two papers (Gurciullo, 2014; Battiston et al., 2016), and has gained the interest of policy makers who perform macroprudential exercises in central banks (Hartman, 2014). The following pages present two key instances of the real-world policy applications already experienced by the framework.

**Measuring the vulnerability of banks.** By being able to simulate distress propagation in networks, the framework may well be harnessed to perform systemic stress-testing exercises, thereby inspecting how vulnerable a financial institution is under a certain macroeconomic scenario. Vulnerability is the potential amount of capital losses that can be incurred by the institution, conditional to a pre-defined scenario. This variety of stress-testing is here performed to calculate the potential capital losses of three banks included in the dataset: HSBC, Santander and Intesa San Paolo. The ensemble of networks relative to 2013 is subject to a global macroeconomic shock, engendering the initial loss of 0.05% in asset value for all banks in the network. The shock propagation is simulated as described in section 4.3, and the relative capital losses incurred by the three selected banks are conglomerated in distributions. With the distributions in hand, it is possible to utilise two typical risk assessment instruments, Value at Risk (VaR) and Conditional Value at Risk (CVaR). VaR measures the maximum potential loss of a financial value over a pre-determined time period within a confidence interval (Linsmeier and Pearson, 2000). A 5% VaR, for instance, is simply the 95% percentile of a distribution. Analytically, it may be defined as

$$VaR_{\alpha} = \inf\{l \in \mathfrak{R} : (P > l) \leq 1 - \alpha\}. \quad (4.10)$$

The CVaR is an extension of VaR, and yields the expected loss conditional to a potential loss. It is essentially the expected value of the segmented distribution starting at a pre-specified  $x\%$  VaR (Rockafellar and Uryasev, 2002).

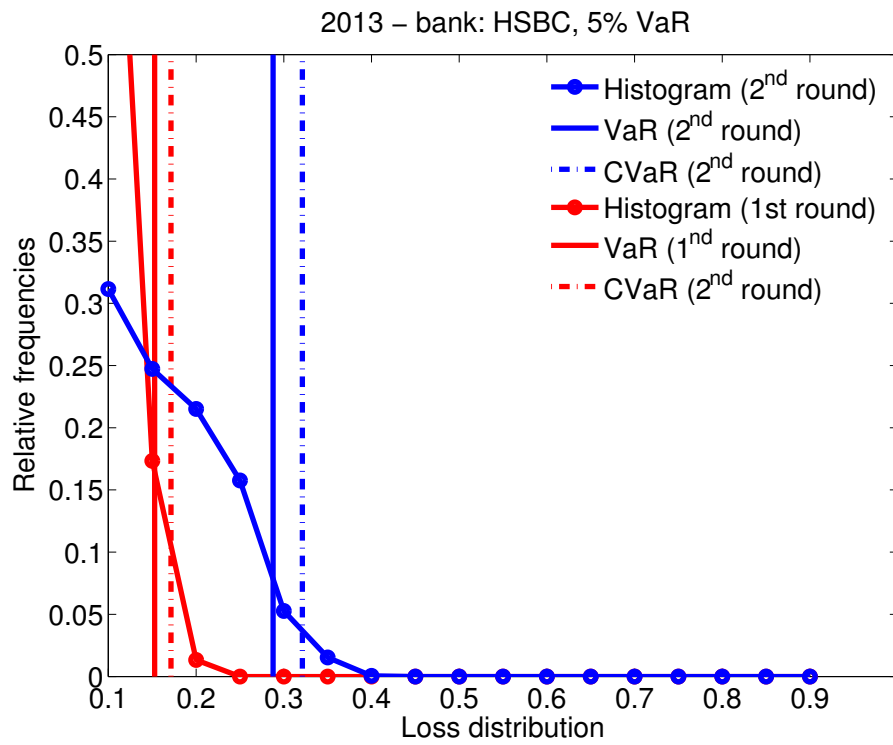


Figure 4.8: Capital loss distribution and 5% VaR, CVaR for HSBC.

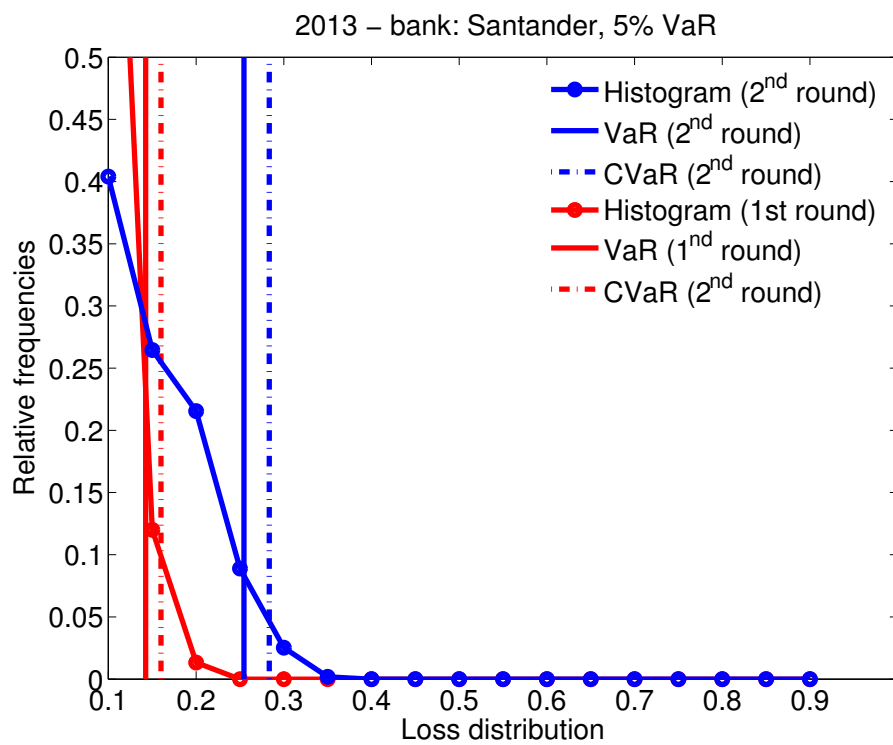
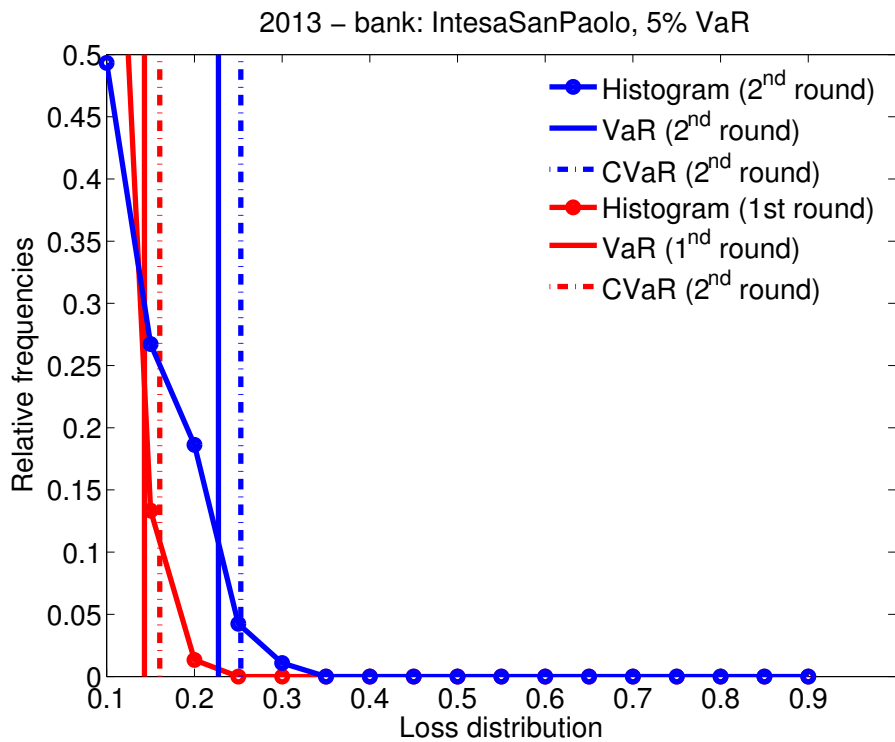



Figure 4.9: Capital loss distribution and 5% VaR, CVaR for Santander.




**Figure 4.10:** Capital loss distribution and 5% VaR, CVaR for Intesa San Paolo.

Figures 4.8, 4.9 and 4.10 show the results of the stress testing exercise. Red lines show the capital loss distribution in the first round, that is, when the shock contagion is not taken into account. The blue lines represent the loss distributions after the contagion occurs. It can be noted that taking into account the network effect does increase the maximum potential capital losses, measured both with 5% VaR and CVaR. The results are easily interpretable. Given the simulated scenario, HSBC has a 5% chance of losing almost 30% of its value in equity as a result of the interconnections with other banks. The expected loss goes above 30% if the CVaR is used. While the figures might not imply the default of the company, they still suggest a quite significant amount of loss which would have been ignored were the network structure of the complex system not considered.

**Identifying systemically influential financial institutions.** A further, immediate application of this NIIF framework instance is to extend the work undertaken in section 4.4 to identify the systemic influence of large institutions. This enables policy-makers to dedicate the appropriate amount of attention to those banks in the position of harming financial stability at greater scales.



(a) 2008



(b) 2013

**Figure 4.11:** Top ten most systemically influential banks, years 2008 and 2013.

An instance of this exercise is shown in figure 4.11. Each of the 183 banks in the ensemble has been assumed to default, and median global capital losses have been evaluated. The top most impactful banks in 2008 and 2013 are reported in the graphs. The first thing to be noticed is that the absolute levels of influence decrease significantly after the crisis. In 2008, the default of the Royal Bank of Scotland could have caused a damage equal to 60% of the total equity of all the institutions in the financial network. In contrast, the default of Unicredit, the most influential bank in 2013, would have led to a loss equal to 13% of the global equity. This is a clear indication that the wiring of the financial markets in 2008 was extremely fragile, as history now teaches. The measure of systemic influence proposed can be adopted as a complementary instrument to the more traditional ranking performed by the Bank of International Settlements and reviewed in section 4.1.

## 4.6 Case study discussion

This Chapter has introduced the use of the NIIF framework with regard to financial systems, more specifically, to tackle the problem of identifying how financial distress reverberates through a European interbank network and how to measure the extent to which one or more institutions are responsible for its propagation. Because of the nature of the data collected, the fitness model has been chosen as a good method to reconstruct an ensemble of interbank lending instances. The ensembles, then, have been subject to the model-based influence metric which can measure the systemic influence of an individual bank in a more sensitive way than another conventional method, the default contagion algorithm. The work presented here has already found room for application in macroprudential policy-making. Two instances of the policy instruments that can be implemented have been presented. The first relates to the possibility of performing stress-testing scenarios in order to investigate the vulnerability of the institutions. By simulating any pre-determined scenario, it is possible to elaborate a loss distribution for each institution, and calculate their Values at Risk or Conditional Values at Risk. It has also been demonstrated how the framework can supply a novel ranking of the most systemically-influential banks. This provides policy-makers with the ability to identify those institutions that, in case of distress, are capable of causing the most damage to the whole system.

This case study application is not immune to limitations, which may impede the validity of its insights. Three are discussed in this paragraph. First, the applica-

tion is sensitive to the assumptions of the fitness model used to generate the network ensemble. In truth, in Battiston et al. (2016) extensive sensitivity analysis has been performed by changing the latent density of the network, probably the most significant assumed input to the model. Interestingly, the stress-testing scenarios, even with different densities, change only from the fourth decimal digit. Yet, users ought to be careful about the latent structure of the network, which may feature multiple cores rather than just one, as in this instance. In a European context, interbank networks could have banks of the same nationality being more densely-connected, thus forming multiple cores around the largest, international, core. Such a structure might significantly change the results of a stress-testing scenario, as the final losses would be more sensitive to the country from which the shock has originated. At the time of writing, there is little data accessible that would enable the construction of a multi-core interbank network. With the right data, the fitness model can be modified by calculating different  $z$  values for countries of different nationalities, thereby generating the networks by having different linking probabilities conditioned to this property.

A further important limitation is the specified form of the flow function. In the case study, it has been opted for a rather simple flow function, where the distance to default is dictated by the magnitude of interbank loans relative to the institution's equity. In the real world, a bank might survive even when its losses are greater than its equity, due to policy or shareholders' actions. This and other factors should be modelled within the flow function, so as to capture a more realistic dynamic of how distress diffuses through the interrelations.

A final consideration relates to the frequency of the raw data used to reconstruct interbank networks. In this instance, annual quantities have been used. In reality, the degree to which a financial distress flows may be dependent on the monthly and daily fluctuations of the interbank relations. Further research should find room for framework application with higher frequency data. The results would be able to yield more timely and pertinent insights to policy-makers, therefore improving the chances of better stability within the financial system.





## Chapter 5

# Second case study: Influence in parliamentary debates

Policy-making is the process by which a government addresses a societal goal or problem, by designing and implementing a solution in the form of a law, regulation, or a programme (Birkland, 2014). Why one particular policy is instantiated rather than another remains an open-ended question in the study of policy processes. From a complex systems perspective, the policy outcome can be seen as the result of both the resources owned by the political system, and the dynamics of the intersecting interests of policy actors (Lempert, 2002). For example, the conditions by which free school meals are introduced in the UK education system depends upon both the capacity of the government of doing so, given budget constraints and goals, as well as the relative influence of those actors who promote or obstruct the policy in the pre-implementation phase (Murphy et al., 2011). The latter side of the policy process is of special interest to political scientists and practitioners. Measuring the degree to which a political actor manages to influence public discourse in institutional settings can provide a significant leap forward in understanding (and, perhaps, predicting) political outcomes, with valuable insights to be used when assessing the accountability and transparency of governmental bodies.

The case study application of this thesis attempts to demonstrate the use of NIIF to measure the extent to which political actors influence parliamentary debates about certain policy issues. More specifically, the case study performs the analysis onto a dataset containing speeches made at the UK House of Commons during the last three governments between 2001 and 2014. For each governmental cycle, six policy themes are identified and taken into account. A novel method is used to measure influence and tests its results with some of the insights yielded by the relevant literature. The chapter is structured into six main sections. The first supplies the necessary background needed to understand the importance of study-

ing policy processes, and briefly reviews some of the key literature on the topic. It also clarifies how this Chapter intends to contribute to it. This is then followed by a section describing the dataset of interest, accompanied by an overview of the preprocessing work performed to extract the six policy themes for each government cycle, and how the corpus is prepared for network reconstruction. Using the transfer entropy method, for each policy theme a single network is constructed using the sequence of relevant parliamentary speeches as inputs. The fourth section introduces the analysis performed on the reconstructed hidden networks. The influence of each parliamentary speaker within a given policy network is measured topologically, by calculating their eigenvector centrality. The distribution of centralities is then evaluated against the role the speaker features in the House of Commons. The test determines whether the most influential actors are those with Cabinet or Office positions, and the hypothesis is not falsified. This provides more ground for theoretical and empirical contributions that have suggested a more influential role for government leaders than backbenchers. The fifth section comments on the practical policy application of the methods shown in the case study. It is suggested that analysing the topology of policy networks can yield empirical ways to assess the accountability of elected House members. Also, a discussion takes place about how the NIIF framework can be used to evaluate the overall concentration of influence in debates, thereby acting as a novel measure for government transparency and power concentration. As the study is only the first application of NIIF in this context, it is not immune to setbacks. These are discussed in the concluding section, together with suggestions for amelioration and further research.

## 5.1 Case study background

Having a grasp of how political actors are in a position to exert some form of influence in government is crucial for the well-being of society. Political systems, albeit often considered secondary to macroeconomic processes, have the capability of determining the future of a collective for generations, for better or worse. While a financial crisis might have an immediate impact lasting no more than a decade, an important political decision can impact severely in the long term. The case of Brexit exemplifies. Despite the significant economic consequences of leaving the European Union, party political factors have proven more significant in defining the final decision which was unexpected by many when the referendum was first proposed (Dagnis Jensen and Snaith, 2016). Just as in financial systems, a system of political actors shares information and resources, with final systemic outcomes that

have a wide-ranging effect on society. In comparison to political systems, though, financial system dynamics of influence are easier to understand. If a bank loses money, it will prompt a loss in other banks as well, given counterparty exposure. In politics the mechanisms of information diffusion are much less explicit and involve both endogenous factors, such as party politics, and exogenous ones, such as media and public opinion (Sabatier and Weible, 2014). Furthermore, the diffusion of information occurs at multiple levels, thus adding complexity to an already complex task. The units of information diffused may be very specific policy proposals, or more abstract policy views, or even ideological positions.

In this first application of NIIF in political systems, a practical and modest stance is taken, with the focus solely on policy themes as the unit of analysis. The underlying assumption throughout the case study is that it is possible to think of political systems as complex collectives of actors who share information about a policy theme, by means of a latent network that determines the structure of inter-actor influence. Furthermore, the case study focuses on a very specific type of political system, parliament, where the components are its members. In the following paragraphs, some of the key efforts in the study of influence in political systems are reviewed. A specific discussion takes place on how parliamentary speeches have already provided a novel source of study of political interest and positions.

**Influence of interest groups in political systems.** Influence in political systems is a concept that can be approached from a great variety of theoretical and methodological angles. Generally, research in the field has occupied itself in investigating an operationalised definition of influence based on clear empirical outcomes, namely voting results and bills passed. In this area, the study of interest groups and their impact on policy-making has received significant attention. By interest group, it is meant an organised or unorganised set of actors with a common political interest, and which uses its resources – be they financial or informational – to drive the policy process towards an outcome deemed favourable (Berry et al., 2015). Examples of interest groups include formal lobbies, spontaneous groups of citizens who petition for a political cause, and intra- or inter-party collectives of political members who are interested in common policy proposals. Clearly, the underlying definition of political system in this area is a broad one, since it includes both official actors, such as MPs, and industry or civil society agents (Denzau and Munger, 1986).

A great amount of research has been dedicated to tracking the impact of financial resources that flow from interest groups to actors who have the ability to action

specific policies. The study by Silberman and Yochum (1980) is probably the first exploratory attempt to track this sort of flow, with regards to the US Congress. A more recent, extremely interesting study of the phenomenon in the United States is presented in Baumgartner et al. (2014). The authors make use of the insights provided by a large sample of interviews on 98 randomly selected cases of lobbying to understand the state of interest group influence in Congress. Their major finding is a profound disjuncture between the concerns of the public, and those of the actors directing financial resources into the public body. The study validates the findings by Gilens and Page (2014), who tested whether the US government was a majoritarian electoral democracy. Their multivariate analysis on variables related to 1,779 policy issues suggesting that the policy process found in the United States appears to be akin to an economic elite oligarchy. In other words, in the US political system, formal political actors and their policy decisions seem to be greatly influenced by exogenous interest groups.

Interest groups and their influence has also been the subject of study also with respect to the European Union legislative process. Michalowitz (2007) introduces a theoretical framework that determines EU interest group influence according to the degree of conflict with the incumbent legislation on the topic of interest, the structural conditions of interest exertion, and the type of interest pursued. The author tests the hypothesis on case studies relating to IT and transport policies, finding that policy outcomes heavily depend upon the legislator's initial interest and technical knowledge. Dür and Mateo (2012) provide a comprehensive map of the variety and scope of interest groups around EU public bodies. Based on the quantitative analysis of 1,417 surveys about organisations in Germany, Ireland and Spain, they find that resource-rich, national groups have greater access to executive bodies such as the European Commission, while citizen-based groups tend to have greater access to EU Parliament members. Overall, their insights suggest that national interests tend to have a significant leverage on EU policy making. The contention contributes to the line of thought stating that the European Union is a multi-level polity system, severely biased by the interests of single nations (Klüver, 2013).

The analysis of interest group influence has provided invaluable information about policy process dynamics, and supplies evidence that enables both scholars and citizens to effectively understand if and to what extent governments actually operate in the public interest rather than of a minority. Interest groups are, however, only part of one level of political systems. It is just as important to identify what actors in the very policy process contribute more to the individual policy idea. This

case study focuses on this level of analysis, which implies that the units of the system are the very individuals who compose the government or legislative bodies. Efforts focusing on these issues are reviewed in the next paragraphs.

**Members' influence in political systems.** Estimating the degree of influence exerted by one or more political actors in a governmental system is a difficult task because of the latency problem. As has been stressed in several sections of this thesis, the structure of influence relations of this type is not observable, at least with current data collection methods. Public policy scholars have thus turned to data which are assumed to act as a manifestation of those latent factors. That which is gaining substantial attention is textual data. Political speeches, statements, and bills contain invaluable information about the actor's stance with respect to policy and ideological concepts, and with natural language processing it is possible to extract them (Grimmer and Stewart, 2013).

Early attempts on this side include Laver and Garry (2000), who use hand-coded text analysis to estimate the policy positions of party manifestos in Britain and Ireland from 1992 to 1997. Their results enable them to obtain a first glimpse of what parties feature policies with a higher salience in the respective parliaments. They also pave the way to the subsequent computerised text analysis of policy positions. More recent research includes Baturo and Mikhaylov (2014), who uncovered the policy position of Russian regional leaders analysing 233 annual gubernatorial addresses. The scholars use the natural language processing methods first exposed in Laver et al. (2003) and Slapin and Proksch (2008) to test how leaders' positions change between 2009 and 2011. The findings suggest a clear pro-Putin stance across the group in 2011, implying that the political figure has attained over time a larger influence within the Russian executive system. Another sophisticated research effort is found in Tsur et al. (2015). The team is able to quantify the degree of control exercised by US parties on policy topics through a methodology based on probabilistic topic modelling (Blei, 2012). The scholars inspect four years of congressional statements, and manage to recover - in an unsupervised fashion - instances of spinning, i.e. the process by which political actors induce a cognitive bias through consistent linkage between a topic and specific context (frame). The study is not only able to measure the degree of party influence on a policy topic that has been subject to political spin, but also functions as a tool for government accountability.

A significant amount of research has also been dedicated to unearth the influence of political actors on final voting outcomes. A landmark study was published

by Snyder Jr and Groseclose (2000), who explored the influence of parties on the voting patterns at the US Congress. The duo constructed an estimation model of the proximity between official party position and individual roll-call voting, on a dataset referring to the years 1871-1998. They found a strong party influence in virtually all years. Party influence appears to be stronger on budget resolutions, tax policy, social security and national debt; it is weaker on moral or religious issues, such as civil rights. Theme-based party influence seem evident also at the European Parliament. Klüver and Spoon (2013) explored 400,000 vote decisions across 1948 roll-call votes, and measure the voting patterns of individual MPs against party-level data coming from the Euromanifestos Project (Schmitt and Wüst, 2012). Their conclusions indicate that European parliamentary members are likely to defect their European Party if the issue at hand is deemed important and in contrast with the national party's interests. National parties, therefore, are relatively influential in the European policy-making process.

The research surrounding the study of political influence has managed to overcome the latency problem in clever and effective ways, shedding some light on the topic, while engendering important practical implications. Nonetheless, what still remains poorly understood is the role of individual political members within the hidden web of influence that permeates the policy making process. Only very recently have attempts on this front been presented to the scientific community. An example is Lum et al. (2013), who use topological data analysis on voting patterns in the US Congress, reconstructing a network of House speakers where the edges determine voting similarity. The structure of the network closely resembles the two-cluster party division between Republican and Democrat. Yet, by taking into account only voting results, the extracted network ignores the potential channels of information diffusion that occur before a bill is passed. This case study attempts to fill the gap by reconstructing the hidden networks of House of Commons speakers, using parliamentary debates. The goal is to quantify the extent to which parliamentary members are able to contribute to a political system surrounding a certain policy theme or idea. The next section concerns itself with clarifying the data on which the analysis is based.

## 5.2 Data and network reconstruction

In order to untangle the latent structure of influence in the UK House of Commons with respect to policy themes, the case study makes use of legislative debates. As expressed by Proksch and Slapin (2012), debates can be seen as the manifestation of the interests of the speakers who carefully prepare them according to strategic considerations related to parties and individual preferences. In this Chapter, a corpus of speeches is used containing the minutes of all debates occurring in the House of Commons over three governments: the Labour government between 2001 and 2005, the Labour administration between 2005 and 2010, and the Conservative government between 2010 and 2014<sup>1</sup>. The corpus is a subset of a much larger one dating back to 1935, extracted by Alexander Herzog from a transparency website known as TheyWorkForYou.com, and has been previously presented in former conjunct work (Gurciullo et al., 2015b,a). Each government mandate is split into parliamentary sessions that last about twelve months, and commence with the Queen’s Speech, which outlines the government agenda for the commencing session. On average, in each session there are about 630 speakers. As table 5.1 shows, speeches vary greatly in length, albeit they tend to be rather short, at around 187 words.

Session ID	Year	Number of Speakers	Number of speeches	Median length of speeches	10th percentile	90th percentile
5301	2001-2002	645	55227	273	28	1572
5302	2002-2003	645	52591	256	30	1330
5303	2003-2004	636	47947	259	34	1400
5304	2004-2005	611	20549	262	31	1379.4
5401	2005-2006	634	64507	254	35	1269
5402	2006-2007	628	40142	264	39	1356.9
5403	2007-2008	630	50175	254	30	1234.6
5404	2008-2009	622	39883	253	35	1322
5405	2009-2010	596	21452	239	36	1230
5502	2012-2013	639	53975	222	44	1104.6
5503	2013-2014	634	46467	225	46	1165
Average		630	47973	187	24	631

**Table 5.1:** UK House of Commons Debates - Summary statistics.

The data is used to extract latent networks of influence for six policy themes in each government session, totalling 18 different networks. To accomplish the task three key preprocessing steps are required. These are outlined below.

**Preprocessing.** The first step of preprocessing involves the ridding of characters and text from the speeches that yield no useful information to the subsequent analysis. After lower-casing the entire corpus, English stopwords were removed.

<sup>1</sup>The dataset time span stops on February 13, 2014, and does not incorporate the final part of David Cameron’s government.

These include overly frequent words such as ‘the’, ‘a’, ‘an’ and so on. The list of stopwords is pre-determined by the module used to perform the task, which is the Natural Language Processing Kit (NLTK) in Python<sup>2</sup>. In addition, all punctuation is removed, whereas digits are kept.

Having cleaned the text, a semi-automatic method was needed to identify six policy themes for each government session. In the context of parliamentary debates, it is assumed that a policy theme is represented by a small set of concurrent words. If a policy proposal refers to, say, the proliferation of nuclear weapons, it is safe to assume that the a speaker would pronounce the words ‘proliferation’, ‘nuclear’, and ‘weapon’ close to one another, and that any variant of interest would most likely contain at least two of the three tokens. Given this assumption, for each government session (not to be confused with the annual parliamentary sessions), it an algorithm was run that ranks quadgrams by frequency. A quadgram is a set of four tokens appearing concurrently in the text<sup>3</sup>. The quadgrams are then clustered by word commonality, that is, if they feature at least three tokens in common with another cluster member. The resulting clusters contain expressions that are therefore likely to deal with the same topic. Out of the several clusters dedicated to a specific topic, six are selected for subsequent analysis. The selection is totally manual, but informed by frequency. In other words, the top six policy themes with the most frequent quadgrams have been chosen. Obviously, further research may well opt for different criteria for selection.

2001-2005	2005-2010	2010-2014
West coast line	West coast line	West coast line
Human Rights Act	Human Rights Act	Syrian humanitarian situation
Agricultural reform	Iraqi war	National Planning Policy Framework
Tax credit	Criminal Justice Act	Youth unemployment
Free school meals	Nuclear power	Small and medium sized enterprises
Al-Qaeda	Carbon emissions	Strategic Defence Security Review

**Table 5.2:** Selected policy themes for each governmental session.

The six policy themes identified at each governmental session are shown in ta-

<sup>2</sup>The documentation is available at [www.nltk.org](http://www.nltk.org)

<sup>3</sup>The reason for choosing quadgrams rather than other n-grams of different length is the degree of specification that the former can yield. A bi- or trigram would most likely result in topics that are too general, while penta- or exagrams would be over-specific (Cavnar et al., 1994). As an example, the bigram ‘human right’ would lead to the identification of all speakers mentioning it, but the bigram can be pronounced in several contexts, thus not allowing a clear extraction of influence among politicians. The pentagram ‘human rights act clause 5’, on the other hand, would select only speakers citing the clause, failing to capture information around the Human Rights Act in general.



ble 5.2. These all refer to topics that have received a substantial amount of attention from the House speakers, assuming that attention is proxied by frequency. Some tend to appear over time, such as the proposal to modernise the West Coast main line, and the Human Rights Act. Others are idiosyncratic and specific to the then-government agenda. In the 2001-2005 parliamentary sessions, the Common Agricultural Policy reform, working families tax credit, free school meals and the threat of Al-Qaeda were mentioned. In the subsequent government, the reconstruction of post-war Iraq was a main policy of interest, together with the Criminal Justice Act reform, the management of carbon emissions, and the proposal to build novel nuclear power stations. As for the incumbent Conservative government, the policy themes selected relate to the Syrian humanitarian situation, the National Planning Policy Framework, proposals to curb youth unemployment, the stimulus of small and medium sized enterprises (SMEs), and the Strategic Defence Security Review (SDSR).

Table 5.3 displays an example of what kind of quadgrams correspond to a policy theme. The table shows one topic for each governmental session, namely, the policy on tax credits for working families, improvements on the West Coast main line, and the Strategic Defence Security Review. For all policy themes, there exists a quadgram that dominates the relative frequency distribution, followed by a number of less frequent variants.

The third and final step of the preprocessing stage concerns the transformation of the raw data into a time series of discrete states that can be fed to the entropy-based network reconstruction method of NIIF. For each policy theme at each government mandate, all House of Commons speakers are filtered in whose speeches include a relevant quadgram. The result is a subset of political actors who have mentioned the issue at least once during their mandate. Given the subset of speakers, for each day of parliamentary activity it is checked whether their speeches contain a relevant quadgram. If they do, the speakers are assigned a state variable equal to 1, 0 otherwise. At the end of the assignment, for each policy theme it is obtained a time series of binary states associated with the MPs who mentioned the issue at least once, akin to the example to explain how the transfer entropy equation 3.18 works. The next subsection shows how networks are generated from the timeseries.

Tax credit	Frequency	West Coast Line	Frequency	Strategic Defence Security Review	Frequency
working families tax credit	0.685633001	west coast main line	0.68796992	strategic defence security review	0.75495
families tax credit disabled	0.058321479	improvements west coast main	0.04511278	defence security review set	0.022277
families tax credit child	0.028449502	upgrading west coast main	0.03383459	review counter terrorism security	0.022277
tax credit working families	0.021337127	upgrade west coast main	0.02255639	defence security review 2015	0.019802
families tax credit new	0.014224751	west coast route modernisation	0.02255639	defence security review 2010	0.019802
families tax credit children	0.014224751	south west coastal path	0.01879699	review strategic defence security	0.012376
families tax credit increased	0.008534851	modernisation west coast main	0.01879699	defence security review hon	0.012376
families tax credit minimum	0.008534851	capacity west coast main	0.01879699	security review comprehensive spending	0.009901
families tax credit does	0.007112376	work west coast main	0.01503759	defence security review make	0.009901
families tax credit measures	0.007112376	east coast west coast	0.01503759	defence security review announced	0.009901
families tax credit tax	0.007112376	problems west coast main	0.0112782	defence security review based	0.009901
families tax credit people	0.0056899	investment west coast main	0.0112782	defence security review comprehensive	0.009901
families tax credit taken	0.0056899	spent west coast main	0.0112782	defence security review defence	0.009901
families tax credit fraud	0.0056899	services west coast main	0.0112782	defence security review published	0.009901
families tax credit make	0.0056899	billion west coast main	0.0112782	defence security review come	0.007426
families tax credit example	0.0056899	west coast main lines	0.0112782	defence security review committed	0.007426
families tax credit just	0.0056899	coast west coast main	0.0112782	defence security review strategic	0.007426
new tax credits families	0.0056899	spent upgrading west coast	0.0112782	defence security review october	0.007426
families tax credit important	0.0056899	north west coast scotland	0.0112782	security review secretary state	0.007426

**Table 5.3:** Quadgrams for selected policy themes, and respective relative frequencies.

**Network reconstruction through transfer entropy.** The transfer entropy method is used to extract a representation of the latent networks of influence between speakers in relation to the selected policy themes. As explained in the methodological Chapter, the transfer entropy between two processes indicates the strength by which one determines the state of the other. In the context of the House of Commons, the transfer entropy between two speakers is a measure of the likelihood that one speaker would mention a policy theme in their speech, provided that the other speaker has previously mentioned it. The timeseries of the case study features a daily time unit, implying that one MP influences another if, given that they have spoken about a policy on day  $t$ , the other mentions it on day  $t + 1$ . It is extremely important to note that this is an assumption. Influence may well take place over shorter time frames, or longer. In the UK House of Commons, the order of speeches is defined in advance by the Speaker of the House (May, 1997), who accommodates the strategic requests of parliamentary members. Given this institutional constraint, it can be deemed a reasonable assumption that speakers can decide to respond to or echo the speech of another speaker a day afterwards. Obviously, the assumption might hold in all contexts, and further research ought to contemplate other timeframes.

Transfer entropy is calculated for each pair of speakers associated with the same policy, and statistically insignificant values are filtered out by means of bootstrapping, as outlined in section 3.1. Performing the task is very computationally intensive, and it was decided to make use the UCL high performance computing Legion clusters, where 36 CPUs have been accessed for parallel computation. The results are 18 different directed weighted networks of members of the House. The direction of the link indicates the flow of the latent influence, while the weight represents the strength of the relation. In the next section the graphs are shown, in conjunction with an exposition of the hypotheses to be tested, along with annexed results.

### 5.3 Analysis settings and results

The inferred graphs of this case study represent a reconstruction of the potential, latent information diffusion networks of House of Commons speakers surrounding certain policy themes. A strong relation from speaker  $A$  to speaker  $B$  implies that, given the input data, speaker  $B$  has repeatedly spoken about the same policy theme as  $A$  in the subsequent time unit. This, in turn, can imply that  $A$  exerts some sort of influence over  $B$ . The influence might not necessarily be due to some intrinsic property of speaker  $A$ , such as the very content of their speeches. It may also be

caused by structural reasons, inherent to the institutional role of *A*, such as party position. The analysis does not distinguish between these two sources of influence. In this section, two hypotheses are tested:

- (i) The degree distributions of the networks of influence approximate an exponential distribution.
- (ii) The most influential actors are those with a current or past Cabinet or Great Office of State position.

The first part of the section explains the motivation behind the hypotheses, and outlines the procedure by which they are tested. The second part discusses the results of the analysis.

**Outline of the analysis.** Hypothesis (i) is informed by recent advances in the study of graphs and information diffusion. In section 2.3 it has already been shown that several research efforts suggest the existence of a power law or exponential distribution of the degrees (i.e. the number of links per node) in a complex social network. The pattern has been found in the diffusion of information through web blogs (Gruhl et al., 2004; Agarwal et al., 2008), social media applications (Lerman and Ghosh, 2010; Jiang et al., 2015), and real-world human interaction (Schlöpfer et al., 2014). In all those instances, there exist a comparably very small number of units that determine the flow of information at a systemic level, thereby controlling its whole state over time. With the data in hand, it is possible to test whether the same pattern can be found in political networks reconstructed from debates. Finding an exponential distribution of the degrees would suggest the presence of elite actors who are significantly more influential than others in setting the policy agenda, providing a novel empirical validation for an elitist view of policy processes (Knoke, 1993). Yet, both the institutional and physical constraints of the system may well prevent such a degree of inequality to take place. In the House of Commons, the Speaker of the House reportedly gives less experienced MPs the opportunity to debate. Restricted time and days dedicated to a specific policy can also function as barriers against overly-influential actors. An answer to this question can pave the way towards new hypotheses on government, decision and influence. If a few exert more influence, does it imply that they also possess greater decision power? Furthermore, it contributes to the general study of information diffusion networks embedded in institutional constraints. So far, the networks studied in network science do not exhibit the type of rules of behaviour dictated in the House of

Commons. Discovering whether the shape of the degree distributions of the case study's networks resemble an exponential signifies a step forward in the field, and supplies novel empirical evidence that can inform a theory of topological behaviour of such networks co-existing with institutional rules.

The hypothesis is assessed by using a Kolmogorov–Smirnov (KS) test that compares the empirical distributions of the 18 policy networks to respective fitted exponential distributions. The KS test is nonparametric, and measures the distance between an empirical distribution given by a sample and the cumulative distribution function of a reference one, which, in this case, is exponential (Chicheportiche and Bouchaud, 2012). The null hypothesis of the test is that the sample has been drawn from an exponential distribution. A large and statistically significant KS test statistic suggests that the hypothesis is rejected.

Hypothesis (ii) is informed by theoretical and empirical studies on the role of legislative debates in government. Relevant literature seems to be sceptical about the idea of parliamentary debates as a persuasive tool in law-making (Austen-Smith, 1990; Proksch and Slapin, 2012). Other studies propose, rather, that debates are harnessed by MPs to differentiate their positions from one another (Martin and Vanberg, 2008), or to secure favourable polls (Kam, 2009). Certainly, debates can serve a great variety of purposes. This case study assesses the idea that actors with a current or past official role in government tend to wield greater leverage in the policy themes with which they confront themselves. This hypothesis is mainly justified by two arguments. First, political actors with a Cabinet or Great Office role are more likely to be active within the political system for a longer time. More experienced actors, it can be argued, are more likely to engage with a policy theme of interest in a more active fashion than less experienced ones, despite the structural organisation of the House of Commons. The second argument contends that speakers with a governmental position may exhibit more perceived authority during debates. This, in turn, would make the probability more likely that other speakers would engage or echo the topic in their speeches.

The hypothesis is tested by means of a robust Ordinary Least Square (OLS) model, where the position of the speaker functions as a predictor, and its topological-based measure of influence (the eigenvector centrality) as the dependent variable. A positive, statistically significant coefficient for the regressor implies that the hypothesis is not rejected. In order to dwell more on the relationship between MP position and influence, it is necessary to compute the probability distribution of the influence metric across all identified policy themes, conditioned to the role


of the speakers. Data collected by Vannoni and John (2015) provide information about the position of each parliamentary speaker in the time frame of interest. This takes the form of a dummy variable assigned to each speaker: the value is 1 if they feature or have featured a position in the Cabinet or in the Great Office of State, 0 otherwise. The eigenvector centralities distribution is conditioned to this very variable.

**Results.** Eighteen different network of speakers are inferred by means of the transfer entropy method, each related to one of the topics listed in table 5.2. Table 5.4 provides some descriptive statistics. All networks appear to be sparse, with an average density of 0.20, implying that about 20% of all possible links actually occur in the graphs. The number of nodes varies significantly, indicating that some policy themes were engaged by a larger portion of the political community. It is also possible to note that for all networks, the mean degree is greater than the median, indicating a skewness of the graphs' degree distributions towards higher values. In other words, there exists a minority of speaker who enjoy a greater number of edges than the rest.

A visual inspection adds validity to the impression that the graphs feature a skewed degree distribution. Figures 5.1, 5.2 and 5.3 show three examples, each from one governmental session. The layout algorithm used to depict the networks locates the most connected nodes at their centres. It is possible, then, to note that the networks of parliamentary speakers possess a centre-periphery structure, perhaps not so different from that encountered in the previous case study. It is now required to test whether the degree distribution skewness is shaped by an exponential trend.


	Years	Density	Number of nodes	Mean degree	Median degree
Al-Qaeda	2001-2005	0.208988	77	0.412548	0.285714
Human Rights Act	2001-2005	0.234709	143	0.466135	0.363636
Agricultural reform	2001-2005	0.233159	102	0.461745	0.343137
Free school meals	2001-2005	0.173362	44	0.338843	0.295455
Tax credit	2001-2005	0.242338	130	0.480947	0.376923
West Coast line	2001-2005	0.196839	88	0.389205	0.272727
Criminal Justice Act	2005-2010	0.237364	110	0.470413	0.340909
Carbon emissions	2005-2010	0.217233	104	0.430288	0.326923
Human Rights Act	2005-2010	0.251196	133	0.498615	0.37594
Iraqi war	2005-2010	0.170599	132	0.338613	0.231061
Nuclear power	2005-2010	0.185756	155	0.369116	0.277419
West Coast line	2005-2010	0.191756	63	0.377425	0.285714
National Planning Policy Framework	2010-2014	0.201235	81	0.3975	0.283951
Strategic Defence Security Review	2010-2014	0.191136	90	0.378025	0.272222
Small and medium sized enterprises	2010-2014	0.175983	325	0.350883	0.261538
Syrian humanitarian situation	2010-2014	0.229649	76	0.453255	0.315789
West Coast line	2010-2014	0.168557	110	0.33405	0.227273
Youth unemployment	2010-2014	0.165271	174	0.328643	0.232759
Average		0.204174	118.7222	0.404236	0.298283

**Table 5.4:** Summary statistics for the inferred networks of speakers.




**Figure 5.1:** Network of House of Commons speakers, Human Rights Act - 2001–2005.





**Figure 5.2:** Network of House of Commons speakers, Nuclear power - 2005–2010.



**Figure 5.3:** Network of House of Commons speakers, Youth unemployment - 2010–2014.


The results of the KS test are displayed in table 5.5. The hypothesis is clearly rejected for all network instances. The values in the first column are the KS statistics, which quantifies the absolute difference between the empirical degree distribution and the fitted exponential one. Values closer to zero indicate that it is likely that the sample distribution has been drawn from an exponential. Clearly, it is not the case with regard to the case study graphs. The magnitude of the values is extremely high, and all are statistically significant. In KS tests, the range of the KS statistics depends on the range of values of the distributions considered. In this case, the degree distributions span from 0 to 1. KS statistic values of 0.50 and above imply, evidently, a clear difference between the actual distributions and the hypothesised ones.

Rejecting hypothesis (i) opens new ground for novel theories on the structure of social networks under institutional constraints. Figure 5.4 shows the degree distribution for two networks for each governmental session. A skewness towards higher values is confirmed, albeit not of an exponential form. Overall, it appears that the


distribution gravitates towards mediocre, yet not minimal levels of influence, with a long tail that implies the presence of more influential actors. Alternative hypotheses for the shape of distribution can include the beta function, a more general form of power distribution. The parameters of the beta function may be determined by the dynamics of information sharing, which can be explored in further research.

	<b>Years</b>	<b>KS statistic</b>	<b>p-value</b>
Al-Qaeda	2001-2005	0.549772	0.000
Human Rights Act	2001-2005	0.545343	0.000
Agricultural reform	2001-2005	0.5433	0.000
Free school meals	2001-2005	0.688305	0.000
Tax credit	2001-2005	0.533116	0.000
West Coast line	2001-2005	0.603899	0.000
Criminal Justice Act	2005-2010	0.536908	0.000
Carbon emissions	2005-2010	0.569837	0.000
Human Rights Act	2005-2010	0.526446	0.000
Iraqi war	2005-2010	0.635627	0.000
Nuclear power	2005-2010	0.63098	0.000
West Coast line	2005-2010	0.611222	0.000
National Planning Policy Framework	2010-2014	0.589643	0.000
Strategic Defence Security Review	2010-2014	0.643283	0.000
Small and medium sized enterprises	2010-2014	0.626579	0.000
Syrian humanitarian situation	2010-2014	0.540718	0.000
West Coast line	2010-2014	0.612655	0.000
Youth unemployment	2010-2014	0.642998	0.000


**Table 5.5:** Kolmogorov–Smirnov test for exponentiality results.



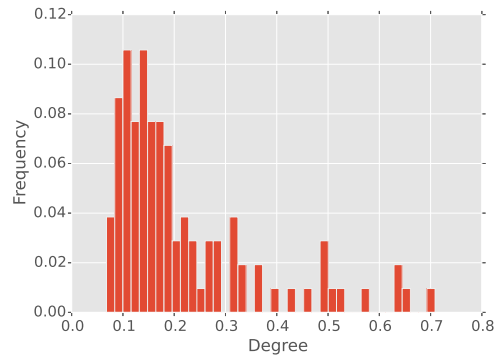
(a) Human rights Act, 2001-2005




(b) Free school meals, 2001-2005



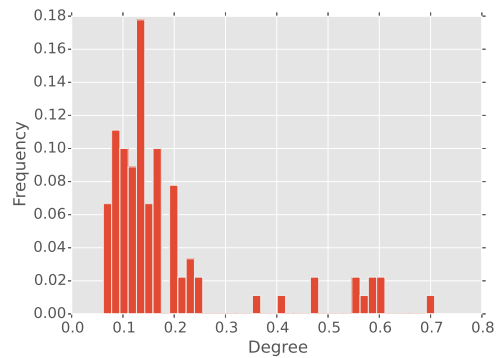
(c) Nuclear power, 2005-2010



(d) Carbon emissions, 2005-2010



(e) Youth unemployment, 2010-2014



(f) Strategic Security Review, 2010-2014

**Figure 5.4:** Degree distributions for six networks of parliamentary speakers.

In order to test hypothesis (ii), the ordinary and weighted eigenvector centralities of nodes in all networks have been evaluated and used as dependent variables for two OLS models, with the position dummy variable as a predictor. Table 5.6 shows the result of the analysis. The hypothesis has not been rejected. In both models, the regressor's coefficient is positive and highly statistically significant. In the case of eigenvector centrality, the coefficient indicates that, on average, its value is greater by 0.311 if the speaker has an official position. The average difference slightly decreases to 0.246 when weighted eigenvector centralities are utilised. Considering that both metrics range between 0 and 1, such a jump in value is huge.


	Eigenvector centrality	Weighted eigenvector centrality
Position	0.311 (0.020)	0.246 (0.017)
t-statistic	15.20	14.22
p-value	0.000	0.000
$N$	2090	2090
$R^2$	0.180	0.179
Root MSE	0.192	0.243

**Table 5.6:** OLS models testing MP position as a predictor of influence.

In order to further inspect the relationship between influence and position, the eigenvector centralities across all 18 graphs are conglomerated to obtain a general distribution of the variable. This has then been conditioned to the binary variable indicating whether a member of the House is a current or past Cabinet or Great Office of State member. The results are shown in the stacked bar graph of figure 5.5. The distribution of the influence measure has been segmented into ten quantiles, in order to better capture the change in probabilities. For lower values of influence, it is more likely that the speaker is not a former or current officer. A shift happens as the value becomes greater, with values greater than 0.9 being dominated by officers.

The pattern is robust to the other measure of influence. Figure 5.6 shows the conditional stacked bar for the distribution of weighted eigenvector centralities, i.e. eigenvectors that also take into account the weight of the links. The dramatic difference in the level of influence between backbenchers and officers is even more emphasised with this variant. Speakers with current or past Cabinet or Office positions are much more likely to feature an influence metric of 0.5 or greater.

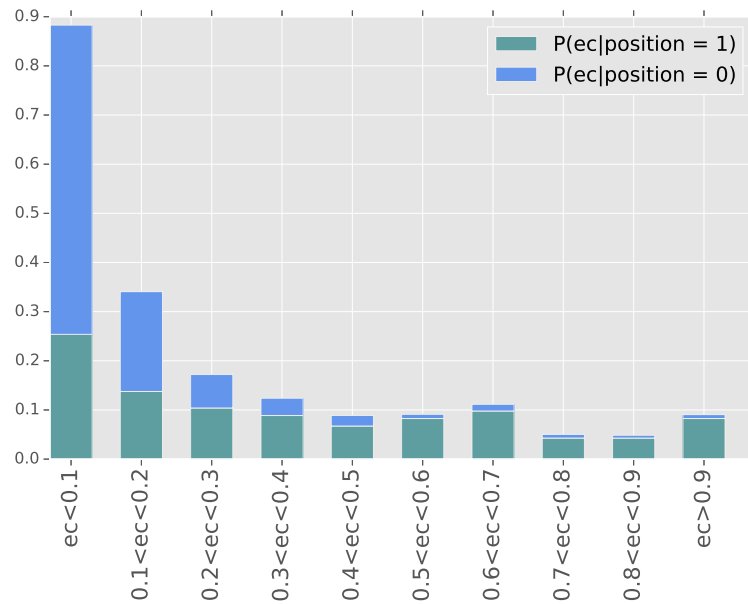
The results of this analysis suggest that individuals with formal positions are likely to exert a much greater leverage in influencing debates about a policy pro-



**Figure 5.5:** Conditional probability distribution of eigenvector centralities.

posal than backbenchers. This adds validity to the notion that underlying factors, such as experience, authority and prestige, could determine the dynamics of information sharing and influence more than the informational content of the debates themselves. As a final consideration, it is also worth noting that while the distribution of influence metrics for backbenchers is heavily skewed towards low values, that of officers appear to be much more uniform. The insight hints to the fact that being an officer does not necessarily lead an individual to engage more in debates; rather, further latent factors are likely to determine that.

This section has tested the hypothesis that the inferred networks of speakers feature an exponential distribution, and that speakers with a Cabinet or Office position are likely to be more influential. The first hypothesis has been disproved. The degree distributions of the 18 networks considered present a heavy tail, yet the shape does not resemble an exponential one. It is advocated that a main reason behind that is the presence of institutional constraints, which prevent the total monopolisation of a policy debate by an elite set of actors. Another concurrent reason might also be the size of the graphs. Power laws and exponential functions generally characterise very large networks, comprising thousands or more nodes (Barabási and Albert, 1999). In small networks, such an asymptotic shape may not be approximated. From this perspective, the case study function as a starting point for the study of the laws governing the shape of small sized networks, which – especially in the context



**Figure 5.6:** Conditional probability distribution of weighted eigenvector centralities.

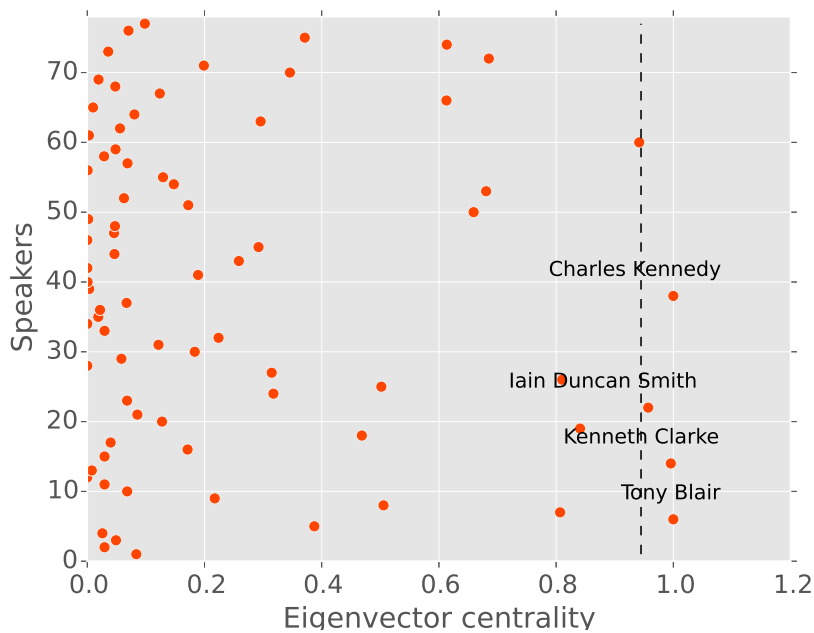
of policy-making - are just as important as large scale ones.

The second hypothesis has not been rejected, since it has been discovered that political actors who benefit from an official position are much more likely to be central in policy debates. The finding contributes to the contention that information diffusion in policy-making and agenda-setting may be dominated more by latent structures of power, authority and prestige rather than the pure content of the information shared. The next section explores further the role of influential speakers by providing an example of policy implementation.

## 5.4 Practical implementations

With a metric of parliamentary members influence in debates, it is possible to devise tools that monitor the role of an individual with regard to a policy topic. This section provides an example. Similar to the Value at Risk instrument shown in the previous Chapter, the most influential actors can be identified by observing whether the values of the respective influence metrics are above a pre-determined threshold percentile. In the three examples that follow, the 95<sup>th</sup> percentile is used.

Figure 5.7 is a scatter plot of the eigenvectors centralities in the graph of speakers whose debates were about Al-Qaeda and terrorism. The most influential actors are those beyond the threshold dashed line. The names are not surprising. Certainly, Tony Blair had the most influential role, given his dedication to the Iraq war and the

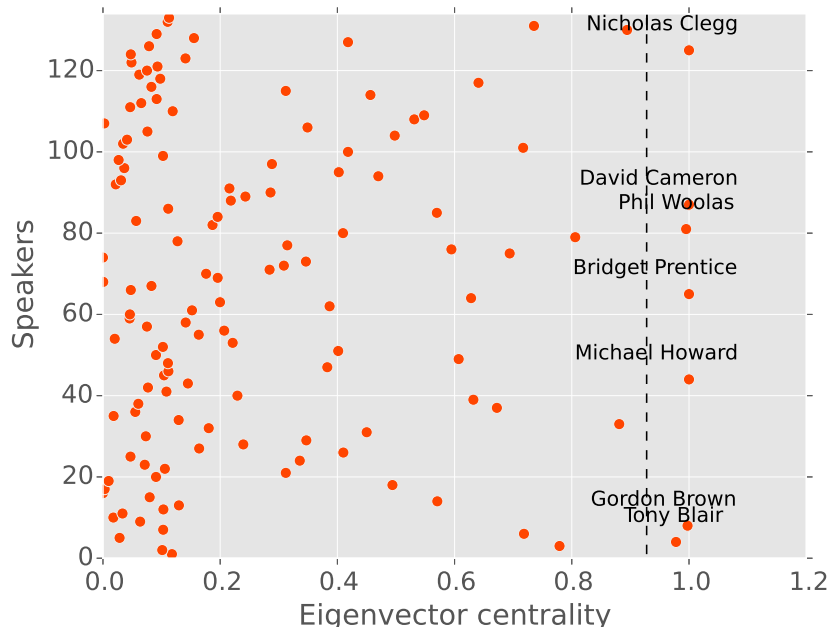


**Figure 5.7:** Most influential actors, Al-Qaeda - 2001-2005.

fight against terrorism alongside George Bush following the 9/11 terrorist attacks. Iain Duncan Smith, then-leader of the opposition, played a very important role in the management of terrorism policy and the war on Iraq (The Guardian, 2002). Charles Kennedy was the leader of the Liberal Democrats and played a huge part in the opposition against the Iraq war and Blair's anti-terrorism policies. In fact, one of his most historic speeches, in which he pleaded with Tony Blair to end the Iraq war and find other ways to fight Al-Qaeda, has recently gained great interest from the media (The Guardian, 2015). Kenneth Clarke is a very experienced Conservative politician who held several positions within government over the years and, like Kennedy, declared a fierce opposition to a war that he deemed 'catastrophic' and not addressing the threat of Islamic terrorism (BBC, 2005).

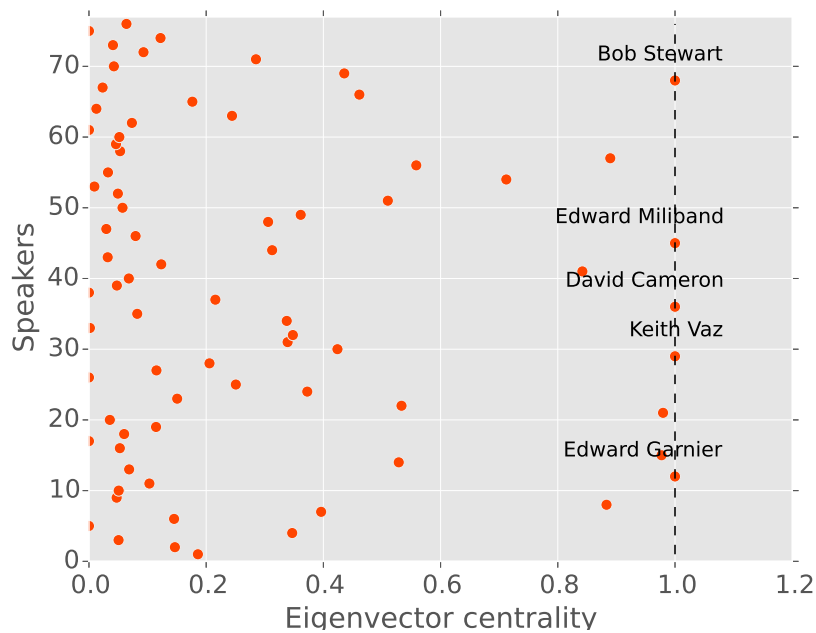
Figure 5.8 provides the same type of scatter plot about the debates surrounding the Human Rights Act, during the second governmental debate taken into account by the case study. Among the actors with the top 5% influence metric values is Phil Woolas, who acted as Minister of State for Borders and Immigration between 2008 and 2010. His mandate required much attention towards asylum seekers, implying that discussions about the implementation of the Human Rights Act were of foremost importance to the politician. Bridget Prentice appears to have had a very influential role in this topic as well. She took the position of Minister of Justice in 2007, and a crucial part of her role involved, just like that of Woolas, immigration and the implementation of rules dictating who was entitled to stay in the United





**Figure 5.8:** Most influential actors, Human Rights Act - 2005-2010.

Kingdom. Michael Howard was a Conservative leader until 2005, and despite his resignation in 2006, he seems to have been very influential on the Human Rights topic. In fact, his views have spread across the media, as in debates he called for a redefinition of the Act (Howard, 2011).



**Figure 5.9:** Most influential actors, Syrian humanitarian crisis - 2010-2014.

The last scatter plot is about the Syrian humanitarian crisis. Bob Stewart is among the most central individuals in the graph. The MP has a distinguished career in the British Army, and his military expertise has often been considered with regard to Syria. Edward Garnier is a Conservative party member who held the position of Solicitor General of England and Wales between 2010 and 2012. He strongly advocated for the bombing of the country in order to prevent further empowerment of Daesh. Keith Vaz, Labour party politician, appears to be extremely influential as well. He chairs the Home Affairs Select Committee wherein one of his duties has been to deal with the resettling of Syrian refugees in British territory (BBC, 2016b).

This section has supplied an example of the application of the case study results to assess the relative influence exerted by a parliamentary member on the discussion of a policy topic. Interestingly, the evaluation of the top 5% political actors may serve both as a transparency tool, to assess their degree of activity, and as an accountability instrument, because interest groups can use the insights to target individuals with a higher likelihood to affect the direction of a policy debate. In the next, final section of this Chapter, a summary of the findings and a discussion of their potential limitations is presented.

## 5.5 Case study discussion

This chapter has shown how the NIIF can be used to uncover the latent structure of influence of political actors, sharing information about policy topics. Eighteen different networks of British MPs have been extracted from UK House of Commons debates by means of transfer entropy, each related to a different policy theme. It has been tested and discovered that the networks do not feature an exponential degree distribution, thus acting as a significant exception to current observable cases of social information diffusion networks. This paves the way for novel hypotheses about small sized networks and their topology. It has been suggested that institutional constraints and the graphs' endogenous dynamics at small scale may function as the underlying causal factors. The chapter has also tested whether the role of a political actor in government determines their level of influence, here measured as their eigenvector centrality. It has been found that individuals with a Cabinet or Great Office position are more likely to feature higher centrality values than those who do not.

The case study presents some limitations that can be addressed with further research efforts. The first limitation concerns the construction of the timeseries functioning as input to the network inference method. It has been opted to use a daily time step, rather than a sub-daily one. Yet, the dynamics of debates may enable speakers to respond to or echo the message delivered by their colleagues over the same debate day. If this is the case, the study has missed a set of important information about the hidden structure of influence in this context. Furthermore, the time step can be contingent to the issue at hand and the debate rules within the studied parliament. What applies to the House of Commons might not apply to the Italian Camera dei Deputati. It appears that in this case study a daily time step worked well enough to provide sensible and valid results, however, further efforts ought to spend more time researching the parliamentary context and perform robustness checks with timeseries of different time steps.

A second limitation might be related to the method used to identify the policy themes, i.e. by the extraction and clustering of quadgrams. A parliamentary speaker can discuss a policy without explicitly mentioning it, perhaps when responding to a previous speaker who initiated the topic. If so, this chunk of information has not been captured by the preprocessing stage, with the potential of creating biases during the network inference stage. Further research can perform robustness studies by coupling the preprocessing stage with a topic model. Topic models are unsupervised machine-learning algorithms grouping concurrent words together, able, as the name suggests, to extract the 'topics' in a corpus (Wallach, 2006). The method may perhaps identify words associated to a policy theme even when this has not been cited explicitly.

A further limitation of the study is conceptual, rather than methodological. To put it in non-technical terms, throughout the chapter the following has essentially been assumed: if a political actor speaks repeatedly after another political actor, and if this pattern is deemed not to be by chance according to the transfer entropy method, then the former is influenced by the latter. In a strict informational sense, this statement is valid. Speaker A sends information generally after speaker B does, and with the use of transfer entropy, this is more than enough to say that the two processes are coupled. If the informational standpoint is relaxed, some difficulties may arise about the kind of influence exerted by actors on one another. Two speakers might respond to each other often and be in opposition. In this case, the influence is not of a persuasive nature, but merely structural, caused by the underlying interest groups within parliament. In its current state, the framework produced by this research is not capable of distinguishing between different concepts of influence, and stops at the strict informational definition. Further research ought to confront

the issue. A potential way forward would be to couple the preprocessing analysis with other natural language processing techniques able to identify the sentiment of a message. This, possibly, allows for a differentiation across different categories of influence.

The next and last case study provides an instance whereby not only influence is impossible to differentiate, but in addition its direction cannot be untangled. Yet, with the use of NIIF it is still viable to extract affinities among the actors in a system. This is shown with regard to the UN member nations.

## Chapter 6

# Third case study: Global political change in the United Nations

Nations are agents in a complex social system with several facets. Economic relations form part of that. The study of the complex web of international trade has allowed us to see the face of current globalisation in all its might and true shape. Serrano et al. (2007) was one of the first attempts to observe the change in the flow of goods internationally between 1960 and 2000, showing a massive increase in the intensity of the exchange, and in the number of participating countries. The network study of world trade provided novel empirical grounds for identifying communities of countries, joined together by similar patterns or roles within the wider system (Fan et al., 2014). Modelling trade imbalances can also point towards structural changes that anticipate distress across nations, such as the 2007/8 financial crisis (Saracco et al., 2016).

Yet, countries not only exchange goods but also ideas and relations of power. As Kant, among others, conjectured (Franke, 2014), the system of nations is epistemologically evolutionary: countries exchange political information, values, which are selected and evolved, both endogenously and according to seemingly exogenous events. From a research perspective a threefold challenge arises here, asking whether it is possible to observe, model, and predict the dynamics of international idea diffusion and mutation, and the political positions of countries with respect to that. In this final case study, the use of NIIF is proposed to tackle the first segment of the challenge. Observing political changes by means of non-structured information can be translated into the problem that has accompanied the thesis so far: there exists a latent network of influence across nations, and it is possible that detecting structural changes within such a network implies a change in global politics. This chapter reconstructs a timeseries of the networks of nations, extracted by means of discrete state neural learning from UN General debate speeches. It demonstrates

that the reconstructed relations can capture the global political change started with the demise of the USSR and the end of the Cold War. The chapter is structured into five sections. The first elaborates the background of the study, reviewing the literature on ideology diffusion and countries' voting patterns. The second section proceeds with a description of the protocol for the preprocessing of the corpus of UN speeches, and establishes the initial parameters of the network inference method. In the third section the hypotheses of the analysis are explained, together with the methods of testing. The results are then shown and commented upon. The fourth section introduces an example of the practical instrument that can be derived from the approach. The tool can track the salience of a political theme across UN member nations over time, thereby functioning as a political thermometer for policy interests. The fifth and final section, as per convention, concludes with a summary of the analysis, as well as comments on its limitations and further research needed.

## 6.1 Case study background

It is doubtless that the rise and fall of political norms and ideologies throughout history have had a massive impact on society. The Enlightenment brought the now taken-for-granted idea that the state is a social contract with the people, building the conceptual foundations of Western democracies (Outram, 2013). Communist ideologies broke through millennial political orders in Asia, justifying revolutions, war and famine. Policies based on those caused the deaths of about 30 million people in Mao Zedong's China between 1958 and 1962 (Ashton et al., 1992). Neoliberalist norms put transition economies in peril. In Brazil, for instance, free-market policies are believed to be responsible for part of the increase in relative poverty and income inequality. Between 1986 and 1999 – the main years in which neoliberalist regulations were implemented – the percentage of people living in poverty increased from 28.2% to 34.1% (Amann and Baer, 2002). A significant amount of scholarship has been dedicated to understanding the diffusion of ideas and norms across international regimes, and the respective roles of political actors – be they states, nongovernmental organisations, or collectives – in this kind of process. Two strands of literature are reviewed in the next paragraphs. The first focuses on the role of political actors in spreading norms and ideas, mainly composed of qualitative research and case study analyses. The second tackles questions of ideas diffusion and influence from a more quantitative perspective, through the use of voting data as the empirical starting point of the research.

**Norm diffusion in the international realm.** The study of the diffusion of norms in political systems is akin to the literature of diffusion of other types of phenomena in contexts such as epidemiology and information theory. Finnemore and Sikkink (1998) provide one of the most accepted theoretical groundings of the phenomenon, proposing the existence of a norm life cycle, dictated by three stages: norm emergence, norm diffusion, and norm internalisation. Interestingly, the authors posit that states and international organisations are the main characters of the second stage, as they have the capability to render an emergent norm legitimate and be implemented. The theory has been followed by several case studies in order to test it. Fernández and Lutter (2013) analyse the diffusion of same sex marriage across Europe between 1988 and 2009 through an event history model (Box-Steffensmeier and Jones, 2004). The authors find that states with a more secular tradition are more likely to adopt the policy and to act as norm promoters in the European Union and United Nations. In Linde (2013), a similar kind of analysis is performed with regards to the diffusion of laws against child death penalty. The scholar observes how after the Second World War, a series of norms for the protection of childhood was promoted heavily by European and Northern American countries which harnessed the international context of the United Nations to push the rest of the world towards an internalisation of their own policy preferences. Grossback et al. (2004) added a further, interesting insight on what makes a country more influential in diffusing a norm internationally: ideological proximity. Two countries with a similar ideological view are more likely to adopt a norm implemented by the other, all things being equal.

The state of the art in the field is currently represented by research efforts implementing computerised approaches. A critical example is Ring (2014) who develops an instance of agent-based modelling (ABM) for international norms diffusion. ABMs are simulations of agents, that iteratively interact with each other and the environment under a set of behavioural rules set by the creator (Bonabeau, 2002). They are heavily used in fields such as economics, ecology, and the study of innovation diffusion (Kiesling et al., 2012; Farmer and Foley, 2009; Grimm and Railsback, 2013), and have found room in political science, especially in the study of political economy (Kollman et al., 2003) and party competition (Laver and Sergenti, 2011). The model in Ring (2014) is anchored by three variables that capture three fundamental aspects of international society: hierarchy, neighbourhood, and identity. It then simulates norm dynamics through four different channels, in accordance with part of the theoretical literature, namely coercion, competition, emulation and learning (Franzese and Hays, 2008). The author defines a number of experiments within a range of different initial parameters, such as population

size and the distribution of the countries' political statuses. His results suggest that both competition and learning enhance an exponential type of norm diffusion, resembling the three-stage norm life cycle described above. Coercion and emulation, instead, characterise a much slower and incremental dynamic of diffusion. The study is a landmark in the field, yet its validity is impeded by the lack of specification of interstate relations. In the model, two countries are connected only if they are neighbours. The assumption is clearly not respected in reality, as alliances, treaties and trade relations are not necessarily dependent on geographical proximity. Subsequent work by Froncek (2015) attempts to fill the gap by explicitly taking into account network structures, and simulating these alongside norm diffusion dynamics. Their results place the centrality of a node in the networks as a key determinant of whether it would be a key player in the diffusion or halting of a norm. While ABMs are providing the very first room for experimentation of this type of phenomena, they tend to lack the insights that would be given by studies using empirical data. The next subsection reviews some of the key working that addresses the issue.

#### **Empirical quantitative approaches to international political preferences.**

Norms entail political preferences as to what kinds of policies are welcomed for implementation. Recently, an increasing amount of quantitative research has turned its attention towards the extraction of political preference and influence as expressed through voting patterns. The underlying assumption is that political actors sharing norms would vote in a similar fashion across a variety of issues, thereby forming clusters around the country-leaders promoting the norm (Epstein and Mershon, 1996). With respect to the United Nations, Kim and Russett (1996) published the first work that untangles political preferences using voting patterns in the UN General Assembly (UNGA). The two scholars analyse three post-Cold War sessions of the General Assembly and compare them with the alignments and issues that characterised sessions during the Cold War. Among their findings, it is shown that Eastern European countries appear to be more in line with Western European ones, forming a cohesive continental bloc. The main division seems to be located across the income spectrum, with rich and poor countries having the most different patterns of voting. Potrafke (2009) attempts to measure the degree of influence exerted by the United States on the UNGA voting patterns of 21 OECD countries. With a multivariate analysis of voting data over the 1984-2005 period, the authors suggest that nations with left-wing governments are less likely to be aligned with the United States. The issue of influence over voting is taken more explicitly by



Dreher and Sturm (2012), who explored the concept in relation to two nongovernmental actors, the International Monetary Fund and the World Bank. Using panel data on 188 countries, they tested whether receiving adjustment projects and larger non-concessional loans from the World Bank, or being part of a non-concessional IMF programme, impact the likelihood of voting in line with G7 countries. Interestingly, only the World Bank appears to exert such influence, while results related to the IMF are deemed to be not statistically significant. The most complete dataset on UNGA voting was presented in Bailey et al. (2015), and is kept up to date on the authors' personal academic web page. The dataset comprises all UNGA roll-call voting from 1946 to the last available date. In their study, they performed a dynamic spatial model to extract a one-dimensional metric, quantifying how close a state is to the US-led liberal order that they assume. This is the current state of the start situation in the quantification of political preference and influence with respect to global politics, despite some methodological issues inherent in the interpretation of the data, such as confusing abstentions and absentee votes (Voeten, 2012).

Most of the research on norm diffusion and influence suffers from two shortcomings. The first is partly conceptual, partly methodological: influence, rather than being extracted, is assumed to exist in relation to certain entities. As can be noted, many studies assume the United States or Western advanced economies to be at the centre of a hidden web of influence. While these kinds of assumptions can be judged appropriate in certain contexts and at certain timeframes, efforts should be made towards a methodological step backwards, allowing the extraction of a latent web of influence without a priori assumptions about the hierarchy of such a system. The second shortcoming relates to the type of empirical data utilised. Voting patterns can certainly provide a valuable overview of ideological commonalities across states, yet, analogously to the national parliamentary context, voting is the final, strategic outcome of a constrained process on specific issues contained in UN Resolutions, not necessarily embodying countries' attitudes towards greater scale policy themes (Schwarz et al., 2014). Textual data can be an alternative source from which to extract political influences, as they may contain a less noisy, clear set of information about the political and ideological position of a nation. The work of Baturu et al. (2016) provides the first analysis of UN member states speeches. They use correspondence analysis, an unsupervised clustering algorithm for discrete state data such as words (Greenacre, 2007), on UN General Debates to determine nations' policy preferences. This case study intends to contribute to this line of work, by inferring a time series of graphs of UN members based on their speeches, and using

the influence metrics contained in NIIF to detect global political changes. The next sections dwell on a description of the data, and how it has been preprocessed to infer the networks.

## 6.2 Data and network reconstruction

This Chapter plans to reconstruct a network of influence across nations over time, according to the information contained in their UN General Debate (UNGA) speeches. The following paragraphs introduce the dataset and outline the pre-processing and network inference stages. It is important to note that in contrast to the previous two case studies, this features undirected rather than directed graphs. Justification for the difference and the resulting implications are discussed.


This paper's dataset comprises all speeches made at the General Debate of the UNGA from 1970 to 2014. The General Assembly is the main deliberative body of the international organisation, comprising all member nations at a given point in time<sup>1</sup>. The General Debate dictates the start of UNGA's activities. Taking place over several days, the General Debate schedules speeches by Heads of States, Prime Ministers, foreign ministers, and representatives of their respective UN delegations. These speeches are used by countries to express their views on current issues in foreign affairs, and are regarded as an invaluable source of information used to understand policy preferences a sort of 'barometer of international opinion', as stated by Smith (2006).

Akin to Baturo et al. (2016), who first introduce the corpus, this research retrieves all General Debate speeches in the cited years from the official online database managed by the UN Bibliographic Information System. Where the original speech was not made in English, the official English translation is collected. All speeches made by representatives of the same country in one year are grouped together. It follows that the unit of analysis is a country-year, each identified by the country ISO Alpha-3 denomination. At each year, a country appears in the corpus if it is a member of the UNGA. Where a country ceases to exist because of a merger (e.g. the German Democratic Republic which merged into the German Federal Republic), the ISO code of the assimilated country disappears. A new identification code instead appears if a country splits. In total, the corpus is composed of 7,310 documents, with an average of 945 unique tokens per country-year. Figure 6.1 shows the number of country-years and mean token frequency over the entire

---

<sup>1</sup><http://www.un.org/en/ga/about/>

time period considered. Interestingly, as the number of UNGA members grows, speeches become shorter, perhaps due to time constraints.



**Figure 6.1:** Number of nations and mean token frequency per document, 1970-2014.

The preprocessing stage to applying discrete state neural learning is, in fact, minimal. Contrary to the previous two network inference approaches observed, this method only requires that the corpus is already structured into the units of analysis, i.e. that each document related to a countryyear is already grouped together. The exact format of the text preprocessing part is based on the judgement of the researcher. Discrete state neural learning can use punctuation as single states, perhaps to also classify how it is used by different actors. In this case study, all text has been converted to lower case, and the punctuation removed. Numbers and English stopwords have been retained. The corpus is tokenised by white space, and all words are stemmed, that is, reduced to their root form. For example, tokens such as ‘economic’ and ‘economy’ are reduced to ‘econ’.

The entire corpus is subject to the learning model, the functioning of which has been explained in section 3.1. Recall that the model transforms each token and document into a vector of length defined by the researcher. Also, as the learning

process for each word occurs in relation to its neighbouring words, a window size must be set. In this analysis, it has been opted to generate vectors of 200 coordinates. In truth, at the moment there is no analytical method to choose vector length, leaving it as an arbitrary value. The value above has been determined reasonable after some trials with vector lengths ranging between 100 and 300. The window size has been set to 10, implying that the model would determine the vector representation of a word by looking at its position in relation to a maximum of 10 tokens situated at one or the other side of that word. That is, currently, the standard set by previous application of the model (Le and Mikolov, 2014). The output of the model is a table with vectors corresponding to each unique word and whole speeches. In other words, there would be a vector for the word, say, 'human', and for the entire speech of the United States in 1987.


In order to reconstruct a dynamic representation of the network of UNGA members, for each year the respective speeches are selected, and a matrix of cosine similarity is constructed for the vector representations of each pair of countries. The result is a yearly time series of adjacency matrices, wherein the links are undirected, and the weights indicate the degree of similarity between the speeches of two countries. The matrix is a complete one, but is filtered to contain only edges with weights equal to or above 0.6<sup>2</sup>. As with the other parameters, there is no analytical, hard method that would justify the choice of this threshold value over another one. Prior to the analysis, the author experimented with several threshold values and determined that cosine similarities with values equal to or greater than 0.6 allow the discarding of most connections which would be deemed feeble by a human observer, while at the same time maintaining a sparse yet well-connected network. Needless to say, further studies may well indulge in robustness studies dedicated to this stage of network inference.

The resulting dynamic network is not directed, meaning that it is not possible to state that one country influences another. Yet, where a country is particularly central in the network, as measured by eigenvector centrality, it is possible to presume that it definitively features some sort of systemic influence, as many other UNGA members conform to its message. This sort of conformity, though, does not engender that countries with high eigenvector centralities exert power onto others. In the previous two case studies, the edges of the network had a clear connotation: financial interdependence and the likelihood that an MP would speak given that another had done so. In this context, the edges include a latent combination of linguistic and


---

<sup>2</sup>Recall that cosine similarity values range between 0 and 1.

semantic similarity. Two countries have high rates of similarity if their speeches are close, both in terms of language style and topics discussed. Unfortunately, it is still not possible to discern the two aspects. Yet, the changes in network topology with respect to some countries can apparently yield extremely valuable insights on the global political situation experienced by UNGA members, as tested in the next section. A visualisation of the evolution of the dynamic network is shown in figures 6.2 and 6.3.



**Figure 6.2:** UNGA General Debate inferred network, 1970.



**Figure 6.3:** UNGA General Debate inferred network, 2010.

### 6.3 Analysis settings and results

Thanks to the latent linguistic and ideological information yielded by UNGA speeches, it is possible to use the network inference and influence network to build a dynamic undirected network of UN member nations, the edges of which indicate affinity both in the policy topics discussed by pairs of countries, and in the style and words adopted. In this semantic political network, it may be reasonable to state that structural changes mirror actual political changes in the international political regime. This section inspects this idea by testing whether structural systemic and node-specific mutations co-occur with the greatest international political disruption experienced in modern history: the demise of the Soviet Union and the end of the Cold War. More specifically, three hypotheses are tested:

- (i) The density of the filtered dynamic network experiences a structural break between the years 1985 and 1992.
- (ii) The centrality of the Soviet Union experiences a structural break between the years 1985 and 1992.
- (iii) The centrality of the United States experiences a structural break between the years 1985 and 1992.

All three hypotheses imply that it is possible to statistically detect a structural change in the time series of the variable extracted from the dynamic network. The range of years is not arbitrary. In 1985 Mikhail Gorbachev's presidential mandate began. According to several international relations scholars, this represented a critical point for the Soviet Union, as the new political leader challenged previous foreign policy approaches with the other super power, and opens the country towards less hostile relations to the West (Lebow and Risse-Kappen, 1995). Gorbachev's years are characterised by a historic series of attempted reforms and economic and political crises, which concluded with the formal dissolution of the USSR in 1992 (Sakwa, 1990). The set of political events associated with the end of the Cold War imply a change of political order. From a more analytical perspective, such a change is nothing but a sudden change of state, which – if its output is measured by a series of dynamic processes – should be detectable as a structural break from past behaviour (Davis et al., 2006).

In the dynamic network inferred in this case study, the global political change should be associated with a structural break in the time series of at least three signals. The first is in the dynamics of the network's density. Density is the number of existing links in relation to the number of all possible links in a graph. With the end of the Cold War, a sudden fall in density is expected. The reason lies with the fact

that the ideological position and topics discussed towards and during the end of the Cold War become less homogeneous among countries. UN agendas become less entrenched with international security topics related to conflicts of states sponsored by the USA and the USSR, and more focused on less polarising topics such as human rights, or global economic issues such as the oil crisis (Roberts and Kingsbury, 2008). The potential, observable result is that UNGA countries become, overall, less semantically and ideologically similar, and therefore less connected to one another. The change would not be characterised by a trend, rather, it should appear as an almost sudden change, fuelled by the abrupt crisis that brought the USSR to its knees in a relatively short amount of time. A similar background process should govern the dynamic process of the eigenvector centrality of the Soviet Union. More specifically, a sudden decrease of the variable is expected, associated with a number of political factors, the most important of which are: Gorbachev's decision to end aid to Soviet allies (Kalinovsky et al., 2011) and the formal dissolution finalised in December 1991 (Walker, 2003). The first reason would disincentivise Soviet allies to show themselves as being ideologically close to their former donor, while the second engenders both the birth of new UNGA members, with agendas very much in contrast with the former centralised power, and the very change of political leadership and vision in post-Cold War Russia. The Russian Federation at the onset of 1992 would most probably exert much less political leverage within the international system. Also, with the demise of the Soviet Union, the topological status of the United States in the dynamic network is expected to have a structural break. Conventional scholarly work goes on to state that the end of a bipolar world would be the start of a monopolar one headed by the USA and its political ideology (Fukuyama, 1989). According to this line of thought, the post-Cold War regime shift would imply an increase in the eigenvector centrality of the country. Yet, this may not be necessarily true: the regime shift might well cause UNGA members to diversify their ideological and semantic positions in such a way as to render the issues of the UN Agenda different from those of the US. This would implicate a different trend for the centrality associated with the US, after the potential structural break.

The three hypotheses are subject to falsification by constructing a test statistic for a structural break in an auto-regressive (AR) model for each variable, without imposing a known break date. This is occurs by combining the test statistics computed for each possible break date in the sample, and selecting the date for which the test statistic is at the highest value, and statistically significant (Wright, 1998). The exact test statistic used for the grid search is a robust version of the Wald test (Andrews, 1993), which accounts for the presence of unknown forms of het-



eroskedasticity. A hypothesis is not rejected if a statistically significant structural break is identified at the proposed range of years. Statistical significance is, as per convention, set beyond a 5% p-value. In order to check for the robustness of the results, AR models of both first and second orders are evaluated<sup>3</sup>, and the tests are applied to both. The results are shown and discussed in the next subsection.

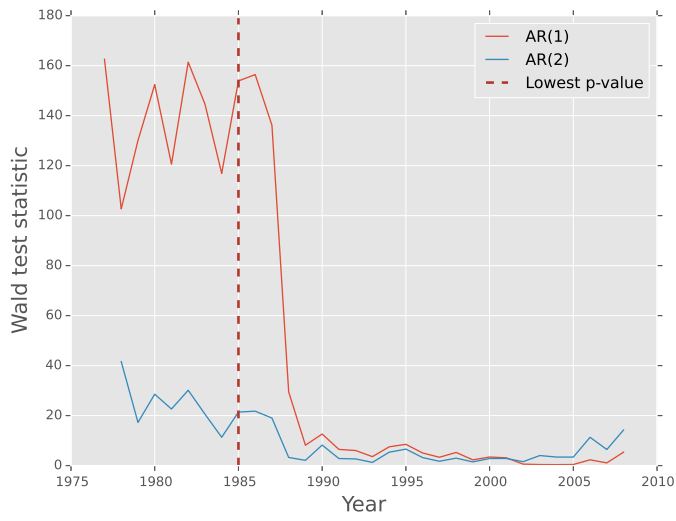
**Results.** The results of the tests for an unknown structural break for each of the three variables is shown in table 6.1. The dynamics of the Wald test statistics are displayed in figure 6.4.

	Year	Test statistic	P-value
Density, AR(1)	1985	153.869	0.000
Density, AR(2)	1985	21.425	0.000
USSR, AR(1)	1992	41.19	0.000
USSR, AR(2)	1992	37.466	0.000
USA, AR(1)	1989	7.538	0.023
USA, AR(2)	1989	16.969	0.000

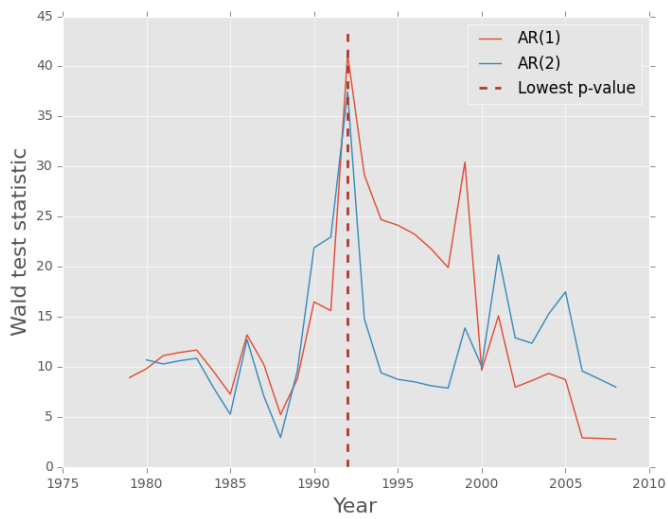
**Table 6.1:** Structural break test statistics and p-values.

---


<sup>3</sup>The order of an auto-regressive model indicates the number of lags used as regressors.



(a) Network density



(b) USSR/Russian Federation



(c) USA

Figure 6.4: Wald test statistics for each year.

Hypothesis (i) is clearly not rejected. In both autoregressive models, the test identifies a structural break in year 1985, with a very high and statistically significant Wald statistic. Figure 6.5 makes the result even clearer. The graph shows an indexed version of the density of the dynamic graph, rescaled to 0 in the first year of observation, 1970. The filtered graph experiences a dramatic decrease in density from 1985, dropping to a minimum value that is 90% less than that measured in the first year. Interestingly, the unfiltered version of the graph (i.e. where all edges with a non-zero value are kept) appears to be a rather stationary process, fluctuating around the same level of density across all years. The implication is extremely important: with the onset of the end of the Cold War, the number of high-level similarities across nations' political interests and semantics – as extrapolated from their UNGA speeches – severely diminishes. A diversification in positions, topics tackled, and language used evidently takes place. The application of the framework in this context, therefore, has not only allowed the unsupervised identification of a global political change, but also provides an empirical case for theories of international relations about a multi-polar world. In contrast to the 'end of history' theory, the latest advances in the field suggest the emergence of a global political regime where several entities act without benefiting from hegemony (Cooper and Flandes, 2013). In this scenario, global political discourse arenas such as the UN General Assembly should not be governed by the interest and priorities of one countries; rather, more variation in interests should be observed (Lesage and Van de Graaf, 2016). Lower network densities, mirroring weaker semantic and ideological connection, are an operationalisation of this theoretical phenomenon, which is worth verifying with further empirical research. Another interesting insight is supplied by figure 6.4(a). Very high Wald test statistics are found in some years before and after 1985, although not with such a high level of statistical significance. This may indicate the presence of more structural breaks, or of a very unstable systematic regime at the onset of a break. Both hypotheses are sound, and will be clarified through further inspection of the timeseries.

Hypothesis (ii) is not rejected either. The eigenvector centrality time series referring to the USSR/Russian Federation is characterised, according to the test, by a structural break in 1992, whose Wald statistic is highly significant both in the first and second order auto-regressive models. The graph in figure 6.6 provides more information about the kind of regime change that occurred in the time series.

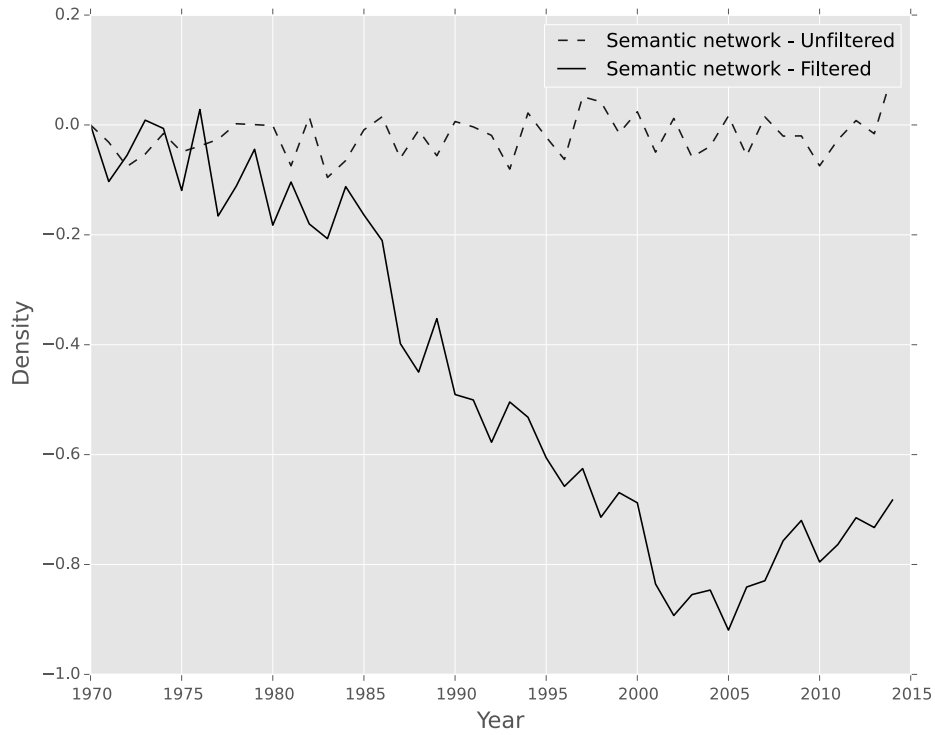


Figure 6.5: Density index of the UNGA dynamic networks, 1970-2014. Base year = 1970.




Figure 6.6: Eigenvector centrality indices for USA and USSR/Russian Federation. Base year = 1970.

The graph shows an indexed version of the centrality, with a base year of 1970. At the onset of the Russian Federation and dismissal of the Soviet Union, the time-series experiences a significant drop, big enough to be considered by the Wald test as a structural change from past patterns. Indeed, the break is so unequivocal that the 1992 Wald statistic surpassed other years' by far in figure 6.4(b). For the former leader of the Soviet Union, the change in regime meant a drastic change in the political discourse confronted, policy agenda and status (Lucas, 2014). Having transformed into a post-Soviet nation, with structural economic problems which then led to a crisis, in the following years the country reasonably enjoyed less ideological and semantic conformity from the rest of the international community than in the Cold War era. The post-1992 timeseries process appears to be highly volatile. Interestingly, the three peaks coincide with important national socio-economic events, respectively the 1998 Russian financial crisis, the 2004 re-election of Vladimir Putin as president, and the 2009 Great Recession and gas dispute with Ukraine. It is possible to interpret the relative increases in the centrality of the Russian Federation in those years as an increment in the policy themes and agenda points shared with other nations, due to the critical events listed. However, further research ought to discern and test for such effects, since the current model is not capable of identifying the semantic content which made Russia more central.

A structural break is also found with regards to the United States. Table 6.1 shows that the test found a break in 1989, which is significant at a 5% threshold value in the two auto-regressive models run. In figure 6.4(c), it is interesting to note that the Wald test spikes in that year in the AR(2), while in the AR(1) it competes with that in 1986, the year after the rise of Gorbachev. The index of centrality with a base year of 1970 is in figure 6.6. Contrary to the conventional theoretical explanation for a regime change previously advanced, US centrality appears to have decreased in comparison to the Cold War trend, although the high volatility of the process after 1989 does not provide a clear and definite answer to the question of the nation's political standing in this system. It is interesting to note high spikes in centrality in the last time steps of the series, although the causes for the phenomenon are unclear. This provides further evidence of the inability of the framework to untangle the underlying semantic and ideological mix that makes up the information forming the network.

In this section, it has been demonstrated that it is possible to detect mutations in the global political order by taking into consideration the sole semantic content of the speeches pronounced at the UNGA General Debate, and associated intercon-

nections. A series of Wald tests for a timeseries structural break at unknown date have successfully identified a change in regime for three different topological measures extracted from the inferred dynamic network, namely the density of the graph, and the eigenvector centralities of both the USSR – then the Russian Federation – and the United States of America. The structural break are found within the time frame hypothesised, and the dynamics of the three graph properties is surprisingly interpretable, despite the inability of the method to discern the exact causes behind the processes – be they rooted in merely political or economic events. The implications of the analysis are extremely important for the study of political systems. The NIIF application in this case study demonstrates that it is possible to detect signals in international politics with the sole use of textual data, paving the way for further investigations with other datasets not yet explored. Its results, furthermore, generate novel empirical insights that may inform international relations theory. The decreasing trend in density observed towards and after the end of the Cold War, for instance, hints at the emergence of a semantically more heterogeneous political regime, where the topics and narratives adopted are more diversified across countries. Furthermore, at least in relation to the Russian Federation, it has been noticed that some changes in the relative value of the influence metric correlate with important political or economic events. Future research can ground novel hypotheses on the dynamics of international political interests upon the empirical data provided by this study. Additional suggestions and ameliorations of the method are provided in section 6.5, accompanied by some concluding remarks.

## 6.4 Practical implementations

In the context of international relations, the network inference and influence framework can be tweaked to yield policy instruments that can be of use to the policy practitioner, as well as policy advocates and active citizenship. This section introduces the topic-related semantic index, a tool that monitors the degree of interest exhibited by one or more countries towards a pre-determined policy theme. The rest of the section outlines its analytical evaluation, and demonstrates its use with regard to four representative policy areas: health, education, nuclear weaponry and Islamic terrorism.

Recall that discrete state neural learning transforms into vector representations not only the set of discrete states (collection of words, for instance) associated with each element of a complex social system, but also each unique discrete state as

well. In this case study, it follows that a 200-element vector of hyper-dimensional coordinates has been evaluated not only for the speeches made by each UNGA member for each year, but also for each unique word contained in the entire corpus. It is, therefore, possible to extract not only the semantic distance across nations, but also between nations and words. This idea is at the base of the topic-related semantic index.

A topic-related semantic index is defined as the deviation in the arithmetic mean cosine similarity of a group of country-years (or simply the cosine similarity, in case of only one country) and a set of tokens deemed to represent a policy theme, with respect to a base year. The index is best explained through an example. Say that it is intended to calculate the topic-related index for country group  $N = [n_1, n_2, n_3, \dots, n_n]$ , with respect to a policy theme represented by the set of tokens  $W = [w_1, w_2, w_3, \dots, w_k]$ . Let year  $y_b$  be the base year, and year  $y_t$  the year of interest. The topic-related semantic index  $I_{y_t}^{N,W}$  is evaluated by

$$I_{y_t}^{N,W} = \frac{\frac{1}{n} \sum_{i=1}^n C_{y_t}(n_i, W) - \frac{1}{n} \sum_{i=1}^n C_{y_b}(n_i, W)}{\left| \frac{1}{n} \sum_{i=1}^n C_{y_b}(n_i, W) \right|}, \quad (6.1)$$

where  $C_{y_t}(n_i, W)$  is the cosine similarity between nation  $i$  and words  $W$  at year  $y_t$ . Note that the cosine similarity with a group of words can be easily performed by taking into account the mean across their respective vector values.  $I^{N,W}$  is real-valued, approaching zero when the mean similarity at a given point in time is close to the base year. Negative values imply a semantic similarity that is less than in the base year, vice versa when positive values are observed.


This implementation evaluates topic-related indices for the entire UNGA membership, with the aim of identifying macro-scale changes in policy preferences. Obviously, any other arbitrary choice of countries is possible. Four indices are calculated, each referring to a different policy theme which has been arbitrarily chosen, and operationalised by a couple of key words, as shown in table 6.2.

Policy theme	Key words
Health	'health', 'sanit'
Education	'educ', 'school'
Nuclear weapons	'nuclear', 'weapon'
Islamic terrorism	'terror', 'islam'


**Table 6.2:** Policy themes considered and respective key words

The choice of tokens is informed by the sense of representativeness that they feature for their respective policy themes. It is, indeed, common sense to think that the vector representations of the words ‘nuclear’ and ‘weapon’ would be central when countries refer to nuclear arsenals. Obviously, this arguably arbitrary method can be substituted with more analytical ones in the next generations of this proof of concept. For instance, key words associated to a policy theme can be identified by means of topic modelling.







(a) Education



(b) Health



(c) Nuclear weapons



(d) Islamic terrorism

**Figure 6.7:** Topic-related semantic indices for education, health, nuclear weapons and Islamic terrorism. Base year = 1995.

Figure 6.7 shows the dynamics of the topic-related semantic indices with regard to the four policy themes selected. They summarise the average semantic similarity of all UNGA members' speeches to the key words that represent such topics. Their dynamics appear to be well explained by historical events. Both the education and health indices feature lower levels until about 1985, the year Mikhail Gorbachev came to power in the USSR. With Gorbachev and the subsequent end of the Cold War, policies at the United Nations started becoming less focused on security issues and more on social and development goals (Simes, 1987). The steady, increasing trend of the indices since then can be explained by more and more focus on the Millennium Development Goals policy efforts, which retain both education and health at their core (Sachs and McArthur, 2005; Griggs et al., 2013).

In a similar vein, the nuclear weapons index experiences an abrupt decline with Gorbachev's rise to power, featuring a slow declining trend since then, thus suggesting that the topic has had much less importance than during the nuclear arms race. Interestingly, the Islamic terror index is more volatile and responsive to large terrorist attacks. Starting in 1979 the year regarded by experts as the commencing point for organised fundamentalist Islamic terror (O'Ballance, 1997) – the index registers high levels, in concomitance with the bombings of US embassies in Lebanon and Kuwait (Quillen, 2002). A new spike occurs in 2001, most likely due to the 9/11 attacks. A rising trend is observed during the last years of the index, explained by the emergent threat from ISIL, which has targeted both civilians and UNESCO sites (Harmanşah, 2015).

This variety of semantic index can also pave the way towards a more quantitative study of political ideologies and their dynamics. Provided that a reliable method for the choice of representative key words is developed, this novel 'science of political ideas' would allow for the investigation of how specific political ideologies gain or lose salience over one or more countries, and what factors underlie such mechanisms. A potential interesting question to tackle is, for instance, the empirical inspection of the diffusion of liberalist ideas (Simmons et al., 2006). Great amounts of theoretical work and case studies have provided valuable insights into the role now possessed by liberalism in the international regime, yet it has not been feasible to understand its microdynamics: do countries become more or less liberal over time? Is there a latent fluctuation dependent on socio-economic or political events? Further application of the framework of this thesis can contribute toward this line of work.

## 6.5 Case study discussion

This final case study has tested the ability of the network and inference framework to yield insights capable of detecting regime changes in international politics. As in the cases of financial institution and parliamentary speakers, it has been assumed that there exists a latent, not directly observable, network of nations which share political ideas and themes, thereby affecting their political interests. Their speeches at the UN General Debate provide a unique source of information capable of manifesting such information, and have therefore been used as the pillars of the network inference, this time performed using discrete state neural learning. The machine learning model evaluated vector representations of the countries' speeches at each year of the corpus (between 1970 and 2014), and the calculation and subsequent filtering of the cosine similarity for each pair of vectors has allowed the creation of a dynamic network of countries. Different from the networks obtained in the previous case study, this one is undirected, implying that it is not possible to distinguish a flow of information from one country to another. Yet, the analysis of the prominence of nations by calculating their eigenvector centralities has still shown to provide interesting information about the relative political position of a country in relation to others. This has been tested by attempting to detect a structural break in the time series of the centralities of the two major players during the Cold War: the Soviet Union then identified as the Russian Federation and the United States of America. A statistical test for a structural break in their eigenvector centralities has found statistically significant regime changes, respectively in 1992 and 1989. In those years, in fact, both countries experienced a sudden decline in their centrality, entailing that other countries began to be less semantically close to the views expressed in their speeches. The shift towards a more heterogeneous state of the world at least with regard to the policy ideas as expressed in the dataset is mirrored by a further structural break of the density of the dynamic graph in 1985, with the rise of Gorbachev as leader of the Soviet Union. From 1985, the graph experienced a decreasing trend in density, which signifies greater semantic distance between countries, and therefore a set of more diverse political ideas and interests shown. These empirical findings appear to be in line with theoretical contributions advocating that the international regime has become polycentric, despite the presence of potentially hegemonic powers such as the United States. This chapter has then provided an extension of the results, by introducing topic-related semantic indices which measure the semantic distance between the textual data assigned to one or more nations, and a set of key words associated to a policy theme. Four indices have been evaluated, showing the closeness of the entire UNGA membership to the themes of education, health, nuclear weapons and Islamic terrorism. Over time,

nations appear to be more semantically related to the first two topics, in line with the efforts made since the end of the Cold War and greater attention paid to the Millennium Development Goals. The topic of nuclear weapons has shown a declined centrality since the rise of Gorbachev, whereas the index of Islamic terror is extremely sensitive to large terrorist attacks.

The validity of this case study application may be challenged by a number of limitations. Two are discussed here. The first relates to the initial parameters that must be arbitrarily set for the discrete state neural learning, namely vector representation window size. The author has experimented with a pre-determined range of parameters, yielding similar results, yet at the current state of the research it is not possible to state whether other, unexplored sets of parameters could have provided different yet still valid results. Being in its infancy, machine learning specialists have not dedicated themselves to elaborating analytical ways to automatically identify the right value for the parameters, according to the data context. Most probably, as in other black-box machine learning models, such a solution would remain infeasible. Further research, though, can provide robustness checks, by replicating the study with different vector sizes and windows.

A further limitation is the non-directionality of the dynamic graph inferred. The reason for the lack of information about the direction of potential influence between nations lies on the nature of the data in hand. General Debate speeches are annual and concurrent, it is therefore not possible to know whether one speech influenced another in the same year. Also, it is not likely that a speech made in a given year would influence one the year after. The time distance is large, and the number of political events and political interests of nations are most probably dependent on other information than speeches made by other countries a year previous. Directional influence can be observed by applying the framework on more frequent data points, such as General Assembly speeches, or speeches at specific UN Commissions. As in the case study of the previous chapter, in this context it would be feasible to obtain directional networks which would help to better identify those actors exerting higher degrees of influence.

Related to the discussion above, this case study is an example of how the concept of influence is very sensitive to the empirical conditions in which a study is located. At its core, influence is about a flow of information between entities which results in the change of behaviour of one of those entities. Yet, it may still be possible to detect a 'soft' version of influence even when the data does not allow for the extraction of information flow directionality. The reason is the very nature of

the concept of the graph. If a node is highly central, with a relatively high number of connections to other nodes, it is reasonable to think that such a node is influential in the latent flow of information taking place. If the node is removed from the graph, certainly the information flow would be severely impaired. Because of this property, the concept of influence elevates to a systemic level: even when it is not now the exact channel of diffusion of the semantic content, a central node remains critical for the very existence of the graph. This case study exemplifies the concept, and can function as a precedent for data contexts with no insights into directionality.



## Chapter 7

# Concluding remarks

Social life is one the least understood instances in which complexity arises. At the same time, it is the one that most closely affects the life of billions of people on Earth. Untangling even a small glimpse, piece by piece, of the properties of social complex systems can result in valuable outcomes for the well-being of humankind. The science of society is a science that unveils not just laws and patterns, but also points of action.

This is, at least, the principle on which this thesis has been cemented. This research has been a humble attempt to shed light onto one of many problems surrounding the study of social complex systems. It has focused on the twofold information diffusion latency problem: it is often impossible to directly observe the network by which social entities share information; by extension, it is often impractical to measure the extent to which a social agent is able to influence the behaviour of other agents or of the whole system. Social complex systems vary greatly in terms of the empirical data that they generate for the benefit of their observers, making it difficult for one method to tackle and solve the latency problem in all possible empirical instances. To address the issue, it has been therefore chosen to create a multi-methodological framework, the Network Inference and Influence Framework (NIIF). As its name implies, NIIF is comprised of two modules; the first dedicated to the reconstruction or estimation of the latent network of social agents using available data, and the second focused on synthesising a metric of the extent to which agents influence each other. The methodological specifications of NIIF are explained in Chapter 3, and involve insights coming from graph theory, information theory and epidemiological science. To demonstrate the versatility of the framework, three apparently diverse case study applications have been presented. They respectively involve an economic and two political systems. While the areas and topics covered may seem disconnected at first sight, they in fact all revolve around the same latency problem, that is, ignorance about the latent network by

which social actors shared information and resources.

In Chapter 4, the latency problem instantiated in the lack of knowledge about the flow of financial credit resources among a system of European banks. The implication was that at times of financial distress, the degree to which a bank or group of banks were capable of harming the stability of the whole system was unknown too. The application of NIIF allowed the reconstruction of an ensemble of financial networks based on aggregate data about the institutions' credit exposure to each other. In turn, it has also measured the degree of damage that would originate by a bank given a financial shock in more sensitive ways than alternative approaches. The banks that were most influential in determine the stability of the financial system could therefore been identified, and the information can be used for stress-testing and other macroprudential policy action.

Chapter 5 has faced the latency problem surrounding political parliamentary systems: it is not known how political agents share information and influence each other in the process of policy debate and making. NIIF has been used to extract representations of parliamentary networks surrounding several topics from House of Common Debates. The networks have been used to investigate the level of influence of parliamentary speakers, discovering that individuals with governmental positions are more likely to have other speakers repeating their message. The practical application of the analysis permits, among other things, to keep track of the political actors who exert more influence in particular policy topics. The instrument can be used to render elected leaders more accountable, and their work more transparent.

In Chapter 6, the latency problem characterised the semantic and ideological information shared across country members of the United Nations. Knowing this form of information flow may allow the detection of changes in political order. In this case study, NIIF has been applied on General Debates' speeches, where UNGA member nations have the opportunity to share their view on key contemporary issues with the international community. The result was a dynamic network of nations, connected through semantic and ideological similarity. Structural change in the network and in the relative influence of the USSR (then Russian Federation) and the USA anticipated and accompanied the end of the Cold War. The framework has also allowed the construction of instruments able to measure the salience of policy themes over time. Ideological change, it has been shown, can be quantified.

In each case study, a discussion of the main limitations has been exposed, reminding the reader that the framework needs further improvement and validity testing. Definitively, a more general limitation is its multi-methodological nature, which could foment some confusion about which specific approaches to use in a



determined context. Hopefully, the explanations in Chapter 3 and the case studies have shed some clarity of the issue.

The framework can be applied to gain further insight into other valuable economic and political systems and phenomena which require a greater policy attention. An example is the link between commodity markets and social instability. Between late 2007 and early 2008 the prices of grain, corn and rice rose by over 100%. A similar spike occurred in the fall of 2010 (Gilbert and Morgan, 2010). The shocks, due mainly to ethanol conversion<sup>1</sup> and investor speculation, acted as a main triggering point for the social dissent in several Arab-speaking countries, leading to the so-called Arab Spring, and the Syrian war and refugee crisis (Lagi et al., 2015). The Arab Spring counts 180,000 people killed, and 6 million displaced (World Bulletin, 2013), while in Syria alone the number of refugees has surpassed 4.5 million (Amnesty International, 2016). The application of transfer entropy on commodity price changes and social unrest data can result in the inference of the underlying network of flow of information from financial systems to social collectives.

Similarly, another interesting area to explore and test is the existence of a nexus between corporate power and political events. Studies such as Vitali et al. (2011) have shown how that a handful of companies appear to have corporate control over the system of international firms, comprising tens of thousands of elements. It is, however, not yet ascertained whether this sort of control exerts some influence in the non-corporate context, particularly political systems. NIIF applied on corporate and political events data could untangle the existence of any information flow influencing the systems, and test its significance.

A further issue that can be confronted with the use of NIIF is shadow banking and its leverage on regulated financial systems. Shadow banking is the term used to refer to financial intermediaries that are not formally regulated by supervisory bodies such as central banks (Gennaioli et al., 2013). The volume of assets exchanged through the industry has been estimated to be about \$60 trillion globally in 2011 (Schwarcz, 2012). Yet it remains an unanswered question whether and to what extent this sector has the capacity to make regulated financial systems unstable. The difficulty in tackling the question mainly revolves around the lack of data. Banking institutions are obliged to provide information – albeit incomplete – to their regulators, whereas shadow banking has no duty of doing so. The network

---

<sup>1</sup>Ethanol conversion is the practice of using crops of corn, wheat or other plants for the production of ethanol biofuel, rather than for food production (Pimentel and Patzek, 2005).

of resource flow between the two financial sectors is therefore latent. Provided the use of incomplete data about shadow banks exposure to securities and commodities, it is possible to use a similar instantiation of NIIF as in Chapter 2. The result would be the estimation of an ensemble of networks of dependency across regulated and unregulated financial firms. The measure of influence of the nodes in the system may result in valuable insights to be adopted by policy makers.

Overall, this thesis has demonstrated that the solution of the latency problem in social complex systems is not only an advancement of knowledge itself; it also enables to supply potential solutions to some of the problems faced in economics and government. The future needs the ingenuity of complexity science.

# Bibliography

Acharya, V. V.

2009. *On the financial regulation of insurance companies*. PhD thesis, NYU Stern School of Business.

Adolfson, M., M. Andersson, J. Lindé, M. Villani, and A. Vredin

2007. Modern forecasting models in action: Improving macroeconomic analyses at central banks. *Riksbank Research Paper Series*, (23).

Afonso, G., A. Kovner, and A. Schoar

2013. Trading partners in the interbank lending market. *FRB of New York Staff Report*, (620).

Agarwal, N., H. Liu, L. Tang, and P. S. Yu

2008. Identifying the influential bloggers in a community. In *Proceedings of the 2008 international conference on web search and data mining*, Pp. 207–218. ACM.

Allen, F., A. Babus, and E. Carletti

2012. Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104(3):519–534.

Allen, F. and D. Gale

2000. Financial contagion. *Journal of political economy*, 108(1):1–33.

Almog, A., T. Squartini, and D. Garlaschelli

2015. A gdp-driven model for the binary and weighted structure of the international trade network. *New Journal of Physics*, 17(1):013009.

Amann, E. and W. Baer

2002. Neoliberalism and its consequences in brazil. *Journal of Latin American Studies*, 34(04):945–959.

Amnesty International

2016. Syria's refugee crisis in numbers.

Anand, K.

2015. The missing links: A global study on uncovering financial network structure from partial data. *Financial Network Theory Conference*.

Andrews, D. W.

1993. Tests for parameter instability and structural change with unknown change point. *Econometrica: Journal of the Econometric Society*, Pp. 821–856.

Anushka, A., E. Jessica, and H. Stewart

2016. Jeremy corbyn preparing for leadership contest. *The Guardian*. (Accessed on 07/27/2016).

Arthur, W. B.

1999. Complexity and the economy. *science*, 284(5411):107–109.

Arthur, W. B.

2014. *Complexity and the Economy*. Oxford University Press.

Ashton, B., K. Hill, A. Piazza, and R. Zeitz

1992. Famine in china, 1958–61. In *The Population of Modern China*, Pp. 225–271. Springer.

Atkinson, T., D. Luttrell, H. Rosenblum, et al.

2013. How bad was it? the costs and consequences of the 2007–09 financial crisis. *Staff Papers*, (Jul).

Austen-Smith, D.

1990. Information transmission in debate. *American Journal of political science*, Pp. 124–152.

Avellaneda, M. and J. Zhu

2001. Distance to default. *Risk*, 14(12):125–129.

Bailey, M. A., A. Strezhnev, and E. Voeten

2015. Estimating dynamic state preferences from united nations voting data. *Journal of Conflict Resolution*, P. 0022002715595700.

Barabási, A.-L.

2013. Network science. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 371(1987):20120375.

Barabási, A.-L. and R. Albert

1999. Emergence of scaling in random networks. *science*, 286(5439):509–512.

- Barringer, S. N., S. R. Eliason, and E. Leahey  
2013. A history of causal analysis in the social sciences. In *Handbook of Causal Analysis for Social Research*, Pp. 9–26. Springer.
- Basel Committee  
2011. Global systemically important banks: assessment methodology and the additional loss absorbency requirement.
- Basel Committee  
2013. Updated assessment methodology and the higher loss absorbency requirement. *Basel Committee on Banking Supervision: Basel, Switzerland*.
- Battiston, S., S. Gurciullo, G. Caldarelli, and M. D’Errico  
2016. Leveraging the network: a stress-test framework based on debtrank. *Available at SSRN 2571218*.
- Battiston, S., M. Puliga, R. Kaushik, P. Tasca, and G. Caldarelli  
2012. Debtrank: Too central to fail? financial networks, the fed and systemic risk. *Scientific reports*, 2.
- Baturo, A., N. Dasandi, and S. Mikhaylov  
2016. The general debate of the united nations and government preferences in world politics. *Article pre-print*.
- Baturo, A. and S. Mikhaylov  
2014. Reading the tea leaves: Medvedev’s presidency through political rhetoric of federal and sub-national actors. *Europe-Asia Studies*, 66(6):969–992.
- Baumgartner, F. R., J. M. Berry, M. Hojnacki, D. C. Kimball, and B. L. Leech  
2014. Money, priorities, and stalemate: How lobbying affects public policy. *Election Law Journal*, 13(1):194–209.
- Baumgartner, F. R. and B. Jones  
1993. Agendas and instability in american politics. *Chicago: University of Chicago Press. Baumgartner Agendas and Instability in American Politics 1993*.
- BBC  
2005. Clarke attacks ‘catastrophic’ war. [http://news.bbc.co.uk/2/hi/uk\\_news/politics/4203296.stm](http://news.bbc.co.uk/2/hi/uk_news/politics/4203296.stm). (Accessed on 08/22/2016).
- BBC  
2016a. Brexit: David cameron to quit after uk votes to leave eu. <http://www.bbc.co.uk/news/uk-politics-36615028>. (Accessed on 07/27/2016).

## BBC

2016b. Keith vaz: Scotland 'leading the way' in resettling refugees. <http://www.bbc.com/news/uk-scotland-36956721>. (Accessed on 08/22/2016).

## Beinhocker, E. D.

2006. *The origin of wealth: Evolution, complexity, and the radical remaking of economics*. Harvard Business Press.

## Berry, J. M., C. Wilcox, et al.

2015. *Interest group society*. Routledge.

## Bevan, G. and R. Robinson

2005. The interplay between economic and political logics: path dependency in health care in England. *Journal of Health Politics, Policy and Law*, 30(1-2):53–78.

## Bianchetti, M., D. Galli, C. Ricci, A. Salvatori, and M. Scaringi

2016. Brexit or bremain? evidence from bubble analysis. *Evidence from Bubble Analysis (June 20, 2016)*.

## Birkland, T. A.

2014. *An introduction to the policy process: Theories, concepts and models of public policy making*. Routledge.

## Blei, D. M.

2012. Probabilistic topic models. *Communications of the ACM*, 55(4):77–84.

## Bonabeau, E.

2002. Agent-based modeling: Methods and techniques for simulating human systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7280–7287.

## Bonacich, P.

2007. Some unique properties of eigenvector centrality. *Social networks*, 29(4):555–564.

## Bond, R. M., C. J. Fariss, J. J. Jones, A. D. Kramer, C. Marlow, J. E. Settle, and J. H. Fowler

2012. A 61-million-person experiment in social influence and political mobilization. *Nature*, 489(7415):295–298.

- Bondy, J. A. and U. S. R. Murty  
1976. *Graph theory with applications*, volume 290. Citeseer.
- Boss, M., H. Elsinger, M. Summer, S. Thurner, et al.  
2004a. An empirical analysis of the network structure of the austrian interbank market. *Oesterreichische Nationalbanks Financial stability Report*, 7:77–87.
- Boss, M., H. Elsinger, M. Summer, and S. Thurner  
2004b. Network topology of the interbank market. *Quantitative Finance*, 4(6):677–684.
- Box-Steffensmeier, J. M. and B. S. Jones  
2004. *Event history modeling: A guide for social scientists*. Cambridge University Press.
- Brämer, P., H. Gischer, et al.  
2012. *Domestic systemically important banks: an indicator-based measurement approach for the Australian banking system*. Univ., Faculty of Economics and Management.
- Brandes, U. and D. Wagner  
2004. Analysis and visualization of social networks. In *Graph drawing software*, Pp. 321–340. Springer.
- Caccioli, F., M. Shrestha, C. Moore, and J. D. Farmer  
2014. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking & Finance*, 46:233–245.
- Caldarelli, G. and M. Catanzaro  
2012. *Networks: A very short introduction*, volume 335. Oxford University Press.
- Caldarelli, G., A. Chessa, F. Pammolli, A. Gabrielli, and M. Puliga  
2013. Reconstructing a credit network. *Nature Physics*, 9(3):125–126.
- Campos, M. and C. Jacobs-Wagner  
2013. Cellular organization of the transfer of genetic information. *Current opinion in microbiology*, 16(2):171–176.
- Carpenter, D. P., K. M. Esterling, and D. M. Lazer  
2004. Friends, brokers, and transitivity: Who informs whom in washington politics? *Journal of Politics*, 66(1):224–246.

Carr, C.

2005. Are german, japanese and anglo-saxon strategic decision styles still divergent in the context of globalization? *Journal of Management Studies*, 42(6):1155–1188.

Cavnar, W. B., J. M. Trenkle, et al.

1994. N-gram-based text categorization. *Ann Arbor MI*, 48113(2):161–175.

Cederman, L.-E.

2001. Back to kant: Reinterpreting the democratic peace as a macrohistorical learning process. In *American Political Science Association*, volume 95, Pp. 15–31. Cambridge Univ Press.

Chen, S.-H., C.-L. Chang, and Y.-R. Du

2012. Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(02):187–219.

Chesters, G.

2004. Global complexity and global civil society. *Voluntas: International Journal of Voluntary and Nonprofit Organizations*, 15(4):323–342.

Chicheportiche, R. and J.-P. Bouchaud

2012. Weighted kolmogorov-smirnov test: Accounting for the tails. *Physical Review E*, 86(4):041115.

Chinazzi, M. and G. Fagiolo

2015. Systemic risk, contagion, and financial networks: A survey. *Contagion, and Financial Networks: A Survey (June 3, 2015)*.

Choudhry, M., E. Marelli, M. Signorelli, et al.

2010. The impact of financial crises on youth unemployment rate. *Quaderni del Dipartimento di Economia, Finanza e Statistica*, 79:2010.

Chumacero, R. A. and K. S. Hebbel

2004. General equilibrium models: an overview. *Documentos de Trabajo (Banco Central de Chile)*, (307):1.

Cingranelli, D. L. and D. L. Richards

1999. Respect for human rights after the end of the cold war. *Journal of Peace Research*, 36(5):511–534.



- Citi, M.  
2013. Eu budgetary dynamics: incremental or punctuated equilibrium? *Journal of European Public Policy*, 20(8):1157–1173.
- Claessens, S., G. DellAriccia, D. Igan, and L. Laeven  
2010. Cross-country experiences and policy implications from the global financial crisis. *Economic Policy*, 25(62):267–293.
- Clemens Jr, W. C.  
2001. *The Baltic transformed: Complexity theory and European security*. Rowman & Littlefield Publishers.
- Clement, P.  
2010. The term 'macroprudential': origins and evolution. *BIS Quarterly Review*, March.
- Cohen, K. J., G. A. Hawawini, S. F. Maier, R. A. Schwartz, and D. K. Whitcomb  
1983. Friction in the trading process and the estimation of systematic risk. *Journal of Financial Economics*, 12(2):263–278.
- Colander, D., H. Föllmer, A. Haas, M. D. Goldberg, K. Juselius, A. Kirman, T. Lux, and B. Sloth  
2009. The financial crisis and the systemic failure of academic economics. *Univ. of Copenhagen Dept. of Economics Discussion Paper*, (09-03).
- Commission, F. C. I. and U. S. F. C. I. Commission  
2011. *The financial crisis inquiry report: Final report of the national commission on the causes of the financial and economic crisis in the United States*. PublicAffairs.
- Conrad, J. and G. Kaul  
1998. An anatomy of trading strategies. *Review of Financial studies*, 11(3):489–519.
- Cooper, A. F. and D. Flemes  
2013. Foreign policy strategies of emerging powers in a multipolar world: An introductory review. *Third World Quarterly*, 34(6):943–962.
- Csermely, P., A. London, L.-Y. Wu, and B. Uzzi  
2013. Structure and dynamics of core/periphery networks. *Journal of Complex Networks*, 1(2):93–123.

- Dagnis Jensen, M. and H. Snaith  
2016. When politics prevails: the political economy of a brexit. *Journal of European Public Policy*, Pp. 1–9.
- Davis, R. A., T. C. M. Lee, and G. A. Rodriguez-Yam  
2006. Structural break estimation for nonstationary time series models. *Journal of the American Statistical Association*, 101(473):223–239.
- De Masi, G., G. Iori, and G. Caldarelli  
2006. Fitness model for the italian interbank money market. *Physical Review E*, 74(6):066112.
- De Walque, G., O. Pierrard, and A. Rouabah  
2010. Financial (in) stability, supervision and liquidity injections: a dynamic general equilibrium approach. *The Economic Journal*, 120(549):1234–1261.
- Denzau, A. T. and M. C. Munger  
1986. Legislators and interest groups: How unorganized interests get represented. *American Political Science Review*, 80(01):89–106.
- Devaney, R. L.  
1992. *A first course in chaotic dynamical systems*. Westview Press.
- Dhingra, S., G. I. Ottaviano, T. Sampson, and J. V. Reenen  
2016. The consequences of brexit for uk trade and living standards.
- Dimpfl, T. and F. J. Peter  
2013. Using transfer entropy to measure information flows between financial markets. *Studies in Nonlinear Dynamics and Econometrics*, 17(1):85–102.
- Dopfer, K.  
1991. Toward a theory of economic institutions: Synergy and path dependency. *Journal of Economic Issues*, 25(2):535–550.
- Dreher, A. and J.-E. Sturm  
2012. Do the imf and the world bank influence voting in the un general assembly? *Public Choice*, 151(1-2):363–397.
- Drehmann, M. and N. Tarashev  
2013. Measuring the systemic importance of interconnected banks. *Journal of Financial Intermediation*, 22(4):586–607.

- Dungey, M., M. Luciani, and D. Veredas  
2012. Ranking systemically important financial institutions.
- Dür, A. and G. Mateo  
2012. Who lobbies the european union? national interest groups in a multilevel polity. *Journal of European public policy*, 19(7):969–987.
- Eisenberg, L. and T. H. Noe  
2001. Systemic risk in financial systems. *Management Science*, 47(2):236–249.
- Epstein, L. and C. Mershon  
1996. Measuring political preferences. *American journal of political science*, Pp. 261–294.
- Erikson, R. S., P. M. Pinto, and K. T. Rader  
2014. Dyadic analysis in international relations: A cautionary tale. *Political Analysis*, 22(4):457–463.
- Fagiolo, G. and A. Roventini  
2012. Macroeconomic policy in dsge and agent-based models. *Revue de l'OFCE*, (5):67–116.
- Fan, Y., S. Ren, H. Cai, and X. Cui  
2014. The state's role and position in international trade: A complex network perspective. *Economic Modelling*, 39:71–81.
- Farina, L. and S. Rinaldi  
2011. *Positive linear systems: theory and applications*, volume 50. John Wiley & Sons.
- Farmer, J. D. and D. Foley  
2009. The economy needs agent-based modelling. *Nature*, 460(7256):685–686.
- Farrell, J.  
2016. Network structure and influence of the climate change counter-movement. *Nature Climate Change*, 6(4):370–374.
- FDIC  
2011. Failed bank list.
- Fernández, J. J. and M. Lutter  
2013. Supranational cultural norms, domestic value orientations and the diffusion of same-sex union rights in europe, 1988–2009. *International Sociology*, 28(1):102–120.

- Fiedler, B. and M. Bukovsky  
2011. The effect of a giant wind farm on precipitation in a regional climate model. *Environmental Research Letters*, 6(4):045101.
- Fienberg, S. E. and M. M. Meyer  
2006. Iterative proportional fitting. *Encyclopedia of Statistical Sciences*.
- Financial Stability Board  
2009. Guidance to assess the systemic importance of financial institutions, markets and instruments: initial considerations. *Report to G20 finance ministers and governors*.
- Finnemore, M. and K. Sikkink  
1998. International norm dynamics and political change. *International organization*, 52(04):887–917.
- Flockhart, T.  
2006. complex socialization: A framework for the study of state socialization. *European Journal of International Relations*, 12(1):89–118.
- Flood, R. L. and E. Carson  
2013. *Dealing with complexity: an introduction to the theory and application of systems science*. Springer Science & Business Media.
- Focardi, S. M. and F. J. Fabozzi  
2013. Fat tails, scaling, and stable laws. *Encyclopedia of Financial Models*.
- Fowler, J. H. and N. A. Christakis  
2010. Cooperative behavior cascades in human social networks. *Proceedings of the National Academy of Sciences*, 107(12):5334–5338.
- Franke, M. F.  
2014. *Global limits: Immanuel Kant, international relations, and critique of world politics*. SUNY Press.
- Franzese, R. J. and J. C. Hays  
2008. Interdependence in comparative politics substance, theory, empirics, substance. *Comparative Political Studies*, 41(4-5):742–780.
- Freixas, X. and B. Parigi  
1998. Contagion and efficiency in gross and net interbank payment systems. *Journal of Financial Intermediation*, 7(1):3–31.

- Fricke, D. and T. Lux  
2015a. Core–periphery structure in the overnight money market: evidence from the e-mid trading platform. *Computational Economics*, 45(3):359–395.
- Fricke, D. and T. Lux  
2015b. On the distribution of links in the interbank network: Evidence from the e-mid overnight money market. *Empirical Economics*, 49(4):1463–1495.
- Froncek, P.  
2015. The effects of network structure on the emergence of norms in adaptive populations. *Journal of Artificial Societies and Social Simulation*, 18(4):14.
- Fukuyama, F.  
1989. The end of history? *The national interest*, (16):3–18.
- Fukuyama, F.  
2011. *The origins of political order: from prehuman times to the French Revolution*. Macmillan.
- Gallagher, R., T. Appenzeller, D. Normile, et al.  
1999. Beyond reductionism. *Science*, 284(5411):79.
- Gangel, M., M. J. Seiler, and A. Collins  
2013. Exploring the foreclosure contagion effect using agent-based modeling. *The Journal of Real Estate Finance and Economics*, 46(2):339–354.
- Gennaioli, N., A. Shleifer, and R. W. Vishny  
2013. A model of shadow banking. *The Journal of Finance*, 68(4):1331–1363.
- Georg, C.-P.  
2013. The effect of the interbank network structure on contagion and common shocks. *Journal of Banking & Finance*, 37(7):2216–2228.
- Gerali, A., S. Neri, L. Sessa, and F. M. Signoretti  
2010. Credit and banking in a dsge model of the euro area. *Journal of Money, Credit and Banking*, 42(s1):107–141.
- Gerlach, M. L.  
1992. *Alliance capitalism: The social organization of Japanese business*. Univ of California Press.
- Geyer, R.  
2003. European integration, the problem of complexity and the revision of theory. *JCMS: Journal of Common Market Studies*, 41(1):15–35.

Gibney, E.

2014. Force of nature gave life its asymmetry.

Gil de Zúñiga, H., N. Jung, and S. Valenzuela

2012. Social media use for news and individuals' social capital, civic engagement and political participation. *Journal of Computer-Mediated Communication*, 17(3):319–336.

Gilbert, C. L. and C. W. Morgan

2010. Food price volatility. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 365(1554):3023–3034.

Gilens, M. and B. I. Page

2014. Testing theories of american politics: Elites, interest groups, and average citizens. *Perspectives on politics*, 12(03):564–581.

Glattfelder, J. B.

2013. The network of global corporate control. In *Decoding Complexity*, Pp. 95–119. Springer.

Goldberg, Y. and O. Levy

2014. word2vec explained: deriving mikolov et al.'s negative-sampling word-embedding method. *arXiv preprint arXiv:1402.3722*.

Goldenfeld, N. and L. P. Kadanoff

1999. Simple lessons from complexity. *science*, 284(5411):87–89.

Gorbachev, M.

2007. The nuclear threat. *Wall Street Journal*, 31.

Greenacre, M.

2007. *Correspondence analysis in practice*. CRC press.

Griggs, D., M. Stafford-Smith, O. Gaffney, J. Rockström, M. C. Öhman, P. Shyam-sundar, W. Steffen, G. Glaser, N. Kanie, and I. Noble

2013. Policy: Sustainable development goals for people and planet. *Nature*, 495(7441):305–307.

Grimm, V. and S. F. Railsback

2013. *Individual-based modeling and ecology*. Princeton university press.

Grimmer, J.

2015. We are all social scientists now: how big data, machine learning, and causal inference work together. *PS: Political Science & Politics*, 48(01):80–83.

- Grimmer, J. and B. M. Stewart  
2013. Text as data: The promise and pitfalls of automatic content analysis methods for political texts. *Political Analysis*, P. mps028.
- Grossback, L. J., S. Nicholson-Crotty, and D. A. Peterson  
2004. Ideology and learning in policy diffusion. *American Politics Research*, 32(5):521–545.
- Gruhl, D., R. Guha, D. Liben-Nowell, and A. Tomkins  
2004. Information diffusion through blogspace. In *Proceedings of the 13th international conference on World Wide Web*, Pp. 491–501. ACM.
- Guille, A., H. Hacid, C. Favre, and D. A. Zighed  
2013. Information diffusion in online social networks: A survey. *ACM SIGMOD Record*, 42(2):17–28.
- Gurciullo, S.  
2014. Stress-testing the system: Financial shock contagion in the realm of uncertainty. *arXiv preprint arXiv:1412.1679*.
- Gurciullo, S., A. Herzog, P. John, and S. Mikhaylov  
2015a. How do policy memes spread? an analysis of the uk house of commons debates. *An Analysis of the UK House of Commons Debates (September 25, 2015)*.
- Gurciullo, S., M. Smallegan, M. Pereda, F. Battiston, A. Patania, S. Poledna, D. Hedblom, B. T. Oztan, A. Herzog, P. John, et al.  
2015b. Complex politics: A quantitative semantic and topological analysis of uk house of commons debates. *arXiv preprint arXiv:1510.03797*.
- Haldane, A. G.  
2013. Rethinking the financial network. In *Fragile stabilität–stabile fragilität*, Pp. 243–278. Springer.
- Haldane, A. G. and R. M. May  
2011. Systemic risk in banking ecosystems. *Nature*, 469(7330):351–355.
- Harmanşah, Ö.  
2015. Isis, heritage, and the spectacles of destruction in the global media. *Near Eastern Archaeology*, 78(3):170–177.
- Harrison, R., K. Nikolov, M. Quinn, G. Ramsay, A. Scott, and R. Thomas  
2005. *The Bank of England quarterly model*. Bank of England London.

Hart, O. and L. Zingales

2011. A new capital regulation for large financial institutions. *American Law and Economics Review*, 13(2):453–490.

Hartman, P.

2014. *Results of the ESCB Macro-prudential Research Network (MaRs)*, number 13-15. European Central Bank.

Hayles, N. K.

1991. *Chaos and order: Complex dynamics in literature and science*. University of Chicago Press.

Heaberlin, B. and S. DeDeo

2016. The evolution of wikipeidias norm network. *Future Internet*, 8(2):14.

Hecht-Nielsen, R.

1989. Theory of the backpropagation neural network. In *Neural Networks, 1989. IJCNN., International Joint Conference on*, Pp. 593–605. IEEE.

Heemeijer, P., C. Hommes, J. Sonnemans, and J. Tuinstra

2009. Price stability and volatility in markets with positive and negative expectations feedback: An experimental investigation. *Journal of Economic dynamics and control*, 33(5):1052–1072.

Henry, A. D., M. Lubell, and M. McCoy

2011. Belief systems and social capital as drivers of policy network structure: The case of california regional planning. *Journal of Public Administration Research and Theory*, 21(3):419–444.

Hoff, P. D. and M. D. Ward

2004. Modeling dependencies in international relations networks. *Political Analysis*, 12(2):160–175.

Hong, Z. and Y. Sun

2000. The butterfly effect and the making of 'ping-pong diplomacy'. *Journal of contemporary China*, 9(25):429–448.

Howard, M.

2011. Michael howard: Parliament must redefine human rights. *The Telegraph*.

Huang, X., I. Vodenska, S. Havlin, and H. E. Stanley

2013. Cascading failures in bi-partite graphs: model for systemic risk propagation. *Scientific reports*, 3.



Hulme, D.

2007. The making of the millennium development goals: Human development meets results-based management in an imperfect world. *Brooks World Poverty Institute Working Paper*, (16).

Ingber, D. E.

1998. The architecture of life. *Scientific American*, 278(1):48–57.

Jallo, D., D. Budai, V. Boginski, B. Goldengorin, and P. M. Pardalos

2013. Network-based representation of stock market dynamics: an application to american and swedish stock markets. In *Models, Algorithms, and Technologies for Network Analysis*, Pp. 93–106. Springer.

Jiang, B., L. Wang, C. Yang, and R. Li

2015. Mining multiplex power-law distributions and retweeting patterns on twitter. In *Fuzzy Systems and Knowledge Discovery (FSKD), 2015 12th International Conference on*, Pp. 1437–1443. IEEE.

Jiang, Y. and Z. Huang

2012. The rich get richer: Preferential attachment in the task allocation of cooperative networked multiagent systems with resource caching. *IEEE Transactions on Systems, Man, and Cybernetics-Part A: Systems and Humans*, 42(5):1040–1052.

Johansen, A. and D. Sornette

2000. The nasdaq crash of april 2000: Yet another example of log-periodicity in a speculative bubble ending in a crash. *The European Physical Journal B-Condensed Matter and Complex Systems*, 17(2):319–328.

Johansen, A. and D. Sornette

2001. Bubbles and anti-bubbles in latin-american, asian and western stock markets: An empirical study. *International Journal of Theoretical and Applied Finance*, 4(06):853–920.

John, P., S. Bevan, and W. Jennings

2014. Party politics and the policy agenda: the case of the united kingdom. *Policy Agenda-Setting Theory: From Policy Theory to a Theory of Politics*, forthcoming.

Kalinovsky, A., S. Radchenko, et al.

2011. *The end of the Cold War and the Third World: new perspectives on regional conflict*. Taylor & Francis.

Kam, C. J.

2009. *Party discipline and parliamentary politics*. Cambridge University Press.

Kauffman, S. A.

1993a. *The origins of order: Self organization and selection in evolution*. Oxford University Press, USA.

Kauffman, S. A.

1993b. *The origins of order: Self organization and selection in evolution*. Oxford University Press, USA.

Kavalski, E.

2007. The fifth debate and the emergence of complex international relations theory: notes on the application of complexity theory to the study of international life. *Cambridge Review of International Affairs*, 20(3):435–454.

Kelly, H.

2013. The power of one wrong tweet. <http://edition.cnn.com/2013/04/23/tech/social-media/tweet-ripple-effect/>. (Accessed on 06/28/2016).

Kennedy, J.

2011. Particle swarm optimization. In *Encyclopedia of machine learning*, Pp. 760–766. Springer.

Kiesling, E., M. Günther, C. Stummer, and L. M. Wakolbinger

2012. Agent-based simulation of innovation diffusion: a review. *Central European Journal of Operations Research*, 20(2):183–230.

Kim, G.-r. and H. M. Markowitz

1989. Investment rules, margin, and market volatility. *The Journal of Portfolio Management*, 16(1):45–52.

Kim, S. Y. and B. Russett

1996. The new politics of voting alignments in the united nations general assembly. *International Organization*, 50(04):629–652.

Kim, Y., S. Park, and S.-H. Yook

2016. The origin of the criticality in meme popularity distribution on complex networks. *Scientific reports*, 6.

- Kindleberger, C. P.  
2000. Manias, panics, and crashes: a history of financial crises. *The Scriblerian and the Kit-Cats*, 32(2):379.
- Klüver, H.  
2013. *Lobbying in the European Union: interest groups, lobbying coalitions, and policy change*. Oxford University Press.
- Klüver, H. and J.-J. Spoon  
2013. Bringing salience back in explaining voting defection in the european parliament. *Party Politics*, P. 1354068813487114.
- Knoke, D.  
1993. Networks of elite structure and decision making. *Sociological Methods & Research*, 22(1):23–45.
- Kollman, K., J. H. Miller, and S. E. Page  
2003. *Computational models in political economy*. MIT Press.
- Korten, D. C.  
2015. *When corporations rule the world*. Berrett-Koehler Publishers.
- La Due Lake, R. and R. Huckfeldt  
1998. Social capital, social networks, and political participation. *Political Psychology*, 19(3):567–584.
- Ladyman, J., J. Lambert, and K. Wiesner  
2013. What is a complex system? *European Journal for Philosophy of Science*, 3(1):33–67.
- Lagi, M., Y. Bar-Yam, K. Z. Bertrand, and Y. Bar-Yam  
2015. Accurate market price formation model with both supply-demand and trend-following for global food prices providing policy recommendations. *Proceedings of the National Academy of Sciences*, 112(45):E6119–E6128.
- Laver, M., K. Benoit, and J. Garry  
2003. Extracting policy positions from political texts using words as data. *American Political Science Review*, 97(02):311–331.
- Laver, M. and J. Garry  
2000. Estimating policy positions from political texts. *American Journal of Political Science*, Pp. 619–634.

- Laver, M. and E. Sergenti  
2011. *Party competition: An agent-based model*. Princeton University Press.
- Lazer, D.  
2011. Networks in political science: Back to the future. *PS: Political Science & Politics*, 44(01):61–68.
- Le, Q. V. and T. Mikolov  
2014. Distributed representations of sentences and documents. In *ICML*, volume 14, Pp. 1188–1196.
- Lebow, R. N.  
1994. The long peace, the end of the cold war, and the failure of realism. *International Organization*, 48(02):249–277.
- Lebow, R. N. and T. Risse-Kappen  
1995. *International relations theory and the end of the Cold War*. Cambridge Univ Press.
- Lehman Brothers, B. o.  
2008. *LEHMAN BROTHERS HOLDINGS INC. ANNOUNCES IT INTENDS TO FILE CHAPTER 11 BANKRUPTCY PETITION*.
- Leifeld, P. and V. Schneider  
2012. Information exchange in policy networks. *American Journal of Political Science*, 56(3):731–744.
- Lempert, R. J.  
2002. A new decision sciences for complex systems. *Proceedings of the National Academy of Sciences*, 99(suppl 3):7309–7313.
- Leontief, W. W.  
1951. The structure of american economy, 1919-1939: an empirical application of equilibrium analysis. Technical report.
- Lerman, K. and R. Ghosh  
2010. Information contagion: An empirical study of the spread of news on digg and twitter social networks. *ICWSM*, 10:90–97.
- Lesage, D. and T. Van de Graaf  
2016. *Global energy governance in a multipolar world*. Routledge.

- Li, W.  
1990. Mutual information functions versus correlation functions. *Journal of statistical physics*, 60(5-6):823–837.
- Linde, R.  
2013. The globalization of childhood: The international diffusion of norms and law against the child death penalty. *European Journal of International Relations*, P. 1354066113475464.
- Lindemann, B. A.  
2014. Case study 4: The united nations (un) and its specialized agencies. In *Cross-Strait Relations and International Organizations*, Pp. 253–286. Springer.
- Linsmeier, T. J. and N. D. Pearson  
2000. Value at risk. *Financial Analysts Journal*, 56(2):47–67.
- Lloyd, S.  
2001. Measures of complexity: a nonexhaustive list. *IEEE Control Systems Magazine*, 21(4):7–8.
- Lorenz, E.  
2000. 7. the butterfly effect. *The chaos avant-garde: Memories of the early days of chaos theory*, 39:91.
- Lucas, E.  
2014. *The New Cold War: Putin's Russia and the Threat to the West*. Macmillan.
- Lum, P., G. Singh, A. Lehman, T. Ishkanov, M. Vejdemo-Johansson, M. Alagappan, J. Carlsson, and G. Carlsson  
2013. Extracting insights from the shape of complex data using topology. *Scientific reports*, 3.
- MacShane, D.  
2015. *Brexit: How Britain Will Leave Europe*. IB Tauris.
- Mandelbrot, B. B.  
1997. The variation of certain speculative prices. In *Fractals and Scaling in Finance*, Pp. 371–418. Springer.
- Martin, L. W. and G. Vanberg  
2008. Coalition government and political communication. *Political Research Quarterly*, 61(3):502–516.

Masciantonio, S.

2015. Identifying and tracking global, eu, and eurozone systemically important banks with public data. *Applied Economics Quarterly*, 61(1):25–64.

May, T. E.

1997. *Erskine May's treatise on the law, privileges, proceedings and usage of Parliament*. Butterworth-Heinemann.

McKelvey, B.

2001. What is complexity science? it is really order-creation science. *Emergence, A Journal of Complexity Issues in Organizations and Management*, 3(1):137–157.

McNamara, R. S.

1986. *Blundering into disaster: Surviving the first century of the nuclear age*.

Michalowicz, I.

2007. What determines influence? assessing conditions for decision-making influence of interest groups in the eu 1. *Journal of European Public Policy*, 14(1):132–151.

Mikolov, T., K. Chen, G. Corrado, and J. Dean

2013a. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

Mikolov, T., I. Sutskever, K. Chen, G. S. Corrado, and J. Dean

2013b. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, Pp. 3111–3119.

Minoiu, C. and J. A. Reyes

2013. A network analysis of global banking: 1978–2010. *Journal of Financial Stability*, 9(2):168–184.

Montagna, M. and T. Lux

2016. Contagion risk in the interbank market: A probabilistic approach to cope with incomplete structural information. *Quantitative Finance*, Pp. 1–20.

Moran, K.

2001. Dynamic general-equilibrium models and why the bank of canada is interested in them. *Bank of Canada Review*, 2000(Winter):3–12.

- Morgan, G. and I. Kubo  
2005. Beyond path dependency? constructing new models for institutional change: the case of capital markets in japan. *Socio-Economic Review*, 3(1):55–82.
- Moshirian, F.  
2012. The future and dynamics of global systemically important banks. *Journal of Banking & Finance*, 36(10):2675–2679.
- Moussa, A.  
2011. *Contagion and systemic risk in financial networks*. Columbia University.
- Murphy, S., G. Moore, K. Tapper, R. Lynch, R. Clarke, L. Raisanen, C. Desousa, and L. Moore  
2011. Free healthy breakfasts in primary schools: a cluster randomised controlled trial of a policy intervention in wales, uk. *Public health nutrition*, 14(02):219–226.
- Musmeci, N., S. Battiston, G. Caldarelli, M. Puliga, and A. Gabrielli  
2013. Bootstrapping topological properties and systemic risk of complex networks using the fitness model. *Journal of Statistical Physics*, 151(3-4):720–734.
- Muzy, J.-F., E. Bacry, and A. Kozhemyak  
2006. Extreme values and fat tails of multifractal fluctuations. *Physical Review E*, 73(6):066114.
- Nicolis, G., I. Prigogine, et al.  
1977. *Self-organization in nonequilibrium systems*, volume 191977. Wiley, New York.
- Nikitinsky, N. S., D. A. Ustalov, and S. A. Shashev  
2015. Introducing government contracts to technology forecasting. *International Journal of Economics and Financial Issues*, 5(3S).
- Nobi, A., S. Lee, D. H. Kim, and J. W. Lee  
2014. Correlation and network topologies in global and local stock indices. *Physics Letters A*, 378(34):2482–2489.
- Norberg, J. and G. S. Cumming  
2013. *Complexity theory for a sustainable future*. Columbia University Press.

Norton, P.

2015. The times guide to the house of commons 2015: The definitive record of britain's historic 2015 general election. *The Journal of Legislative Studies*, 21(4):588–589.

O'Ballance, E.

1997. *Islamic Fundamentalist Terrorism, 1979-95: The Iranian Connection*. Macmillan.

Oneal, J. R. and B. M. Russett

2015. The kantian peace: The pacific benefits of democracy, interdependence, and international organizations, 1885–1992. In *Bruce M. Russett: Pioneer in the Scientific and Normative Study of War, Peace, and Policy*, Pp. 74–108. Springer.

Osinski, J., K. Seal, and M. L. Hoogduin

2013. *Macprudential and microprudential policies: toward cohabitation*, number 13-15. International Monetary Fund.

Ottosson, J. and L. Magnusson

1997. Evolutionary economics and path dependence.

Outram, D.

2013. *The enlightenment*. Cambridge University Press.

Owen, J. M.

1994. How liberalism produces democratic peace. *International security*, 19(2):87–125.

Page, L., S. Brin, R. Motwani, and T. Winograd

1999. The pagerank citation ranking: bringing order to the web.

Palla, G., A.-L. Barabási, and T. Vicsek

2007. Quantifying social group evolution. *Nature*, 446(7136):664–667.

Palmer, R. G., W. B. Arthur, J. H. Holland, B. LeBaron, and P. Tayler

1994. Artificial economic life: a simple model of a stockmarket. *Physica D: Nonlinear Phenomena*, 75(1):264–274.

Pimentel, D. and T. W. Patzek

2005. Ethanol production using corn, switchgrass, and wood; biodiesel production using soybean and sunflower. *Natural resources research*, 14(1):65–76.



Potrafke, N.

2009. Does government ideology influence political alignment with the us? an empirical analysis of voting in the un general assembly. *The Review of International Organizations*, 4(3):245–268.

Preis, T., D. Y. Kenett, H. E. Stanley, D. Helbing, and E. Ben-Jacob

2012. Quantifying the behavior of stock correlations under market stress. *Scientific reports*, 2.

Proksch, S.-O. and J. B. Slapin

2012. Institutional foundations of legislative speech. *American Journal of Political Science*, 56(3):520–537.

Quack, M.

2012. Molecular parity violation and chirality: the asymmetry of life and the symmetry violations in physics. In *Quantum Systems in Chemistry and Physics*, Pp. 47–76. Springer.

Quillen, C.

2002. A historical analysis of mass casualty bombers. *Studies in Conflict and Terrorism*, 25(5):279–292.

Rasche, A., S. Waddock, and M. McIntosh

2013. The united nations global compact retrospect and prospect. *Business & Society*, 52(1):6–30.

Reh, C., A. Héritier, E. Bressanelli, and C. Koop

2013. The informal politics of legislation explaining secluded decision making in the european union. *Comparative Political Studies*, 46(9):1112–1142.

Reinhart, C. M. and K. S. Rogoff

2009. The aftermath of financial crises. Technical report, National Bureau of Economic Research.

Reinhart, C. M. and K. S. Rogoff

2013. Banking crises: an equal opportunity menace. *Journal of Banking & Finance*, 37(11):4557–4573.

Rhee, Y. P.

2000. Complex systems approach to the study of politics. *Systems research and behavioral science*, 17(6):487.

Rimer, M. and T. Martinez

2004. Softprop: softmax neural network backpropagation learning. In *Neural Networks, 2004. Proceedings. 2004 IEEE International Joint Conference on*, volume 2, Pp. 979–983. IEEE.

Rind, D.

1999. Complexity and climate. *science*, 284(5411):105–107.

Ring, J.

. An agent-based model of international norm diffusion. *PhD Dissertation, University of Iowa*.

Roberts, A. and B. Kingsbury

2008. The un's roles in international relations. *Nankai Journal (Philosophy, Literature and Social Science Edition)*, 5:003.

Rochet, J.-C. and J. Tirole

1996. Interbank lending and systemic risk. *Journal of Money, credit and Banking*, 28(4):733–762.

Rockafellar, R. T. and S. Uryasev

2002. Conditional value-at-risk for general loss distributions. *Journal of banking & finance*, 26(7):1443–1471.

Rose, A. K. and M. M. Spiegel

2012. Cross-country causes and consequences of the 2008 crisis: early warning. *Japan and the World Economy*, 24(1):1–16.

Roukny, T., H. Bersini, H. Pirotte, G. Caldarelli, and S. Battiston

2013. Default cascades in complex networks: Topology and systemic risk. *Scientific reports*, 3.

Ruddell, B. L., N. A. Brunsell, and P. Stoy

2013. Applying information theory in the geosciences to quantify process uncertainty, feedback, scale. *Eos, Transactions American Geophysical Union*, 94(5):56–56.

Sabatier, P. A. and C. Weible

2014. *Theories of the policy process*. Westview Press.

Sachs, J. D. and J. W. McArthur

2005. The millennium project: a plan for meeting the millennium development goals. *The Lancet*, 365(9456):347–353.

- Sakwa, R.  
1990. *Gorbachev and his Reforms, 1985-1990*. Prentice Hall.
- Saracco, F., R. Di Clemente, A. Gabrielli, and T. Squartini  
2016. Detecting early signs of the 2007–2008 crisis in the world trade. *Scientific Reports*, 6.
- Sawyer, R. K.  
2005. *Social emergence: Societies as complex systems*. Cambridge University Press.
- Sbordone, A. M., A. Tambalotti, K. Rao, and K. J. Walsh  
2010. Policy analysis using dsge models: an introduction. *Economic policy review*, 16(2).
- Schakel, A. M. and B. J. Wilson  
2015. Measuring word significance using distributed representations of words. *arXiv preprint arXiv:1508.02297*.
- Scheffer, M.  
2009. *Critical transitions in nature and society*. Princeton University Press.
- Schläpfer, M., L. M. Bettencourt, S. Grauwin, M. Raschke, R. Claxton, Z. Smoreda, G. B. West, and C. Ratti  
2014. The scaling of human interactions with city size. *Journal of the Royal Society Interface*, 11(98):20130789.
- Schmitt, H. and A. M. Wüst  
2012. Euromanifestos project (emp) 1979–2004. *Cologne: GESIS Data Archive*.
- Schreiber, T.  
2000. Measuring information transfer. *Physical review letters*, 85(2):461.
- Schwarcz, S. L.  
2008. Systemic risk. *Geo. LJ*, 97:193.
- Schwarcz, S. L.  
2012. Regulating shadow banking. *Review of Banking and Financial Law*, 31(1).
- Schwarz, D., D. Traber, and K. Benoit  
2014. Estimating intra-party preferences: comparing speeches to votes. *Political Science Research and Methods*, Pp. 1–18.

Seeley, J. R.

1949. The net of reciprocal influence. a problem in treating sociometric data. *Canadian Journal of Experimental Psychology*, 3:234.

Sengupta, A.

2006. Chaos, nonlinearity, complexity: A unified perspective. In *Chaos, Nonlinearity, Complexity*, Pp. 270–352. Springer.

Serrano, M. Á., M. Boguñá, and A. Vespignani

2007. Patterns of dominant flows in the world trade web. *Journal of Economic Interaction and Coordination*, 2(2):111–124.

Servedio, V. D., G. Caldarelli, and P. Butta

2004. Vertex intrinsic fitness: How to produce arbitrary scale-free networks. *Physical Review E*, 70(5):056126.

Shachar, R. and B. Nalebuff

1999. Follow the leader: Theory and evidence on political participation. *American Economic Review*, Pp. 525–547.

Shahbazi, M., J. R. Barr, V. Hristidis, and N. N. Srinivasan

2016. Estimation of the investability of real estate properties through text analysis. In *2016 IEEE Tenth International Conference on Semantic Computing (ICSC)*, Pp. 301–306. IEEE.

Shannon, C. E.

1949. Communication in the presence of noise. *Proceedings of the IRE*, 37(1):10–21.

Shively, W. P.

2011. *Power & choice: An introduction to political science*. McGraw-Hill.

Silberman, J. and G. Yochum

1980. The market for special interest campaign funds: An exploratory approach. *Public Choice*, 35(1):75–83.

Simes, D. K.

1987. Gorbachev: A new foreign policy? *Foreign Affairs*, 65(3):484.

Simmons, B. A., F. Dobbin, and G. Garrett

2006. Introduction: The international diffusion of liberalism. *International Organization*, 60(04):781–810.

- Slapin, J. B. and S.-O. Proksch  
2008. A scaling model for estimating time-series party positions from texts. *American Journal of Political Science*, 52(3):705–722.
- Smith, C. B.  
2006. Politics and process at the united nations. *The Global Dance*. Boulder/London: Lynne Rienner.
- Snyder, D. and E. L. Kick  
1979. Structural position in the world system and economic growth, 1955-1970: A multiple-network analysis of transnational interactions. *American journal of Sociology*, Pp. 1096–1126.
- Snyder Jr, J. M. and T. Groseclose  
2000. Estimating party influence in congressional roll-call voting. *American Journal of Political Science*, Pp. 193–211.
- Solé, R. V. and J. Bascompte  
2006. *Self-Organization in Complex Ecosystems.(MPB-42)*, volume 42. Princeton University Press.
- Sornette, D.  
2009. *Why stock markets crash: critical events in complex financial systems*. Princeton University Press.
- Sornette, D., R. Woodard, and W.-X. Zhou  
2009. The 2006–2008 oil bubble: Evidence of speculation, and prediction. *Physica A: Statistical Mechanics and its Applications*, 388(8):1571–1576.
- Speel, H.-C.  
1997. A memetic analysis of policy making. *Journal of Memetics-Evolutionary Models of Information Transmission*, 1(2).
- Spizzirri, L.  
2011. Justification and application of eigenvector centrality. *Algebra in Geography: Eigenvectors of Network*.
- Stanley, H. E., L. A. N. Amaral, D. Canning, P. Gopikrishnan, Y. Lee, and Y. Liu  
1999. Econophysics: Can physicists contribute to the science of economics? *Physica A: Statistical Mechanics and its Applications*, 269(1):156–169.
- Starr, R. M.  
2011. *General equilibrium theory: An introduction*. Cambridge University Press.

Stern, G. H. and R. J. Feldman

2004. *Too big to fail: The hazards of bank bailouts*. Brookings Institution Press.

Stiglitz, J. E.

1975. Incentives, risk, and information: notes towards a theory of hierarchy. *The Bell Journal of Economics*, Pp. 552–579.

Sydow, J., G. Schreyögg, and J. Koch

2005. Organizational paths: Path dependency and beyond.

The Guardian

2002. Party leaders call for iraq debate. <http://www.theguardian.com/politics/2002/sep/05/foreignpolicy.uk>. (Accessed on 08/22/2016).

The Guardian

2015. Charles kennedy's 2003 speech on his party's anti-iraq war stance. <http://www.theguardian.com/politics/2015/jun/02/charles-kennedys-2003-speech-on-his-partys-anti-iraq-war-stance>. (Accessed on 08/22/2016).

The Guardian

2016a. Eu referendum full results find out how your area voted. <http://www.theguardian.com/politics/ng-interactive/2016/jun/23/eu-referendum-live-results-and-analysis>. (Accessed on 07/27/2016).

The Guardian

2016b. Nigel Farage resigns as ukip leader after 'achieving political ambition' of brexit. <http://www.theguardian.com/politics/2016/jul/04/nigel-farage-resigns-as-ukip-leader>. (Accessed on 07/27/2016).

Theraulaz, G., E. Bonabeau, and J.-L. Deneubourg

1995. Self-organization of hierarchies in animal societies: the case of the primitively eusocial wasppolistes dominuluschrist. *Journal of theoretical Biology*, 174(3):313–323.

Tichy, N. M., M. L. Tushman, and C. Fombrun

1979. Social network analysis for organizations. *Academy of management review*, 4(4):507–519.

- Tovar, C. E.  
2009. Dsge models and central banks. *Economics: The Open-Access, Open-Assessment E-Journal*, 3.
- Tsur, O., D. Calacci, and D. Lazer  
2015. A frame of mind: Using statistical models for detection of framing and agenda setting campaigns. In *Proc. of ACL*.
- Upper, C.  
2007. Using counterfactual simulations to assess the danger of contagion in interbank markets.
- Upper, C. and A. Worms  
2004. Estimating bilateral exposures in the german interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827–849.
- Valenzuela, S., Y. Kim, and H. G. de Zúñiga  
2012. Social networks that matter: Exploring the role of political discussion for online political participation. *International journal of public opinion research*, 24(2):163–184.
- Valverde, S. and R. V. Solé  
2007. Self-organization versus hierarchy in open-source social networks. *Physical Review E*, 76(4):046118.
- Van Deth, J. W.  
2001. Studying political participation: towards a theory of everything. In *joint sessions of workshops of the European consortium for political research, Grenoble*, Pp. 6–11.
- Van Horne, J. C., M. E. Blume, and I. Friend  
1975. The asset structure of individual portfolios and some implications for utility functions. *The Journal of Finance*, 30(2):585–603.
- van Lelyveld, I. et al.  
2014. Finding the core: Network structure in interbank markets. *Journal of Banking & Finance*, 49:27–40.
- Vannoni, M. and P. John  
2015. Proto stars, dwarfs and giants: A career path analysis of british mps. *Dwarfs and Giants: A Career Path Analysis of British MPs (October 31, 2015)*.

- Vitali, S., J. B. Glattfelder, and S. Battiston  
2011. The network of global corporate control. *PloS one*, 6(10):e25995.
- Vizgunov, A., B. Goldengorin, V. Kalyagin, A. Koldanov, P. Koldanov, and P. Pardalos  
2014. Network approach for the russian stock market. *Computational Management Science*, 11(1-2):45–55.
- Voeten, E.  
2012. Data and analyses of voting in the un general assembly. *Available at SSRN 2111149*.
- Vovk, V., H. Papadopoulos, and A. Gammerman  
2015. *Measures of Complexity*. Springer.
- Walker, E. W.  
2003. *Dissolution: Sovereignty and the breakup of the Soviet Union*. Rowman & Littlefield.
- Walker, J. C., P. Hays, and J. F. Kasting  
1981. A negative feedback mechanism for the long-term stabilization of earth's surface temperature. *Journal of Geophysical Research: Oceans*, 86(C10):9776–9782.
- Wallach, H. M.  
2006. Topic modeling: beyond bag-of-words. In *Proceedings of the 23rd international conference on Machine learning*, Pp. 977–984. ACM.
- Wang, X., A. Mueen, H. Ding, G. Trajcevski, P. Scheuermann, and E. Keogh  
2013. Experimental comparison of representation methods and distance measures for time series data. *Data Mining and Knowledge Discovery*, 26(2):275–309.
- Wasserman, S. and K. Faust  
1994. *Social network analysis: Methods and applications*, volume 8. Cambridge university press.
- Weistroffer, C., B. Speyer, S. Kaiser, and T. Mayer  
2011. Identifying systemically important financial institutions (sifis). *Deutsche Bank Research, Current Issues*.
- West, G.  
2013. Wisdom in numbers. *Scientific American*, 308(5):14–14.



- Whitesides, G. M. and R. F. Ismagilov  
1999. Complexity in chemistry. *science*, 284(5411):89–92.
- Wibral, M., R. Vicente, and M. Lindner  
2014. Transfer entropy in neuroscience. In *Directed Information Measures in Neuroscience*, Pp. 3–36. Springer.
- Wickens, M.  
2011. *Macroeconomic theory: a dynamic general equilibrium approach*. Princeton University Press.
- Wiebrecht, F.  
2013. *Democratic Peace Theory*. GRIN Verlag.
- World Bulletin, N.  
2013. 180,000 killed, 6 million displaced in arab spring.
- Wray, C. M. and S. R. Bishop  
2016. A financial market model incorporating herd behaviour. *PloS one*, 11(3):e0151790.
- Wright, J. H.  
1998. Testing for a structural break at unknown date with long-memory disturbances. *Journal of Time Series Analysis*, 19(3):369–376.
- Yates, F. E.  
2012. *Self-organizing systems: The emergence of order*. Springer Science & Business Media.
- Zeigler, B., H. Praehofer, and T. Kim  
2013. Theory of modeling and simulation: integrating discrete event and continuous complex dynamic systems. 2000.
- Zhou, C.  
2009. Are banks too big to fail? measuring systemic importance of financial institutions. *Measuring Systemic Importance of Financial Institutions (December 1, 2009)*.