

T. R. Hurd
Department of Mathematics and Statistics
McMaster University
Hamilton, ON, Canada
L8S 4K1

Contagion! The Spread of Systemic Risk in Financial Networks

June 15, 2015

Springer

Preface

This slim volume logs the development of a cascade of contagious ideas that has occupied my space, time and mind in recent years. There was a clear triggering event that occurred in April 2009. Late in that month, Michael Lynch and his colleagues at MITACS Canada brought together a host of scientists, mathematicians and finance industry participants for three days to brainstorm about underlying causes of the ongoing financial crisis and how mathematical thinking could be brought to bear on it. My role there was as gadfly to provoke discussion on a special topic no one at the meeting was very aware of, namely financial systemic risk.

Since that event introduced me to the subject, I have had many opportunities to present to a diversity of audiences an evolving view of how the architecture of the financial system can be described in terms of network science, and how such a network formulation can be made amenable to a certain type of mathematical analysis. This book is not intended to be a definitive work on the subject of financial systemic risk, and does not try to represent a broad consensus. Instead, it is a personal attempt to crystallize the early results of research that focuses on the basic modelling structure that ensures some kind of mathematical tractability, while allowing a great deal of both reality and complexity in the actual finance network specification. I owe a debt of thanks to a great number of people who have listened, commented, and added new nodes to this complex network of ideas, too many to list here in this preface.

My McMaster colleague, Matheus Grasselli, was instrumental in many ways, not least in providing the original impetus to write this SpringerBrief. Nizar Touzi encouraged and supported me in my first attempt at delivering a minicourse on Systemic Risk. The scope of this minicourse grew over time: Jorge Zubelli hosted me for an extended period at IMPA, where I delivered another version; Peter Spreij arranged a session for me to speak at the Winter School on Financial Mathematics in Luntenen; James Gleeson provided me with multiple invitations to Limerick. The Fields Institute for Research in Mathematical Sciences gave me encouragement and organized multiple events relevant to my work. The Global Risk Institute for Financial Services, in particular Michel Maila and Catherine Lubochinsky, have provided substantial financial and moral support for this research. I give my hearty thanks to

Mario Wüthrich and Paul Embrechts who hosted my extended stay at ETH Zürich in 2014 where I was extremely fortunate to be able to deliver a Nachdiplom lecture series based on the material contained in this book. Finally, to my wife, Rita Bertoldi, I offer my affectionate acknowledgment of her patient support throughout my lengthy exposure to this dangerous contagion.

Contents

1	Systemic Risk Basics	9
1.1	The Nature of this Book	10
1.2	What is Systemic Risk?	11
1.2.1	Defining SR	12
1.2.2	Haldane's 2009 Speech	13
1.2.3	A Lesson from Network Science: The Sandpile Model	15
1.3	Channels of Systemic Risk	16
1.4	Capital Structure of a Bank	18
1.4.1	Bank Asset Classes	19
1.4.2	Debt and Liabilities	20
1.4.3	Equity	22
1.5	Regulatory Capital and Constraints	22
2	Static Cascade Models	25
2.1	Bank Balance Sheets	27
2.2	The Eisenberg-Noe 2001 Model	29
2.2.1	Reduced Form E-N Cascade Mechanism	34
2.2.2	Clearing Algorithm	35
2.3	The Gai-Kapadia 2010 Default Model	36
2.4	Liquidity Cascades	37
2.4.1	Gai-Kapadia 2010 Liquidity Cascade Model	38
2.4.2	The Liquidity Model of S. H. Lee 2013	39
2.4.3	Generalized Liquidity Cascades	40
2.5	Asset Fire Sales	41
2.5.1	Fire Sales of One Asset	42
2.5.2	Fire Sales of Many Assets	45
2.6	Random Financial Networks	48
3	Random Graph Models	51
3.1	Definitions and Basic Results	52
3.1.1	Random Graphs	54

3.2	Configuration Random Graphs	55
3.3	Preferential Attachment and Detachment	57
3.3.1	Preferential Attachment Models	57
3.4	Inhomogeneous Random Graphs	62
3.5	Assortative Configuration Graphs	68
3.5.1	Finite Assortative Configuration Graphs	69
3.5.2	Asymptotic Configuration Probabilities	73
3.5.3	Approximate ACG Simulation	77
3.6	Measures of Network Topology	79
3.6.1	Connectivity	79
3.6.2	Measures from the Degree Distributions	80
3.6.3	Centrality Measures	80
3.6.4	Clustering Coefficients	81
3.6.5	Connectivity and Connected Components	82
4	Percolation and Cascades	85
4.1	Branching Processes	86
4.2	Percolation on Configuration Graphs	87
4.3	Site Percolation	90
4.4	Bootstrap Percolation	91
4.5	Watt's 2002 Model of Global Cascades	92
4.5.1	The Framework	93
4.5.2	The Main Result	95
4.5.3	The Cascade Condition	97
4.5.4	The Frequency of Global Cascades	98
4.6	Numerical Experiments on the Watts Model	99
4.7	Extensions of the Watts Model	102
4.8	Dependence in Random Financial Networks	103
4.8.1	The LTI Property	105
4.8.2	Ramifications of the LTI Property	106
5	Zero Recovery Default Cascades	107
5.1	The Gai-Kapadia Model	108
5.1.1	Shocks and the Solvency Condition	109
5.2	The GK Model with Random Link Weights	112
5.3	Proof of Theorem 12	115
5.4	Measures of Cascade Impact	118
5.5	The Cascade Condition	120
5.5.1	Frequency and Size of Global Cascades	123
5.6	Testing and Using the Cascade Mapping	125
5.7	Default Cascades with Asymmetric Shocks	126
5.8	Cascade Computations	130
5.9	Numerical Experiments	131
5.9.1	Experiment 1: Benchmark Gai-Kapadia Model	132
5.9.2	Experiment 2: Assortative Networks	133

Contents	7
5.9.3 Experiment 3: Random Buffers and Link Weights	134
6 Future Directions for Cascade Models	137
A Background Material	141
A.1 Special Notation	141
A.2 The Discrete Fourier Transform	142
References	143
Index	149

Chapter 1

Systemic Risk Basics

*Annual income twenty pounds, annual expenditure nineteen nineteen six, result happiness. Annual income twenty pounds, annual expenditure twenty pounds ought and six, result misery. The blossom is blighted, the leaf is withered, the god of day goes down upon the dreary scene, and—and, in short, you are for ever floored.*¹

Abstract Attempts to define systemic risk are summarized and found to be deficient in various respects. In this introductory chapter, after considering some of the salient features of financial crises in the past, we focus on the key characteristics of banks, their balance sheets and how they are regulated.

Bankruptcy! Mr. Micawber, David Copperfield's debt-ridden sometime mentor, knew first hand the difference between surplus and deficit, between happiness and the debtors' prison. In Dickens' fictional universe, and perhaps even in the real world of Victorian England, a small businessman's unpaid debts were never overlooked but were always a trigger leading him and his family to the unmitigated misery of the poorhouse. On the other hand, the bigger players, the aristocrats and upper middle classes, were treated more delicately, and were usually able to find a comfortable escape.

For people, firms, and in particular banks, bankruptcy in modern times is more complicated yet still retains the flavour of the olden days. When a bank fails, the rich financiers responsible for its collapse and the collateral damage it inflicts, often walk away from the wreckage with their bonuses and compensation packages intact. Occasionally, a particularly egregious case arises where a scapegoat is needed: then a middle rank banker is identified who takes the bullet for the disaster. A cynic might say that despite the dictates of Basel I, II, III ...∞, bank executives are still free to take excessive risks with their company, receiving a rich fraction of any upside while insulating themselves from any possible disaster they might cause.

As we learn afresh during every large scale financial crisis, society at large pays the ultimate costs when banks fail. Spiking unemployment leads to the poverty of

¹ CHARLES DICKENS, *David Copperfield*, Chapter 12, p. 185 (1950). First published 1849–1850.

the less well-to-do, while salary freezes and imploded pension plans lead to belt-tightening and delayed retirement for the better-off. Those at the top of the pile, even those responsible, seem to do just fine. Banks are bailed out by central banks and governments, their debts taken over and paid by the taxpayers.

If anything is different since the crisis of 2007-08, perhaps it is the wide recognition of the fact that society needs to find ways and rules to ensure that the responsible parties pay the downside costs of bank failure. New ideas on bank resolution, including contingent capital and bail-in regulation, aim to force the financial stakeholders, not the central bank, to pay much higher fractions of the costs of failure: banks' creditors, bondholders and equity investors should in the future be forced to take their fair share of losses. Bank failures might then be better anticipated, prepared for and managed to reduce more catastrophic threats of serial collapse of the financial system.

1.1 The Nature of this Book

The title of this book, "Contagion! The Spread of Systemic Risk in Financial Networks", suggests that financial contagion is analogous to the spread of disease, and that damaging financial crises may be better understood by bringing to bear ideas that have been developed to understand the breakdown of other complex systems in our world. It also suggests that the aim of systemic risk management is similar to a primary aim of epidemiology, namely to identify situations when contagion danger is high, and then make targeted interventions to damp out the risk.²

This book is intended to be two things, a timely summary of a growing body of systemic risk research as well as a unified mathematical framework for the primary channels that can transmit damaging shocks through financial systems. Much of its contents are new, not having appeared previously in published journals. It aims to serve as a coherent guide of equal interest to quantitative finance practitioners, financial regulators and a broad range of academics doing network research, including economists, physicists, applied mathematicians and computer scientists.

The introductory Chapter One develops the setting for the problem of systemic risk in financial networks. It provides a brief survey of how people view financial crises and systemic risk, a look at the type of business banks and other financial institutions deal with, and some of what we know about the state of actual financial networks, both in terms of stylized features and datasets that have been studied. From Chapter Two onwards, we delve more deeply into the mechanics of the interactions between banking counterparties. Chapter Two puts a sharp focus on the type of bank behaviour that can negatively impact the functioning of the entire system, by surveying and classifying the range of economic cascade models that have

² Interestingly, I found on Wikipedia that epidemiology has a code of nine principles, called the "Bradford Hill criteria", that should be considered to help assess evidence of a causal relationship between an incidence and a consequence. Perhaps, researchers can codify an analogous set of principles for assessing systemic risk.

been proposed in recent years. We will make the important discovery that similar mathematical models can describe a multiplicity of financial cascade channels. Given the intrinsic opacity of financial institutions and their interconnections, we eventually argue for the usefulness of what we call *random financial networks* that provide a statistical representation of networks of banks and their balance sheets. The design of this concept, fleshed out in subsequent chapters, reflects the type of models that network science has already developed in other domains, allowing us to bring their understanding to bear on our new problem. Chapter Three therefore provides the mathematical underpinning we need for systemic risk, by surveying, developing and adapting the theory of random graphs to describe the *skeleton* structure at the heart of the financial network. Four distinct classes of random graphs, each characterized by its specific stochastic construction algorithm, are described in detail. Chapter Four is devoted to understanding the relation between the Watts 2002 model of information cascades and the concept of bootstrap percolation in random networks. The Watts model provides us with a template for more specific financial networks, and therefore we analyze it in detail from first principles. We shall learn that its properties can be best understood in terms of the mathematics of *percolation* which addresses the size distribution of connected network components. Chapter Five returns to the simplest possible financial network model, the zero recovery cascade mechanism due to Gai and Kapadia, and develops a purely analytical approach based on an assumption called *locally treelike independence*. This theory provides us with a computational methodology that is independent of and complementary to the usual Monte Carlo simulation techniques used everywhere in network science.

Do there exist classes of mathematical systemic risk models that provide a degree of realism, but at the same time are sufficiently tractable that all critical parameters can be varied at will and resulting network characteristics computed? Can these model systems be tested for their systemic susceptibility? Are the mathematical conclusions robust and relevant to the real world of financial crisis regulation? Ultimately we hope this book will persuade the reader that the answer to these questions is an emphatic “YES”!

1.2 What is Systemic Risk?

First it is helpful to identify what systemic risk is not. For example, Duffie and Singleton [31] identify five categories of risk faced by financial institutions: (i) market risk: the risk of unexpected changes in market prices; (ii) credit risk: the risk of changes in value due to unexpected changes in credit quality, in particular if a counterparty defaults on one of their contractual obligations; (iii) liquidity risk: the risk that costs of adjusting financial positions may increase substantially; (iv) operational risk: the risk that fraud, errors or other operational failures lead to loss in value; (v) systemic risk: the risk of market wide illiquidity or chain reaction defaults. The first four risk categories are for the most part focussed on individual institutions, while the topic of this book must be the market wide risks of the fifth category.

Kaufman and Scott [50], John B. Taylor [69] and others all seem to agree that the concept of systemic risk must comprise at least three ingredients. First, a triggering event. Second, the propagation of shocks through the financial system. And third, significant impact of the crisis on the macroeconomy. Possible triggers might come from outside the financial system, for example a terrorist attack that physically harms the system. Or triggers might come internally, for example as the surprise spontaneous failure of a major institution within the system. Propagation of shocks may be through direct linkages between banks or indirectly, such as through the impact on the asset holdings of many banks caused by the forced sales of a few banks or through a crisis of confidence. The impact of systemic crises on the macroeconomy may take many forms: on the money supply, on the supply of credit, on major market indices, on interest rates, and ultimately on the production economy and the level of employment.

As Admati and Hellwig [2] have argued, ambiguity in the definition of systemic risk implies that mitigation of systemic risk might mean different things to different people. One approach might seek to reduce impact on the financial system, whereas a different approach might instead try to mitigate the damage to the economy at large. These aims do not necessarily coincide: the demise of Lehman Bros. illustrates that key components of the financial system might be sacrificed to save the larger economy during a severe crisis.

1.2.1 Defining SR

The economics literature has used the term *systemic risk* in the context of financial systems for many years. Nonetheless, Kaufman and Scott, Taylor and many others argue that there is as yet no generally accepted definition of the concept, and furthermore, that without an agreed definition, it may be pointless and indeed dangerous to implement public policy that explicitly aims to reduce systemic risk. To see that there is as yet no consensus definition over the years, consider the following examples of definitions proposed in the past.

1. Mishkin 1995 [57]: “the likelihood of a sudden, usually unexpected, event that disrupts information in financial markets, making them unable to effectively channel funds to those parties with the most productive investment opportunities.”
2. Kaufman 1995 [49] “The probability that cumulative losses will accrue from an event that sets in motion a series of successive losses along a chain of institutions or markets comprising a system. . . . That is, systemic risk is the risk of a chain reaction of falling interconnected dominos.”
3. Bank for International Settlements 1994 [34] “ the risk that the failure of a participant to meet its contractual obligations may in turn cause other participants to default with a chain reaction leading to broader financial difficulties.”
4. Board of Governors of the Federal Reserve System 2001 [63] “In the payments system, systemic risk may occur if an institution participating on a private large- dollar payments network were unable or unwilling to settle its net debt position. If such a settlement

failure occurred, the institution's creditors on the network might also be unable to settle their commitments. Serious repercussions could, as a result, spread to other participants in the private network, to other depository institutions not participating in the network, and to the nonfinancial economy generally."

In the light of the 2007-08 financial crisis, the above style of definitions, deficient as they are in several respects, can be seen to miss or be vague about one key attribute of any systemic crisis, namely that it also causes damage outside the network, through its failure to efficiently perform its key function of providing liquidity, credit and services. S. L. Schwarcz' definition [65] of systemic risk explicitly includes this important aspect:

The risk that (i) an economic shock such as market or institutional failure triggers (through a panic or otherwise) either (X) the failure of a chain of markets or institutions or (Y) a chain of significant losses to financial institutions, (ii) resulting in increases in the cost of capital or decreases in its availability, often evidenced by substantial financial-market price volatility.

While the Schwarcz definition is hardly elegant in its phrasing, we will accept it provisionally as the closest thing we have to a concise definition of the spirit of systemic risk.

However, if this definition captures much of the spirit of systemic risk, it fails to address how to measure or quantify the level of systemic risk, and how it might be distributed over the network. Much of current research on systemic risk is dedicated to defining measures of systemic risk and identifying where it is concentrated. Some of the important concepts are *counterparty value at risk (CoVaR)* introduced by Brunnermeier and Pedersen [20]; and *systemic expected shortfall* introduced by Acharya, Pedersen, Philippon, and Richardson [1]. For a recent and comprehensive review of these and many other systemic risk measures, please see [11].

1.2.2 Haldane's 2009 Speech

In 2009, in the aftermath of the crisis, Andrew G. Haldane, Executive Director of Financial Stability at the Bank of England, gave a provocative and visionary talk, entitled "Rethinking the Financial Network" [41]. In this brilliant summary of the nature of networks, he compares the 2002 SARS epidemic to the 2008 collapse of Lehman Bros, with the aim to inspire efforts to better understand the nature of systemic risk. For a very broad free thinking overview, we can't do better than summarize the high points of his speech.

In these two examples of contagion events he identifies the following pattern:

- an external event strikes;
- panic ensues and the complex system seizes up;
- *collateral damage* is wide and deep;
- in hindsight, the trigger event was modest;
- during the event itself, dynamics was chaotic.

He claims this type of pattern is a manifestation of any complex adaptive system, and should be the target where we need to direct our attention.

So, in more detail, what went wrong with the financial network in 2008? Haldane identifies two contributing trends: increasing complexity and decreasing diversity. In real world networks these two trends are observed to lead to fragility, and ring alarm bells for ecologists, engineers, geologists. Figure 1.1 illustrates how the global financial network has grown in complexity. Highly connected, heterogeneous networks may be *robust yet fragile*, by which he means that they may be resistant to average or typical shocks, yet highly susceptible to an attack that targets a highly connected or dominant node. In such networks, connections that we think of as shock absorbers may turn out to act as shock amplifiers during a crisis. There may be a sharp *tipping point* that separates normal behaviour from a crisis regime. Thus, a network with a fat-tailed *degree distribution* (i.e. where there is a significant number of highly connected nodes) may be robust to random shocks while vulnerable to shocks that preferentially target these highly connected nodes.




Fig. 1.1 The global financial network in 1985 (left) and 2005 (right). Here line thickness denotes link strength as fraction of total GDP. (figure taken from Haldane [41].)

In both of Haldane’s examples of contagion events, agents exhibit a variety of behavioural responses that create feedback and influence the stability of the network. In epidemics, two classic responses, “hide” or “flee”, may prevail and the virulence of the event is highly dependent on which behaviour dominates. In a financial crisis, two likely responses of banks are to hoard liquidity or to sell assets. Both responses are rational, but both make the systemic problem worse. Massive government intervention to provide liquidity and restore capital to banks in a timely manner may be needed in order to curtail systemic events.

Financial networks generate chains of claims and at times of stress, these chains can amplify uncertainties about true counterparty exposures. In good times, coun-

terparty risk is known to be small, and thus “Knightian” uncertainty³ is small, and in such times we might expect that stability will improve with connectivity. In bad times, counterparty risk can be large and highly uncertain, due to the complicated web and the nature of the links: we then expect stability to decline with connectivity. Financial innovation, particularly *securitization*, created additional instability. As CDOs, MBSs, RMBSs and similar high dimensional products proliferated internationally, they dramatically expanded the size and scope of the precrisis bubble (see [66]). The structure of these contracts was opaque not transparent. They dramatically increased the connectedness and complexity of the network, and moreover adverse selection made them hard to evaluate. As Haldane wrote:

Haldane 2009 [41]: “With no time to read the small-print, the instruments were instead devoured whole. Food poisoning and a lengthy loss of appetite have been the predictable consequences.”

In ecosystems, many instances have been observed that show that biodiversity tends to improve stability. On the other hand, Haldane argues that during the *Great Moderation* prior to 2007, financial diversity has been reduced. Pursuit of returns led many major players, including global banks, insurance companies and hedge funds, to follow similar strategies leading to averaged portfolio correlations in excess of 90% during 2004-2007. Moreover, risk management regulation following Basel II led to similar risk management strategies for banks. As a result of such trends, bank balance sheets became increasingly homogeneous. Finance became almost a monoculture, and thus vulnerable to *viral infection*.

What one learns from Haldane’s analysis is that networks arising in ecology, engineering, the internet, and in finance, are complex and adaptive. Such networks are in a sense *robust yet fragile*. He asks “what properties of the financial network most influence stability?” and expresses the hope that the key determinants for financial stability can be deduced from studies of other types of networks.

1.2.3 A Lesson from Network Science: The Sandpile Model

Is there more specific guidance to understanding systemic risk that comes from other branches of the science of complex adapted systems? Consider the following thought experiment, first proposed by Bak, Tang and Wiesenfeld [7]. A very slow trickle of sand is allowed to fall in the middle of a large circular table. How do we expect the system to evolve? The growing accumulation of sand forms a pile on the table and our common experience tells us that the steepness of the pile cannot exceed a certain critical slope that depends on the microscopic and statistical properties of the sand. As more sand is added, the sandpile, still near its critical slope,

³ In *Knightian* terms, *uncertainty* describes modelling situations where probabilities cannot plausibly be assigned to outcomes. On the other hand, *risk* describes situations where uncertainty can be adequately captured in a probability distribution.

eventually expands to cover the entire surface of the table. Having reached this maximal extent, the properties of the system take on a new character. On average, as sand is added near the centre of the table, an equal amount of sand must fall off the edge. The interesting thing however is the probabilistic nature of the likelihood of n grains falling off, for each single grain added: BTW's assertion, vindicated since by experiments, is roughly that the frequency for between N and $2N$ grains to fall is twice the frequency as for $2N$ to $4N$ grains. In other words, it is a power law or scale-invariant distribution similar to the Gutenberg-Richter frequency law for earthquakes, that carries the implication that disturbances of unbounded size can be triggered by a very small event. They coined the term *self-organized criticality*, or "SOC", for this type of phenomenon, and boldly claimed that such large scale driven systems have an innate tendency to build into a steady state that exhibits power law statistics that are *universal*, or insensitive to the microscopic details of the system.

Self-organized criticality has also been invoked to explain the widespread observation of fat-tailed *Pareto distributions* in economic contexts, such as the size of cities, the distribution of wealth, and the distribution of firm sizes. Network scientists are thus not surprised to see evidence of Pareto tails in the size and connectivity of financial networks, with large, highly connected *hub* banks that form a core within a periphery of intermediate and small banks.

It might sound naive to assert that something like sand piling is happening in financial systems. However, as Minsky wrote in [56], "Stability—even of an expansion—is destabilizing in that more adventuresome financing of investment pays off to the leaders, and others follow." Perhaps the financial system is like a sand pile near its maximal size, where unbounded disturbances are possible. The *Minsky moment* when a financial bubble bursts might then be analogous to one of these large scale disturbances. Adrian and Shin [4] provide a possible explanation. They demonstrate that in the 10 year period leading up to the 2007-08 crisis, financial institutions exhibited strongly pro cyclical investment strategies: as asset prices rose during the prolonged period of stability, so did the balance sheets and leverage ratios of banks, showing that they pursued ever more adventurous strategies. Eventually, as the financial system approached a critical state with little government oversight, only small triggers were needed to create the inevitable collapse.

1.3 Channels of Systemic Risk

Systemic contagion that causes the failure or impairment of a large number of banks will in reality always manifest itself through a multitude of different channels, with spillover or domino effects from one to another. In the language of network science, financial networks are multiplex, meaning there are interbank links of many different types, and a contagious event that starts with one type of link will likely quickly infect all other types of links. Nonetheless, it is important to identify the basic types of shock mechanisms that we expect to find activated during a financial crisis, either as the primary cause, or else as the result of spillover effects stemming from the

initial shock. For an in-depth discussion of various channels of systemic risk, and in particular, contagion, please see [29].

Asset Correlation: Different banks tend to hold common assets in their portfolios. Haldane [41] has argued that banks' asset portfolios became increasingly similar during the Great Moderation, making them more and more susceptible to correlated asset shocks that can be considered as a channel of systemic risk. In 2007, most large banks around the world held significant positions in the US sub-prime mortgage market. The prolonged drawdown of US housing prices in that year acted as a huge downward asset shock that exposed the vulnerability of most banks' asset portfolios. Such systemic events undermine the health of the system, in the same way that famine impairs the health of a community without exhibiting the amplification factors that characterize contagion, but make it vulnerable to other types of contagion.

Default Contagion: Bank deposits held in other banks can be considered as a form of interbank lending, but banking in modern times has dramatically expanded the range of interbank exposures. There is a multitude of linkage types between bank counterparties that range well beyond traditional interbank lending, to include swaps, derivatives and other securitized assets. At any moment, banks can at least in principle identify their exposures to all other banks and they also work hard to identify their expected potential exposure over different future time horizons. When a bank becomes insolvent, if it is not bailed out by a government agency, it will be forced into bankruptcy. Its creditors, including other banks, will then experience severe losses given this default, possibly losing close to 100% of their total exposure in the short term aftermath. Such shocks to creditor banks' interbank assets at the time of default of a debtor bank are the channel for *default contagion*. If left unchecked by government intervention, such shocks can in principle chain together like dominos to create a *default cascade*. Default cascades can only happen when interbank exposures are a high fraction of lending banks' equity, and [70] provides evidence that this was the case in Europe before and during the crisis, when many banks' interbank exposures exceeded their capital by factors of 5 or more.

Liquidity Contagion: *Funding illiquidity* is the situation of a bank with insufficient access to short term borrowing. Such banks, being short of cash or other liquid assets, will adopt a variety of strategies that can be considered as shrinking their balance sheets. They will try to access the repo markets for untapped sources of collateralized borrowing. They will refuse to rollover short term loans and repo lending to other counterparties. When banks respond to funding illiquidity by curtailing a large fraction of their interbank lending, the resulting funding shocks to other banks are the channel for *liquidity contagion* in the system.

Market Illiquidity and Asset Fire Sales: As [4] discussed, in good times banks tend to create upward asset price spirals by increasing their leverage through large scale asset purchasing. This pushes up prices, creating the illusion of even better times ahead and further increases in leverage. As they also discuss, the reverse is true in bad times. This tendency for distressed banks to sell assets into a depressed

market creates the contagion mechanism known as an *asset fire sale*. A fire sale cascade proceeds through a double step mechanism: first, asset sales by distressed banks decreases prices, then marking-to-market leads to losses by other banks holding these assets.

Other Channels: Many authors have identified further channels for systemic risk. *Rollover risk* is the name given to the effect that lenders to a bank may fail to renew or “roll-over” short term debt. [60] models this effect as a coordination game played by the lenders to a single counterparty: a player that perceives that other players are likely not to rollover their debt, will be much more likely not to rollover their debt. Such a decision may be due either to a lending bank’s assessment of the health of the counterparty (which was termed *structural uncertainty*), or to that bank’s assessment of the lending behaviour of other banks (termed *strategic uncertainty*). [6] extend this picture to a network setting by considering a multitude of simultaneous coordination games, leading to runs in the interbank network. In [39], it is argued that the 2008 crisis was largely a crisis of confidence in the repo market that led to a drying up of banks’ funding opportunities. In normal times, the repo market provides a huge source of short term funding for banks that is *information insensitive* in the sense that the lender has little incentive to be concerned about the health of its counterparty. During the crisis however, lenders became *information sensitive* and questioned the quality of counterparties and their underlying collateral. Consequently, they began to demand large increases in repo *haircuts*. In other words, they demanded collateral valued well above the loan amounts, and in consequence dramatically decreased the availability of repo funding at a time it was most needed. This effect was contagious: banks that raised haircuts imposed funding illiquidity on their counterparties, leading to further questioning of the quality of counterparties and their collateral.

1.4 Capital Structure of a Bank

Banks around the globe form a diverse family of firms, spanning a huge range of sizes and types. In addition to traditional retail and investment banks, financial network models need eventually to include a whole host of shadow banking institutions, including hedge funds, pension and investment funds, savings and credit institutions, and so on. As our systemic models evolve, we will include in our system more and more components of the wider production and retail economy. It will clearly be impossible to capture here in a brief overview the full range of holdings and liabilities that such institutions might have, and that should in principle be understood.

Quarterly financial reports of a firm’s balance sheet offer a snapshot at a moment in time of the details of a firm’s capital structure, that is, the valuation of the totality of their assets and liabilities. It is helpful to imagine these balance sheet entries as existing at all times, even if not observed by the market. One can imagine that the bank management maintains its accounting books, updating them daily, and only

once a quarter makes them available to the public. Regulators, on the other hand, have the power to examine the books of any bank, at any moment.

In studies of the “financial system”, it is important to carefully define the boundary between the interior and exterior of the system. Correspondingly, for systemic analysis, assets and liabilities will always be separated into intra- and extra-network components. The available classes of assets relevant in banking is extremely diverse, but from a systemic risk perspective their most important characteristics are in the following dimensions: duration or maturity, credit quality, interest rate and liquidity. Similarly, liabilities primary characteristics are maturity, interest rate, and seniority.

One important insight to keep in mind when considering capital structure is the formal duality between assets and liabilities. Almost any asset is someone else’s liability and in many cases where a story can be told of asset side contagion, an analogous story can be told of liability side contagion.

1.4.1 Bank Asset Classes

This section gives a schematic overview of some of the most important basic forms that bank assets take, and how these assets may be controlled. The key attributes of bank assets are duration or maturity, credit quality and liquidity.

Loan portfolio: The main business line of a bank, like any other firm, is to invest in endeavours for which it has a competitive advantage. Such *irreversible projects* are by their nature illiquid, and fail to recoup their full value when sold. Indeed, their mark-to-market value may not even be well-defined, and we should suppose that when liquidated suddenly are sold well below their book value. For banks, this business line, the *bank book*, consists of a vast portfolio of loans and mortgages of all maturities, to a wide variety of counterparties ranging from the retail sector, to small and medium enterprises (SMEs) and major corporates. Far from receding in importance since the crisis, [48] shows that real estate lending (i.e. mortgages) in particular accounts for an ever increasing percentage of bank assets. They comment: “Mortgage credit has risen dramatically as a share of banks’ balance sheets from about one third at the beginning of the 20th century to about two thirds today” and suggest that real estate bubbles will continue to provide a dominant component of future systemic risk events.

As mentioned above, in systemic risk analysis, assets placed within the financial system, called *interbank assets* are always distinguished from *external assets* placed outside the system.

Over-the-counter securities: Bonds, derivatives and swap contracts between banks are to a large extent negotiated and transacted in the OTC markets, sometimes bilaterally, but increasingly within a *central clearing party* (CCP) . Some of these exposures fluctuate rapidly in time, both in magnitude and in sign, and may or may not be collateralized by margin accounts to reduce counterparty risk. Between a pair of financial counterparties there may be many contracts, each with positive and

negative exposures. To reduce risk, counterparties often negotiate a bilateral *master netting agreement* (MNA), subject to the regulations stipulated by ISDA, that allows them to offset exposures of opposite signs. Entering into an MNA is a costly endeavour, and thus existence of an MNA is an indication of a strong network connection between two banks. Methods of counterparty risk management, reviewed for example in [28], is a particular flavour of credit risk management that has developed rapidly since the crisis. As part of this methodology, banks now routinely forecast the *potential future exposure*, or PFE, for all their important counterparties. This is a high quantile (often 95%) of the positive part of their exposure to the given counterparty on a given date in the near future. The PFE of one bank to another comes close to what we mean when we assign a weight to an interbank link: in the event that one bank suddenly defaults, PFE is a pessimistic estimate of losses to which its counterparties are exposed.

An important subclass of OTC securities are *total return swaps* (TRS) that exchange the random returns on an underlying asset for a fixed periodic payment. An example of a TRS is the credit default swap (CDS) that exchanges fixed quarterly payments for a large payment at the moment the underlier defaults, providing a form of default insurance. From a network perspective, complex and poorly understood effects are bound to arise when the underlier is itself part of the system, as would be the case of a CDS written on the default of another bank.

Cash and market securities: In addition to cash, the firm may hold other securities for which there is a liquid market, and low or moderate transaction costs. Some examples would be the money-market account that pays the over-night rate, stocks, T-bills, or exchange traded derivatives. In the case of banks, a new aspect of the Basel III regulatory framework requires active and prudent liquidity management which means that a fraction of assets must be held in a portfolio of cash and market securities that can be easily liquidated when the bank needs to meet its short term obligations in a timely manner.

Reverse repo assets: A repo-lending bank (see the next section on repos) receives collateral assets known as reverse repos. Such assets can be “rehypothecated”, which means they can themselves be used as collateral for repo borrowing.

Other assets: Of lesser importance are a range of further asset classes such as real estate, accounts receivable, and the like.

1.4.2 Debt and Liabilities

Deposits: A large fraction of the debt of a traditional bank is in the form of deposits made by both institutional investors and small retail investors. Since there are many of them, with a diverse range of maturities, the collective of deposits can be thought of as a multiple of an asset that pays a constant dividend rate (for banks, we will assume it is less than the risk free rate). One important class of wholesale depositor is short-term money-market funds. Their widespread meltdown in early stages

of the last financial crisis played an important contagion role. Small depositors are typically protected by deposit insurance in the event of the bank's default, while institutional depositors have no such protection. Banks in particular seek protection through collateralization. Uncollateralized lending between banks takes other forms: certificates of deposit and bankers' acceptances are variants banks use to lend to each other.

Bonds: Like most large firms, banks issue bonds as a primary means to raise long term debt capital, each bond being characterized by its notional amount, maturity and coupon rate, plus a variety of additional attributes. Sometimes a bank's bonds differ in seniority, meaning that in the event of default, the most senior bonds are paid in full before junior bonds. Typically, the firm cannot retire existing bonds or issue new bonds quickly.

Market securities Hedge funds and investment banks have the characteristic that they often take large short positions in market securities, such as stocks and derivatives. To a lesser extent, commercial banks also hold short positions for hedging and risk management reasons. Short positions can be thought of as holding negative amounts in market securities.

Collateralized Loans (Repos): Short for *repurchase agreements*, repos are an important class of collateralized short term debt issued between banks and other institutional investors. Typically, money is borrowed for a short term, often overnight, at an interest rate r called the *repo rate*. They are backed by assets (*repo assets*) whose value exceeds the loan amount by a percentage h called the *haircut*. The haircut reflects the liquidation value of the collateral in the event the money is not repaid. This haircut thus compensates the lender, also called the asset buyer, for the counterparty risk inherent in the contract when the borrower (or asset seller) defaults. To illustrate how repos are used, suppose a bank has an excess \$1 in cash. Then they might undertake an overnight repo of \$1 with another bank for collateral valued at $\$1/(1-h)$. The next day the contract is closed by *repurchase* of the collateral for the price $\$(1+r)$. While the lending bank holds the collateral, they may also elect to use it to finance a second repo with another counterparty: we then say the collateral has been *rehypothecated*.

Hybrid capital: This term denotes parts of a firm's funding that possess both equity and debt features. Like equity, it may pay tax-deductible dividends, and like debt it maintains seniority over equity in the event of default. There is now active interest in forms of hybrid capital issued by banks such as *contingent convertible bonds* (COCOs) that behave like bonds as long as the firm is healthy, but provide additional equity cushioning for the bank when its balance sheets weaken.

Other Liabilities: Accounts payable are analogous to accounts receivable. Investment banks act as prime brokers and hold their clients' securities in trust. They often act on the right to borrow such securities to be used as collateral for purchasing further assets. Such a debt can be understood as similar to a collateralized loan.

1.4.3 Equity

Equity, defined to be the value of a firm's assets minus its liabilities, what it owns minus what it owes, represents the total value conveyed by ownership of the firm. For a publicly owned firm, ownership is divided into shares that have an observable fluctuating market price: in this case, total market capitalization, the share price times the number of shares outstanding, is a market value of equity. For privately held firms, equity does not have a transparent market value: valuation of privately held firms can only be gleaned through investigation of the firm's accounts. Firms, especially banks, are typically highly leveraged, meaning their assets are a large multiple of their equity. In such situations, equity, being a small difference of large positive numbers, is inherently difficult to estimate and this uncertainty is reflected in the high volatility of the stock price.

Limited liability is the principle, applying to almost all publicly held companies, that share owners are never required to make additional payments. In the event the firm ceases to operate, the shareholders are not held responsible for unpaid liabilities. We can say this means equity can never be negative, and is zero at the moment the firm ceases to operate. This is called bankruptcy, and limited liability is a central principle in bankruptcy law worldwide.

Firms return profits to their shareholders two ways, either through the payment of regular dividends, or through the increase in share price when the value of firm equity rises. Share issuance and its opposite, the share buyback, are two further financing strategies firms can adopt. A firm near to default may attempt a share issuance, hoping that new investors will view the decline in its fortunes as a temporary effect before the firm's return to health.

1.5 Regulatory Capital and Constraints

Largely as a result of lessons hard learned during the crisis, the world is currently moving quickly beyond a Basel II regulatory regime that can be characterized as *microprudential* in emphasis. This means Basel II regulations are imposed bank by bank without taking into account the dependence of risks between banks. For example, the *capital adequacy ratio* (CAR) which stipulates

$$\text{Risk-weighted Assets} \leq 12.5 \times \text{Total Capital}$$

is the main requirement of Basel II, and is based only on the individual bank's balance sheets. The importance of network effects is recognized at the core of Basel III in measures that are *macroprudential* in nature, meaning they try to account for the network and the interconnectivity of risks between banks.

An example of macroprudential regulation is the new requirement by Basel III that banks must report their large exposures to individual counterparties or groups of counterparties, both financial and non-financial. This is clear recognition that large

interbank linkages are systemically important during a crisis. It has also become clear that the fixed regulatory capital ratios of Basel II were *procyclical* and can dangerously amplify the swings of the credit cycle. When the financial system enters a contraction phase, capital buffers of some banks will be squeezed below the regulatory minimum, leading naturally to fire sales and further asset shocks. Basel III seeks to ward off this tendency by making the capital requirements countercyclical. During normal periods, capital requirements have a surcharge which can be removed as the system begins to contract to provide banks with more flexibility. Yet another example of macroprudential regulation is that Basel III now recognizes the existence of *SIFIs* (for *systemically important financial institutions*), also called *G-SIBs* (for *global systemically important banks*), and subjects them to a regulatory capital surcharge that will hopefully make them more resilient. Clearly, the identification of SIFIs must be grounded in well established systemic risk theory, and the SIFIs themselves will demand a theoretical basis for what is to them a punitive measure.

The Basel II capital adequacy ratio, although it has been strengthened in Basel III, still leads to distortions banking balance sheets through its use of *risk weights*. For example, the risk weight for sovereign bonds of OECD countries remains zero, meaning these assets require no offsetting equity capital, allowing banks to operate with unsustainably high leverage ratios. Basel III provides a counterbalance to the CAR by requiring an additional constraint on bank leverage. Liquidity risk is for the first time explicitly addressed in Basel III through the implementation of two new regulatory ratios. The *Liquidity Coverage Ratio* (LCR) ensures that every bank's liquid assets will be sufficient to cover an extreme stress scenario that includes a 30 day run off of its short-term liabilities. The *Net Stable Funding Ratio* (NSFR) similarly seeks to ensure that enough long term (greater than one year) funding will be available to cover a stress scenario that hits the bank's long term assets. Considering the Lucas critique [53](p. 41) that "any change in policy will systematically alter the structure of econometric models", we must expect that the systemic ramifications of the LCR and NSFR will be subtle and far-reaching.

Critics, notably Haldane [42], argue that the Basel III regulatory framework has become excessively complex, and that the regulatory community must do an about-turn in strategy and operate by a smaller, less-detailed rulebook. Haldane writes: "As you do not fight fire with fire, you do not fight complexity with complexity. Because complexity generates uncertainty, not risk, it requires a regulatory response grounded in simplicity, not complexity."

Our task now is to begin exploring the channels of cascading systemic risks in detail, with an aim to closing the gap in understanding that exists between knowing bank-to-bank interactions and knowing how a large ensemble of banks will behave.

Chapter 2

Static Cascade Models

Abstract Distinct financial network effects such as default contagion and liquidity hoarding share the characteristic that they are transmitted between banks by direct contact through their interbank exposures. It follows they should be modelled starting from a common framework. In this chapter, a number of different cascade mechanisms are developed and shown to have common features. Banks maintain safety buffers in normal times, but these may be weakened or fail during a crisis. Banks react to such stresses by making large adjustments to their balance sheets, sending shocks to their counterparties that have the potential to trigger further stresses. The eventual extent of a crisis can be described mathematically as a fixed point or equilibrium of a cascade mapping. Asset fire sales can be modelled analogously as a form of single buffer contagion, while more complex effects arise when banks' behaviour is determined by multiple buffers. In the end, we propose a definition of a random financial network.

Keywords: Cascade mechanism, default and liquidity buffers, fixed point equations, cascade equilibrium, asset fire sales, random financial network.

Contagion, meaning the transmission of a disease by direct or indirect contact, is an appropriate term for the damaging effects that can be transmitted through the interbank network, and which are the essential point of interest in this book. This chapter will focus on the development of a number of frameworks for stylized static cascades that model contagion that can arise in hypothetical financial networks. The essential picture will be to treat banks and their behaviour as determined by their balance sheets, and to concentrate on effects that can be transmitted between banks through the contracts they exchange.

We begin in the traditional fashion by considering first how the default of one bank might directly impact other banks, leading to what we call a *default cascade*. In such situations, the default of any bank that occurs when its assets are insufficient to cover its liabilities, will have a large impact on each of its creditor banks. These default shocks might then build up and cause further bank defaults, and trigger new shocks, creating a default cascade. In an analogous way, part of the recovery strategy of an illiquid bank may be to call its loans to its *debtor* banks, thereby inflicting

liquidity shocks on these banks. In view of this formal similarity, it is interesting to consider how the same basic cascade framework, reinterpreted, can also model such a simple *liquidity cascade*. A third type of contagious effect, called asset fire sales, occurs when large scale selling of an asset class by a stressed bank can lead to price shocks that in turn negatively impact other banks holding the same asset class. Finally, one can investigate more complex models of double or higher dimensional cascades that combine different contagion transmission mechanisms, each with some of these basic characteristics.

Our primary purpose is therefore to focus on contagion effects that are transmitted through interbank linkages. These are genuinely network effects that cannot be understood within traditional financial mathematics modelling. The models we consider are simple and stylistic and are intended to capture the essence of contagion rather than specific details. By studying the schematics of financial cascades, we can hope to connect with more developed and better understood domains of network science where analogous mechanisms have been well-studied, allowing us to carry over their intuition and rigorous results. In later chapters, we will attach these aspects of systemic risk firmly to the rich mathematical framework that has been developed to understand cascades and contagion in network science. At some basic level, cascades of failing banks are not so different from certain phenomena that are seen in other complex networks, such as the failure of power grids, the spread of memes in society or cascades of nerve impulses in neural networks.

A model of *static default cascades* starts with a set of banks characterized by their balance sheets that are interconnected through directed links representing simple interbank loans. The balance sheets contain the information of the nominal (or book) values of all the relevant securities held as assets or liabilities of each bank. Insolvency, or default, of a bank arises when its assets are insufficient to cover their debt liabilities. It is natural to assume that a defaulted bank can no longer repay the full nominal values of its debt. The solvency of banks are then determined by the reduced, or “mark-to-market”, values of interbank assets.

In this book, we use the word “static” to describe a cascade whose end result is a deterministic function of the initial balance sheets and exposures. We shall assume that such cascades proceed from a possibly random initial configuration through a *cascade mechanism* that generates a series of deterministic steps until a steady state or *cascade equilibrium* is reached. At the outset of a default cascade, a subset of banks, called the *seed*, are taken to be defaulted due to some triggering event coming from outside the system. In each cascade step, banks update the mark-to-market valuation of their assets to assess their own solvency, and any newly insolvent bank is identified. The cascade equilibrium is reached at the step when no further banks default. In a static cascade, it is assumed that balance sheets do not experience cash flows, such as interest or dividend payments, in any cascade step. Each cascade step can be viewed as a mapping from the collection of marked-to-market bank balance sheets to itself, depending on the nominal bank balance sheets. The state of the network at the end of the cascade is then necessarily a fixed point of this cascade mapping.

The well-known Eisenberg-Noe model [32] is the prototype of a static cascade model. As we shall review in this chapter, the systemic risk literature contains many variations, that differ in the types of securities and exposures considered, and in the way defaulted assets are marked-to-market. Underlying all such cascade models are some interesting mathematical questions. Is the cascade mapping well-defined, and does it necessarily have an economically acceptable fixed point? If there is a fixed point, is it unique? When the fixed point is unique, many pathological behaviours can be ruled out. On the other hand, when there are many fixed points, determining their stability properties may be a challenging problem. The main goal in this chapter is to create a framework flexible enough to cover the spectrum of known static cascade models.

2.1 Bank Balance Sheets

The financial cascade framework of Eisenberg and Noe [32] begins with a financial system assumed to consist of N “banks”, labelled by $v = 1, 2, \dots, N := [N]$ (which may include non-regulated, non-deposit taking, leveraged institutions such as hedge funds, or other regulated financial institutions such as insurance companies). Their balance sheets can be characterized schematically like this:

Assets	Liabilities
external assets \bar{Y}	external debt \bar{D}
interbank assets \bar{Z}	interbank debt \bar{X}
	equity E

Table 2.1 An over-simplified bank balance sheet.

At the outset, the entries in these banks’ balance sheets refer to *nominal values* of assets and liabilities, and give the aggregated values of contracts, valued as if all banks are solvent. Nominal values, denoted by upper case letters with bars, can also be considered *book values* or *face values*. Assets and liabilities are also decomposed into interbank and external quantities depending on whether the loan or debt counterparty is a bank or not. Banks and institutions such as foreign banks that are not part of the system under analysis are deemed to be part of the exterior, and their exposures are included as part of the external debts and assets.

Definition 1. The *nominal value of assets* of bank v at any time consists of *nominal external assets* denoted by \bar{Y}_v , plus *nominal interbank assets* \bar{Z}_v . The *nominal value of liabilities* of the bank includes *nominal external debt* \bar{D}_v and *nominal interbank debt* \bar{X}_v . The bank’s *nominal equity* is defined by $\bar{E}_v = \bar{Y}_v + \bar{Z}_v - \bar{D}_v - \bar{X}_v$. The *nominal exposure* of bank w to bank v , that is the amount v owes w , is denoted by $\bar{\Omega}_{vw}$. We define interbank loan fractions to be $\bar{\Pi}_{vw} = \bar{\Omega}_{vw}/\bar{X}_v$ as long as $\bar{X}_v > 0$. Interbank assets and liabilities satisfy the constraints:

$$\bar{Z}_v = \sum_w \bar{\Omega}_{wv}, \quad \bar{X}_v = \sum_w \bar{\Omega}_{vw}, \quad \sum_v \bar{Z}_v = \sum_v \bar{X}_v, \quad \bar{\Omega}_{vv} = 0.$$

The stylized balance sheets of all banks in the network can be combined into a table, as shown in Table 2.2.

	1	2	...	N	\bar{X}	\bar{D}	\bar{E}
1	0	$\bar{\Pi}_{12}\bar{X}_1$...	$\bar{\Pi}_{1N}\bar{X}_1$	\bar{X}_1	\bar{D}_1	\bar{E}_1
2	$\bar{\Pi}_{21}\bar{X}_2$	0	...	$\bar{\Pi}_{2N}\bar{X}_2$	\bar{X}_2	\bar{D}_2	\bar{E}_2
...
N	$\bar{\Pi}_{N1}\bar{X}_N$	$\bar{\Pi}_{N2}\bar{X}_N$...	0	\bar{X}_N	\bar{D}_N	\bar{E}_N
Z	\bar{Z}_1	\bar{Z}_2	...	\bar{Z}_N			
Y	\bar{Y}_1	\bar{Y}_2	...	\bar{Y}_N			

Table 2.2 The matrix of interbank exposures contains the values $\bar{\Omega}_{vw} = \bar{\Pi}_{vw}\bar{X}_v$. The first N rows of this table represent different banks' liabilities and the first N columns represent their assets.

When a bank v is known to be insolvent or defaulted, certain contracts connected to v will naturally be valued at less than their nominal values. These *actual* or *mark-to-market* values, denoted by symbols Z, Y, X, D, E, Ω without upper bars, typically decrease during the steps of the cascade.

Economic cascade models usually invoke the notion of limited liability, and define a firm to be defaulted when its mark-to-market equity is non-positive. This means its aggregated assets are insufficient to pay the aggregated debt. At the onset of the cascade, initially defaulted banks are therefore identified as those banks whose nominal equity \bar{E} is non-positive. As a static default cascade progresses, the market value of equity of banks generally decreases, potentially leading to secondary defaulted banks.

Definition 2. A *defaulted bank* is a bank with $E \leq 0$. A *solvent bank* is a bank with $E > 0$.

The nominal amounts $\bar{Y}, \bar{Z}, \bar{D}, \bar{X}, \bar{\Omega}$ are used to determine the relative claims by creditors in the event a debtor bank defaults. The precise mechanism by which defaulted claims are valued is an important modelling decision that distinguishes different approaches.

Later in the chapter we will study cascade models of illiquidity, in which the impairment banks suffer is not default, but rather insufficient funding. We can call such banks *stressed*. The static cascade models we shall discuss all share a number of timing assumptions, which we describe as if the cascade proceeds in daily steps.

Assumption 1. 1. Prior to the cascade, all banks are in the normal state, not insolvent.

2. The crisis commences on day 0 triggered by the default of one or more banks;
3. Balance sheets are recomputed daily on a mark-to-market basis;
4. Banks respond daily on the basis of their newly computed balance sheets;

5. All external cash flows, interest payments, and external asset price changes are ignored throughout the cascade.

2.2 The Eisenberg-Noe 2001 Model

This famous model of default contagion makes three additional assumptions that determine a precise default resolution mechanism¹:

Assumption 2. 1. External debt is senior to interbank debt and all interbank debt is of equal seniority; 2. There are no losses due to bankruptcy charges; and 3. $\bar{Y}_v \geq \bar{D}_v$.

In tandem with the limited liability assumption, the first of these means that on default, a bank's equity is valued at zero, and its external debt must be paid in full before any of its interbank debt is paid. A variant of this assumption is discussed at the end of this section. The third assumption means that defaulted banks always have sufficient external assets to pay their external debt in full.

The no-bankruptcy costs assumption is very important, and somewhat optimistic in the context of systemic risk. It has the strong consequence that when the system is viewed as a whole, no system-wide equity is lost during the crisis. It is easy to see that the system equity, defined as total assets minus total liabilities, is independent of the payments within the inter banking sector:

$$\bar{E}_{sys} = \sum_v (\bar{Y}_v + \bar{Z}_v - \bar{D}_v - \bar{X}_v) = \sum_v (\bar{Y}_v - \bar{D}_v).$$

Let us suppose the banking network, previously in equilibrium with no defaulted banks, experiences a catastrophic event, such as the discovery of a major fraud in a bank or a system wide event, whereby the nominal assets of some banks suddenly contract. If one or more banks are then found to be in such a state of *primary default*, they are assumed to be quickly liquidated, and any proceeds go to pay off these banks' creditors, in order of seniority. We let $p_v^{(n)}, v \in [N]$ denote the (mark-to-market) amount available to pay v 's internal debt at the end of the n th step of the cascade, and $\mathbf{p}^{(n)} = [p_1^{(n)}, \dots, p_N^{(n)}]$. By Assumption 2, the value $p_v^{(n)}$ is split amongst the creditor banks of v in proportion to the fractions $\bar{\Pi}_{vw} = \bar{Q}_{vw}/\bar{X}_v$ (when $\bar{X}_v = 0$, we define $\bar{\Pi}_{vw} = 0, w = 1, \dots, N$). Therefore, at step $n+1$ of the cascade, every bank w values its interbank assets as $Z_w^{(n+1)} = \sum_v \bar{\Pi}_{vw} p_v^{(n)}$. Since by assumption there are no bankruptcy charges, we find

$$p_v^{(n+1)} = F_v^{(EN)}(\mathbf{p}^{(n)}); F_v^{(EN)}(\mathbf{p}) := \min(\bar{X}_v, \bar{Y}_v + \sum_w \bar{\Pi}_{vw} p_w - \bar{D}_v), v = [N]. \quad (2.1)$$

Now we let $\mathbf{p} = [p_1, \dots, p_N]$ denote the vector of banks' internal debt values at the end of the cascade. This *clearing vector* satisfies the *clearing condition*, or *fixed*

¹ The original paper is expressed in terms of the differences $\bar{e}_v := \bar{Y}_v - \bar{D}_v$.

point condition:

$$\mathbf{p} = F^{(EN)}(\mathbf{p}). \quad (2.2)$$

For any clearing vector \mathbf{p} , the equity of bank v at the end of the cascade is given by

$$E_v = (\bar{Y}_v + \sum_w \bar{\Pi}_{wv} p_w - \bar{D}_v - \bar{X}_v)^+ = \bar{Y}_v + \sum_w \bar{\Pi}_{wv} p_w - \bar{D}_v - p_v, \quad (2.3)$$

where one can see that the second equality holds because bankruptcy charges are zero.

The first main theorem of Eisenberg and Noe, written using a vector and matrix notation described in Appendix A.1, is the following:

Theorem 1. *Corresponding to every financial system $(\bar{Y}, \bar{Z}, \bar{D}, \bar{X}, \bar{\Omega})$ satisfying Assumption 2,*

1. *There exists a greatest and a least clearing vector \mathbf{p}^+ and \mathbf{p}^- .*
2. *Under all clearing vectors, the value of the equity at each node is the same, that is, if \mathbf{p}' and \mathbf{p}'' are any two clearing vectors,*

$$(\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}' - \bar{D} - \bar{X})^+ = (\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}'' - \bar{D} - \bar{X})^+.$$

Proof: Part (1) of the Theorem follows by the Knaster-Tarski Fixed Point Theorem² once we verify certain characteristics of the mapping F^{EN} . We note that F^{EN} maps the hyperinterval $[\mathbf{0}, \bar{X}] := \{x \in \mathbb{R}^N : 0 \leq x_v \leq \bar{X}_v\}$ into itself. We also note that it is monotonic: $x \leq y$ implies $F^{EN}(x) \leq F^{EN}(y)$. Finally, note that $[\mathbf{0}, \bar{X}]$ is a complete lattice. We therefore conclude that the set of clearing vectors, being the fixed points of the mapping F^{EN} , is a complete lattice, hence nonempty, and with maximum and minimum elements \mathbf{p}^+ and \mathbf{p}^- .

Part (2) follows by showing that for any clearing vector \mathbf{p}' ,

$$(\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}' - \bar{D} - \bar{X})^+ = (\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}^+ - \bar{D} - \bar{X})^+.$$

By monotonicity, $\mathbf{p}' \leq \mathbf{p}^+$ implies

$$(\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}' - \bar{D} - \bar{X})^+ \leq (\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}^+ - \bar{D} - \bar{X})^+$$

and equivalently

$$\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}' - \bar{D} - p' \leq \bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}^+ - \bar{D} - p^+.$$

Then we take the inner product of this equation with the vector $\mathbf{1} = [1, \dots, 1]$ and use the fact that $\mathbf{1} \cdot \bar{\Pi}^T = \mathbf{1}$ to show that

$$\mathbf{1} \cdot (\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}' - \bar{D} - p') = \mathbf{1} \cdot (\bar{Y} + \bar{\Pi}^T \cdot \mathbf{p}^+ - \bar{D} - p^+)$$

² A statement and proof of this result can be found at <http://en.wikipedia.org/wiki/Knaster-Tarski-theorem>

and the result is proved. Note there is a glitch in this proof because when $\bar{X}_v = 0$ we have defined $\bar{\Pi}_{wv}^T = 0$, invalidating the condition $\mathbf{1} \cdot \bar{\Pi}^T = \mathbf{1}$: however for this v , $p_v = 0$, and hence it is still true that $\mathbf{1} \cdot \bar{\Pi}^T \cdot p = \mathbf{1} \cdot p$. \square

Remark 1. We note that by Part 2 of the Theorem, $\bar{\Pi}^T \cdot p' - p' = \bar{\Pi}^T \cdot p'' - p''$. Also, note that Part (2) depends crucially on the assumption of no-bankruptcy charges.

Finding natural necessary and sufficient conditions on E-N networks to ensure the uniqueness of the clearing vector proves to be more challenging. Theorem 2 of [32] gives a sufficient condition, and some economic intuition that underlies it. We will address this question in a slightly more general setting shortly, but in the meantime, the following is an example of non-uniqueness of the clearing vector.

Remark 2. The $N = 4$ bank network shown in Figure 2.1 with

$$\bar{\Omega} = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 2 & 0 \\ 3 & 0 & 0 & 0 \\ 1 & 1 & 1 & 0 \end{pmatrix}$$

and $\bar{Y} - \bar{D} = 0$ has multiple clearing vectors $p = \lambda[1, 1, 1, 0]$ with $\lambda \in [0, 1]$.




Fig. 2.1 A simple network with an irreducible in-graph consisting of the nodes $\{1, 2, 3\}$.

Including debt \bar{Y}' with equal seniority to interbank debt can be handled easily by introducing a fictitious source bank $v = 0$. However, despite a contrary claim made in [32], it is not trivial to extend the above results to networks which allow $\bar{Y} < \bar{D}$: The extension to this general case, which we now describe, seems to have first been done by [30]. By writing $\bar{Y} - \bar{D} = (\bar{Y} - \bar{D})^+ - (\bar{Y} - \bar{D})^-$ one can now express the clearing condition in terms of a vector $q = [q_1, \dots, q_N]^T$ where q_v denotes the amount bank v has available to pay both the excess external debt $(\bar{Y}_v - \bar{D}_v)^-$ and the interbank debt \bar{X}_v :

$$q = \min((\bar{Y} - \bar{D})^+ + \bar{\Pi}^T \cdot p, (\bar{Y} - \bar{D})^- + \bar{X})$$

$$p = (q - (\bar{Y} - \bar{D})^-)^+.$$

This can be written more compactly in terms of p alone:

$$p = \min(\bar{X}, \max(\bar{Y} - \bar{D} + \bar{\Pi}^T \cdot p, 0)) . \quad (2.4)$$

We present here a complete characterization of the clearing vector set in the E-N model without the condition $\bar{Y} \geq \bar{D}$. First, we identify groups of banks called *in-subgraphs* (these are essentially the same as the *surplus sets* in [32]), that do not lend outside their group.

Definition 3.

1. In-subgraph: any subgraph $\mathcal{M} \subset \mathcal{N}$ (possibly $\mathcal{M} = \mathcal{N}$) with no out-links to its complement \mathcal{M}^c .
2. Irreducible in-subgraph: an in-subgraph $\mathcal{M} \subset \mathcal{N}$ with at least 2 nodes that does not contain a smaller in-subgraph.

Having identified all the irreducible in-subgraphs of a given network, we can simplify the following discussion by removing all singleton *in-banks* (i.e. that do not lend to other banks). Any such bank v will have $p_v = \bar{X}_v = 0$ and its state has no effect on other banks. We then consider the exposure matrix $\bar{\Pi}^T$ restricted to the *reduced network* $\tilde{\mathcal{N}}$ without such in-banks, which may then be strictly substochastic³ in some columns.

The theorem that characterizes the set of clearing vectors in the E-N model is:

Theorem 2. *Let the reduced network $\tilde{\mathcal{N}}$ have exactly $K \geq 0$ non-overlapping irreducible in-graphs $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_K$, and a decomposition $\tilde{\mathcal{N}} = \mathcal{M}_0 \cup (\cup_{k=1}^K \mathcal{M}_k)$. Let the possible fixed points be written in block form $p = [p_0^*, p_1^*, \dots, p_K^*]^T$. Then:*

1. *In case $K = 0$, the clearing vector $p^* = p_0^*$ is unique;*
2. *In case there are exactly $K \geq 1$ (non-overlapping) irreducible in-graphs $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_K$, then the multiplicity of the clearing vectors is characterized as follows: p_0^* is unique, and each p_k^* is either unique or of the form $p_k^* = \alpha_k v_k$ where the vector v_k is a 1-eigenvector of the matrix $P_k \cdot \bar{\Pi}^T \cdot P_k$ projected onto \mathcal{M}_k and $\alpha_k \in [0, \bar{\alpha}_k]$ for some $\bar{\alpha}_k$. The precise form of each p_k^* is shown in the proof.*

Remark 3. This theorem demonstrates that non-uniqueness in this model is a highly non-generic property: only extremely special arrangements of the network lead to multiplicity of solutions.

The proof of the theorem involves extending the following lemma to the most general kind of E-N fixed point equation.

Lemma 1. *Consider the system*

$$p = \min(X, \max(Y + \Pi \cdot p, 0)) \quad (2.5)$$

where Π is substochastic, X is a positive vector and Y is arbitrary.

³ A stochastic (substochastic) matrix has non-negative entries and columns that sum to 1 (respectively, ≤ 1).

1. If Π is strictly substochastic and has spectral radius less than one, then there is a unique fixed point.
2. If Π is stochastic and irreducible, with y the unique eigenvector such that $\Pi \cdot y = y$ and $\mathbf{1} \cdot y = 1$, then one of two possible cases holds:
 - a. If $\mathbf{1} \cdot Y = 0$, there is a one-parameter family of solutions that have the form $p^* = \lambda y$, $\lambda \in [0, \lambda_{\max}]$.
 - b. If $\mathbf{1} \cdot Y \neq 0$, there is a unique fixed point, of which at least one component is either 0 or X .

Proof of Lemma: In Part 1, uniqueness is guaranteed because $I - \Pi$ has an explicit inverse given by the convergent matrix power series $(I - \Pi)^{-1} = \sum_{k=0}^{\infty} \Pi^k$.

Under the conditions of Part 2, by dotting (2.5) with the vector $\mathbf{1}$, we see that the system $p = Y + \Pi \cdot p$ has a solution if and only if $\mathbf{1} \cdot Y = 0$. If $\mathbf{1} \cdot Y = 0$ one can check that case 2(a) of the lemma holds. If $\mathbf{1} \cdot Y \neq 0$, at least one component of any fixed point of (2.5) must be X or 0. Substituting in this component value, and reducing the system by one dimension now leads to a new fixed point equation of the same form (2.5) but with the matrix $\bar{\Pi}$ strictly substochastic with spectral radius less than one. Such systems have a unique fixed point by part 1. \square

Proof of Theorem: We must now deal with the system (2.5) when Π is neither strictly substochastic nor stochastic and irreducible. In general, it is easy to see that \mathcal{N} , itself an in-subgraph, contains within it a maximal number of irreducible non-overlapping in-subgraphs $\mathcal{M}_1, \mathcal{M}_2, \dots, \mathcal{M}_K$, $K \geq 0$ plus the possibility of additional single non-lending nodes with $\bar{X}_v = 0$. As discussed above, as a first step we can eliminate all non-lending nodes and consider the reduced piece-wise linear fixed point problem on the subgraph $\tilde{\mathcal{N}}$. The case when $\tilde{\mathcal{N}}$ has no irreducible in-subgraphs, i.e. $K = 0$, has a unique clearing vector because then, by Part 1 of the Lemma, $\bar{\Pi}^T$ must be substochastic with spectral radius less than one.

If $K > 0$, we decompose into $\tilde{\mathcal{N}} = \mathcal{M}_0 \cup (\cup_{k=1}^K \mathcal{M}_k)$, and after reordering nodes write the matrix $\bar{\Pi}^T$ in block form with respect to this decomposition:

$$\bar{\Pi}^T = \begin{pmatrix} A & 0 & \dots & 0 \\ B_1 & \Pi_1 & 0 & 0 \\ \vdots & 0 & \ddots & 0 \\ B_K & 0 & 0 & \Pi_K \end{pmatrix}$$

with the column sums less than or equal to one. We shall characterize the possible fixed points written in block form $p = [p_0^*, p_1^*, \dots, p_K^*]^T$. First we note that \mathcal{M}_0 is itself an out-subgraph, and A is easily seen to be strictly substochastic with spectral radius strictly less than 1. Therefore the fixed point equation for p_0^* is a closed equation, that has a unique solution by part 1 of Lemma 1. Each of the remaining pieces of p^* , $p_k^*, k \geq 1$, is a solution of a piecewise linear equation in the following form:

$$p_k^* = \min(\bar{X}_k, \max(B_k \cdot p_0^* + \bar{Y}_k - \bar{D}_k + \Pi_k \cdot p_k^*, 0)) .$$

Now we note that each Π_k is an irreducible stochastic matrix, and by Part 2 of Lemma 1, p_k^* is unique if $\mathbf{1}_k \cdot (B_k \cdot p_0^* + \bar{Y}_k - \bar{D}_k) \neq 0$ and a point in a one-dimensional interval if $\mathbf{1}_k \cdot (B_k \cdot p_0^* + \bar{Y}_k - \bar{D}_k) = 0$. \square

2.2.1 Reduced Form E-N Cascade Mechanism

In models such as this, cascades of defaults arise when primary defaults trigger further losses to the remaining banks. Theorem 1 proves the existence of an “equilibrium” clearing vector, which is usually unique, that gives the end result of cascades in the E-N framework. Typically, however, different balance sheet specifications lead to identical cascades, and we can characterize the cascade mechanism and resultant clearing vectors in terms of a reduced set of balance sheet data. It turns out that the key information to track is something we call the *default buffer*, which extends the notion of equity. We define the initial default buffer $\Delta_v^{(0)}$ of bank v to be its nominal equity:

$$\Delta_v^{(0)} := \bar{E}_v = \bar{Y}_v + \sum_w \bar{\Omega}_{wv} - \bar{D}_v - \bar{X}_v. \quad (2.6)$$

As before, define $p_v^{(n)}$ to be the amount available to pay \bar{X}_v at the end of cascade step n , initialized at $n = 0$ to be $p_v^{(0)} = \bar{X}_v$. Introduce the normalized “threshold function” h that maps the extended real line $[-\infty, \infty]$ to the unit interval $[0, 1]$:

$$h(x) = (x + 1)^+ - x^+ = \max(0, \min(x + 1, 1)). \quad (2.7)$$

Then, as one can easily verify, the result of the n th step of E-N cascade is expressible as

$$\begin{cases} p_v^{(n)} = \bar{X}_v h(\Delta_v^{(n-1)} / \bar{X}_v) \\ q_v^{(n)} = ((\bar{Y}_v - \bar{D}_v)^- + \bar{X}_v) h(\Delta_v^{(n-1)} / ((\bar{Y}_v - \bar{D}_v)^- + \bar{X}_v)) \\ \Delta_v^{(n)} = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} (1 - p_w^{(n)} / \bar{X}_w) \\ \quad = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} (1 - h(\Delta_w^{(n-1)} / \bar{X}_w)). \end{cases} \quad (2.8)$$

The mark-to-market equity is the positive part of the default buffer, $E_v^{(n)} = (\Delta_v^{(n)})^+$, and default of bank v occurs at the first step that $\Delta_v^{(n)} \leq 0$. As $n \rightarrow \infty$, the monotone decreasing sequence $p^{(n)}$ converges to the maximal fixed point p^+ . We can now see clearly that in each step, the cascade mapping boils down to a vector-valued function $p^{(n-1)} \mapsto p^{(n)} = F(p^{(n-1)} | \bar{\Delta}, \bar{\Omega})$

$$F_v(p) = \bar{X}_v h \left(\bar{\Delta}_v / \bar{X}_v - \sum_w \bar{\Pi}_{wv} (1 - p_w^{(n-1)} / \bar{X}_w) \right)$$

that depends parametrically only on the initial equity buffers $\bar{\Delta}$ and the interbank exposures $\bar{\Omega}$. Thus, the resulting maximal clearing vector p^+ is a well-defined function of the reduced balance sheet data

$$p^+ = G^+(\bar{\Delta}, \bar{\Omega}).$$

If instead of starting the cascade at the initial value $p_v^{(0)} = \bar{X}_v$, we had begun with $p_v^{(0)} = 0$, we would obtain a monotone *increasing* sequence $p^{(n)}$ that converges to the minimal fixed point $p^- := G^-(\bar{\Delta}, \bar{\Omega})$.

The scaled variable Δ/\bar{X} (which if $\bar{X} = 0$ we can define to be $\text{sgn}(\Delta) \times \infty$), or alternatively $\Delta/((\bar{Y} - \bar{D})^- + \bar{X})$, has the interpretation of a bank's *distance-to-default*, and the threshold function h determines both the fractional recovery on interbank debt and on total debt when Δ is negative. Different threshold functions h are an important characteristic of different cascade models.

To determine the amount of external debt that bank v eventually repays, one obviously also needs the value of $(\bar{Y}_v - \bar{D}_v)$:

$$q_v^+ = ((\bar{Y}_v - \bar{D}_v)^- + \bar{X}_v) h(\Delta_v^+ / ((\bar{Y}_v - \bar{D}_v)^- + \bar{X}_v)); \quad \Delta_v^+ = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} (1 - p_w^+ / \bar{X}_w).$$

2.2.2 Clearing Algorithm

The previous algorithm, while financially natural, typically takes an infinite number of steps to converge to the maximal fixed point p^+ , even for finite N . The following algorithm resolves the E-N cascade to the fixed point p^* in at most $2N$ iterations by constructing an increasing sequence of defaulted banks $A^k \cup B^k, k = 0, 1, \dots$

1. **Step 0** Determine the primary defaults by writing a disjoint union $[N] := \{1, \dots, N\} = A^0 \cup B^0 \cup C^0$ where

$$\begin{aligned} A^0 &= \{v \mid Y_v + \bar{Z}_v - \bar{D}_v < 0\} \\ B^0 &= \{v \mid Y_v + \bar{Z}_v - \bar{D}_v - \bar{X}_v < 0\} \setminus A^0 \\ C^0 &= [N] \setminus (A^0 \cup B^0). \end{aligned}$$

2. **Step $k, k = 1, 2, \dots$** Solve, if possible, the $|B^{k-1}|$ dimensional system of equations⁴:

$$p_v = Y_v - \bar{D}_v + \sum_{w \in C^{k-1}} \bar{\Gamma}_{wv} \bar{X}_w + \sum_{w \in B^{k-1}} \bar{\Gamma}_{wv} p_w, \quad v \in B^{k-1}$$

and define result to be p^{k*} . Define a new decomposition

⁴ Remember that if B^{k-1} is an in-subgraph, the restriction to B^{k-1} of $I - \bar{\Gamma}$ is not invertible

$$\begin{aligned}
A^k &= A^{k-1} \cup \{v \in B^{k-1} | p_v^{k*} \leq 0\} \\
B^k &= (B^{k-1} \setminus A^k) \cup \{v \in C^{k-1} | Y_v - \bar{D}_v + \sum_{w \in C^{k-1}} \bar{\Pi}_{wv} \bar{X}_w + \sum_{w \in B^{k-1}} \bar{\Pi}_{wv} p_w^{k*} \leq \bar{X}_v\} \\
C^k &= [N] \setminus (A^k \cup B^k)
\end{aligned}$$

and correspondingly

$$p_v^k = \begin{cases} 0 & v \in A^k \\ Y_v + \sum_{w \in C^k} \bar{\Pi}_{wv} \bar{X}_w + \sum_{w \in B^k} \bar{\Pi}_{wv} p_w^{k*} - \bar{D}_v & v \in B^k \\ \bar{X}_v & v \in C^k. \end{cases} \quad (2.9)$$

If $A^k = A^{k-1}$ and $B^k = B^{k-1}$, then halt the algorithm and set $A^* = A^k, B^* = B^k, \mathbf{p}^* = \mathbf{p}^{k*}$. Otherwise perform step $k+1$.

2.3 The Gai-Kapadia 2010 Default Model

The threshold function $h(x)$ for the E-N 2001 model encodes a *soft* type of default in which the interbank debt of defaulted banks with $x \sim 0$ recover almost all their value. In their 2010 paper [38], Gai and Kapadia offer a model with very *hard defaults*: interbank debt on defaulted banks recover zero value. They justify their zero recovery assumption with the statement⁵: “This assumption is likely to be realistic in the midst of a crisis: in the immediate aftermath of a default, the recovery rate and the timing of recovery will be highly uncertain and banks’ funders are likely to assume the worst-case scenario.”

While the main results of their paper concern random financial networks, the underlying cascade mechanism is precisely of the E-N type, but with a *zero-recovery* threshold function

$$\tilde{h}(x) = \mathbf{1}(x \leq 0), \quad (2.10)$$

Exactly as in Section 2.1.2, the cascade can be defined by the sequence of vectors $\mathbf{p}^{(n)}$ and buffers $\Delta^{(n)}$ satisfying equations (2.6) and (2.8), with h replaced by \tilde{h} , that is,

$$p_v^{(n)} = \bar{X}_v \tilde{h}(\Delta_v^{(n-1)}), \quad (2.11)$$

$$\Delta_v^{(n)} = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} (1 - \tilde{h}(\Delta_w^{(n-1)})). \quad (2.12)$$

The clearing vector condition is simply

$$\mathbf{p} = \bar{\mathbf{X}} \mathbf{1}(\bar{\mathbf{Y}} + \bar{\mathbf{\Pi}}^T \cdot \mathbf{p} - \bar{\mathbf{D}} - \bar{\mathbf{X}} \geq 0),$$

or equivalently,

⁵ [38], footnote 9.

$$\Delta = \bar{\Delta} - \sum_w \bar{\Omega}_w \cdot (1 - \tilde{h}(\Delta_w)) .$$

Note that the second version of the clearing vector condition depends only on the reduced balance sheet information $(\bar{\Delta}, \bar{\Omega})$, which is a special simplifying feature of the zero recovery framework.

The uniqueness of clearing vectors \mathbf{p}^* has not been carefully studied for this model. Nonetheless, we conjecture that a result similar to Theorem 2 still holds, in which non-uniqueness arises only connected to irreducible in-graphs. Now, when there are one or more irreducible in-graphs, the multiplicity of clearing vectors is parametrized by discrete values of λ . For example, the $N = 4$ bank network with $\bar{\Omega}$ as in the previous examples, and with $\bar{Y} = [0, 0, 0, 4]$, $\bar{D} = [1, 1, 1, 0]$ has exactly two clearing vectors of the form $\mathbf{p}^* = \lambda [1, 1, 1, 0]^T$ with $\lambda = 0$ or $\lambda = 1$.

The zero-recovery assumption implies there is a large cost to the system at the time any bank defaults, which gives us a simple measure of *cascade impact*:

$$\text{CI} = \sum_v \bar{X}_v \mathbf{1}(v \in \mathcal{D}^*) . \quad (2.13)$$

We are also interested in *default correlation*, especially for counterparty pairs (w, v) where w owes to v .

Recall that bankruptcy charges are ruled out in the E-N 2001 model and are maximal in the GK model. We can interpolate between these two extreme cases with a single parameter $\tau \in [0, 1]$ that represents the fraction of interbank debts that are paid as bankruptcy charges at the time any bank defaults. The cascade mapping is again given by equations (2.6) and (2.8), now with h replaced by the interpolated threshold function $h^{(\tau)}(x) = (1 - \tau)h(\frac{x}{1-\tau}) + \tau\tilde{h}(x)$. It seems to be true that for any $\tau > 0$, the multiplicity of clearing vectors will coincide with those that arise in case $\tau = 1$.

2.4 Liquidity Cascades

We have just learned that default cascades can be characterized by shocks that are transmitted *downstream* from defaulting banks to the asset side of their creditor banks' balance sheets. Funding illiquidity cascades are a systemic phenomenon that occur when banks curtail lending to each other on a large scale. In such a cascade, shocks are naturally transmitted *upstream*, from creditor banks to their debtors. A fundamental treatment of the connection between funding liquidity and market liquidity can be found in [20], who propose an economic model that provides a picture of how the funding constraints a bank must satisfy impact the market portfolio, that is external assets, that the bank holds. In this paper, one finds the idea that when a bank's capital is reduced to below a threshold at which a funding liquidity constraint becomes binding, that bank will reduce its assets by a discontinuous amount, and experience a discontinuous increase in its margin and collateral requirements. It is natural to assume that at this threshold, the bank will also reduce its interbank lending

by a discontinuous amount. This picture provides the seed of a cascade mechanism that is transmitted through the interbank network, from creditors to debtors. In view of this, it should not be a surprise that our schematic default cascade models, when reversed and reinterpreted, can turn into models of funding illiquidity cascades.

To our knowledge, the first paper to introduce a schematic model of funding illiquidity cascades is a companion paper [37] to the default model by Gai-Kapadia in [38]. It, and a related model by Gai, Haldane and Kapadia [36] aim to provide a cascade mechanism of *liquidity hoarding* that can describe the observed freezing of liquidity that was a key feature of the period during the great financial crisis.

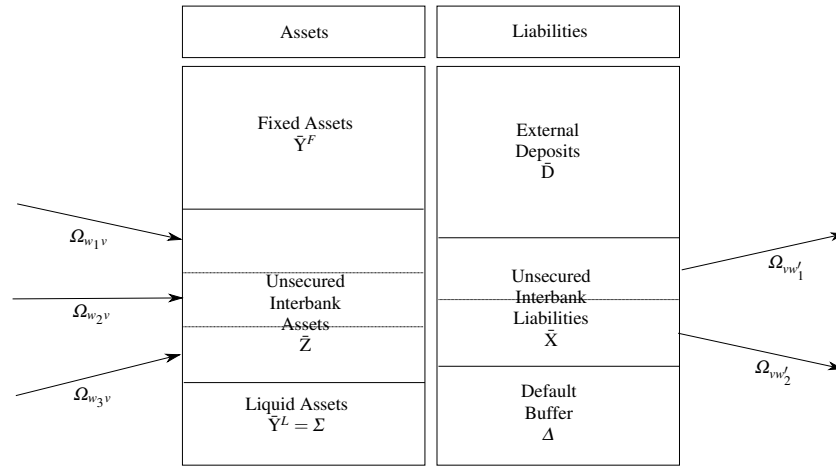


Fig. 2.2 The stylized balance sheet of a bank v with in-degree $j_v = 3$ and out-degree $k_v = 2$. Banks w_1, w_2, w_3 are debtors of v while w'_1, w'_2 are creditors.

2.4.1 Gai-Kapadia 2010 Liquidity Cascade Model

This systemic risk model aims to account for the observation that starting in August 2007 and continuing until after September 2008, interbank lending froze around the world as banks hoarded cash and curtailed lending. As Gai and Kapadia explain, during the build up of credit risk prior to 2007, some banks that held insufficiently liquid assets might have faced funding liquidity difficulties. Such banks would be expected to move to a more liquid position by hoarding liquidity, in particular by reducing their interbank lending almost entirely.

What would a counterparty bank of such a hoarding bank do? Of course they might seek funding elsewhere, but alternatively they might themselves elect to become liquidity hoarders by preemptively curtailing interbank lending. In such a way, a liquidity cascade can ensue.

In their model they make assumptions about a network where prior to the crisis, all banks hold assets and liabilities as shown in Figure 2.2. External fixed assets \bar{Y}_v^F correspond to the bank book of loans to the economy at large, interbank assets \bar{Z}_v are assumed to be short term unsecured loans to other banks while the liquid assets \bar{Y}_v^L correspond to cash or liquid securities such as treasury bonds. When positive, the liquid assets \bar{Y}_v^L are used as a *stress buffer* $\bar{\Sigma}_v$ from which to pay liabilities as they arise. In analogy to the default buffers, Σ_v can become negative: such a bank is called a *stressed bank*. On the liability side we have as before external debt \bar{D}_v , interbank debt \bar{X}_v and the equity or default buffer $\bar{\Delta}_v$. The nominal interbank exposure of v to w is represented by $\bar{\Omega}_{vw}$ and hence there are the constraints $\bar{Z}_v = \sum_w \bar{\Omega}_{vw}$ and $\bar{X}_v = \sum_w \bar{\Omega}_{vw}$.

- Assumption 3.** 1. At step 0 of the cascade, one or more banks experience funding liquidity shocks or *stress shocks* that make their *stress buffers* $\Sigma_v^{(0)} = \bar{\Sigma}_v \leq 0$ go negative.
2. Banks respond at the moment they become stressed by preemptively hoarding a fixed fraction λ of interbank lending. This sends a stress shock of magnitude $\lambda \bar{\Omega}_{vw}$ to each of the debtor banks w of v .
3. At each step $n \geq 0$ of the crisis, bank v pays any interbank liabilities $\lambda \bar{\Omega}_{vw}$ that have been recalled by newly stressed banks w .

These simple behavioural rules lead to a cascade mechanism (CM) that can be expressed entirely in the recursive updating of the stress buffers of all banks. Given the collection of stress buffers $(\Sigma_v^{(n-1)})$ at step $n - 1$ of the cascade, the updated stress buffers are given by

$$\Sigma_v^{(n)} = \bar{\Sigma}_v - \lambda \sum_w \bar{\Omega}_{vw} (1 - \tilde{h}(\Sigma_w^{(n-1)})) . \quad (2.14)$$

In this proposed cascade mechanism, the parameter λ represents the average strength of banks' collective stress response.

It should not be a surprise that (2.14) is identical in form to (2.12), but with shocks going upstream from creditors to debtors instead of downstream from debtor banks to creditors. The pair of models [38] and [37] by Gai and Kapadia is a first instance of a formal symmetry of financial networks under interchange of assets and liabilities.

2.4.2 The Liquidity Model of S. H. Lee 2013

As another illustration of how liquidity cascades are the mirror image of default cascades, we now show how a simple liquidity cascade model proposed by S. H. Lee [51] is formally identical to a version of the E-N cascade. This model is based on banks with balance sheets as shown Table 2.3.

To put this model into an E-N form while preserving the labelling of Lee's paper, we introduce a fictitious "source" bank labelled by 0 that borrows from but does not

Assets	external illiquid assets \bar{Y}_v^F	liquid assets \bar{Y}_v^L	liquid interbank assets $\sum_w \bar{\Omega}_{vw}$
Liabilities	external deposits \bar{D}_v	interbank debt $\sum_w \bar{\Omega}_{vw}$	equity \bar{E}_v

Table 2.3 The stylized balance sheet of the Lee liquidity model.

lend to other banks and arbitrarily define $\bar{Z}_0 = \bar{D}_0 = 0$. We now label the interbank exposures

$$\bar{\Omega}_{vw} = \begin{cases} b_{vw} & v, w \neq 0 \\ q_w & v = 0, w \neq 0 \\ 0 & w = 0. \end{cases}$$

As before, let $\bar{Z}_v = \sum_w \bar{\Omega}_{vw}$, $\bar{X}_v = \sum_w \bar{\Omega}_{vw}$ and identify $\bar{Y}_v^F = z_v$, $\bar{Y}_v^L = \bar{\Omega}_{0v}$, $\bar{D}_v = d_v$.

At time 0, each bank experiences deposit withdrawals (which we can call a liquidity shock) $\Delta d_v \geq 0$. These withdrawals are paid immediately by each bank v in order of seniority: first from the liquid interbank assets \bar{Z}_v (which includes $\bar{\Omega}_{0v}$) until these are depleted, and then by selling fixed external assets \bar{Y}_v^F . Let us define the notional liquidity buffer to be $\Sigma_v^{(0)} = \bar{\Sigma}_v = -\Delta d_v$, then at the n th step of the liquidity cascade each buffer $\Sigma_v^{(n)}$, which is the negative of bank v 's total liquidity needs $\ell_v^{(n)}$, will have accumulated shocks as follows

$$\Sigma_v^{(n)} = \bar{\Sigma}_v - \sum_w \bar{\Omega}_{vw} \left(1 - h(\Sigma_w^{(n-1)} / \bar{Z}_w) \right).$$

This equation reveals a formal identity with (2.8), provided \bar{Z} and \bar{X} are interchanged, and the exposures are reversed. However, in the Lee model, we begin with all buffers $\Sigma_v^{(0)} \leq 0$ except when $v = 0$. For completeness, at step n we can define $p_v^{(n)} = \bar{Z}_v h(\Sigma_v^{(n-1)} / \bar{Z}_v)$, the amount of interbank assets remaining. The amount of fixed assets sold after n steps is $\ell_v^{(n)} - \bar{Z}_v + p_v^{(n)} = -\Sigma_v^{(n)} - \bar{Z}_v + p_v^{(n)}$.

2.4.3 Generalized Liquidity Cascades

The liquidity cascade model of S.H. Lee supposes that deposit withdrawals are funded in equal proportion by interbank assets and liquid assets. A reasonable alternative picture is that each bank keeps a ‘‘first line’’ reserve of liquid external assets (or simply ‘‘cash’’) \bar{Y}^L to absorb liquidity shocks. We think of this as the stress buffer, labelled by Σ , to be kept positive during normal banking business. However, when the stress buffer becomes zero or negative, we will call the bank ‘‘stressed’’, and the bank then meets further withdrawals by liquidating first interbank assets \bar{Z} , and finally the illiquid fixed assets \bar{Y}^F .

As for the Lee model, we add a fictitious sink bank $v = 0$ to represent external agents that borrow amounts $\bar{\Omega}_{0v}$. In terms of liquidation priority, these external loans will be considered a component of a bank's interbank assets: $\bar{Z}_v = \sum_{w=0}^N \bar{\Omega}_{vw}$.

Let us suppose that just prior to a withdrawal shock that hits any or all of the banks, the banks' balance sheets are given by notional amounts $(\bar{Y}^F, \bar{Z}, \bar{Y}^L, \bar{D}, \bar{X}, \bar{E}, \bar{\Omega})$. At the onset of the liquidity crisis, all banks are hit by withdrawal shocks ΔD_v that reduce the initial stress buffers $\Sigma_v^{(0)} = \bar{Y}_v^L - \Delta D_v$ of at least some banks to below zero, making them stressed. Stressed banks then liquidate assets first from \bar{Z} , inflicting additional liquidity shocks to their debtor banks' liabilities, and then from \bar{Y}^F . A stressed bank that has depleted all of \bar{Z} will be called "illiquid", and must sell external fixed assets \bar{Y}^F in order to survive.

Let $p_v^{(n)}$ be the amount of bank v 's interbank assets remaining after n steps of the liquidity cascade, starting with $p_v^{(0)} = \bar{Z}_v$. Stressed and illiquid banks will be those with $p_v^{(n)} < \bar{Z}_v$ while normal, unstressed banks have $p_v^{(n)} = \bar{Z}_v$. We also define $\Sigma_v^{(n)}$ to be the stress buffer after n steps. Assuming each stressed bank liquidates exactly enough additional interbank assets at each step to meet the additional liquidity shocks, the update rule is

$$p_v^{(n)} = \max(0, \min(\bar{Z}_v, (D_v - \Delta D_v) - \bar{Y}^F + \sum_w \bar{\Omega}_{vw} (p_w^{(n-1)} / \bar{Z}_w)) . \quad (2.15)$$

We note that

$$\Sigma_v^{(n)} = \Sigma_v^{(0)} - \sum_w \bar{\Omega}_{vw} (1 - p_w^{(n-1)} / \bar{Z}_w) , \quad (2.16)$$

and that (2.15) can be written

$$p_v^{(n)} = \bar{Z}_v h(\Sigma_v^{(n-1)} / \bar{Z}_v)$$

with the threshold function h of (2.7) used before.

Comparison of these equations with (2.8) reveals that our model is now precisely equivalent to the extended E-N 2001 model, with the role of assets and liabilities, and stress and default buffers, interchanged: $\bar{Y}^F \leftrightarrow \bar{D}, \bar{Z} \leftrightarrow \bar{X}, \bar{Y}^L \leftrightarrow \bar{E}, \Delta \leftrightarrow \Sigma$. As we expect, we recover the Lee model simply by taking $\bar{Y}_v^L = 0$, which also has the effect of making all the banks initially stressed since the initial stress buffers are $\Sigma_v^{(0)} = -\Delta D_v \leq 0$.

2.5 Asset Fire Sales

Certainly one of the basic triggers of financial crises is when a major asset class held by many banks is suddenly beset by bad news, resulting in a major devaluation shock that hits these banks. In our language we could identify this as a situation in which the external assets Y_v held by many banks exhibit a sharp one-time decline. If this *asset correlation shock* is sufficient to cause the default of some banks, we face

the possibility of a pure default cascade of the same nature as we have described already.

Of a different nature are *asset fire sales*, the name given to the crisis effect in which banks under stress (of which there will be many during a crisis) react by selling external assets on a large scale. As described in detail in the 2005 paper by Cifuentes, Ferrucci and Shin [25], an asset fire sale creates a negative feedback loop in the financial network: large scale selling of an asset class by banks leads to strong downward pressure on the asset price, which leads to market-to-market losses by all banks holding that asset, to which they respond by selling this and other assets.

Of course, small and medium scale versions of such selling spirals are an everyday occurrence in financial markets, sometimes leading to an asset correlation shock. In the present context, we will focus on large scale selling spirals that form during and as a result of the crisis and are essential amplifiers of financial distress. Our aim in this section is to provide a simple stylized modelling framework that highlights the network cascade aspects of this mechanism.

2.5.1 Fire Sales of One Asset

The basic network picture of asset fire sales is most clearly explained by the CFS model of Cifuentes, Ferrucci and Shin [25]. The baseline CFS model consists of a network of N banks with balance sheets with the same components $(\bar{Y}^F, \bar{Z}, \bar{Y}^L, \bar{D}, \bar{X}, \bar{Q})$ as shown in Table 2.3 for the Lee 2013 model [51]. Since liquidity and solvency are both considered in this model, it can be regarded as a generalization of the Eisenberg-Noe model. In the one asset model, all banks hold their fixed assets in the same security, which we might view as the market portfolio. We set the initial price of the asset to be $\bar{p} = p^{(0)} = 1$, of which each bank v holds $s_v^{(0)} = \bar{Y}_v^F$ units.

The essential new feature is to include a *capital adequacy ratio* (CAR) as an additional regulatory constraint: for some fixed regulatory value r^* (say 7%), the bank must maintain the lower bound

$$\frac{\Delta_v}{Y_v^F + Y_v^L + Z_v} \geq r^* . \quad (2.17)$$

As soon as this condition is violated, one assumes that the bank is compelled to restore the condition by selling, first the liquid assets, then the fixed illiquid assets. In contrast to the Gai-Kapadia and Lee liquidity cascade models who explicitly assume that interbank assets are fully callable instantaneously and are sold before fixed assets, [25] assumes that “interbank loans normally cannot be expected to be recalled early in the event of default of the lender.”⁶ Based on this assumption, the solvency condition assumed in [25] differs as well: a bank must be liquidated in the

⁶ It is interesting that these modelling frameworks make essentially contradictory assumptions at this point. The assumption of [25] removes the need to consider how interbank assets are liquidated.

event that

$$\frac{\Delta_v}{Z_v} < r^* . \quad (2.18)$$

This inequality indicates that even the selling of all non-interbank assets is insufficient to restore the CAR condition, which implies the bank is insolvent and must be fully liquidated. Finally, [25] assume that in the event of insolvency the defaulted interbank assets are distributed at face value proportionally among the bank's creditors, and the bank ceases to function. The picture to have is that such a bank can never achieve the CAR condition, and hence must be terminated, even if its equity buffer may still be positive.

The remaining important economic assumption determines the impact on prices when assets are sold. When boiled down to its essentials, the assumption of [25] is of an inelastic supply curve and a downward sloping inverse demand function $d^{-1}(\cdot)$, such that the asset price is $p = d^{-1}(s)$ when $s = \sum_v s_v$ is the aggregated amount sold. For the inverse demand function, they work with the family of exponential functions $d^{-1}(s) = e^{-\alpha s}$ for a specific value of α .

We deviate from the original model in our treatment of cash: in line with the logic of the GK and Lee liquidity models, we consider cash as exactly equivalent to other liquid assets and remove both from the denominator of the CAR so that (2.17) is replaced by

$$\frac{\Delta_v}{Y_v^F + Z_v} \geq r^* . \quad (2.19)$$

Then Y^L has the interpretation of a liquidity buffer, and is used until depleted to pay liabilities before any fixed assets are sold. Finally, it is assumed that external deposits have equal seniority to interbank debt: one then introduces a non-borrowing fictitious bank $v = 0$ that lends $\bar{D}_w := \bar{Q}_{w0}$ to each bank w .

The crisis unfolds starting at step $n = 0$ from an initial balance sheet configuration $(\bar{Y}^F, \bar{Z}, \bar{Y}^L, \bar{X}, \bar{Q})$ with equity buffers

$$\Delta^{(0)} = \bar{\Delta} = \bar{Y}^F + \bar{Z} + \bar{Y}^L - \bar{X} \quad (2.20)$$

in which at least one bank is found to be in violation of its CAR bound. Recall we set $p^{(0)} = 1$ so the number of fixed assets is $s^{(0)} = \bar{Y}^F$. Then, recursively for $n = 1, 2, \dots$ the balance sheets of each bank $(Y^{F(n)}, Z^{(n)}, Y^{L(n)}, \Delta^{(n)})$ and the asset price $p^{(n)}$ are updated according to the following steps:

1. Each bank v adjusts its fixed asset holdings by selling⁷

$$\delta s_v = \min(s_v^{(n-1)}, \max(0, s_v^{(n-1)} + Z_v^{(n-1)} / p^{(n-1)} - \Delta_v^{(n-1)} / (r^* p^{(n-1)})) \quad (2.21)$$

units at the price $p^{(n-1)}$. Note that $\delta s_v = 0$ corresponds to a compliant bank, while $\delta s_v = s_v^{(n-1)}$ corresponds to an insolvent bank. When $\delta s_v > 0$, the sale of

⁷ This is a slight modification of [25] who assume these units are sold at an n dependent equilibrium price somewhat lower than $p^{(n-1)}$.

the fixed asset increases the bank's liquid assets to $Y_v^{L(n)} = Y_v^{L(n-1)} + \delta s_v p^{(n-1)}$.

We also remember to update the number of shares to $s_v^{(n)} = s_v^{(n-1)} - \delta s_v$.

2. In case $s_v^{(n)} = 0$, bank v is now insolvent, and must be liquidated in the manner described above and the mark-to-market value of its debt adjusted: $X_v^{(n)} = \bar{X}_v + \min(0, \Delta_v^{(n-1)})$.
3. After all banks have completed their asset sales, it is assumed that the market price moves downwards according to the aggregate amount sold so $p^{(n)} = d^{(-1)}(\sum_v (s_v^{(0)} - s_v^{(n)}))$ and that the interbank assets are updated to account for new default losses, $Z_v^{(n)} = \sum_w \bar{\Gamma}_{wv} X_w^{(n)}$.
4. Finally, corresponding to these changes, the updated equity buffer of bank v becomes

$$\Delta_v^{(n)} = s_v^{(n)} p^{(n)} + Y_v^{L(n)} + \sum_w \bar{\Gamma}_{wv} X_w^{(n)} - \bar{X}_v. \quad (2.22)$$

Just as the E-N framework could be simplified into a reduced form cascade mapping by focussing on the default buffers $\Delta_v^{(n)}$, it turns out the above recursion simplifies in a very similar way if we focus on the pairs $\Delta_v^{(n)}, s_v^{(n)}$. The key fact to recognize is that once the bank becomes noncompliant, it can never become compliant, nor can a defaulted bank recover. Having seen this, one can easily verify that the result of the n -th step of the CFS cascade is given by

$$\begin{cases} X_v^{(n)} = \bar{X}_v h(\Delta_v^{(n-1)} / \bar{X}_v) \\ \Delta_v^{(n)} = \bar{\Delta}_v - \sum_{m=1}^n (p^{(m-1)} - p^{(m)}) s_v^{(m)} - \sum_w \bar{\Omega}_{wv} (1 - h(\Delta_w^{(n-1)} / \bar{X}_w)) \\ s_v^{(n)} = \max(0, \min(s_v^{(0)}, \frac{1}{p^{(n-1)}} [\frac{\Delta_v^{(n-1)}}{r^*} - \sum_w \bar{\Omega}_{wv} h(\Delta_w^{(n-1)} / \bar{X}_w)]) \end{cases} \quad (2.23)$$

together with the equation for the price

$$p^{(n)} = d^{(-1)} \left(\sum_v (s_v^{(0)} - s_v^{(n)}) \right).$$

The third of these equations corresponds to the trichotomy of possibilities of the bank being compliant, noncompliant but solvent, and insolvent.

By comparison to the E-N cascade dynamics given by (2.8), we note that the key effect of the fire sale on the cascade is to provide additional price shocks that further reduce the default buffers, thereby amplifying the contagion step. In case the price impact is omitted by assuming $d^{(-1)}(\cdot) = 1$, one reproduces the E-N model.

An interesting special case of the model emerges if we set the compliancy parameter $r^* = 0$ and interpret the third equation of (2.23) to mean

$$s_v^{(n)} = s_v^{(0)} \mathbf{1}(\Delta_v^{(n-1)} > 0). \quad (2.24)$$

We then interpret the breach of this regulatory condition as the insolvency of the bank in two senses: first, the bank has insufficient assets to repay its liabilities; sec-

ond, the bank has insufficient equity to support its assets and thus needs to liquidate all fixed assets. Thus the bank must sell all fixed assets at this moment. However, as in the E-N model, the bank continues to lose money after it defaults, further eroding the value of $X_v^{(n)}$. Thus, when $r^* = 0$ this model looks very similar to the E-N 2001 model, albeit with a GK-like condition to determine the amount of fixed assets sold.

The simplification $r^* = 0$ also provides a simpler setting to address a different question: how do fire sales create contagion when banks hold different portfolios of a multiplicity of assets. As it turns out, this variation has been studied already, in a paper [21] we shall now describe.

2.5.2 Fire Sales of Many Assets

A model due to Caccioli et al [21] addresses exactly this question: how do fire sales create contagion when banks hold different portfolios of a multiplicity of assets? Their paper considers a variation of the CFS approach that takes $r^* = 0$ and M assets, and assumes moreover, that the interbank sector is set to zero: $Z_v = X_v = 0$. One can describe their model in terms of a bipartite graph with nodes of two colours, blue nodes that represent banks and red nodes that represent non-banking assets, and links connecting banks to assets when the bank has a significant holding of that asset. Figure 2.3 shows a typical network with 5 banks and 4 assets.




Fig. 2.3 A bipartite graph with 5 banks (blue nodes) co-owning 4 assets (red nodes).

In this section we consider a multi-asset fire sale model with bank balance sheets $(\bar{Y}^F, \bar{Y}^L, \bar{D})$ where the interbank sector is set to zero $\bar{Z}_v = \bar{X}_v = 0$. The fixed assets are portfolios of the M assets, and the CAR condition is

$$\frac{\Delta}{Y^F} \geq r^* \quad (2.25)$$

with $r^* \geq 0$. When $r^* = 0$, this reduces to the [21] model. The cascade picture that arises from these assumptions is that banks are forced to liquidate some fixed assets when the compliance condition (2.25) no longer holds, and all their fixed assets at the moment of their default. Forced asset sales drive down the prices of assets (in other words, causing a shock to be transmitted along each link from the defaulted bank to the assets it holds). Then, any other bank with a link to any of these depreciated assets will suffer a mark-to-market loss transmitted along that link. If after such losses some banks are still found to be non-compliant or insolvent, the cascade continues to the next step.

We label the assets by $a \in \mathcal{M} := \{1, 2, \dots, M\} := [M]$ and let the total initial market capitalization of asset a be labelled \bar{A}_a . Without loss of generality, we can assume the initial share prices are $\bar{p}_a = 1$. The fixed (illiquid) external assets of bank v can then be written

$$\bar{Y}_v^F = \sum_a \bar{s}_{av} . \quad (2.26)$$

where \bar{s}_{av} denotes the initial number of shares of asset a owned by bank v .

If after n cascade steps each bank v holds $s_{av}^{(n)}$ shares of asset a , the asset share price will have been driven down to the value

$$p_a(n) = d_a^{(-1)} (\bar{A}_a^{-1} \sum_v (s_{av}^{(0)} - s_{av}^{(n)})) \quad (2.27)$$

determined by the total amount of selling and the inverse demand function.

It is necessary to specify a portfolio selection rule that determines in which proportion banks choose to sell (and buy) assets. We make a very natural assumption that banks follow a *fixed-mix strategy* that keeps the dollar amounts in different assets in proportion to the initial ratios

$$\bar{\phi}_{av} := \frac{\bar{s}_{av}}{\bar{Y}_v^F} \quad (2.28)$$

at all times. Just as in the single asset model, banks also manage the total fixed assets $Y_v^{F(n)}$ close to a target $\tilde{Y}_v^{F(n)}$ which depends on the state of the bank. Thus each cascade step n , banks behave depending on whether they are compliant, non-compliant but solvent, or newly insolvent:

- Compliant banks rebalance their portfolios, with the target that assumes the prices have not changed, that is $\tilde{Y}_v^{F(n)} = \sum_a s_{av}^{(n-1)} p_a^{(n-1)}$ with the updated fixed-mix ratios $s_{av}^{(n)} = \bar{\phi}_{av} \tilde{Y}_v^{F(n)} / p_a^{(n-1)}$. Note that the fixed-mix strategy requires buying assets, i.e. $s_{av}^{(n)} > s_{av}^{(n-1)}$ that have just dropped in price, and selling assets that have just gained.

- Non-compliant solvent banks have the target $\tilde{Y}_v^{F(n)} = r^{-1} \Delta_v^{(n-1)}$ with the updated fixed-mix ratios $s_{av}^{(n)} = \bar{\phi}_{av} \tilde{Y}_v^{F(n)} / p_a^{(n-1)}$. This amounts to a net selling of assets, although since rebalancing still happens, some assets may be bought.
- Newly insolvent banks sell all assets, so $s_{av}^{(n)} = 0$.

Based on these assumptions, recursively for $n = 1, 2, \dots$ the balance sheets of each bank $(Y^{F(n)}, Y^{L(n)}, \Delta^{(n)})$, the portfolio allocations $s_{av}^{(n)}$ and asset prices $p_a^{(n)}$ are updated according to the following steps, starting from the initial values $Y^{F(0)} = \tilde{Y}^F, Y^{L(0)} = \tilde{Y}^L, \Delta^{(0)} = \tilde{Y}^F + \tilde{Y}^L - \bar{D}$:

1. The bank adjusts its asset holdings by picking the target

$$\tilde{Y}_v^{F(n)} = \min(Y_v^{F(n-1)}, \max(0, \Delta_v^{(n-1)} / r^*)) \quad (2.29)$$

and the new share holdings

$$s_{av}^{(n)} = \bar{\phi}_{av} \tilde{Y}_v^{F(n)} / p_a^{(n-1)}. \quad (2.30)$$

The liquid assets increase:

$$Y_v^{L(n)} = Y_v^{L(n-1)} + \sum_a (s_{av}^{(n-1)} - s_{av}^{(n)}) p_a^{(n-1)}. \quad (2.31)$$

2. After all banks have completed their asset sales, it is assumed that the market prices move downwards according to the aggregate fraction sold, so $p_a^{(n)} = d_a^{(-1)} (\bar{A}_a^{-1} \sum_v (s_{av}^{(0)} - s_{av}^{(n)}))$.
3. Finally, corresponding to these price changes, the updated equity buffer and fixed assets of bank v decrease:

$$\Delta_v^{(n)} = \Delta_v^{(n-1)} - \sum_a s_{av}^{(n)} (p_a^{(n-1)} - p_a^{(n)}), \quad (2.32)$$

$$Y^{F(n)} = \sum_a s_{av}^{(n)} p_a^{(n)}. \quad (2.33)$$

We observe the formal similarity between this model and the original E-N 2001 model, in the sense that ultimately the cascade dynamics is Markovian in the buffer variables $\Delta_v^{(n)}$ attached to blue nodes and asset prices $p_a^{(n)}$ attached to the red nodes. In addition, the dynamics depends explicitly on the initial balance sheets only through the fixed-mix ratios $\bar{\phi}_{av}$, analogous to the edge weights $\bar{\Omega}_{wv}$. The inverse demand functions d_a^{-1} play a role similar to the h, \bar{h} functions in our default cascade models. Therefore the asset prices behave as if they were ‘‘asset buffers’’. One can easily verify that the mapping from step $n - 1$ to n is monotonic in the collection of buffer variables Δ_v, p_a and bounded below, and thus its iterates converge to a fixed point.

2.6 Random Financial Networks

This chapter has explored stylistic features of various types of shocks that can be transmitted through the interbank exposures (or, in one case, through bank-to-asset exposures), and how they might potentially cascade into large scale disruptions of the sort seen during the 2007-08 financial crisis. We have seen how to build these shock channels into a variety of different financial network models.

The real world financial systems in most countries are of course far from behaving like these models. Bank balance sheets are hugely complex. Interbank exposure data are never publicly available, and in many countries nonexistent even for central regulators. Sometimes, the only way to infer exposures is indirectly, for example, through bank payment system data as done in [35]. Interbank exposures are of a diversity of types and known to change rapidly day to day. In a large jurisdiction like the US, the banking sector is highly heterogeneous, and the systemic impact due to the idiosyncrasies of individual banks will likely overwhelm anything one might predict from their average properties.

Nonetheless, a large and rapidly growing web of economics and complex systems research continues to address the real world implications of theoretical network cascades. The conceptual tools that we will explore in the remainder of this book come partly from the experience gained by modelling large complex networks that arise in other areas of science.

The central theme of this book is that something fundamental about financial systems can be learned by studying very large stochastic networks. There are at least three important reasons why stochastic network models are a good way to approach studying real world financial systems. The first and most fundamental reason comes from over a century of statistical mechanics theory, which has discovered that the macroscopic properties of matter, for example crystals, are determined by the ensemble averages of the deterministic dynamics of constituent microscopic particles. Even a completely known deterministic system, if it is large enough, can be well described by the average properties of the system. From this fact we can expect that for large N , an E- N model with fully specified parameters will behave as if it were a stochastic model with averaged characteristics.

The second important reason is that in a concrete sense, the true networks are to be thought of as stochastic at any moment in time. The balance sheets of banks, between reporting dates, are not observed even in principle. Moreover, they change so quickly that last week's values, if they were known, will have only limited correlation with this week's values. This fact is especially true for the interbank exposures that provide the main channels for transmitting network cascades: even in jurisdictions where they are reported to central regulators, they comprise a diversity of different securities, including derivatives and swaps whose mark-to-market valuations fluctuate dramatically on intraday time-scales.

A third important reason is that a hypothetical financial system, with all balance sheets completely known, will be hit constantly by random shocks from the outside, stochastic world. A deterministic system, subjected to a generic random shock, becomes after only one cascade step a fully stochastic system.

For all these reasons, and more, the next chapters will consider the stochastic nature of financial networks, and the extent to which the large scale properties of cascades might be predictable from models of their stochastic properties. From now on in this book, our various contagion channels will usually take their dynamics within the framework of “random financial networks”, defined provisionally as follows:

Definition 4. A *random financial network* or RFN is a random object representing a possible state of the financial network at an instant in time. It consists of three layers of mathematical structure. At the base structural level, the *skeleton* is a random directed graph $(\mathcal{N}, \mathcal{E})$ whose nodes \mathcal{N} represent “banks” and whose edges or links represent the presence of a non-negligible “interbank exposure” between a debtor bank and its creditor bank. Conditioned on a realization of the skeleton, the second structural layer is a collection of random *balance sheets*, one for each bank. In our simple models this is usually a coarse-grained description, listing for example the amounts $(\bar{Y}_v, \bar{Z}_v, \bar{D}_v, \bar{X}_v)$ for each $v \in \mathcal{N}$ as in Section 2.1. Finally, conditioned on a realization of the skeleton and balance sheets, the third level is a collection of positive random *exposures* $\bar{\Omega}_\ell$ for each link $\ell = (w, v) \in \mathcal{E}$. The interbank assets and liabilities are constrained to equal the aggregated exposures:

$$\bar{Z}_v = \sum_w \bar{\Omega}_{wv}, \quad \bar{X}_v = \sum_w \bar{\Omega}_{vw}. \quad (2.34)$$

Typically in cascade scenarios, we consider the RFN at the instant that a crisis triggering event unexpectedly occurs. We will combine the choice of RFN with the choice of a cascade mechanism such as the ones described in this chapter to describe what happens next. Depending on the cascade mechanism, only reduced balance sheet information in the form of buffer random variables $\bar{\Delta}_v$ and exposures $\bar{\Omega}_\ell$ is needed to follow the progress of the cascade. In that case, and we can work with a minimal parametrization of the RFN by the quadruple $(\mathcal{N}, \mathcal{E}, \bar{\Delta}, \bar{\Omega})$.

This schematic definition will prove to be acceptable for the description of simple contagion models. But more than that, it scales conceptually to much more complex settings. Nodes may have additional attributes or “types” beyond their connectivity to represent a more diverse class of economic entities. Links might have extended meaning where the random variables $\bar{\Omega}_\ell$ take vector values which represent different categories of exposures. We also recognize that even a very simple RFN is a complicated random variable of enormous dimension. Before proceeding to any analysis of network dynamics, the distributions of these collections of random variables must be fully specified. We will proceed stepwise, first focussing in the next chapter on characterizing possible models for the skeleton. In subsequent chapters we will consider how to specify the random structure of balance sheets and interbank exposures. Our optimistic view that something meaningful can be learned about systemic risk in the real world through the study of schematic or stylized RFNs is derived from our collective experience in other fields of complex disordered stochastic systems, rooted in the theory of statistical mechanics.

Chapter 3

Random Graph Models

Abstract The network of interbank counterparty relationships, or “skeleton”, can be described as a random graph that acts as the medium through which financial contagion is propagated. The basic properties are developed for a number of promising families of random graph constructions including the configuration models, preferential attachment models, and the inhomogeneous random graphs. A new extension, called the assortative configuration model, is proposed. The main results of this chapter are theorems describing the large graph asymptotics of this new assortative configuration model, including a proof of the locally tree-like property. Finally, measures of network structure are briefly surveyed.

Keywords: Skeleton, counterparty network, configuration graph, preferential attachment, inhomogeneous random graph, assortativity, random graph simulation, large graph asymptotic, network topology, locally tree-like.

The “skeleton” of a random financial network or RFN at a moment in time is most simply thought of as the directed random graph that indicates which pairs of banks are deemed to have a significant counterparty relationship, where the arrow on each link points from debtor to creditor. Later on, we will consider alternative, more complex, specifications of skeletons, such as to replace directed links by two-dimensional undirected links whose two components represent the directed exposures in either direction. In this chapter we survey what the random graph literature can tell us about models that will be useful in describing the skeleton. We begin with definitions, then a description of four rich and useful families of random graphs.

Random graph theory is a general framework mathematicians and physicists have developed to capture the most salient and basic attributes of what have been come to be called *complex adapted systems*. Random graphs form the lowest structural layer of networks: general categories of “networks” can be built by adding further layers of structure on top of random graphs. The double adjective “complex adapted” has taken on a higher meaning in recent years to describe the nature of a system whose totality is “more than the sum of its parts”. In other words, complex adaptivity refers to the property that key responses and behaviours of a system emerge

in a way that cannot be anticipated from the microscopic interactions, in analogy to the way consciousness is presumed to emerge from the collective interaction of the brain's myriad neurons. The list of categories of complex adapted systems in the real world is ever-growing, and questions about such systems reach beyond our depths of understanding. The theme of the present book is to explore the profound “complex adapted” nature of the world's financial systems, using the tools of network science, as it extends random graph theory.

3.1 Definitions and Basic Results

In this section, we provide some standard graph theoretic definitions and develop an efficient notation for what will follow. Since in this book we are most interested in directed graphs rather than undirected graphs, our definitions are in that setting and the term “graph” will have that meaning. Undirected graphs fit in easily as a subcategory of the directed case.

- Definition 5. 1.** For any $N \geq 1$, the collection of *directed graphs* on N nodes is denoted $\mathcal{G}(N)$. We consider that the set of *nodes* \mathcal{N} is numbered by integers, i.e. $\mathcal{N} = \{1, \dots, N\} := [N]$. Then $g \in \mathcal{G}(N)$, a graph on N nodes, is a pair $(\mathcal{N}, \mathcal{E})$ where the set of edges is a subset $\mathcal{E} \subset \mathcal{N} \times \mathcal{N}$ and each element $\ell \in \mathcal{E}$ is an ordered pair $\ell = (v, w)$ called an *edge* or *link*. We often label links by integers $\ell \in \{1, \dots, E\} := [E]$ where $E = |\mathcal{E}|$. Normally, “self-edges” with $v = w$ are excluded from \mathcal{E} , that is, $\mathcal{E} \subset \mathcal{N} \times \mathcal{N} \setminus \text{diag}$.
2. A given graph $g \in \mathcal{G}(N)$ can be represented by its $N \times N$ *adjacency matrix* $M(g)$ with components

$$M_{vw}(g) = \begin{cases} 1 & \text{if } (v, w) \in g \\ 0 & \text{if } (v, w) \in \mathcal{N} \times \mathcal{N} \setminus g \end{cases} .$$

3. The *in-degree* $\deg^-(v)$ and *out-degree* $\deg^+(v)$ of a node v are

$$\deg^-(v) = \sum_w M_{wv}(g), \quad \deg^+(v) = \sum_w M_{vw}(g) .$$

4. A node $v \in \mathcal{N}$ has *node type* (j, k) if its in-degree is $\deg^-(v) = j$ and its out-degree is $\deg^+(v) = k$; the node set partitions into node types, $\mathcal{N} = \cup_{jk} \mathcal{N}_{jk}$. We shall write $k_v = k, j_v = j$ for any $v \in \mathcal{N}_{jk}$ and allow degrees to be any non-negative integer.
5. An edge $\ell = (v, w) \in \mathcal{E} = \cup_{kj} \mathcal{E}_{kj}$ is said to have *edge type* (k, j) with in-degree j and out-degree k if it is an out-edge of a node v with out-degree $k_v = k$ and an in-edge of a node w with in-degree $j_w = j$. We shall write $\deg^+(\ell) = k_\ell = k$ and $\deg^-(\ell) = j_\ell = j$ whenever $\ell \in \mathcal{E}_{kj}$.
6. For completeness, we define an undirected graph to be any directed graph g for which $M(g)$ is symmetric.

Remark 4. We use the term “type” for nodes and edges to denote characteristics that potentially extend beyond degree characteristics to include other characteristics such as the size, kind and location of a bank, or type of security. In this book, we narrow the definition of edge-type of $\ell = (v, w)$ to the pair (k_v, j_w) rather than the quadruple $(j_v, k_v; j_w, k_w)$ which may seem more natural: this is a choice that can be relaxed without difficulty and is made purely to reduce model complexity.

The standard visualization of a graph g on N nodes is to plot nodes as “dots” with labels $v \in \mathcal{N}$, and any edge (v, w) as an arrow pointing “downstream” from node v to node w . In our systemic risk application, such an arrow signifies that bank v is a debtor of bank w and the in-degree $\deg^-(w)$ is the number of banks in debt to w , in other words the existence of the edge (v, w) means “ v owes w ”. Figure 3.1 illustrates the labelling of types of nodes and edges.




Fig. 3.1 A type (3,2) debtor bank that owes to a type (3,4) creditor bank through a type (2,3) link.

There are obviously constraints on the collections of node type $(j_v, k_v)_{v \in \mathcal{N}}$ and edge type $(k_\ell, j_\ell)_{\ell \in \mathcal{E}}$ if they derive from a graph. If we compute the total number of edges $E = |\mathcal{E}|$, the number of edges with $k_\ell = k$ and the number of edges with $j_\ell = j$ we find three conditions:

$$\begin{aligned}
 E &:= |\mathcal{E}| = \sum_v k_v = \sum_v j_v \\
 e_k^+ &:= |\mathcal{E} \cap \{k_\ell = k\}| = \sum_\ell \mathbf{1}(k_\ell = k) = \sum_v k \mathbf{1}(k_v = k) \\
 e_j^- &:= |\mathcal{E} \cap \{j_\ell = j\}| = \sum_\ell \mathbf{1}(j_\ell = j) = \sum_v j \mathbf{1}(j_v = j). \quad (3.1)
 \end{aligned}$$

It is useful to define some further graph theoretic objects and notation in terms of the adjacency matrix $M(g)$:

1. The *in-neighbourhood* of a node v is the set $\mathcal{N}_v^- := \{w \in \mathcal{N} | M_{wv}(g) = 1\}$ and the *out-neighbourhood* of v is the set $\mathcal{N}_v^+ := \{w \in \mathcal{N} | M_{vw}(g) = 1\}$.

2. We write \mathcal{E}_v^+ (or \mathcal{E}_v^-) for the set of out-edges (respectively, in-edges) of a given node v and v_ℓ^+ (or v_ℓ^-) for the node for which ℓ is an out-edge (respectively, in-edge).
3. Similarly, we have second-order neighbourhoods $\mathcal{N}_v^{--}, \mathcal{N}_v^{-+}, \mathcal{N}_v^{+-}, \mathcal{N}_v^{++}$ with the obvious definitions. Second and higher order neighbours can be determined directly from the powers of M and M^\top . For example, $w \in \mathcal{N}_v^{-+}$ whenever $(M^\top M)_{wv} \geq 1$.
4. We will always write j, j', j'', j_1 , etc. to refer to in-degrees while k, k', k'', k_1 , etc. refer to out-degrees.

Our financial network models typically have a sparse adjacency matrix $M(g)$ when N is large, meaning that the number of edges is a small fraction of the $N(N-1)$ potential edges. This reflects the fact that bank counterparty relationships are expensive to build and maintain, and thus \mathcal{N}_v^+ and \mathcal{N}_v^- typically contain relatively few nodes even in a very large network.

3.1.1 Random Graphs

Random graphs are simply probability distributions on the sets $\mathcal{G}(N)$:

Definition 6. 1. A random graph of size N is a probability distribution \mathbb{P} on the finite set $\mathcal{G}(N)$. When the size N is itself random, the probability distribution \mathbb{P} is on the countable infinite set $\mathcal{G} := \cup_N \mathcal{G}(N)$. Normally, we also suppose that \mathbb{P} is invariant under permutations of the N node labels.

2. For random graphs, we define the *node-type distribution* to have probabilities $P_{jk} = \mathbb{P}[v \in \mathcal{N}_{jk}]$ and the *edge-type distribution* to have probabilities $Q_{kj} = \mathbb{P}[\ell \in \mathcal{E}_{kj}]$.

P and Q can be viewed as bivariate distributions on the natural numbers, with marginals $P_k^+ = \sum_j P_{jk}, P_j^- = \sum_k P_{jk}$ and $Q_k^+ = \sum_j Q_{kj}, Q_j^- = \sum_k Q_{kj}$. Edge and node type distributions cannot be chosen independently however, but must be consistent with the fact that they derive from actual graphs. The following additional assumptions impose that equations (3.1) hold in expectation, and convenient finite moment conditions.

Assumption 4. The P distribution has finite first and second moments

$$z := \sum_{jk} j P_{jk} = \sum_{jk} k P_{jk}; \quad \sigma_K^2 := \sum_k k^2 P_k^+ - z^2 < \infty; \quad \sigma_J^2 := \sum_j j^2 P_j^- - z^2 < \infty \quad (3.2)$$

and Q is “consistent” with P :

$$Q_k^+ = k P_k^+ / z, \quad Q_j^- = j P_j^- / z \quad \forall k, j. \quad (3.3)$$

A number of random graph construction algorithms have been proposed in the literature, motivated by the desire to create families of graphs that match the types and measures of network topology that have been observed in nature and society. The remainder of this chapter reviews the properties of the random graph constructions that seem most closely related to the types of networks observed in financial systems. The textbook “Random Graphs and Complex Networks” by van der Hofstad [71] provides a much more complete and up-to-date review of the entire subject.

In the analysis to follow, asymptotic results are typically expressed in terms of convergence of random variables in probability, defined as:

Definition 7. A sequence $\{X_n\}_{n \geq 1}$ of random variables is said to *converge in probability* to a random variable X , written $\lim_{n \rightarrow \infty} X_n \stackrel{P}{=} X$ or $X_n \xrightarrow{P} X$, if for any $\varepsilon > 0$

$$\mathbb{P}[|X_n - X| > \varepsilon] \rightarrow 0.$$

We also recall further standard notation for asymptotics of sequences of real numbers $\{x_n\}_{n \geq 1}$, $\{y_n\}_{n \geq 1}$ and random variables $\{X_n\}_{n \geq 1}$:

1. Landau’s “little oh”: $x_n = o(1)$ means $x_n \rightarrow 0$; $x_n = o(y_n)$ means $x_n/y_n = o(1)$;
2. Landau’s “big oh”: $x_n = O(y_n)$ means there is $N > 0$ such that x_n/y_n is bounded for $n \geq N$;
3. $x_n \sim y_n$ means $x_n/y_n \rightarrow 1$;
4. $X_n = o(y_n)$ means $X_n/y_n \xrightarrow{P} 0$.

3.2 Configuration Random Graphs

In their classic paper [33], Erdős and Renyi introduced the undirected model $G(N, M)$ that consists of N nodes and a random subset of exactly M edges chosen uniformly from the collection of $\binom{N}{M}$ possible such edge subsets. This model can be regarded as the M th step of a random graph process that starts with N nodes and no edges, and adds edges one at a time selected uniformly randomly from the set of available undirected edges. Gilbert’s random graph model $G(N, p)$, which takes N nodes and selects each possible edge independently with probability $p = z/(N-1)$, has mean degree z and similar large N asymptotics provided $M = zN/2$. In fact, it was proved by [14] and [58] that the undirected Erdős-Renyi graph $G(N, zN/2)$ and $G(N, p_N)$ with probability $p_N = z/(N-1)$ both converge in probability to the same model as $N \rightarrow \infty$ for all $z \in \mathbb{R}_+$. Because of their popularity, the two models $G(N, p) \sim G(N, zN/2)$ have come to be known as “the” random graph. Since the degree distribution of $G(N, p)$ is $\text{Bin}(N-1, p) \sim_{N \rightarrow \infty} \text{Pois}(z)$, this is also called the *Poisson graph* model.

Both the above constructions have obvious directed graph analogues: henceforth we use the notation $G(N, M)$ and $G(N, p)$ to denote the directed graph models. In the directed Gilbert $G(N, p)$ model, each possible directed edge selection is an in-

dependent Bernoulli trial, and thus the adjacency matrix $M(g)$ is extremely easy to simulate in Matlab:

```
M=( rand(N,N) < p );
diag(M)=0;
```

Another undirected graph construction of interest is the random r -regular model with $r \geq 3$, that draws uniformly from the set of r -regular graphs on N nodes, that is, graphs for which each node has degree exactly r . This model is a particular case of a general class called “undirected configuration graphs” which takes as data an arbitrary degree distribution $P = \{P_k\}$; similarly “directed configuration graphs” take as data an arbitrary bi-degree distribution $P = \{P_{jk}\}$. When $P \sim \text{Pois}(z) \times \text{Pois}(z)$, we have an example of the configuration graph, that turns out to be asymptotic for large N to both the directed Erdős-Renyi and Gilbert models.

The well known directed configuration multigraph model introduced by Bollobas [13] with general degree distribution $P = \{P_{jk}\}_{j,k=0,1,\dots}$ and size N is constructed by the following random algorithm:

1. Draw a sequence of N node-degree pairs $(j_1, k_1), \dots, (j_N, k_N)$ independently from P , and accept the draw if and only if it is feasible, i.e. $\sum_{n \in [N]} (j_n - k_n) = 0$. Label the n th node with k_n *out-stubs* (picture this as a half-edge with an out-arrow) and j_n *in-stubs*.
2. While there remain available unpaired stubs, select (according to any rule, whether random or deterministic) any unpaired out-stub and pair it with an in-stub selected uniformly amongst unpaired in-stubs. Each resulting pair of stubs is a directed edge of the multigraph.

The algorithm leads to objects with self-loops and multiple edges, which are usually called multigraphs rather than graphs. Only multigraphs that are free of self-loops and multiple edges, a condition called “simple”, are considered to be graphs. For the most part, we do not care over much about the distinction, because we will find that the density of self-loops and multiple edges goes to zero as $N \rightarrow \infty$. In fact, Janson [46] has proved in the undirected case that the probability for a multigraph to be simple is bounded away from zero for well-behaved sequences $(g_N)_{N>0}$ of size N graphs with given P .

Exact simulation of the adjacency matrix in the configuration model with general P is problematic because the feasibility condition met in the first step occurs only with asymptotic frequency $\sim \frac{\sigma}{\sqrt{2\pi N}}$, which is vanishingly small for large graphs. For this reason, practical Monte Carlo implementations use some type of rewiring or *clipping* to adjust each infeasible draw of node-degree pairs. We shall return to address this issue in detail in Section 3.5.

Because of the uniformity of the matching in step 2 of the construction, the edge-type distribution of the resultant random graph is

$$Q_{kj} = \frac{jkP_k^+P_j^-}{z^2} = Q_k^+Q_j^- \quad (3.4)$$

which we call the *independent edge condition*. For many reasons, financial and otherwise, we are interested in the more general situation when (3.4) is not true, called *assortativity*. It turns out that we have something new to say about how to construct assortative configuration graphs, which we will postpone until Section 3.5.

We will find that this extended class of assortative configuration graphs encompasses all reasonable type distributions (P, Q) and that it has special properties that make it suitable for exact analytical results. However, given that (P, Q) is freely specifiable, the model offers no explanation or justification for any particular choice of type distribution, nor any understanding of the dynamical formation of the network. On the other hand, real world networks such as the internet, world wide web and some would say financial networks, are often observed to have degree distributions consistent with Pareto (power law) tails, which means they are dominated by highly connected “hubs”. As first proposed by Barabasi and Albert [9], networks with such a structure arise from a natural *preferential attachment* (PA) property of their growth dynamics. Under the analogy with the sandpile model, while configuration graphs might be adequate to describe a critical network with fat-tailed type distribution, PA models give a hint how it managed to get into this state. The next class of random graph models shows how natural rules for growing networks lead to Pareto degree distributions.

3.3 Preferential Attachment and Detachment

In 1999, Barabási and Albert [9] introduced the first in a wide class of models of network growth, now known as *preferential attachment models* (PA), that have the property that their degree distributions have power law tails, and are thus (at least approximately) *scale free*. The important feature of this model class is that it provides an explanation of power law tails and hub-dominated networks through the plausible hypothesis that naturally occurring networks grow in size by adding nodes and edges that connect preferentially to nodes of high connectivity.

3.3.1 Preferential Attachment Models

We describe here a class of scale-free *directed* graphs first considered in [15] who compute the Pareto exponents for the marginal degree distributions, with more details due to Hadrien de March and Simon Leguil. The PA network growth process starts at time t_0 with an initial graph G_{t_0} containing t_0 edges, and evolves step by step according to three types of moves so that at step t , the current graph $G(t) = (\mathcal{N}(t), \mathcal{E}(t))$ contains t edges and a random number $N(t)$ of nodes. These three move types follow *preferential attachment* rules:

1. **α -move**: With probability α , a new node v is created along with an edge pointing from v to an existing node w .

2. **β -move**: With probability β , a directed edge is added between two existing nodes v and w .
3. **γ -move**: With probability γ , a new node v is created along with an edge going from an existing node w to v .

In each of these moves, existing nodes are *preferentially* chosen. Letting $j_{\tilde{v}}$ and $k_{\tilde{v}}$ denote the in-degree and out-degree of a potential node \tilde{v} , then for two constants ζ^- and ζ^+ :

1. The potential node \tilde{v} of an α or β move will be selected as the existing in-node w with probability

$$\mathbb{P}(w = \tilde{v}) = \frac{j_{\tilde{v}} + \zeta^-}{\sum_{v \in \mathcal{N}} (j_v + \zeta^-)} = \frac{j_{\tilde{v}} + \zeta^-}{t + \zeta^- N(t)} .$$

2. The potential node \tilde{v} of a γ or β move will be selected as the existing out-node v with probability

$$\mathbb{P}(v = \tilde{v}) = \frac{k_{\tilde{v}} + \zeta^+}{\sum_{v \in \mathcal{N}} (k_v + \zeta^+)} = \frac{k_{\tilde{v}} + \zeta^+}{t + \zeta^+ N(t)} .$$

The four independent parameters α , β , ζ^- and ζ^+ will be chosen subject to some particular conditions:

- $\alpha + \gamma = 1 - \beta$, controls the growth of the network and the mean degree z , since at time t , $\mathbb{E}[N(t)] = N(t_0) + (\alpha + \gamma)(t - t_0)$ and for times $t \rightarrow \infty$

$$\mathbb{E} \left[\frac{t}{N(t)} \right] \xrightarrow{\text{P}} z := \frac{1}{\alpha + \gamma} . \quad (3.5)$$

- As [15] first showed, for given α , β , $\gamma = 1 - \alpha - \beta$, the parameters ζ^- and ζ^+ determine the Pareto exponents of the marginal in- and out-degree distributions as well as the in-degree (out-degree) distribution conditioned on k (respectively j).

The next Proposition and Corollary proved by Hadrien de March and Simon Leguil¹ extend the original results of [15] by using a general technique based on the Wormald Theorem [73]. For each value of n , $\{G^{(n)}(t)\}_{t=n}^{\infty}$, $G^{(n)}(t) = (\mathcal{N}^{(n)}(t), \mathcal{E}^{(n)}(t))$ will be a directed random graph sequence constructed by the preferential attachment rules starting from an initial random graph $G^{(n)}(t_0)$ set at $t_0 = n$. For each $t \geq n$

$$X_k^+(t) := \sum_{v \in \mathcal{N}(t)} \mathbf{1}(k(v) = k) , \quad (3.6)$$

$$X_j^-(t) := \sum_{v \in \mathcal{N}(t)} \mathbf{1}(j(v) = j) , \quad (3.7)$$

$$X_{jk}(t) := \sum_{v \in \mathcal{N}(t)} \mathbf{1}(k(v) = k, j(v) = j) , \quad (3.8)$$

¹ Private communication

denote the random number of nodes with out-degree k , in-degree j and bi-degree jk in $G(t)$ respectively.

Let $\{P_{0,jk}\}_{j,k \geq 0}$ be a node type distribution such that $\sum_{j,k} P_{0,jk} = 1$ and $\sum_{j,k} k P_{0,jk} = \sum_{j,k} j P_{0,jk} = (\alpha + \gamma)^{-1}$. The initial random graph sequence $G^{(n)}(t_0)$ is selected to have the following asymptotics as $n \rightarrow \infty$:

$$\frac{|\mathcal{N}^{(n)}(n)|}{n} \stackrel{\text{P}}{=} \alpha + \gamma, \quad \frac{|\mathcal{E}^{(n)}(n)|}{n} \stackrel{\text{P}}{=} 1, \quad \frac{X_{jk}^{(n)}(n)}{n} \stackrel{\text{P}}{=} P_{0,jk}. \quad (3.9)$$

For example, this random graph sequence can be taken to be a configuration graph sequence.

The implication of the Wormald Theorem when applied to this random graph construction is that for large n the stochastic processes $\{X_k^+(t), X_j^-(t), X_{jk}(t)\}$ are determined by certain systems of ordinary differential equations (ODEs). We first focus on the collection $\{X_k^+(t), D^+(t)\}_{k \geq 0}$, $D^+(t) = \sum_k (k + \zeta^+) X_k^+(t)$. Some thought about the PA rules leads one to deduce that it must satisfy Markovian conditional expectation formulas:

$$\mathbb{E} [X_k^+(t+1) - X_k^+(t) | G(0), \dots, G(t)] = f_k(X^+(t)/t, D^+(t)/t) \quad (3.10)$$

$$\mathbb{E} [D^+(t+1) - D^+(t) | G(0), \dots, G(t)] = F(X^+(t)/t, D^+(t)/t). \quad (3.11)$$

Here the functions f_k, F are given by

$$\begin{aligned} f_k(z, d) &= (\beta + \gamma) \frac{\mathbf{1}(k \geq 1) A_{k-1} z_{k-1}^+ - A_k z_k^+}{d} + \alpha \mathbf{1}(k=1) + \gamma \mathbf{1}(k=0), \\ F(z, d) &= 1 + (\alpha + \gamma) \zeta^+ \end{aligned} \quad (3.12)$$

with the constants $A_k^+ = k + \zeta^+$. The relevant system for determining the large n asymptotics of this stochastic system is the infinite dimensional collection of functions $(z_k^+(s), d^+(s)), k \geq 0$ defined for $s \geq 1$ that satisfy the ODEs:

$$\begin{aligned} \dot{z}_k^+ &= f_k(z^+, d^+), \quad z_k^+(1) = P_{0,k}^+, \\ \dot{d}^+ &= F(z^+, d^+), \quad d^+(1) = 1 + (\alpha + \gamma) \zeta^+. \end{aligned} \quad (3.13)$$

It turns out that this triangular system of ODEs has an explicit solution:

$$z_k^+(s) = P_k^+ s + \sum_{k'=0}^k d^{k,k'} s^{-c^+ A_{k'}^+} \quad (3.14)$$

where the collection $d_{k,k'}$ are constants that depend on the initial conditions $P_{0,k}^+$. The constants P_k^+ however, are independent of the initial conditions and satisfy the recursion formula:

$$\begin{aligned}
P_k^+ &= \frac{A_{k-1}^+ P_{k-1}^+ + \mathbf{1}(k=1)\alpha/c^+}{A_k^+ + 1/c^+}, \quad k \geq 1 \\
P_0^+ &= \frac{\gamma}{A_0^+ + 1/c^+}.
\end{aligned} \tag{3.15}$$

Invoking the Wormald Theorem, taking care to truncate the finite n systems to the variables $k \in [K(n)]$ for a suitable sequence $K(n) = O(n)$, leads to the desired result.

Proposition 1. *For any $C > 1$, the following asymptotic formula holds with high probability uniformly for $n \leq t \leq Cn$:*

$$\frac{X_k^+(t)}{n} \stackrel{\mathbb{P}}{=} z_k^+(t/n) + o(1) \tag{3.16}$$

where the functions $z_k^+(s)$ are defined for $s \geq 1$ by the ODEs (3.13). The asymptotic marginal out-degree distribution $(P_k)_{k \geq 0}$ is given in terms of the Gamma function $\Gamma(x) := \int_0^\infty e^{-t} t^{x-1} dt$ by the formula

$$P_k^+ := \frac{\Gamma(k + \zeta^+) \Gamma(2 + \zeta^+ + 1/c^+)}{\Gamma(k + 1 + \zeta^+ + 1/c^+) \Gamma(1 + \zeta^+)} P_1^+, \quad k > 1 \tag{3.17}$$

$$P_0^+ = \frac{\gamma}{A_0^+ + 1/c^+}, \quad P_1^+ = \frac{A_0^+ P_0 + \alpha/c^+}{A_1^+ + 1/c^+}. \tag{3.18}$$

The constants P_0^+, P_1^+ normalize the probabilities and $c^+ = \frac{\beta + \gamma}{1 + (\alpha + \gamma)\zeta^+}$.

Using Stirling's formula, $\Gamma(x+1) \underset{x \rightarrow \infty}{\sim} e^{x \log x - x + 1/2} \sqrt{2\pi}$, one finds that $P_k^+ \sim C^+ k^{-1-1/c^+}$ as $k \rightarrow \infty$ where $C^+ := \frac{\Gamma(2 + \zeta^+ + 1/c^+)}{\Gamma(1 + \zeta^+)} P_1^+$. That is, the out-degree distribution has a Pareto tail of exponent $1 + 1/c^+ > 2$, as first shown by [15].

The extension of this result to the remaining processes $\{X_j^-(t), X_{jk}(t)\}$ is given by

Corollary 3. *For any $C > 1$, the following asymptotic formulas hold with high probability uniformly for $n \leq t \leq Cn$:*

$$\frac{X_j^-(t)}{n} \stackrel{\mathbb{P}}{=} z_j^-(t/n) + o(1), \tag{3.19}$$

$$\frac{X_{jk}(t)}{n} \stackrel{\mathbb{P}}{=} z_{jk}(t/n) + o(1), \tag{3.20}$$

where the functions $(z_j^-(s), z_{jk}(s))_{j,k \geq 0}$ satisfy a similar system of ODEs to (3.13). The asymptotic degree distributions are given by the formulas

$$P_j^-(t) = \frac{\Gamma(j + \zeta^-)\Gamma(2 + \zeta^- + 1/c^-)}{\Gamma(j + 1 + \zeta^- + 1/c^-)\Gamma(1 + \zeta^-)} P_1^-, \quad j > 1 \quad (3.21)$$

$$P_0^- = \frac{\alpha}{A_0^- + 1/c^-}, \quad P_1^- = \frac{A_0^- P_0^- + \gamma/c^-}{A_1^- + 1/c^-} \quad (3.22)$$

with $A_j^- = j + \zeta^-$ and

$$P_{j,k} = \frac{\Gamma(j + \zeta^-)\Gamma(k + \zeta^+)}{j!k!\Gamma(1 + \zeta^-)\Gamma(1 + \zeta^+)} \left(\alpha \zeta^- k \int_0^1 t^{c^- \zeta^- + c^+ (1 + \zeta^+)} (1 - t^{c^-})^j (1 - t^{c^+})^{k-1} dt \right. \\ \left. + \gamma \zeta^+ j \int_0^1 t^{c^- (1 + \zeta^-) + c^+ \zeta^+} (1 - t^{c^-})^{j-1} (1 - t^{c^+})^k dt \right) \quad (3.23)$$

$$\text{Here, } c^+ = \frac{\beta + \gamma}{1 + (\alpha + \gamma)\zeta^+}, \quad c^- = \frac{\beta + \alpha}{1 + (\alpha + \gamma)\zeta^-}.$$

It is probably more helpful to note that formula (3.23) for $P_{j,k}$ is the unique solution of the following recursion:

$$(1 + c^-(j + \zeta^-) + c^+(k + \zeta^+))P_{j,k} = c^-(j - 1 + \zeta^-)P_{j-1,k}\mathbf{1}(j > 0) \quad (3.24) \\ + c^+(k - 1 + \zeta^+)P_{j,k-1}\mathbf{1}(k > 0) + \alpha\mathbf{1}(j = 0, k = 1) + \gamma\mathbf{1}(j = 1, k = 0).$$

In particular, formula (3.24) enables us to independently verify the statement of [15] that, for fixed $k \geq 1$,

$$P_{j,k} \underset{j \rightarrow \infty}{\sim} C_k j^{-1-1/c^- - c^+/c^-} (\zeta^+ + \mathbf{1}(\gamma\zeta^+ = 0))$$

and for fixed $j \geq 1$

$$P_{j,k} \underset{k \rightarrow \infty}{\sim} D_j k^{-1-1/c^+ - c^-/c^+} (\zeta^- + \mathbf{1}(\gamma\zeta^- = 0)),$$

where $C_k, D_j, j, k > 0$ are positive constants.

An important aspect of PA models is that networks up to a million or more nodes can be simulated efficiently following techniques developed in [74]. The above asymptotic formulas, and many more general ones, can be easily checked in such simulation experiments.

A consequence of taking power law tails $P_k^+ \sim k^{-\tau}$ seriously in networks is that we must address the essentially unbounded nature of node and edge degrees, and in particular the fact that higher moments of the P and Q will be infinite:

$$\mathbb{E}[k_v^\mu] = \sum_k k^\mu P_k^+ \begin{cases} = \infty, & \mu \geq \tau - 1 \\ < \infty, & \mu < \tau - 1 \end{cases} \quad (3.25)$$

$$\mathbb{E}[k_\ell^\mu] = \sum_k k^\mu Q_k^+ \begin{cases} = \infty, & \mu \geq \tau - 2 \\ < \infty, & \mu < \tau - 2 \end{cases}. \quad (3.26)$$

Since $\tau < 3$ is commonly observed to arise in PA models of real networks, we must face the fact that even the first moments of Q are sometimes infinite. The paper [52] discusses some of the implications in measuring assortativity in such models.

While this class of PA growth models has certain attractive features, the rules do not in any way reflect how a real financial network evolves. We would like to capture the fact that banks make and break counterparty connections far more frequently than new banks are created. In fact, we should better imagine that the number of banks remains roughly constant (except perhaps during financial crises) while the number (and strength) of interbank links fluctuates relatively rapidly. A class of simple models of this type obtained by replacing the α and γ moves just discussed with a fourth type of move that deletes an edge. In the simplest specification, we set $\alpha = \gamma = 0$ and introduce two parameters ξ^-, ξ^+ and the move:

4. **δ -move:** With probability $\delta = 1 - \beta$, a random existing edge ℓ of type k_ℓ, j_ℓ is selected with probability proportional to $\xi^- k_\ell + \xi^+ j_\ell + 1$ and deleted.

Since now nodes are never created or destroyed, $N(t) = N(t_0)$. However, the number of edges $E(t)$ is a random walk with mean $E(t_0) + (\beta - \delta)(t - t_0)$ and variance $4\beta\delta(t - t_0)$. If we select $\beta = \delta = 0.5$, we can predict that the network should evolve into a steady state random graph.

While such preferential attachment and detachment moves apparently lead to interesting classes of skeleton graphs that match well the stylized facts of observed financial networks, for example those identified in [10], their deeper properties have not yet been developed in full. For this reason we move on to some two further classes of potential skeletons whose properties are easier to understand.

3.4 Inhomogeneous Random Graphs

Another class of random graphs that are of potential interest in modelling large scale financial networks is generalization of the Erdős-Renyi random graph known as *inhomogeneous random graphs* (IRG) or *generalized random graphs* (GRG). This class originates in papers [24] and [18] and has been studied in generality in [16]. Although we are most interested in directed graphs, for simplicity, we present here the discussion for undirected graphs. For a much more detailed treatment of the properties of this class, please see the textbook [71].

While the ER graph (or rather the Gilbert graph) selects each potential edge (v, w) independently with a fixed probability $p = z/(N - 1)$, we now select conditionally independently, with probabilities $p_{vw} = p_{wv}$ that depend on the random “type” of the nodes v and w . We suppose that the collection $(u_v)_{v \in [N]}$ of node types are identical independent non-negative random variables with cumulative distribution function $F : \mathbb{R}_+ \rightarrow [0, 1]$. Then, conditioned on the node types u_v, u_w the edge probabilities are defined to be

$$p_{vw} = \frac{\kappa(u_v, u_w)}{1 + \kappa(u_v, u_w)}$$

where $\kappa : \mathbb{R}_+ \times \mathbb{R}_+ \rightarrow \mathbb{R}_+$ is symmetric and nondecreasing in both arguments. An important example is when $\kappa(u, u') = uu'$ which leads to a variant of the Chung-Lu model [24]. Another simple example is when $u_v \in \{0, 1\}$ with $\mathbb{P}[u_v = 0] = p$, which we can think of giving a network with two types of banks. Then κ is a two-by-two symmetric matrix with non-negative entries

$$\kappa = \begin{pmatrix} \kappa_{00} & \kappa_{01} \\ \kappa_{01} & \kappa_{11} \end{pmatrix}.$$

When $\kappa_{00} = \kappa_{11} = 0$, then this specification gives a bipartite Erdős-Renyi graph.

One big advantage of the GRG model is that it is almost as easy to simulate as the Gilbert model. In fact, to generate the adjacency matrix M of a graph with N nodes, one can follow these steps:

1. Simulate the iid random variables $u_v, v \in [N]$ from the distribution F ;
2. Compute the upper-triangular N by N matrix A with the v, w entry given by $\frac{\kappa(u_v, u_w)}{1 + \kappa(u_v, u_w)}$;
3. Generate an upper-triangular matrix B with each entry iid uniform on $[0, 1]$;
4. Let the ones of the adjacency matrix M be put where $B \leq A$.

Conditioned on the values of u , the independent Bernoulli variables $X_{vw} = \mathbf{1}((v, w) \in g), 1 \leq v < w \leq N$ have values $x_{vw} \in \{0, 1\}$ with probabilities

$$\mathbb{P}[X_{vw} = x_{vw}, 1 \leq v < w \leq N | (u_v), v \in [N]] = \prod_{1 \leq v < w \leq N} \frac{\kappa(u_v, u_w)^{x_{vw}}}{1 + \kappa(u_v, u_w)}.$$

Since the total probability is 1, we have the identity

$$\sum_x \prod_{1 \leq v < w \leq N} \kappa(u_v, u_w)^{x_{vw}} = \prod_{1 \leq v < w \leq N} (1 + \kappa(u_v, u_w)) \quad (3.27)$$

where the sum is over all possible configurations $x = (x_{vw}), x_{vw} \in \{0, 1\}$. The degree of node v is $d_v(x) = \sum_{w < v} x_{vw} + \sum_{w > v} x_{vw}$ where we define $x_{vw} = x_{wv}$ for $w < v$. As a first step to investigate the nature of the underlying node degrees, it is useful to compute the conditional joint probability generating function

$$\Psi((t_v), v \in [N] | (u_v), v \in [N]) := \mathbb{E}[\prod_{v \in [N]} (t_v)^{d_v(x)} | (u_v), v \in [N]]. \quad (3.28)$$

To put this function into a more amenable form, we use:

Lemma 2. For any vector $x = (x_{vw})_{1 \leq v < w \leq N} \in \{0, 1\}^{N(N-1)/2}$,

$$\prod_{v=1}^N (t_v)^{d_v(x)} = \prod_{1 \leq v < w \leq N} (t_v t_w)^{x_{vw}}.$$

This and (3.27) allow us to express Ψ as:

$$\begin{aligned}
\mathbb{E} \left[\prod_{v \in [N]} (t_v)^{d_v(x)} \middle| (u_v), v \in [N] \right] &= \mathbb{E} \left[\prod_{1 \leq v < w \leq N} (t_v t_w)^{x_{vw}} \middle| (u_v), v \in [N] \right] \\
&= \frac{\sum_x \prod_{1 \leq v < w \leq N} (t_v t_w \kappa(u_v, u_w))^{x_{vw}}}{\prod_{1 \leq v < w \leq N} (1 + \kappa(u_v, u_w))} \\
&= \prod_{1 \leq v < w \leq N} \frac{1 + t_v t_w \kappa(u_v, u_w)}{1 + \kappa(u_v, u_w)}. \quad (3.29)
\end{aligned}$$

Proof of Lemma: Since for each $v \in [N]$, $d_v(x) = \sum_{w < v} x_{vw} + \sum_{w > v} x_{vw}$ where $x_{vw} = x_{wv}$, we have

$$\begin{aligned}
\prod_{v=1}^N (t_v)^{d_v(x)} &= \prod_v \left[\prod_{1 \leq w < v} t_v^{x_{vw}} \times \prod_{v < w \leq N} t_v^{x_{vw}} \right] \\
&= \left[\prod_{1 \leq w < v \leq N} t_v^{x_{vw}} \right] \times \left[\prod_{1 \leq v < w \leq N} t_v^{x_{vw}} \right] \\
&= \left[\prod_{1 \leq v < w \leq N} t_w^{x_{vw}} \right] \times \left[\prod_{1 \leq v < w \leq N} t_v^{x_{vw}} \right] = \prod_{1 \leq v < w \leq N} (t_v t_w)^{x_{vw}}.
\end{aligned}$$

□

To obtain an interesting asymptotic behaviour of the model as $N \rightarrow \infty$, it turns out to be natural to assume that the probability function $\kappa := \kappa^{(N)}$ scales with N :

$$\kappa^{(N)}(u, u') = (N-1)^{-1} \kappa(u, u') \quad (3.30)$$

and that each u has a fixed cumulative distribution function $F : \mathbb{R}_+ \rightarrow [0, 1]$. We also suppose that for some $\alpha > 0$

$$\|\kappa\|_{1+\alpha, F} := \int_{\mathbb{R}_+^2} |\kappa(u, u')|^{1+\alpha} dF(u) dF(u') < \infty. \quad (3.31)$$

Under these AminC, the following theorem provides us with a full description of the asymptotic joint degree distribution of any finite random collection of nodes:

Theorem 4. *In the GRG, assume that the scaling (3.30) and bound (3.31) hold for some $\alpha > 0$. Then, as $N \rightarrow \infty$,*

1. *The univariate generating function is given by*

$$\Psi^{(N)}(t_1) = \mathbb{E}[e^{(t_1-1)G^{-1}(U)}] (1 + o(1))$$

where U is F distributed, and

$$G^{-1}(u) = \int_0^\infty \kappa(u, u') dF(u'). \quad (3.32)$$

This implies that the degree d_1 converges in distribution to a mixture of Poisson random variables with random parameter λ having the mixing cumulative distribution function $F(G(\lambda))$.

2. For any fixed integer $M > 1$, the joint degrees $d_v, v \in [M]$ converge in distribution to an independent collection of identical random variables.

Part (1) tells us that the degree of a randomly selected node converges in distribution as $N \rightarrow \infty$ to a random variable whose probability distribution is given by $(P_k)_{k \geq 0}$ where

$$P_k = \int_{\mathbb{R}_+} \frac{e^{-\lambda} \lambda^k}{k!} dF(G(\lambda))$$

The mean node-degree is $z = \int_{\mathbb{R}_+} \lambda dF(G(\lambda))$. Part (2) then ensures the asymptotic independence of any finite collection of degrees.

We can see from the theorem that the model parametrization by an arbitrary pair (κ, F) contains redundant information. By defining

$$\tilde{\kappa}(v, v') = \kappa(G(v), G(v')), \quad \tilde{F}(v) = F(G(v)) \quad (3.33)$$

one finds that the pair $(\tilde{\kappa}, \tilde{F})$ leads to the same model:

$$\Psi^{(N)}(t_1) = \mathbb{E}[e^{(t_1-1)V}] (1 + o(1))$$

where V is \tilde{F} distributed, and

$$v = \int_0^\infty \tilde{\kappa}(v, v') d\tilde{F}(v'). \quad (3.34)$$

Without loss of generality therefore, one can take a pair (κ, F) such that (3.32) holds with G the identity mapping, under which condition the Poisson mixing distribution turns out to be F .

A rule of thumb says that a mixed Poisson distribution with an unbounded distribution of the mixing variable inherits the tail distribution of the mixing variable. Thus, if F has a Pareto tail of order $\tau - 1$ for some $\tau > 2$, Theorem 4 applies with $\alpha < \tau - 2$, and leads to a fat-tailed degree distribution with the same order. Note further that many potential integer degree distributions are not Poisson mixtures, for example any distribution with finite support.

Sketch Proof of Theorem: Part (1). For fixed N we compute Ψ by an intermediate conditioning on the random variables u_v and use (3.28) to write

$$\Psi((t_v), v \in [N]) = \mathbb{E} \left[\prod_{1 \leq v < w \leq N} \frac{1 + t_v t_w \kappa^{(N)}(u_v, u_w)}{1 + \kappa^{(N)}(u_v, u_w)} \right].$$

Now putting $t_v = 1$ for $v > 1$ leads to a lot of cancellation of numerator and denominator factors, leading to

$$\Psi(t_1) = \mathbb{E}^{(N)} \left[\prod_{2 \leq v \leq N} \frac{1 + t_1 \kappa^{(N)}(u_1, u_v)}{1 + \kappa^{(N)}(u_1, u_v)} \right] = \mathbb{E} \left[(\psi^{(N)}(t_1, u_1))^{N-1} \right] \quad (3.35)$$

where

$$\psi^{(N)}(t, u) = \int_{\mathbb{R}_+} \frac{1 + t \kappa^{(N)}(u, u')}{1 + \kappa^{(N)}(u, u')} dF(u').$$

This function can be written

$$\psi^{(N)}(t, u) = \int_{\mathbb{R}_+} \frac{1 + t(N-1)^{-1} \kappa(u, u')}{1 + (N-1)^{-1} \kappa(u, u')} dF(u').$$

Now, for any $\alpha > 0$ and every $x \geq 0$, there is $C(\alpha)$ such that

$$\frac{1 + tx}{1 + x} = 1 + (t-1)x + R(x)$$

where the remainder is bounded $|R(x)| \leq |t-1|C(\alpha)x^{1+\alpha}$. Using this bound and the definition of G^{-1} we find

$$\begin{aligned} \psi^{(N)}(t, u) &= \int_{\mathbb{R}_+} (1 + (t-1)(N-1)^{-1} \kappa(u, u')) dF(u') + \tilde{R} \\ &= 1 + (t-1)(N-1)^{-1} G^{-1}(u) + \tilde{R}. \end{aligned}$$

By the bound (3.31) on κ , the remainder \tilde{R} is bounded by $(N-1)^{-1-\alpha}$ times a function with bounded $L^{1+\alpha}$ -norm. Now, by a standard limiting argument, as $N \rightarrow \infty$,

$$[1 + (t-1)(N-1)^{-1} G^{-1}(u) + \tilde{R}]^{N-1} = e^{(t-1)G^{-1}(u)} (1 + O(n^{-\alpha}))$$

which leads to the desired result.

The proof of part (2) is similar, and left as an exercise. \square

We can go further and investigate the *shifted* bivariate distribution of edge degrees $(k_\ell - 1, k'_\ell - 1)$ by computing the expectation $e = \mathbb{E}^{(N)}[t_1^{d_1-1} t_2^{d_2-1} | (1, 2) \in g]$ under the parametrization with G equal to the identity mapping. Following exactly the same steps as above, we find the expression

$$\begin{aligned} e &= \left(\mathbb{P}^{(N)}[(1, 2) \in g] \right)^{-1} \\ &\times \mathbb{E}^{(N)} \left[\frac{\kappa^{(N)}(u_1, u_2)}{1 + \kappa^{(N)}(u_1, u_2)} \prod_{v \geq 3} \left(\frac{1 + t_1 \kappa^{(N)}(u_1, u_v)}{1 + \kappa^{(N)}(u_1, u_v)} \right) \left(\frac{1 + t_2 \kappa^{(N)}(u_2, u_v)}{1 + \kappa^{(N)}(u_2, u_v)} \right) \right]. \end{aligned}$$

The same logic that leads to Theorem 4 leads to

$$e = (\mathbb{E}[\kappa(u_1, u_2)])^{-1} \mathbb{E} \left[\kappa(u_1, u_2) e^{(t_1-1)u_1} e^{(t_2-1)u_2} \right] (1 + o(1)). \quad (3.36)$$

In the Chung-Lu class of models where $\kappa(u_1, u_2) = u_1 u_2$, (3.36) implies that asymptotically, the edge degree distribution is the independent case

$$Q_{kk'} = \frac{kk'P_k P_{k'}}{z^2}.$$

In general there is correlation between the edge degrees, i.e. the graph is assortative, where the edge-type distribution Q equals a bivariate mixture of independent Poisson random variables, shifted up by the vector $(1, 1)$. To verify this statement one needs to check that Q defined by

$$Q_{kk'} = \int_0^\infty \int_0^\infty \text{Pois}(v, k-1) \text{Pois}(v', k'-1) dG(v, v'), \quad \forall k, k' \geq 1 \quad (3.37)$$

with

$$dG(v, v') := z^{-1} \kappa(v, v') dF(v) dF(v') \quad (3.38)$$

is consistent with

$$e = \sum_{kk'} t_1^{k-1} t_2^{k'-1} Q_{kk'}.$$

As a useful variation of the IRG framework, one can introduce a more abstract node type space \mathcal{A} with probability distribution df . and a mapping $u : \mathcal{A} \rightarrow \mathbb{R}_+$. If we define the probability of a type a and type b node to wire together to be

$$\frac{\kappa(u(a), u(b))}{1 + \kappa(u(a), u(b))}$$

where κ and df are consistent:

$$u(a) = \int_{\mathcal{A}} \kappa(u(a), u(b)) df(b), \quad \forall a \in \mathcal{A}$$

then $u(a)$ will be the average degree of a type a node. The next example illustrates this kind of network construction.

Example 1. (Three Bank Types) Consider a financial network with small, medium and large banks whose network fractions are $f = (f_a)_{a=1,2,3} = [0.80, 0.15, 0.05]$ and whose conditional average degrees are $u = (u_a)_{a=1,2,3} = [2, 31/3, 51/2]$. The mean degree is $z = f \cdot u = 4.35$. One connectivity kernel that satisfies the consistency condition $\sum_{b=1}^3 \kappa_{ab} f_b = z^{-1} u_a$ is:

$$\kappa^{(1)} = \begin{bmatrix} 0 & \frac{1}{15} & \frac{1}{5} \\ \frac{1}{15} & \frac{1}{5} & \frac{2}{5} \\ \frac{1}{5} & \frac{2}{5} & \frac{7}{10} \end{bmatrix}.$$

It is easy to verify that this choice leads to the consistent negatively assortative pair of degree distributions

$$Q_{kk'}^{(1)} = z^{-1} \sum_{a,b=1}^3 \kappa_{ab}^{(1)} \text{Pois}(u_a, k-1) \text{Pois}(u_b, k'-1) f_a f_b \quad (3.39)$$

$$P_k = \sum_{a=1}^3 \text{Pois}(u_a, k) f_a . \quad (3.40)$$

The independent edge type distribution $Q^{(2)}$ arises by taking $\kappa_{ab}^{(2)} = z^{-1} u_a u_b$. To simulate a random network with N banks that has such large N asymptotics, one draws N bank types a_1, \dots, a_N from the f distribution. Then for each pair of banks $v < w \in [N]$ one creates an undirected edge between them independently with probability $\frac{\kappa_{a_v, a_w}}{N-1+\kappa_{a_v, a_w}}$.

To summarize, the undirected IRG construction parametrized by pairs (κ, F) that satisfy (3.32) with G equal to the identity, defines a rich class that generalizes the Erdős-Renyi graph, is straightforward to simulate on a computer, and retains a tractable large N asymptotic structure. Both the asymptotic node type and edge type distributions can be fully characterized as mixtures of Poisson random variables. The natural question is what the IRG class has to do with the configuration graph model. It has been proven in [18] that the subclass of Chung-Lu models is asymptotic to a subclass of non-assortative *simple* configuration graphs. In general, it seems that the IRG models forms a class that overlaps substantially with the so-called *assortative configuration model* that will be developed in the next section. In any case, the directed version of the IRG class outlined in this section is a promising foundation for modelling financial networks, and certainly much of the technology we develop later on in this book for configuration graphs is extendible to the IRG class. However, we now return to the configuration graph model and consider how one can modify the construction of Section 3.2 to account for assortativity.

3.5 Assortative Configuration Graphs

It has been often observed in financial networks (and as it happens, also the world wide web) that they are highly *disassortative*, or as we prefer, *negatively assortative* (see for example [68] and [10]). This refers to the property that any bank's counterparties (i.e. their graph neighbours) have a tendency to be banks of an opposite character. For example, it is observed that small banks tend to lend preferentially to large banks rather than other small banks. On the other hand, social networks are commonly observed to have positive assortativity. Structural characteristics such as degree distribution and assortativity are felt by some (see [54]) to be highly relevant to the propagation of contagion in networks but the nature of such relationships is far from clear.

The aim of this section is to put a firm theoretical foundation under the class of configuration graphs with arbitrary node degree distribution P and edge degree distribution Q taken to satisfy Assumption 4 of Section 3.1. The class of config-

uration graphs with general Q has not been well studied previously, and we will generalize some of the classic large N asymptotic results known to be true for the nonassortative configuration graph construction described in Section 3.2. At the end of the section, we will also provide justification of a new approximate Monte Carlo simulation algorithm for assortative configuration graphs.

3.5.1 Finite Assortative Configuration Graphs

The “exact” assortative configuration multigraph of size N parametrized by the node-edge degree distribution pair (P, Q) is constructed by the following random algorithm:

1. Draw a sequence of N node-degree pairs $X = ((j_1, k_1), \dots, (j_N, k_N))$ independently from P , and accept the draw if and only if it is feasible, i.e. $\sum_{n \in [N]} j_n = \sum_{n \in [N]} k_n$, and this defines the number of edges E that will result. Label the n th node with k_n *out-stubs* (picture this as a half-edge with an out-arrow) and j_n *in-stubs*. Define the partial sums $u_j^- = \sum_n \mathbf{1}(j_n = j)$, $u_k^+ = \sum_n \mathbf{1}(k_n = k)$, $u_{jk} = \sum_n \mathbf{1}(j_n = j, k_n = k)$ and the number of out- and in-stubs $e_k^+ = k u_k^+$, $e_j^- = j u_j^-$.
2. Conditioned on X , the result of Step 1, choose an arbitrary ordering ℓ^- and ℓ^+ of the E in-stubs and E out-stubs. For each permutation $\sigma \in S[E]$ of the set $[E]$, select the matching sequence or wiring W of “edges” $\ell = (\ell^- = \ell, \ell^+ = \sigma(\ell))$, labelled by $\ell \in [E]$, with probability weighted by the factor

$$\prod_{\ell \in [E]} Q_{k_{\sigma(\ell)} j_{\ell}}. \quad (3.41)$$

Define the number of type (k, j) edges to be $e_{kj} = \sum_{\ell \in [E]} \mathbf{1}(k_{\ell} = k, j_{\ell} = j)$.

Intuitively, since step 1 leads to a product probability measure subject to a single linear constraint that is true in expectation, it should lead to node degree independence for large N . This intuition can be verified using the Fourier transform of the P measure:

$$\Phi_P(a, b) := \sum_{j, k} P_{jk} e^{i(a_j + b_k)}, \quad (a, b) \in [0, 2\pi]^2 \quad (3.42)$$

which under Assumption 4 is a twice differentiable periodic function. For any finite M , as $N \rightarrow \infty$, it can be shown that

$$\mathbb{E} \left[\prod_{n \in [M]} e^{i(a_n j_n + b_n k_n)} \mid \sum_{n \in [N]} (j_n - k_n) = 0 \right] \rightarrow \prod_{n \in [M]} \Phi_P(a_n, b_n) \quad (3.43)$$

which implies that the first M pairs $(j_n, k_n)_{n \in [M]}$ converge in distribution to the correct product measure. It is also true that the empirical node type distributions converge in probability to their means:

$$N^{-1} u_j^- \xrightarrow{P} P_j^-, \quad N^{-1} u_k^+ \xrightarrow{P} P_k^+, \quad N^{-1} u_{jk} \xrightarrow{P} P_{jk}. \quad (3.44)$$

Similar logic suggests that since the matching weights in step 2 define a product probability measure conditional on a set of linear constraints that are true in expectation, it should lead to edge degree independence in the large N limit. We define

$$\Phi_Q(a, b) := \sum_{k, j} Q_{kj} e^{i(a_j + bk)}, \quad (a, b) \in [0, 2\pi]^2 \quad (3.45)$$

which under Assumption 4 is a once differentiable periodic function. For any finite L , convergence in distribution of the degrees of the first L edges of the wiring sequence to the target distribution follows from

$$\mathbb{E} \left[\prod_{\ell \in [L]} e^{i(a_\ell j_\ell + b_\ell k_\ell)} \right] \rightarrow \prod_{\ell \in [L]} \Phi_Q(a_\ell, b_\ell) \quad (3.46)$$

Eventually we will also need convergence in probability of the empirical density of type (k, j) edges: for every (k, j) ,

$$N^{-1} E \xrightarrow{P} z, \quad E^{-1} e_{kj} \xrightarrow{P} Q_{kj}, \quad E^{-1} e_j^- \xrightarrow{P} Q_j^-, \quad E^{-1} e_k^+ \xrightarrow{P} Q_k^+. \quad (3.47)$$

It appears that Assumption 4 is sufficient for both (3.46), (3.47) to be true in general, but the detailed proof of this fact is not yet available.

The next Proposition says that the probability of any wiring sequence $W = (\ell \in [E])$ in step 2 depends only on the set of quantities (e_{kj}) where for each k, j , $e_{kj} := |\{\ell \in [E] | \ell \in \mathcal{E}_{kj}\}|$. In other words, the random collection $e = (e_{kj})$, subject to the constraints

$$e_j^- = \sum_k e_{kj}, \quad e_k^+ = \sum_j e_{kj}, \quad (3.48)$$

is a sufficient statistic for the probability of a labelled multigraph W conditioned on X .

Proposition 2. *In the assortative configuration graph construction for finite N with probabilities P, Q , the probability of any wiring sequence $W = (\ell \in [E])$ conditioned on the $X = (j_i, k_i), i \in [N]$, is:*

$$\mathbb{P}[W|X] = C^{-1} \prod_{kj} (Q_{kj})^{e_{kj}}, \quad (3.49)$$

$$C = C(e^-, e^+) = \sum_e \prod_{kj} \frac{(Q_{kj})^{e_{kj}}}{e_{kj}!} \prod_j (e_j^-!) \prod_k (e_k^+!), \quad (3.50)$$

where the sum in (3.50) is over collections $e = (e_{kj})$ satisfying the constraints (3.48). Moreover the number of wirings W consistent with a collection $e = (e_{kj})$ is given by

$$|\{W : e(W) = e\}| = \frac{(\prod_j e_j^-!) (\prod_k e_k^+!)}{\prod_{kj} e_{kj}!}. \quad (3.51)$$

Proof of Proposition: The denominator of (3.49) is $C = \sum_{\sigma \in S(E)} \prod_{l \in [E]} Q_{k_{\sigma(l)} j_l}$, from which we verify (3.50) by induction on E . Assuming (3.50) is true for $E-1$, we can verify the inductive step for E :

$$\begin{aligned} C &= \sum_{\tilde{k}} \sum_{\sigma \in S(E)} \mathbf{1}(k_{\sigma(E)} = \tilde{k}) \prod_{l \in [E]} Q_{k_{\sigma(l)} j_l} \\ &= \sum_{\tilde{k}} e_{\tilde{k}}^+ Q_{\tilde{k} j_E} \sum_{\sigma' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_l} \\ &= \sum_{\tilde{k}} e_{\tilde{k}}^+ Q_{\tilde{k} j_E} \sum_{e'} \prod_{k_j} \frac{(Q_{k_j})^{e'_{k_j}}}{e'_{k_j}!} \prod_j (e_j'^-!) \prod_k (e_k'^+!). \end{aligned}$$

Here, $e'_{k_j} = e_{k_j} - \mathbf{1}(k = \tilde{k}, j = j_E)$, $e_j'^- = e_j^- - \mathbf{1}(j = j_E)$, $e_k'^+ = e_k^+ - \mathbf{1}(k = \tilde{k})$. After noting cancellations that occur in the last formula, and re-indexing the collection e' one finds

$$\begin{aligned} C &= \sum_{\tilde{k}} \sum_{e'} \frac{e_{\tilde{k} j_E}^-}{e_{j_E}^-} \prod_{k_j} \frac{(Q_{k_j})^{e_{k_j}}}{e_{k_j}!} \prod_j (e_j^-!) \prod_k (e_k^+!) \\ &= \sum_e \left(\sum_{\tilde{k}} \frac{e_{\tilde{k} j_E}^-}{e_{j_E}^-} \right) \prod_{k_j} \frac{(Q_{k_j})^{e_{k_j}}}{e_{k_j}!} \prod_j (e_j^-!) \prod_k (e_k^+!) \\ &= \sum_e \prod_{k_j} \frac{(Q_{k_j})^{e_{k_j}}}{e_{k_j}!} \prod_j (e_j^-!) \prod_k (e_k^+!) \end{aligned}$$

which is the desired result. Finally, (3.51) follows easily from (3.50). \square

Now we use similar logic to prove further natural properties of the construction. A basic property of the ACG construction is the permutation symmetry in both in- and out-stubs. As one easily checks, it allows the denominator C in (3.49) to be written in various ways:

$$C = \sum_{\tilde{k}} e_{\tilde{k}}^+ Q_{\tilde{k} j_E} \sum_{\sigma' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_l} \quad (3.52)$$

$$= \sum_{\tilde{j}} e_{\tilde{j}}^- Q_{k(E) \tilde{j}} \sum_{\sigma' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\sigma'(l)}} \quad (3.53)$$

$$= E^{-1} \sum_{\tilde{k}, \tilde{j}} e_{\tilde{k}}^+ e_{\tilde{j}}^- Q_{\tilde{k} \tilde{j}} \sum_{\sigma' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_l} \quad (3.54)$$

$$(3.55)$$

In the sums over σ' , one needs to reorder (arbitrarily) the remaining in or out-stubs. These identities can also be invoked at any subsequent step of the wiring process, allowing one to select at each step arbitrarily from either the available in- or out-stubs.

Proposition 3. *In the assortative configuration graph construction for finite N with probabilities P, Q that satisfy Assumption 4:*

1. The probability p of the last edge of any wiring sequence $W = (\ell \in [E])$ having type \tilde{k}, \tilde{j} , conditioned on the $X = (j_i, k_i), i \in [N]$ is

$$p = \mathbb{E}[e_{\tilde{k}\tilde{j}} | X] / E ;$$

2. For each \tilde{k}, \tilde{j} in the support of Q ,

$$\mathbb{E}[e_{\tilde{k}\tilde{j}} | X] / E = Q_{\tilde{k}\tilde{j}}(1 + o(1))$$

with high probability as $N \rightarrow \infty$.

Proof of Proposition: To prove these results, we need to symmetrize the wiring over both the in- and out-stub sequences. For part (1), following a similar notation and logic to the previous proof,

$$\begin{aligned} p &= \frac{1}{C(e^-, e^+)E!} \sum_{\sigma, \tilde{\sigma} \in S(E)} \mathbf{1}(k_{\sigma(E)} = \tilde{k}, j_{\tilde{\sigma}(E)} = \tilde{j}) \prod_{l \in [E]} Q_{k_{\sigma(l)} j_{\tilde{\sigma}(l)}} \\ &= \frac{1}{C(e^-, e^+)E!} e_k^+ e_j^- Q_{\tilde{k}\tilde{j}} \sum_{\sigma', \tilde{\sigma}' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\tilde{\sigma}'(l)}} \\ &= \frac{(E-1)!}{C(e^-, e^+)E!} e_k^+ e_j^- Q_{\tilde{k}\tilde{j}} \sum_{\sigma'} \prod_{k_j} \frac{(Q_{k_j})^{e_{k_j}^+}}{e_{k_j}^+!} \prod_j (e_j^-!) \prod_k (e_k^+!) \\ &= \frac{1}{C(e^-, e^+)E} \sum_e e_{\tilde{k}\tilde{j}} \prod_{k_j} \frac{(Q_{k_j})^{e_{k_j}^+}}{e_{k_j}^+!} \prod_j (e_j^-!) \prod_k (e_k^+!) . \end{aligned}$$

For part (2), as long as $Q_{\tilde{j}}^- \neq 0$, using (3.44) shows it is equivalent to prove

$$p' = \frac{1}{E} \mathbb{E}[e_{\tilde{k}\tilde{j}} Q_{\tilde{j}}^- - Q_{\tilde{k}\tilde{j}} e_{\tilde{j}}^- | X] = o(1)$$

for each \tilde{k}, \tilde{j} . Using part (1), this equals

$$\begin{aligned} p' &= \frac{1}{cE!} \sum_{\tilde{k}'} \sum_{\sigma, \tilde{\sigma} \in S(E)} \\ &\quad \times [\mathbf{1}(k_{\sigma(E)} = \tilde{k}, j_{\tilde{\sigma}(E)} = \tilde{j}) Q_{\tilde{k}'\tilde{j}} - \mathbf{1}(k_{\sigma(E)} = \tilde{k}', j_{\tilde{\sigma}(E)} = \tilde{j}) Q_{\tilde{k}\tilde{j}}] \prod_{l \in [E]} Q_{k_{\sigma(l)} j_{\tilde{\sigma}(l)}} \\ &= \frac{1}{cE!} \sum_{\tilde{k}'} Q_{\tilde{k}'\tilde{j}} Q_{\tilde{k}\tilde{j}} e_{\tilde{j}}^- \\ &\quad \times \left[e_{\tilde{k}}^+ \sum_{\sigma', \tilde{\sigma}' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\tilde{\sigma}'(l)}} - e_{\tilde{k}'}^+ \sum_{\sigma', \tilde{\sigma}' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\tilde{\sigma}'(l)}} \right] . \end{aligned}$$

In this last expression, it is critical to understand that despite the identical notation (which has become quite cumbersome) the collection of out-stubs is different in the two terms: the difference amounts to exactly one substitution of a \tilde{k} -stub with a \tilde{k}' -

stub. The next step is to note that by Assumption 4, $e_k^+ = EQ_k^+(1 + o(1))$, which implies

$$p' = o(1) + \frac{E}{cE!} \sum_{\bar{k}'} Q_{\bar{k}', \bar{j}} Q_{\bar{k}, \bar{j}} e_{\bar{j}}^- \\ \times \sum_{\bar{j}} \left[Q_{\bar{k}, \bar{j}} \sum_{\sigma', \bar{\sigma}' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\bar{\sigma}'(l)}} - Q_{\bar{k}', \bar{j}} \sum_{\sigma', \bar{\sigma}' \in S(E-1)} \prod_{l \in [E-1]} Q_{k_{\sigma'(l)} j_{\bar{\sigma}'(l)}} \right].$$

Now, because of the structure of the imposed symmetrization of remaining $E - 1$ out-stubs in the two terms, we can observe an exact cancellation of all terms, which finishes the proof. \square

The reasoning underlying these two propositions tells us that wiring probabilities are independent of ordering, and the actual wiring order can be chosen flexibly. Moreover, we have made sufficient assumptions that imply that the large N asymptotics of the empirical node and edge degree distributions agree with the target (P, Q) distributions. In the next section, we will use these facts to compute detailed probabilities of more complex configurations.

3.5.2 Asymptotic Configuration Probabilities

It will turn out that the most important objects we need to understand to derive a rigorous treatment of cascade mappings on random financial networks are random connected subgraphs of the skeleton, labelled by their degree types. We shall call such objects *configurations*, and the goal of this section is to determine their asymptotic probability distributions as $N \rightarrow \infty$.

Let us first consider what it means in the (P, Q) configuration model with size N to draw a random pair of vertices v_1, v_2 that happen to have a link, i.e. $v_2 \in \mathcal{N}_{v_1}^-$. Following the ACG algorithm, first we construct the feasible bidegree sequence $X = (j_i, k_i), i \in [N]$. To be clear, we label these corresponding nodes by $v_i = i$. Then, conditioned on X , we construct a random Q -wiring sequence $W = (\ell = (v_\ell^+, v_\ell^-))_{\ell \in [E]}$ with $E = \sum_i k_i = \sum_i j_i$ edges. By an abuse of notation, we label their edge degrees by $k_\ell = k_{v_\ell^+}, j_\ell = j_{v_\ell^-}$.

To compute the joint probability $p = \mathbb{P}[v_i \in \mathcal{N}_{j_i, k_i}, i = 1, 2 | v_2 \in \mathcal{N}_{v_1}^-, X]$, we first condition on the bidegree sequence X for any finite N . The event in question is that the first link in the wiring sequence W attaches to nodes of the required degrees $(j_1, k_1), (j_2, k_2)$. The fraction $j_1 u_{j_1 k_1} / e_{j_1}^-$ of available j_1 -stubs and a fraction $k_2 u_{j_2 k_2} / e_{k_2}^+$ available k_2 -stubs have the correct property. Now, by Proposition 3, the conditional probability for the configuration is therefore

$$p = j_1 u_{j_1 k_1} k_2 u_{j_2 k_2} \frac{\mathbb{E}[e_{k_2 j_1} | X]}{E e_{k_2}^+ e_{j_1}^-} \quad (3.56)$$

which has a limit as $N \rightarrow \infty$:

$$p \rightarrow \frac{j_1 k_2 P_{j_1 k_1} P_{j_2 k_2} Q_{k_2 j_1}}{z^2 Q_k^+ Q_j^-}. \quad (3.57)$$

By summing over k_1, k_2, j_2 , one gets the important formula

$$\mathbb{P}[j_{v_1} = j_1 \mid v_2 \in \mathcal{N}_{v_1}^-] = \sum_{k_2} Q_{k_2 j_1} = Q_{j_1}^-. \quad (3.58)$$

Warning! Note the meaning of this last equation: the in-degree distribution of v , conditioned on v being in-connected to another node, is not equal to the marginal in-degree distribution P_j^- as one might naively expect, but rather Q_j^- . Failure to appreciate this distinction is a major confusion in configuration graph computations.

One can also compute the correct asymptotic expression for p , informally, by successive conditioning:

$$p = \mathbb{P}[v_i \in \mathcal{N}_{j_i k_i}, i = 1, 2 \mid v_2 \in \mathcal{N}_{v_1}^-] \quad (3.59)$$

$$= \mathbb{P}[v_1 \in \mathcal{N}_{j_1 k_1} \mid v_2 \in \mathcal{N}_{v_1}^- \cap \mathcal{N}_{j_2 k_2}] \mathbb{P}[v_2 \in \mathcal{N}_{j_2 k_2} \mid v_2 \in \mathcal{N}_{v_1}^-] \quad (3.60)$$

$$= P_{k_1 | j_1} Q_{j_1 | k_2} P_{j_2 | k_2} Q_{k_2}^+ = \frac{P_{j_1 k_1} P_{j_2 k_2} Q_{k_2 j_1}}{P_{k_2}^+ P_{j_1}^-} \quad (3.61)$$

where we introduce conditional degree probabilities $P_{k|j} = P_{jk}/P_j^-$ etc.

One subtlety in this argument is that occasionally in the above matching algorithm, the first edge forms a self-loop, i.e. $v_1 = v_2$. We can compute the probability of this event, jointly with fixing the degree of v_1 , as follows:

$$\tilde{p} := \mathbb{E}[v_1 = v_2, v_1 \in \mathcal{N}_{jk} \mid v_2 \in \mathcal{N}_{v_1}^- \mid X] = \left(\frac{jk u_{jk}}{e_j^- e_k^+} \right) \frac{\mathbb{E}[e_{kj} \mid X]}{E}.$$

As $N \rightarrow \infty$ this goes to zero, however $N\tilde{p}$ approaches a finite value:

$$N\tilde{p} \xrightarrow{P} \frac{jk P_{jk} Q_{kj}}{z^2 Q_k^+ Q_j^-} \quad (3.62)$$

which says that the relative fraction of edges being self loops is the asymptotically small $\sum_{jk} \frac{jk P_{jk} Q_{kj}}{z^2 Q_k^+ Q_j^-}$. In fact, one can go further with this type of analysis, and prove an asymptotic result that the total number of self loops in the multigraph converges in probability to a Poisson random variable with finite parameter

$$\lambda = \sum_{jk} \frac{jk P_{jk} Q_{kj}}{z^2 Q_k^+ Q_j^-}. \quad (3.63)$$

The probability of a general *configuration* g , that is, a finite connected subgraph with $L = |\mathcal{E}(g)|$ edges labelled by degree types, is best understood as a growth

process that adds one edge at a time. There are two alternatives: processes that start from a random edge or processes that start with a random node. For our purposes, it is most useful to consider a growth process that starts conditionally on an initial node, or *root*, $v \in \mathcal{N}_{jk}$ to which edges are added one at a time.

A tree configuration g will arise if the ℓ th new edge always attaches a new node w_ℓ to an existing node v_ℓ : ℓ is an *in-edge* if it points into v_ℓ , that is if $\ell \in \mathcal{E}_v^-$, or an *out-edge* if it points out of v_ℓ . For a tree g rooted at v , the direction of each edge is unambiguous. If ever ℓ connects two existing nodes, then the resultant configuration will have a new cycle and is no longer a tree, and the direction of some edges will be ambiguous.

Consider the probability of the configuration with one edge obtained by attaching a node $w_1 \in \mathcal{N}_{j_1 k_1}$ pointing to the root v , within a configuration graph $(\mathcal{N}, \mathcal{E})$ of size $|\mathcal{N}| = N$. We follow the usual logic and apply the ACG wiring step, step 2, with $E = \sum_k e_k^+ = \sum_j e_j^-$ edges, conditioned on the finite node type sequence X . The event $W(1)$ that the first k -stub is one of the $k_1 u_{j_1 k_1}$ available k -stubs from a type $(j_1 k_1)$ node has conditional probability

$$\mathbb{P}[W(1)|X] := C^{-1} \sum_{\sigma \in S(E)} \mathbf{1}(W(1)) \prod_{\ell \in [E]} Q_{k(\sigma(\ell), j(\ell))} \quad (3.64)$$

. Using Proposition 3, the result can be expressed as

$$\frac{k_1 u_{j_1 k_1}}{e_{k_1}^+} \mathbb{E}[e_{k_1 j}^- / e_j^- | X]$$

which has the limit

$$\mathbb{P}[W(1)|X] \xrightarrow{P} P_{j_1 | k_1} Q_{k_1 | j} . \quad (3.65)$$

On the other hand, when one creates a new out-edge by attaching a new node w_1 pointing from the root v , one applies the identity (3.53), leading to the reversed formula

$$\mathbb{P}[W(1)|X] = \frac{j_1 u_{j_1 k_1}}{e_{j_1}^-} \mathbb{E}[e_{k j_1}^+ / e_k^+ | X] \xrightarrow{P} P_{k_1 | j_1} Q_{j_1 | k} . \quad (3.66)$$

When one realizes that one can choose to use either (3.52) or (3.53) at each step, each new wire can be chosen to match any specific available in- (or out-) stub to a random available out- (or in-) stub. Thus it is intuitive that a new node w of type (j', k') attaching to an existing node v of type (j, k) will always bring in an additional factor which is either $P_{j' | k'} Q_{k' | j}$ for an in-edge or $P_{k' | j'} Q_{j' | k}$ for an in-edge. Some further thought also suggests that adding a wire between existing nodes will always lead to a factor of 0 in the $N \rightarrow \infty$ limit. Therefore, the main takeaway result for the ACG model, which gives the asymptotic probability of general tree configurations together with the fact that configuration with cycles vanish in probability in the limit, should be intuitively clear.

In the previous section, we remarked that while it appears that Assumption 4 for the ACG construction is sufficient for both large N limits (3.46), (3.47) to be true

in general, a detailed proof of this fact is not yet available. To be on the safe side in what follows we strengthen Assumption 4 to the following:

Assumption 5. The pair (P, Q) satisfies Assumption 4, and in addition (3.46), (3.47) hold as limits as $N \rightarrow \infty$.

Theorem 5. Consider the ACG sequence with (P, Q) that satisfies Assumption 5. For any fixed finite configuration g rooted to $v \in \mathcal{N}_{jk}$, with M added nodes and $L \geq M$ edges, labelled by the node bidegree sequence $(j_m, k_m)_{m \in [M]}$, the joint probability $p = \mathbb{P}[w_m \in \mathcal{N}_{j_m k_m}, m \in [M], g | v \in \mathcal{N}_{jk}, X]$ conditioned on X converges with high probability as $N \rightarrow \infty$:

$$p \xrightarrow{P} \prod_{m \in [M], \text{out-edge}} P_{k_m | j_m} Q_{j_m | k_m'} \prod_{m \in [M], \text{in-edge}} P_{j_m | k_m} Q_{k_m | j_m'}, \quad \text{if } g \text{ is a tree} \quad (3.67)$$

$$p = O(N^{M-L}), \quad \text{if } g \text{ has cycles.} \quad (3.68)$$

For trees, the ℓ th edge has $m = \ell$, and $m' \in \{0\} \cup [\ell - 1]$ numbers the node to which w_ℓ attaches, with $m' = 0$ meaning the root v .

Proof of Theorem: The tree configuration g with its ordered sequence of edges completely determines the sequence of edge types. Based on this information, define the sequences $u_{jk}(\ell), e_k^+(\ell), e_j^-(\ell), \ell \in [L]$ that count the number of available nodes and stubs at step ℓ . Similarly, a given collection $e_{jk} := e_{jk}(0)$ leads to a sequence $e_{kj}(\ell)$. Based on this information, the following identity for the conditional probability of the event $W(L)$ that the L -edge tree g occurs can be verified by induction on L for each fixed N and node type sequence X :

$$\begin{aligned} \mathbb{P}[W(L) | v \in \mathcal{N}_{jk}, X] &= C^{-1} \prod_{\ell \text{ out}} j_\ell u_{j_\ell k_\ell}(\ell) \prod_{\ell \text{ in}} k_\ell u_{j_\ell k_\ell}(\ell) \\ &\times \sum_e \prod_{j' k'} \frac{Q_{k' j'}^{e_{k' j'}}}{e_{k' j'}(L)!} \prod_{j'} e_{j'}^-(L)! \prod_{k'} e_{k'}^+(L)! \quad (3.69) \end{aligned}$$

This formula can be rewritten as

$$\begin{aligned} \mathbb{P}[W(L) | v \in \mathcal{N}_{jk}, X] &= \prod_{\ell \text{ out}} j_\ell u_{j_\ell k_\ell}(\ell) \prod_{\ell \text{ in}} k_\ell u_{j_\ell k_\ell}(\ell) \prod_{j'} \frac{e_{j'}^-(L)!}{e_{j'}^-!} \prod_{k'} \frac{e_{k'}^+(L)!}{e_{k'}^+!} \\ &\times \mathbb{E} \left[\prod_{j' k'} \frac{e_{k' j'}!}{e_{k' j'}(L)!} \mid v \in \mathcal{N}_{jk}, X \right]. \quad (3.70) \end{aligned}$$

Now one can apply the large N limits (3.44) and (3.47) to verify the following facts:

$$\mathbb{E}[E^{-L} \prod_{j,k'} \frac{e_{k'j'}!}{e_{k'j'}(L)!} | v \in \mathcal{N}_{jk}, X] \xrightarrow{P} \prod_{\ell, \text{out}} Q_{k_{\ell}j_{\ell}} \prod_{\ell, \text{in}} Q_{k_{\ell}j_{\ell}} \quad (3.71)$$

$$E^{2L} \prod_{j'} \frac{e_{j'}^-(L)!}{e_{j'}^-!} \prod_{k'} \frac{e_{k'}^+(L)!}{e_{k'}^+!} \xrightarrow{P} \prod_{\ell, \text{out}} \frac{1}{Q_{k_{\ell}j_{\ell}}^+ Q_{j_{\ell}}^-} \prod_{\ell, \text{in}} \frac{1}{Q_{k_{\ell}j_{\ell}}^+ Q_{j_{\ell}}^-} \quad (3.72)$$

$$E^{-L} \prod_{\ell, \text{out}} j_{\ell} u_{j_{\ell}k_{\ell}}(\ell) \prod_{\ell, \text{in}} k_{\ell} u_{j_{\ell}k_{\ell}}(\ell) \xrightarrow{P} \prod_{\ell, \text{out}} \frac{j_{\ell} P_{j_{\ell}k_{\ell}}}{z} \prod_{\ell, \text{in}} \frac{k_{\ell} P_{j_{\ell}k_{\ell}}}{z} \quad (3.73)$$

Applying these limits to the factors in (3.70) yields (3.67).

To prove (3.68) for configurations with cycles, first consider the edge that creates the first cycle for which it is straightforward to demonstrate that the probability is $O(N^{-1})$. Similarly, each additional edge that creates a new cycle introduces an additional $O(N^{-1})$ factor. \square

Remark 5. 1. Formula (3.68) shows clearly what is meant by saying that configuration graphs are *locally tree-like* as $N \rightarrow \infty$. It means the number of occurrences of any fixed finite size graph g with cycles embedded within a configuration graph of size N remains bounded with high probability as $N \rightarrow \infty$.

2. Even more interesting is that (3.67) shows that large configuration graphs exhibit a strict type of conditional independence. Selection of any root node v of the tree graph g splits it into two (possibly empty) trees g_1, g_2 with node bidegrees $(j_m, k_m), m \in [M_1]$ and $(j_m, k_m), m \in [M_1 + M_2] \setminus [M_1]$ where $M = M_1 + M_2$. When we condition on the bidegree of v , (3.67) shows that the remaining degrees form independent families:

$$\begin{aligned} \mathbb{P}[w_m \in \mathcal{N}_{j_m k_m}, m \in [M], g | X, v \in \mathcal{N}_{jk}] &= \mathbb{P}[w_m \in \mathcal{N}_{j_m k_m}, m \in [M_1], g_1 | X, v \in \mathcal{N}_{jk}] \\ &\times \mathbb{P}[w_m \in \mathcal{N}_{j_m k_m}, m \in [M_1 + M_2] \setminus [M_1], g_2 | X, v \in \mathcal{N}_{jk}]. \end{aligned} \quad (3.74)$$

We call this deep property of the general configuration graph the *locally tree-like independence property* (LTI property). It provides the key to unravelling cascade dynamics in large configuration graphs.

3.5.3 Approximate ACG Simulation

As we observed in Section 3.2, the first step of the configuration graph construction draws a sequence $(j_i, k_i)_{i \in [N]}$ of node types that is iid with the correct distribution P , but is only feasible, $\sum_i (k_i - j_i) = 0$, with small probability. Step 2 of the exact ACG algorithm in Section 3.5.1 is even less feasible in practice. Practical simulation algorithms address the first problem by ‘‘clipping’’ the drawn node bidegree sequence when the discrepancy $D = D_N := \sum_i (k_i - j_i)$ is not too large, meaning it is adjusted by a small amount to make it feasible, without making a large change in the joint distribution. The following clipping method generalizes the method introduced by

[23] who verify that the effect of clipping vanishes with high probability as $N \rightarrow \infty$. The problem of Step 2 is overcome by an approximate sequential wiring algorithm.

The *approximate assortative configuration simulation algorithm* for multigraphs of size N , parametrized by the node-edge degree distribution pair (P, Q) that satisfy Assumption 5, involves choosing a suitable threshold $T = T(N)$ and modifying the steps identified in Section 3.5.1:

1. Draw a sequence of N node-degree pairs $X = ((j_1, k_1), \dots, (j_N, k_N))$ independently from P , and accept the draw if and only if $0 < |D| \leq T(N)$. When the sequence $(j_i, k_i)_{i \in [N]}$ is accepted, the sequence is adjusted by adding a few stubs, either in- or out- as needed. First draw a random subset $\sigma \subset \mathcal{N}$ of size $|D|$ with uniform probability $\binom{N}{|D|}^{-1}$, and then define the feasible sequence $\tilde{X} = (\tilde{j}_i, \tilde{k}_i)_{i \in [N]}$ by adjusting the degree types for $i \in \sigma$ as follows:

$$\tilde{j}_i = j_i + \xi_i^-; \quad \xi_i^- = \mathbf{1}(i \in \sigma, D > 0) \quad (3.75)$$

$$\tilde{k}_i = k_i + \xi_i^+; \quad \xi_i^+ = \mathbf{1}(i \in \sigma, D < 0). \quad (3.76)$$

2. Conditioned on \tilde{X} , the result of Step 1, randomly wire together available in and out stubs *sequentially*, with suitable weights, to produce the sequence of edges W . At each $\ell = 1, 2, \dots, E$, match from available in-stubs and out-stubs weighted according to their degrees j, k by

$$C^{-1}(\ell) \frac{Q_{kj}}{Q_k^+ Q_j^-}. \quad (3.77)$$

In terms of the bivariate random process $(e_j^-(\ell), e_k^+(\ell))$ with initial values $(e_j^-(1), e_k^+(1)) = (e_j^-, e_k^+)$ that at each ℓ counts the number of available degree j in-stubs and degree k out-stubs, the ℓ dependent normalization factor $C(\ell)$ is given by:

$$C(\ell) = \sum_{jk} e_j^-(\ell) e_k^+(\ell) \frac{Q_{kj}}{Q_k^+ Q_j^-}. \quad (3.78)$$

The main results now guarantee important convergence properties that show that for large N the clipping and sequential wiring defects get vanishingly small. The first theorem, as proved by Chen and Olvera-Cravioto [23], addresses the clipping step and shows that: (1) the acceptance probability goes to 1; (2) for any fixed finite M , the joint distribution of the random degrees $(\tilde{j}_i, \tilde{k}_i)_{i \in [M]}$ converges in distribution to the target of M i.i.d. bivariate P -distributed random variables $(\hat{j}_i, \hat{k}_i)_{i \in [M]}$ as $N \rightarrow \infty$; and (3) the empirical node bidegree probabilities converge (in probability) to the target probabilities P_{jk} .

Theorem 6. Fix $\delta \in (0, 1/2)$, and for each N let the threshold be $T(N) = N^{1/2+\delta}$. Then:

1. The acceptance probability $\mathbb{P}[\mathcal{D}_N] \rightarrow 1$ as $N \rightarrow \infty$;
2. For any fixed finite M , Λ , and bounded function $f : (\mathbb{Z}_+ \times \mathbb{Z}_+)^M \rightarrow [-\Lambda, \Lambda]$

$$|\mathbb{E}[f((\tilde{j}_i, \tilde{k}_i)_{i=1, \dots, M})] - \mathbb{E}[f((\hat{j}_i, \hat{k}_i)_{i=1, \dots, M})]| \rightarrow 0 ; \quad (3.79)$$

3. The following limits in probability hold:

$$\frac{1}{N} \tilde{u}_{jk} \xrightarrow{P} P_{jk}, \quad \frac{1}{N} \tilde{u}_k^+ \xrightarrow{P} P_k^+, \quad \frac{1}{N} \tilde{u}_j^- \xrightarrow{P} P_j^-. \quad (3.80)$$

The main theorem for Step 2 was proved first by F. Pourbabaee [64] and states that the empirical distribution of edges in the sequence are with high probability clustered near the values of Q , and moreover, that for any fixed L , the joint degree distribution of the first L edges of the matching sequence converges (in distribution) to an independent sequence of identical Q distributed random variables.

Theorem 7. *In the approximate assortative configuration graph construction with probabilities P, Q , the following convergence properties hold as $N \rightarrow \infty$.*

1. *The fraction of type (k, j) edges in the matching sequence $(k_\ell, j_\ell)_{\ell \in [E]}$ concentrates with high probability around the nominal edge distribution Q_{kj} :*

$$\frac{e_{kj}}{E} = Q_{kj} + o(1). \quad (3.81)$$

2. *For any fixed finite number L , the first L edges $\ell \in [L]$ have degree sequence $(k_\ell, j_\ell)_{\ell \in [L]}$ that converges in distribution to $(\hat{k}_\ell, \hat{j}_\ell)_{\ell \in [L]}$, an independent sequence of identical Q distributed random variables.*

Taken together, these two theorems show that when the approximate version of Steps 1 and 2 are combined, the simulation algorithm generates finite assortative configuration graphs, consistent with the target distributions P, Q asymptotically for N large enough.

3.6 Measures of Network Topology

The term *network topology* is commonly used to refer to a wide variety of characteristics observed in the random graphs underlying large scale networks, both synthetic and in the real world. Measures of network topology typically amount to a summary statistic that tells us something important about the way nodes connect or about the relative importance of different node and edge types.

3.6.1 Connectivity

Overall *connectivity* of the network can be measured by the fraction of the number of actual directed links to the number of potential links. Thus, in a finite network with N nodes and E directed edges, the connectivity is given by $C = \frac{E}{N(N-1)}$. Since

our networks are typically sparse enough that $C \rightarrow 0$ for large N , we often focus instead on the mean degree $z = E/N$.

3.6.2 Measures from the Degree Distributions

In addition to measuring moments of the degree distribution P_{jk} , such as the mean (in- and out-) degree $z = \sum_{j,k} jP_{jk} = \sum_{j,k} kP_{jk}$, network practitioners also focus on the tail properties. *Tail exponents* are defined by large graph limits:

$$\alpha^\pm = -\limsup_{j \rightarrow \infty} \frac{\log P_j^\pm}{\log j}.$$

Finite tail exponents are indicators of what are called *Pareto tails*, and signal the existence of non-negligible numbers of *hubs*, or highly connected nodes that can be significant focal points for systemic risk. Log-log plots capture the characteristic tail exponents as the negative of the slope of the best fit line, above a certain cutoff level. Clauset et al [26] provides the definitive statistical inference method for determining Pareto tail exponents for random samples from a distribution with power law tail.

3.6.3 Centrality Measures

Centrality measures require the full adjacency matrix $M(g)$, and aim to decide the relative importance of nodes. At their heart they rest on the fact that the k th power of the adjacency matrix provides the number of (directed) k -step paths between nodes. Different centrality measures typically formalize the idea that important nodes are those that have important neighbours, by summing over these paths with different weights. For directed graphs where we need to distinguish forward paths from backward paths, centrality measures come in several versions.

1. *Degree centrality*: For undirected graphs, this is simply the degree of the node. For directed graphs, one refers to the in-degree and out-degree centralities.
2. *Eigenvalue centrality*: By the Perron-Frobenius theorem for non-negative matrices, there exists a non-negative right-eigenvector of M , and this eigenvector is associated with the maximal eigenvalue λ (which is necessarily positive). If M is irreducible, or equivalently if the directed graph g is strongly connected, then there is a unique right-eigenvector, call it $\mathbf{x}^+ = [x_1^+, \dots, x_N^+]$, normalized so that $\sum_i x_i^+ = 1$. The component for node i is positive, $x_i^+ > 0$, and is called its forward eigenvalue centrality measure. Conversely, the backward eigenvalue centrality measures x_i^- derive from the maximal left-eigenvector of M . Both measures can be easily computed by power iteration using the formulas:

$$x_i^+ = \lambda^{-1} \sum_j M_{ij} x_j^+; \quad x_i^- = \lambda^{-1} \sum_j M_{ji} x_j^-.$$

In more generality, one can apply the same measure to a weighted adjacency matrix, where the link weights are arbitrary non-negative values.

3. *Katz centrality*: This is a parametric family of measures that generalizes both degree and eigenvalue centrality by penalizing long paths with an attenuation factor. For any $\alpha \in (0, \lambda^{-1})$, we define the forward and backward Katz-centrality indices by

$$x_i^{\alpha,+} = \sum_{k=1}^{\infty} \sum_j \alpha^k (M^k)_{ij}; \quad x_i^{\alpha,-} = \sum_{k=1}^{\infty} \sum_j \alpha^k (M^k)_{ji}.$$

We can see that $x_i^{\alpha,+} = \sum_j \alpha M_{ij} (x_j^{\alpha,+} + 1)$ from which it can be shown that the eigenvalue centrality x_i^+ is proportional to the limit of $x_i^{\alpha,+}$ as α approaches λ^{-1} from below.

4. *Betweenness centrality*: This measure differs from the others in considering only *shortest paths* between nodes. For two nodes $v \neq v'$, we define $\sigma_{vv'}$ to be the number of shortest directed paths from v to v' ; for any third node $w \neq v, v'$, $\sigma_{vv'}(w)$ is the number of shortest directed paths from v to v' that go through w . Then the betweenness centrality of node w is defined by the formula

$$b_w = \sum_{v,v' \neq w} \frac{\sigma_{vv'}(w)}{\sigma_{vv'}}.$$

3.6.4 Clustering Coefficients

Clustering in social networks refers to the propensity of the friends of our friends to be our friends. In a general setting it means the likelihood that a connected triple of nodes forms a triangle. It is usually measured by the *clustering coefficient*, the ratio of the number of triangles to connected triples,

$$C(g) = \frac{3 \times (\text{number of triangles})}{(\text{number of connected triples})},$$

where the factor of 3, corresponding to the number of connected triples in a triangle, ensures that $C \in [0, 1]$.

In directed networks, an even more basic node of clustering is *reflexivity*, that refers to the fraction of the number of node pairs that have a *reflexive* pair of edges (i.e. an edge pointing in both directions) to the total number of directed edges. For triples of nodes, accounting for the direction of edges means there are two different kinds of triangles in directed graphs, and three different connected triples. One can use the results of Section 3.5.2 to compute asymptotic values for all these ratios in the ACG model.

3.6.5 Connectivity and Connected Components

Contagion can only propagate across connected components in a network, and thus the sizes of connected subgraphs of a network are additional measures of its susceptibility to large scale contagion. A *strongly connected component* (SCC) of a network is a subgraph each of whose nodes are connected to any other node by a path of downstream edges, and which is maximal in the sense that it is not properly contained in a larger SCC. For any SCC, one defines its *in-component* to be the maximal subgraph of nodes that are downstream connected to the SCC. Similarly, its *out-component* is the maximal subgraph of nodes that are upstream-connected to the SCC. A *weakly connected component* (WCC) of a network is a maximal subgraph each of whose nodes are connected to any other nodes by a path consisting of undirected edges. A WCC may contain a multitude of side branches that are not downstream or upstream connected to the SCC: pieces called *tendrils* and *tubes* form subgraphs that are not connected to the SCC but are either upstream connected to the in-component or downstream connected to the out-component, or both. In undirected graphs, there is no distinction between strong and weak connectivity, and thus any WCC is an SCC.

In random graph models, the probability distribution of sizes of connected components is a topic of interest. When we consider infinite random graphs, a critical question is whether there is strongly connected component that is infinite or even a positive fraction of the entire network. When this happens in a random graph model, the infinite SCC is typically unique, and we call it the *giant strongly connected component* (GSCC). Its associated giant in- and out-components are called G-IN and G-OUT, and GWCC denotes the *giant weakly connected component*. Clearly we have the inclusions:

$$\text{GSCC} = \text{G-IN} \cap \text{G-OUT} \subset \text{G-IN} \cup \text{G-OUT} \subset \text{GWCC}$$

The complement of the GWCC falls into a disjoint collection of finite sized weakly-connected pieces. Figure 3.2 shows a schematic “bow-tie” rendering of the various connected components of a typical directed network, in this instance the World Wide Web as it appeared in 1999.

It has been found that the existence of a giant connected component in a large network is the single most relevant and important distinguishing characteristic that determines its susceptibility to domino-like cascades. The next chapter develops this intuitive observation into a theory called *percolation on random graphs*, and explores its relation to cascade modelling.




Fig. 3.2 The connected components of the World Wide Web in 1999. (Source: [19].)

Chapter 4

Percolation and Cascades

Abstract The right kind of connectivity turns out to be both necessary and sufficient for large scale cascades to propagate in a network. After outlining percolation theory on random graphs, we develop an idea known as “bootstrap percolation” that proves to be the precise concept needed for unravelling and understanding the growth of simple network cascades. These principles are illustrated by the famous Watts model of information cascades.

Keywords: Graph connectivity, branching process, bond and site percolation, bootstrap percolation, vulnerable edge, vulnerable cluster, Watts’ cascade model, with-out regarding property, local tree-like independence.

Before we consider fundamental questions of cascading shocks on random networks, we can answer simpler questions about whether or not the network is highly connected. Obviously cascades cannot propagate if the network isn’t sufficiently connected, but what is far from obvious is that the right kind of connectivity is also a sufficient condition for the possibility of large scale cascading to occur. Fortunately, there is a rich and beautiful theory on the connectivity of networks known as *percolation*, and as it will turn out, this idea of percolation actually permeates our problems of cascades. In the first section of this chapter, we touch on the highlights of percolation theory on random configuration graphs the theory of which we have carefully developed in the previous chapter. In this setting, exact formulas can be obtained, whose proofs, which we will sketch, give detailed insight into the nature of percolation.

On a given undirected graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$, we can consider its connected components, called “clusters”, and order them in decreasing size from largest to smallest. If the graph is infinite, then the largest cluster, which we denote by \mathcal{C} , has the possibility to itself be infinite, or somewhat more strongly, to consist of a positive fraction of the whole graph. When this occurs, we say the graph has a *giant cluster* or *giant component*. Percolation theory begins by considering a finite size random graph model, and aims to show that there will be a large cluster exceeding a certain size with high probability whenever there exists a giant cluster in the infinite graph limit.

The strongest results on percolation on random graphs apply to the family of configuration graphs, because when a configuration graph is large enough and the density of edges is not too great, it has the “locally treelike” property that with high probability the graph has few short cycles or closed loops. The “locally treelike” property provides a key to obtaining exact analytical formulas for percolation. Our intuition about percolation on configuration graphs stems because when we start from a random node, and consider growing its connected component one neighbourhood at a time, the process looks like a growing tree or pure branching process, of the type known as a Galton-Watson process. The potential for this graph component to become large is measured by the potential for the associated GW process to be unbounded. This relation to Galton-Watson processes can be established rigorously in the asymptotic regime as the number of nodes goes to infinity. Therefore before considering percolation, we review the essential properties of branching processes.

4.1 Branching Processes

A *Galton-Watson process*, the simplest class of *branching process*, describes a population evolving in time, starting with a single individual in the 0th generation. In the n th (non-overlapping) generation, let there be Z_n individuals. Each individual i of the n th generation, $n \geq 0$, is assumed to produce a random number $X_{n,i}$ of children or offspring, each drawn independently from an identical distribution X on the non-negative integers.

The central question to answer about a GW process is whether the population will ultimately survive or go extinct. The answer is best expressed in terms of the generating function of the X distribution:

$$g(s) := \mathbb{E}s^X = \sum_{k \geq 0} P_k s^k, \quad P_k = \mathbb{P}[X = k].$$

We also need the generating functions $H_n(s) := \mathbb{E}s^{Z_n}$ for $n \geq 0$, and the fundamental dynamic identity:

$$Z_n = \sum_{i=1}^{Z_{n-1}} X_{i,n}. \quad (4.1)$$

Using iterated conditioning, one computes

$$H_n(s) = \mathbb{E}s^{Z_n} = \mathbb{E} \left[\mathbb{E} \left[\prod_{i=1}^{Z_{n-1}} s^{X_{n,i}} \mid Z_{n-1} \right] \right] = \mathbb{E} \left[\prod_{i=1}^{Z_{n-1}} \mathbb{E} \left[s^{X_{n,i}} \mid Z_{n-1} \right] \right] \quad (4.2)$$

$$= \mathbb{E} \left[(g(s))^{Z_{n-1}} \right] = H_{n-1}(g(s)) \quad (4.3)$$

where the third equality follows from the mutual conditional independence property of the two index collection $X_{n,i}$. This composition identity, when iterated, leads to the essential formula:

$$H_n = \underbrace{g \circ g \cdots \circ g}_{n \text{ factors}} = g \circ H_{n-1}. \quad (4.4)$$

The extinction probability is $\eta := \mathbb{P}[\exists n : Z_n = 0]$ and for each n , define $\eta_n = \mathbb{P}[Z_n = 0] = H_n(0)$. Since $\{Z_{n-1} = 0\} \subset \{Z_n = 0\}$, the sequence η_n is increasing, and converges to η . Since $\eta_n = H_n(0)$,

$$\eta_n = g(H_{n-1}(0)) = g(\eta_{n-1}).$$

Now like any generating function, g satisfies the properties $g \geq 0, g' \geq 0, g'' \geq 0$ and has $g(1) = 1$. It is thus continuous, increasing and convex on $[0, 1]$. Recall also that $g'(1) = \mathbb{E}X$ and $g''(1) = \mathbb{E}X(X-1)$, one or both of which may be infinite. Therefore, by continuity,

$$\eta = \lim_{n \rightarrow \infty} \eta_n = \lim_{n \rightarrow \infty} g(\eta_{n-1}) = g\left(\lim_{n \rightarrow \infty} \eta_{n-1}\right) = g(\eta),$$

that is, $\eta \in [0, 1]$ is a fixed point of g . By the convexity of g , there can be at most 2 fixed points. By induction, one can easily verify that if ψ is any fixed point of g , then $\eta_n \leq \psi$ for all n , and hence $\eta \leq \psi$: Note that $\eta_0 = P_0 \leq \psi$, and if $\eta_{n-1} \leq \psi$ then $\eta_n = g(\eta_{n-1}) \leq g(\psi) = \psi$. Thus η is the smallest fixed point of g on $[0, 1]$. This analysis leads to the following conclusions:

Theorem 8. *The extinction probability $\eta \in [0, 1]$ is the smallest fixed point of g .*

1. *If $\mathbb{E}X > 1$, then $\eta < 1$, which says that with positive probability $1 - \eta$ the population will survive forever.*
2. *If $\mathbb{E}X \leq 1$, then apart from a trivial exception, $\eta = 1$ and the population becomes extinct almost surely. The single trivial exception is if $\mathbb{E}X = 1$ and $g''(1) = 0$, which implies $\mathbb{P}[X = 1] = 1$ and that the population remains 1 for all time.*

Case (1), when survival is possible, is called the *supercritical* case. The case of almost sure extinction subdivides: case $\mathbb{E}X < 1$ is called *subcritical*, and the case $\mathbb{E}X = 1$ and $g''(1) > 0$ is called *critical*.

From our comments about the relation between percolation and GW processes, it should not be a surprise that in the next section we will find that the possibility of a giant cluster boils down to conditions similar to those given in this theorem.

4.2 Percolation on Configuration Graphs

The paper of Janson [46] considers percolation on undirected configuration (multi)graphs for each finite N based on a general degree sequence $\mathbf{d} = (d_v)_{v=1}^N$ with an even number of stubs $\sum_v d_v = 2E$ (note that usually we neglect to show explicit N dependence). In what follows, we consider asymptotics of the sequence of degree sequences as $N \rightarrow \infty$ and use a number of further conventions: *w.h.p.* means “with probability tending to 1 as $N \rightarrow \infty$ ”; \xrightarrow{P} means “convergence in probability as

$N \rightarrow \infty$ ”; the symbols o_P and O_P are used in the standard way. The finite random configuration multigraph model is denoted by $G^*(N, \mathbf{d})$ and $G(N, \mathbf{d})$ denotes $G^*(N, \mathbf{d})$ under the condition that the multigraph is simple, that is that the multigraph is in fact a graph.

Assumption 6. The sequence of undirected multigraphs $G^*(N, \mathbf{d})$ is *well-behaved* in the sense that there exists a probability distribution $(P_k)_{k=0,1,\dots}$ over nonnegative integers such that:

- The empirical degree density converges in distribution:

$$N^{-1} \sum_v \mathbf{1}(d_v = k) \xrightarrow{P} P_k .$$

- The mean degree is finite and positive: $N^{-1} \sum_v \sum_k k \mathbf{1}(d_v = k) \xrightarrow{P} z$ where $z := \sum_{k=1}^{\infty} k P_k < \infty$.

For undirected random graphs, the independent edge probability distribution is $Q_{kk'} = (1 + \delta_{kk'}) Q_k Q_{k'}$ where $Q_k = \frac{k P_k}{z}$ has a nice interpretation as the “size-biased” distribution of d_v : if we select an arbitrary edge $\ell = (v, w)$, Q_k denotes the probability that either of the nodes attached to ℓ has $k - 1$ remaining edges. In other words,

$$Q_k = \mathbb{P}[k_w = k | w \in \mathcal{N}_v] . \quad (4.5)$$

This fact is relevant to understanding the growth of successive neighbourhoods of a randomly selected node. The degree of v , or the number of neighbours of v , has PMF P . However, neglecting the possibility of cycles, we can see that each neighbour of v has k new neighbours with probability $P_k^* := Q_{k+1}$. Therefore, the growth of neighbourhoods of a random node v , that is the growth of the cluster containing v , approximately follows a GW process whose zeroth generation has offspring probability governed by P while for each successive generation, the offspring distribution is given by P^* .

If we let $g(x) = \sum_k P_k x^k$ be the generating function of P then $g^*(x) = g'(x)/z$ is the generating function of P^* . Now, from Theorem 8 for GW processes, the extinction probability ξ^* of any node other than the root node is determined by the condition $\xi = g^*(\xi)$. If there is a fixed point $\xi^* < 1$ then the GW process is supercritical and non-extinction occurs with positive probability $1 - \xi^*$, otherwise $\xi^* = 1$ is the unique fixed point and extinction occurs almost surely. The key insight is that the same condition determines whether the random graph is supercritical or not. The following is a refinement of the main theorem of [59] due to [47], which asserts the precise conditions under which the random graph has a giant cluster (i.e. it is supercritical).

Theorem 9. Consider the random multigraph sequence $G^*(N, \mathbf{d})$ satisfying Assumption 6. Let $g(x)$ be its generating function and let \mathcal{C} be the largest cluster. Then the empirical probabilities of a random node or edge being in \mathcal{C} are governed by the following asymptotic properties:

1. If $\sum_k k(k-2) P_k > 0$, then there is a unique $\xi \in (0, 1)$ such that $g^*(\xi) = \xi$ and

$$\mathbb{P}[v \in \mathcal{C}] \xrightarrow{P} 1 - g(\xi) > 0, \quad (4.6)$$

$$\mathbb{P}[v \in \mathcal{C} \cap \mathcal{N}_k] \xrightarrow{P} P_k(1 - \xi^k), \text{ for every } k \geq 0, \quad (4.7)$$

$$\mathbb{P}[\ell \in \mathcal{C}] \xrightarrow{P} (1 - \xi^2) > 0. \quad (4.8)$$

2. If $\sum_k k(k-2) P_k \leq 0$, then unless $P_2 = 1$, $\mathbb{P}[v \in \mathcal{C}] \xrightarrow{P} 0$ and $\mathbb{P}[\ell \in \mathcal{C}] \xrightarrow{P} 0$.

3. In the trivial special case when $P_2 = 1$, then $\sum_k k(k-2) P_k = 0$ and $\mathbb{P}[v \in \mathcal{C}] \xrightarrow{P} 1$ and $\mathbb{P}[\ell \in \mathcal{C}] \xrightarrow{P} 1$.

This theorem can be phrased in w.h.p. terms: for example, in part one means that with high probability each $G^*(N, \mathbf{d})$ has a giant cluster if $\sum_k k(k-2) P_k > 0$. The result can be understood in an intuitive way based on the asymptotic locally tree-like property of large configuration graphs expressed in equation (3.74) that says the outgoing edges of any node connect to random subgraphs that can be treated as independent. We next introduce a useful definition that will enable us to interpret ξ as a probability of a specific event.

Definition 8. Suppose a node property \mathcal{P} is *local*, meaning the condition $w \in \mathcal{P}$ is determined by certain conditions on nearest neighbours $v \in \mathcal{N}_w$. Then, for any directed edge (v, w) , we say that w satisfies the local property \mathcal{P} *without regarding* v , and write “ $w \in \mathcal{P}$ WOR v ”, if the property is determined by these conditions on all nearest neighbours $v' \in \mathcal{N}_w$ excluding v .

This can be illustrated in our problem because connectedness is a local node property. Thus, for any directed edge (v, w) ,

$$\{w \in \mathcal{C}^c \text{ WOR } v\} = \{(\mathcal{N}_w \setminus v) \cap \mathcal{C} = \emptyset\}$$

whereas

$$\{w \in \mathcal{C}^c\} = \{\mathcal{N}_w \cap \mathcal{C} = \emptyset\}.$$

Now we can analyse the probability $a = \mathbb{P}[v \in \mathcal{C}^c]$ in terms of WOR probabilities. If v has degree k , then $v \in \mathcal{C}^c$ is a local property equivalent to $w \in \mathcal{C}^c$ WOR v for each of the k neighbours $w \in \mathcal{N}_v$. Moreover, by the LT property, these k events form a mutually independent collection. Thus, if we define

$$\xi = \mathbb{P}[w \in \mathcal{C}^c \text{ WOR } v | w \in \mathcal{N}_v] \quad (4.9)$$

then

$$\mathbb{P}[v \in \mathcal{C}^c] = \sum_k P_k \xi^k := g(\xi). \quad (4.10)$$

We can verify that $\xi = g^*(\xi)$ by very similar logic. If w has degree k and $w \in \mathcal{N}_v$ then $w \in \mathcal{C}^c$ WOR v is equivalent to $w' \in \mathcal{C}^c$ WOR w for each of the remaining $k-1$ neighbours $w' \in \mathcal{N}_w \setminus v$, and we find

$$\xi = \sum_k \xi^{k-1} \mathbb{P}[w \in \mathcal{N}_k | w \in \mathcal{N}_v] = \sum_k \xi^{k-1} Q_k = g^*(\xi). \quad (4.11)$$

Since g^* is itself the generating function of the P^* distribution, the discussion leading to Theorem 8 applies here, and we find the three cases for percolation, supercritical, critical and subcritical, correspond to the cases $g^{*'}(1) > 1$, $g^{*'}(1) = 1$, and $g^{*'}(1) < 1$. We also have

$$\mathbb{P}[v \in \mathcal{C}^c, v \in \mathcal{N}_k] = \mathbb{P}[v \in \mathcal{C}^c | k_v = k] P_k = P_k \xi^k$$

which verifies (4.7). Finally, for $\ell = (v, w)$, one can easily see that $\ell \in \mathcal{C}^c$ means both $w \in \mathcal{C}^c$ WOR v and $v \in \mathcal{C}^c$ WOR w . (4.8) follows since these two events are independent and have probability ξ . We see that the fixed point ξ of the equation $g^*(\xi) = \xi$ can be interpreted as a “without regarding” probability.

Janson [46] has proven that a configuration multigraph is simple and hence a graph with positive probability. It follows that the above multigraph result that holds with high probability (*w.h.p.*) also holds *w.h.p.* for graphs.

4.3 Site Percolation

Site percolation on a random graph amounts to asking about the connected clusters of subgraphs created by the deletion of random nodes and their edges, and it has been found that the previous theorem extends beautifully to this more general setting. If we delete nodes v (and their incident edges) of a configuration multigraph independently with probabilities $1 - \delta_k$ determined by their degree k , it can be shown that the resultant subgraph is also a configuration multigraph with a new degree distribution P' . The following theorem due to [45] gives the rigorous statement:

Theorem 10. *Consider site percolation with deletion probabilities $1 - \delta_k \in [0, 1]$, on the random configuration multigraph sequence $G^*(N, \mathbf{d})$ satisfying Assumption 6. Then the subgraph has a giant cluster \mathcal{C} with high probability if and only if*

$$\sum_k k(k-1) \delta_k P_k > z := \sum_k k P_k. \quad (4.12)$$

1. If (4.12) holds then there is a unique $\xi \in (0, 1)$ such that

$$\sum_k k \delta_k P_k (1 - \xi^{k-1}) = z(1 - \xi) \quad (4.13)$$

and then

$$\mathbb{P}[v \in \mathcal{C}] \xrightarrow{\text{P}} \sum_{k \geq 1} \delta_k P_k (1 - \xi^k) > 0, \quad (4.14)$$

$$\mathbb{P}[\ell \in \mathcal{C}] \xrightarrow{\text{P}} \frac{2}{z} \left[(1 - \xi) \sum_k k \delta_k P_k - \frac{1}{2} (1 - \xi)^2 \sum_k k P_k \right]. \quad (4.15)$$

2. If (4.12) does not hold, $\mathbb{P}[v \in \mathcal{C}] \xrightarrow{\text{P}} 0$ and $\mathbb{P}[\ell \in \mathcal{C}] \xrightarrow{\text{P}} 0$.

Why is this apparently special result relevant to our problem of financial cascades? The detailed answer will be revealed shortly, but in the meantime, we can explain that the key to understanding a cascade in a financial network is to focus on banks (or more generally, edges) that are “vulnerable” in the sense that they default if only one debtor bank defaults, and to delete all other banks (or edges). The resultant network of vulnerable banks has a giant in-cluster $G - \text{IN}_V$ if and only if the analogue of (4.12) for directed graphs holds. In this case, if any bank owing money to a bank in this giant vulnerable in-cluster were to default, then all of the giant strongly connected vulnerable cluster GSCC_V , and thereafter the giant vulnerable out-cluster $G - \text{OUT}_V$, must inevitably eventually default. The result will be a *global cascade*. The size of the global cascade will clearly be at least as big as $G - \text{OUT}_V$. However, it may be much larger since in general some banks that are not vulnerable to a single debtor’s default are vulnerable to two or more debtor defaults. Although such secondary defaults may seem unlikely in a locally-tree-like network (since they are impossible on a tree), they are likely to happen in a large LTI network if $G - \text{OUT}_V$ is itself a large fraction of the network. This fact illustrates why we use the term “locally tree-like”: Even if there are few short cycles in a large LTI network, there are many large cycles that connect the network on a large scale, and allow secondary defaults to occur starting from a single seed default.

To understand this better, we now follow the idea that each bank has a threshold number of defaulted neighbours that when exceeded implies its own default, which is reminiscent of a concept that goes by the name of *bootstrap percolation*.

4.4 Bootstrap Percolation

We have just seen that the theory of site percolation can give an indication of the conditions under which a global cascade can occur in any regular lattice or random network. When any node adjacent to the giant vulnerable in-cluster, then the entire giant strongly connected vulnerable cluster will be triggered. However, the resultant global cascade may be much larger in extent than this, because of the possibility that less vulnerable nodes may eventually succumb if they are susceptible when more than one neighbour is triggered. *Bootstrap percolation* considers the general problem of dynamical cascades on random networks where each *susceptible* (or *inactive*) node has a threshold r and will *succumb* or become *activated* when the number of its succumbed (activated) neighbours exceeds r . Furthermore it is assumed that nodes stay activated once becoming activated.

In general, given a set of nodes $v \in \mathcal{N}$ (which may be a graph or a regular lattice) and a *threshold function* $r : v \in \mathcal{N} \rightarrow \{0, 1, \dots\}$, then bootstrap percolation is defined by setting $\mathcal{A}_0 = \{v \in \mathcal{N} \mid r(v) = 0\}$ and for $t \geq 1$,

$$\mathcal{A}_t = \{v \in \mathcal{N} \mid r(v) \leq |\mathcal{N}_v \cap \mathcal{A}_{t-1}|\}.$$

One can say that the threshold model *percolates* if the closure $\mathcal{A} = \cup_{t \geq 0} \mathcal{A}_t$ grows linearly with \mathcal{N} as $|\mathcal{N}| \rightarrow \infty$.

The name *bootstrap percolation* was introduced in a paper [22] in a statistical physics context to denote this type of dynamic percolation theory, and the subject has had a rich development since then. The recent paper [8] focusses on a version of bootstrap percolation on the random regular graph, and contains results and references that are quite relevant to financial cascades.

Our investigations will now follow a similar type of process as we next consider the simplest cascade model, that is essentially bootstrap percolation on the undirected Poisson random graph. The Watts Cascade Model introduced in [72] has been the inspiration for much subsequent work on financial cascades, and provides a prototype for the analytical techniques we shall develop in the remainder of the book.

4.5 Watt's 2002 Model of Global Cascades

This classic paper [72] at the heart of network science considers an undirected random graph that models a social network, with nodes representing people linked to their friends. Each individual is assigned a random threshold and is deemed to adopt a specific new technology as soon as the number of their adopting friends exceeds this threshold. The model addresses the question “how can one understand how social contacts influence people to adopt a new technology, such as an iPhone or Android phone?”. The mathematical analysis of the model focuses on the transmission of “adoption” shocks over the friendship links of the network that represent how someone who adopts influences their social contacts to adopt. It determines conditions under which these shocks accumulate and create a large scale *adoption cascade*, wherein the product eventually gains a large share of the market. This simple cascade model will serve as a template for studying financial cascades such as the propagation of defaults and liquidity shocks. Much of the systemic risk modelling that follows in the remainder of this book will be based on the ideas underlying this basic construction, and so we provide here a complete description of the basic Watts' model.

4.5.1 The Framework

The framework of the Watts model will serve as a template for more complicated financial cascade models that will follow. It consists of two layers: the *skeleton* graph of individuals linked to their friends, and an assignment of random values to people that represents the threshold number of friends having adopted the technology that will trigger that person to adopt. Amongst these nodes are the *early adopters* or *seed nodes* whose thresholds are zero. We don't need a third layer for exposures since friendship links are all trivially set to have weights equal to 1. More precisely,

1. The skeleton graph is the random undirected Gilbert graph model $G(N, p)$ of Section 3.2 with edge probability p and mean degree $z = (N - 1)p$. Here $N \in \{1, 2, \dots\}$ denotes the finite number of people and we introduce the node-degree distribution $P = (P_k)_{k=0,1,\dots}$ with Binomial probabilities $P_k = \mathbb{P}[v \in \mathcal{N}_k] = \text{Bin}(N - 1, p, k)$.¹ Recall that $\text{Bin}(N - 1, z/(N - 1)) \xrightarrow{D} \text{Pois}(z)$ as $N \rightarrow \infty$.
2. The thresholds are random integer variables $\bar{\Delta}_v \geq 0$ meaning v will adopt when at least $\bar{\Delta}_v$ of her friends have adopted. Conditioned on the skeleton, the collection $\{\bar{\Delta}_v\}$ is assumed to be independent, and distributed depending on the degree k_v , that is,

$$\mathbb{P}[\bar{\Delta}_v = x | v \in \mathcal{N}_k] = d_k(x), \quad k, x \geq 0$$

for discrete probability distributions $d_k(\cdot)$ parametrized by k . Let the cumulative probability distributions be given by

$$D_k(x) := \mathbb{P}[\bar{\Delta}_v \leq x | v \in \mathcal{N}_k] = \sum_{y=0}^x d_k(y), \quad k, x \geq 0.$$

3. For each $n = 0, 1, \dots$, we let \mathcal{D}_n denote the set of nodes that have adopted after n steps of the adoption cascade. These sets are defined inductively by $\mathcal{D}_0 = \{v \in \mathcal{N} : \bar{\Delta}_v = 0\}$ and

$$\mathcal{D}_n = \{v \in \mathcal{N} : \bar{\Delta}_v \leq |\mathcal{N}_v \cap \mathcal{D}_{n-1}|\} \quad (4.16)$$

for $n \geq 1$. Their conditional probabilities are defined to be

$$p_k^{(n)} := \mathbb{P}[v \in \mathcal{D}_n | v \in \mathcal{N}_k]. \quad (4.17)$$

A randomly selected degree k node will be an early adopter or seed with probability $p_k^{(0)} := d_k(0)$. Also, observe that the sequence $\{\mathcal{D}_n\}_{n \geq 0}$ is increasing and converges to \mathcal{D}_∞ , the set of nodes that eventually adopt.

In [72], the benchmark threshold specification was $d_k(x) = \rho \mathbf{1}(x = 0) + (1 - \rho) \mathbf{1}(x = \lceil 0.18k \rceil)$ which means early adopters are selected independently with uniform probability $p_k^{(0)} = \rho$ and the remaining nodes adopt when at least a fraction $\phi = 0.18$ of their friends have adopted.

¹ For any $0 \leq k \leq N$ and $p \in [0, 1]$, the Binomial probability $\text{Bin}(N, p, k) = \binom{N}{k} p^k (1 - p)^{N-k}$ is the probability of exactly k successes in N independent Bernoulli trials each with success probability p .

The main results about the Watts model derive from a property based on the *without regarding* concept introduced in Definition 8 in Section 4.2. This property implies that a source of feedback that can be expected in general cascade models is not present in the Watts model. Let D_v^n be the indicator function for the node set \mathcal{D}_n , and $\tilde{D}_{v,w}^n$ be the indicator for the set of directed edges (v, w) such that $v \in \mathcal{D}_n$ WOR w . That is, with initial conditions $D_v^{-1} = 0$, it holds that for $n \geq 0$:

$$D_v^n = \mathbf{1}(v \in \mathcal{D}_n) = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w'}^{n-1} \right) \quad (4.18)$$

$$\tilde{D}_{v,w}^n = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w'}^{n-1} \mathbf{1}(w' \neq w) \right). \quad (4.19)$$

The next theorem shows that as intuitive as (4.19) seems to be for defining the condition $v \in \mathcal{D}_n$ WOR w , it is in fact natural to replace it by a self-consistent definition that is not quite equivalent: Let $D_{v,w}^{-1} = 0$ and for $n \geq 0$:

$$D_{v,w}^n = \mathbf{1}(v \in \mathcal{D}_n \text{ WOR } w) = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w',v}^{n-1} \mathbf{1}(w' \neq w) \right). \quad (4.20)$$

Proposition 4. *[The WOR property of the Watts model] Let the Watts model be specified by $(\mathcal{N}, \mathcal{E}, \{\bar{\Delta}_v\})$ and the sequences $\{D_v^n, \tilde{D}_{v,w}^n, D_{v,w}^n\}_{n=-1,0,1,\dots}$ defined by the recursive equations (4.18), (4.19), (4.20) with the initial conditions $D_v^{-1}, D_{v,w}^{-1}, \tilde{D}_{v,w}^{-1} = 0$. Then for all $n \geq 0$ and $(v, w) \in \mathcal{E}$*

$$D_v^n = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w',v}^{n-1} \right) \quad (4.21)$$

$$= \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} \tilde{D}_{w',v}^{n-1} \right). \quad (4.22)$$

In general, $D_{v,w}^n \leq \tilde{D}_{v,w}^n$, with strict inequality only occurring if $D_w^{n-1} = 1$.

The last part of the theorem says that adoption shocks $D_{v,w}^n$ and $\tilde{D}_{v,w}^n$ transmitted to w can only differ when their impact is inconsequential because w has already adopted.

Proof: To prove (4.21) we introduce additional variables \tilde{D}_v^n defined recursively for $n \geq -1$ by $\tilde{D}_v^{-1} = 0$ and

$$\tilde{D}_v^n = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w',v}^{n-1} \right). \quad (4.23)$$

We show by induction on n that

$$\tilde{D}_v^n = D_v^n \quad (4.24)$$

for all n, v . The proof is based on the monotonicity in n of these recursive mappings, that is $D_v^{n-1} \leq D_v^n$ etcetera. First, note that (4.24) is true for $n = -1, 0$. Then note that by monotonicity $\tilde{D}_v^n \leq D_v^n$ for all n, v .

Now assume there is a minimal $n \geq 1$ and v such that $0 = \tilde{D}_v^n < D_v^n = 1$. Parsing the defining conditions leads to the implication that $\sum_{w \in \mathcal{N}_v} D_{w,v}^{n-1} < \sum_{w \in \mathcal{N}_v} D_w^{n-1}$. Since n is minimal, this in turn implies $D_{w,v}^{n-1} < D_w^{n-1} = \tilde{D}_w^{n-1}$ for some $w \in \mathcal{N}_v$. Thus $D_{v,w}^{n-2} = 1$. This means $D_v^{n-2} \geq D_{v,w}^{n-2} = 1$ while $\tilde{D}_v^{n-2} \leq \tilde{D}_v^n = 0$. But by the minimality of n , we must have that $D_v^{n-2} = \tilde{D}_v^{n-2}$, which is a contradiction. We conclude the non-existence of a minimal $n \geq 2$ and v such that $0 = \tilde{D}_v^n < D_v^n = 1$. Thus (4.21) follows.

Since $D_{v,w}^n \leq \tilde{D}_{v,w}^n \leq \tilde{D}_v^n = D_v^n$ it must be the case that

$$\begin{aligned} D_v^n &= \tilde{D}_v^n = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w',v}^{n-1} \right) \\ &\leq \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} \tilde{D}_{w',v}^{n-1} \right) \leq \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} D_{w'}^{n-1} \right) = D_v^n \end{aligned}$$

which proves (4.22). Finally, to prove the last part of the theorem, suppose $D_{v,w}^n = 0, \tilde{D}_{v,w}^n = 1$. In this case $\tilde{D}_v^n = D_v^n \geq \tilde{D}_{v,w}^n = 1$ as well. Hence it must be that $D_{w,v}^{n-1} = 1$ and thus $D_w^{n-1} = 1$. \square

4.5.2 The Main Result

To state the core result in the Watts model, we reintroduce the conditional *without regarding* probability:

$$\hat{p}_k^{(n)} := \mathbb{P}[w \in \mathcal{D}_n \text{ WOR } v | w \in \mathcal{N}_k \cap \mathcal{N}_v] \quad (4.25)$$

and make use of (4.20) and (4.21) to derive a pair of inductive formulas for the collection of probabilities $\{p_k^{(n)} := \mathbb{P}[v \in \mathcal{D}_n | v \in \mathcal{N}_k]\}$ and $(\hat{p}_k^{(n)})$ for $n \geq 1$ in terms of the WOR-probabilities $(\hat{p}_k^{(n-1)})$, with the initial conditions $\hat{p}_k^{(0)} := d_k(0)$. These two formulas are valid only asymptotically in the limit as the network size N goes to infinity, while keeping the probability data $(P_k, d_k(0))$ fixed. There are also implied finite moment conditions on these distributions. We state the result, and then give a heuristic proof that will provide a guide to further results.

Theorem 11. *Consider the Watts model in the limit as $N \rightarrow \infty$, with fixed mean degree $z > 0$ and with adoption threshold distribution functions $d_k(\cdot), D_k(\cdot)$ for $k \geq 0$. The initial adoption probabilities are $p_k^{(0)} = \hat{p}_k^{(0)} = d_k(0)$. Then:*

1. The collections $(p_k^{(n)})$ and $(\hat{p}_k^{(n)})$ for $n \geq 1$ are given by the recursion formulas

$$p_k^{(n)} = G_k(\hat{p}^{(n-1)}) := \sum_{x=0}^k D_k(x) \text{Bin}(k, \hat{\pi}^{(n-1)}, x) \quad (4.26)$$

$$\hat{p}_k^{(n)} = \hat{G}_k(\hat{p}^{(n-1)}) := \sum_{x=0}^{k-1} D_k(x) \text{Bin}(k-1, \hat{\pi}^{(n-1)}, x) \quad (4.27)$$

where $\hat{\pi}^{(n-1)} = \mathbb{P}[w \in \mathcal{D}_{n-1} \text{ WOR } v | w \in \cap \mathcal{N}_v]$ is the unconditional probability of w adopting at step $n-1$ WOR v for any link (w, v) .

2. The probability $\hat{\pi}^{(n)}$ is given by a scalar mapping $\hat{\pi}^{(n)} = G(\hat{\pi}^{(n-1)})$ where

$$G(\pi) = \sum_k \sum_{x=0}^{k-1} \frac{k P_k}{z} D_k(x) \text{Bin}(k-1, \pi, x). \quad (4.28)$$

The mapping $G: [0, 1] \rightarrow [0, 1]$ is continuous, monotonically increasing and has $G(0) = \hat{\pi}^{(0)}$. Therefore the sequence $\hat{\pi}^{(n)}$ converges to the least fixed point $\pi^* \in [0, 1]$ with $\pi^* = G(\pi^*)$.

Sketch Proof: The proof of the theorem is based on two properties of the model. The first is the LT property of the skeleton as long as N is sufficiently large, and the second is the conditional independence of the thresholds $\bar{\Delta}$, conditioned on the skeleton. Therefore, adoption shocks coming into a node v are always asymptotically independent as $N \rightarrow \infty$, and this collection of shocks is independent of the threshold $\bar{\Delta}_v$.

By the original definition of the set \mathcal{D}_n ,

$$p_k^{(n)} = \mathbb{P}\left[\bar{\Delta}_w \leq \sum_{w' \in \mathcal{N}_w} \mathbf{1}(w' \in \mathcal{D}_{n-1}) | w \in \mathcal{N}_k\right] \quad (4.29)$$

One might try to argue that conditioned on $w \in \mathcal{N}_k$, the events $\{w' \in \mathcal{D}_{n-1}\}$ over nodes $w' \in \mathcal{N}_w$ are mutually independent in the limit $N \rightarrow \infty$ because of the locally tree-like property that becomes exact as $N \rightarrow \infty$, and independent of $\bar{\Delta}_v$ because of the further independence assumption. But this is erroneous: because the links are bi-directional, each $\{w' \in \mathcal{D}_{n-1}\}$ is in fact dependent on $\{w \in \mathcal{D}_{n-2}\}$ and hence on $\bar{\Delta}_w$. However, by (4.20) each $\{w' \in \mathcal{D}_{n-1} \text{ WOR } w\}$ is conditionally independent of the state of $\bar{\Delta}_w$. Using (4.21), (4.29) can be rewritten

$$p_k^{(n)} = \mathbb{P}\left[\bar{\Delta}_w \leq \sum_{w' \in \mathcal{N}_w} \mathbf{1}(w' \in \mathcal{D}_{n-1} \text{ WOR } w) | w \in \mathcal{N}_k\right] \quad (4.30)$$

where the terms in the sum are k independent Bernoulli random variables. Furthermore, because of the independent edge condition, namely $Q_{k'|k} = Q_{k'}$, the Bernoulli probabilities are independent of the value k_w , and thus are each $\hat{\pi}^{(n-1)}$. This leads to equation (4.26). Similarly, using (4.20) we can compute that for random links (v, w)

$$\begin{aligned}\hat{p}_k^{(n)} &:= \mathbb{P}[w \in \mathcal{D}_n \text{ WOR } v | w \in \mathcal{N}_v \cap \mathcal{N}_k] \\ &= \mathbb{P}\left[\bar{\Delta}_w \leq \sum_{w' \in \mathcal{N}_w \setminus v} \mathbf{1}(w' \in \mathcal{D}_{n-1} \text{ WOR } w) | w \in \mathcal{N}_v \cap \mathcal{N}_k\right].\end{aligned}$$

where now there are $k-1$ independent $\text{Bern}(\hat{\pi}^{(n-1)})$ random variables in the sum, leading to (4.27). This verifies part (1) of the theorem.

For part (2), we note that

$$\hat{\pi}^{(n)} = \sum_k \mathbb{P}[w \in \mathcal{D}_n \text{ WOR } v | w \in \mathcal{N}_v \cap \mathcal{N}_k] \mathbb{P}[k_w = k | w \in \mathcal{N}_v] \quad (4.31)$$

Since $\mathbb{P}[k_w = k | w \in \mathcal{N}_v] = Q_k = \frac{kP_k}{z}$ from (4.5), we see that

$$\hat{\pi}^{(n)} = \sum_k \mathbb{P}[w \in \mathcal{D}_n \text{ WOR } v | w \in \mathcal{N}_v \cap \mathcal{N}_k] Q_k \quad (4.32)$$

which leads to (4.32). Finally, the argument why $\lim_{n \rightarrow \infty} \hat{\pi}^{(n)} = \pi^*$ is the smallest fixed point of G is precisely the same as the argument given for the extinction probability η for percolation in Section 4.2. \square

Reviewing for a moment Section 4.4, we can see that the main theorem of the Watts model and its proof are a template for what we can expect in variations of bootstrap percolation dynamics.

4.5.3 The Cascade Condition

In this version of the Watts model, we notice that two initial seed probability distributions $(p_k^{(0)})$ giving the same scalar value $\hat{\pi}^{(0)}$ lead to the same sequence $(p_k^{(n)})$ for $n \geq 1$. So we can consider equally weighted schemes where initial seeds are uniformly random with $p_k^{(0)} = \hat{\pi}^{(0)}$. Then, essentially, the cascade is the iterated scalar mapping G that converges to the fixed point π^* , and the probability that a degree k node eventually defaults is

$$p_k^{(\infty)} = G_k(\pi^*)$$

Let us consider initial seed probabilities $p_k^{(0)} = \hat{\pi}^{(0)} = \varepsilon$ and small $\varepsilon > 0$. We can write $D_k(x) = \varepsilon + (1 - \varepsilon)\bar{D}_k(x)$ where $\bar{D}_k(x) = \mathbb{P}[\bar{\Delta}_v \leq x | v \in \mathcal{N}_k, \bar{\Delta}_v \neq 0]$ is the threshold CDF conditioned on not being an early adopter and consider the fixed point of $G(\cdot; \varepsilon, \bar{D})$ as a function of ε for fixed \bar{D} . The most important question to ask is whether the fixed point $\pi^*(\varepsilon)$ is of order ε or of order 1 as $\varepsilon \rightarrow 0$. In other words, what is the ‘‘cascade condition’’ that determines if an infinitesimally small seed fraction will grow to a large-scale cascade? It turns out this depends on the derivative $\partial G(0; \varepsilon, \bar{D})/\partial x$ at $\varepsilon = 0$, which is easy to calculate:

$$\frac{\partial G(0;0,\tilde{D})}{\partial \pi} = \sum_k \sum_{x=0}^{k-1} \frac{kP_k}{z} \tilde{D}_k(x) \binom{k-1}{x} [\mathbf{1}(x=1) - (k-1)\mathbf{1}(x=0)] = \sum_k \frac{k(k-1)P_k d_k(1)}{z} \quad (4.33)$$

Proposition 5. *Consider the standard Watts model.*

1. *If the cascade condition $\partial G(0;0,\tilde{D})/\partial \pi > 1$ is true, then there is $\bar{\varepsilon} > 0$ such that $|\pi^*(\varepsilon)| > \bar{\varepsilon}$ for all $\varepsilon > 0$. That is, under this condition, an initial seed with a positive fraction of nodes will almost surely trigger a cascade fraction larger than $\bar{\varepsilon}$.*
2. *If $\partial G(0;0,\tilde{D})/\partial \pi < 1$ then there is $\bar{\varepsilon} > 0$ and C such that for all $0 < \varepsilon < \bar{\varepsilon}$, $|\pi^*(\varepsilon)| \leq C\varepsilon$. That is, this network will almost surely not exhibit large scale cascades for any initial seed with fractional size less than $\bar{\varepsilon}$.*

We can interpret this condition by comparing (4.33) to the main result Proposition 10 for site percolation. We see that the above cascade condition is identical to the site percolation condition of (4.12), if we look at the connectivity of the subgraph obtained by deleting all sites except those with $\bar{\Delta} \leq 1$. This connection to site percolation becomes even clearer in the next subsection.

4.5.4 The Frequency of Global Cascades

A necessary condition for a single seed node to trigger a very large cascade is that some of its neighbours are *vulnerable* in the sense that they are susceptible to adopt given that only one of their neighbours adopts. If there are few cycles in the graph, then any potential large scale cascade must first grow through vulnerable nodes, and only when the cascade is sufficiently developed will less vulnerable nodes begin to adopt.

Since we consider the model with a single seed, and we are working formally in an infinite network, we suppose $d_k(0) = 0$ for all k . Then the above picture is formalized by considering the set of vulnerable nodes $\mathcal{V} \subset \mathcal{N}$ defined by the condition $\bar{\Delta}_v \leq 1$, which as we recall has probability $d_k(1)$ if $k_v = k$. We consider whether or not \mathcal{V} has a giant cluster \mathcal{C}_V . Supposing the cluster \mathcal{C}_V is giant, then the condition $v \in \mathcal{C}_V^c$ means either $v \notin \mathcal{V}$ or $v \in \mathcal{V}$ and all neighbours $w \in \mathcal{N}_v$ are not in \mathcal{C}_V without regarding v . Let $\xi = \mathbb{P}[w \in \mathcal{C}_V^c \text{ WOR } v | w \in \mathcal{N}_v]$. Then, following the intuitive logic outlined in Section 4,

$$\mathbb{P}[v \in \mathcal{C}_V^c] = \sum_k P_k \left[(1 - d_k(1)) + d_k(1)\xi^k \right]. \quad (4.34)$$

Moreover, following the same logic,

$$\xi := \mathbb{P}[w \in \mathcal{C}_V^c \text{ WOR } v | w \in \mathcal{N}_v] = \sum_k Q_k \left[(1 - d_k(1)) + d_k(1)\xi^{k-1} \right] := f(\xi). \quad (4.35)$$

Now the function f maps $[0, 1]$ to itself, is increasing and convex, and has $f(0) > 0$ and $f(1) = 1$. As illustrated by Figure 4.1, for (4.35) to have a non-trivial fixed point $\xi < 1$, it is necessary and sufficient that $f'(1) > 1$, that is,

$$\sum_k k(k-1)P_k d_k(1) > z$$

which we recognize as the cascade condition from the previous subsection. When $\sum_k k(k-1)P_k d_k(1) \leq z$, the argument can be reversed and one finds that there can be no giant vulnerable cluster.




Fig. 4.1 Two possible fixed point configurations are illustrated. The supercritical case is found for the green curve showing the function $f(\xi) = e^{-3(\xi-1)/2}$, which has a non-trivial fixed point $\xi < 1$. The blue curve showing $f(\xi) = e^{-2(\xi-1)/3}$ has only the trivial fixed point at $\xi = 1$, and corresponds to a sub-critical random graph.

For a global cascade to arise from a single seed, it is necessary and sufficient for the random seed to have at least one neighbour in the giant vulnerable cluster \mathcal{C}_V , which occurs with frequency

$$f = \sum_k P_k (1 - \xi^k). \quad (4.36)$$

In this case, a cascade of at least the size of \mathcal{C}_V will result. However, since the \mathcal{C}_V is a positive fraction of the entire network, a positive fraction of adoptions by less vulnerable nodes is possible, and therefore the extent of the global cascade may be substantially larger than \mathcal{C}_V .

4.6 Numerical Experiments on the Watts Model

We illustrate the Watts Cascade model by reproducing some of the numerics of the original specification:

- The skeleton is an undirected $G(N, p)$ (Gilbert) random graph with $N = 10000$ and edge probability $p = z/(N - 1)$, chosen with the mean node degree $z \in [0, 10] \ll N$ so that the graph is sparse.
- The threshold random variables $\bar{\Delta}_v$ are proportional to the node degree k_v , with an initial seed of adopters chosen uniformly with probability $p^{(0)} \ll 1$. That is, for the default value $\phi_* = 0.18$,

$$\bar{\Delta}_v = \phi_* k_v \mathbf{1}(v \notin \mathcal{D}_0), \quad \mathbb{P}[v \in \mathcal{D}_0] = p^{(0)}$$

- Link weights are equal to one: $\bar{\Omega}_{wv} = 1$.

Since $N = 10000$ is large, and the $G(N, p)$ model is a configuration graph, we expect Theorem 11 to yield the probability of eventual default $p_k^{(\infty)} = G_k(\pi^*)$ with only small finite size error effects. To make a fair comparison between the finite network computed by Monte Carlo simulation and the infinite network computed analytically, it is important to understand how to generate the initial seed. In the infinite network, any positive initial seed probability generates an infinite set of early adopters. When the density of adopters $p^{(0)}$ is small, each seed generates a “mini-cascade” far from the others, and a large scale cascade will be generated if at least one seed succeeds in generating a large cascade. This will happen almost surely if each single seed has a positive probability of generating a large cascade, that is when the cascade condition of Theorem 5 is true. Moreover, in this case percolation theory suggests that the fractional size of the large scale cascade, that is, the expected default probability, will exceed the fractional size of the giant vulnerable cluster \mathcal{C}_V (recall that on undirected networks, the various giant clusters coincide: $\mathcal{C}_V = G - \text{IN}_V = \text{GSCC}_V = G - \text{OUT}_V$). On the other hand, the probability of a “mini-cascade” growing to a global cascade is zero when there is no giant vulnerable cluster: This defines the non-cascading region of parameter space, where almost surely no global cascade will arise. A quantitative verification of these predictions will follow from the cascade mapping of Theorem 11.

In a Monte Carlo realization of the finite network, we can assert that a single random adoption seed will grow to a large fraction of the network, i.e. it triggers a global cascade, only if it lands on or beside a giant vulnerable cluster \mathcal{C}_V . This will happen with probability f given by (4.36). On the other hand, taking an intermediate seed fraction, say $100/10000$, will generate 100 roughly independent “mini-cascades”, each of which has a probability f of becoming a large scale cascade taking up a significant fraction of the network. Suppose the probability f of hitting \mathcal{C}_V is not too small. Then the probability of a large scale cascade will be about $1 - (1 - f)^{100}$, which will be close to one if $f \gtrsim 0.01$.

Experiment 1: We generated $N_{\text{sim}} = 50$ realizations of the network, each with 50 randomly generated seed nodes and compared the results to the analytic formula for the eventual adoption probability $p^{(\infty)} = \sum_k P_k p_k^{(\infty)}$ from Theorem 11 with initial seed probability $\pi^{(0)} = 0.005$. Figure 4.2(a) shows the results.

In addition to the close agreement between the Monte Carlo and analytical results, we observe in this graph the well known “contagion window” for z in the

approximate range $[1, 7]$. The upper and lower thresholds correspond to the values of the mean degree for which the connected vulnerable out-component \mathcal{C}_V becomes “giant”, that is a finite fraction of the whole network. The contagion window determined by the analytical formula becomes exact as $\pi^{(0)} \rightarrow 0$. The transition near the upper threshold shown in the Monte Carlo result corresponds to the region where the default frequency $f \ll 0.01$ and thus in most realizations none of the initial seeds hits \mathcal{C}_V . \mathcal{C}_V is small in this parameter region because only links pointing to nodes with degree 5 or less are vulnerable.

Experiment 2: To make the correct comparison between the cascade frequency f computed analytically by (4.36), and the Monte Carlo simulations, we should generate only a single seed in our Monte Carlo simulations, and count the number of simulations that lead to a cascade that can be considered “large” or “global”. In a 10000 node network, we consider that a cascade of more than 50 adoptions is a global cascade. Figure 4.2(b) shows the frequency at which single seeds trigger a global cascade of at least 50 nodes in the finite network, and the infinite network frequency given by (4.36).




Fig. 4.2 These two graphs show (a) the mean fractional cascade size, and (b), the fractional extended vulnerable cluster size, in the benchmark Watts model with $\phi_* = 0.18$, as a function of mean degree z , as computed using the large N analytics of Theorem 11 (blue curve) and by Monte Carlo simulation (red crosses). In Figure (a), the Monte Carlo computation involved $N_{\text{sim}} = 50$ realizations of the $N = 10000$ node graph, each with an initial seed generated by selecting nodes independently with probability 0.5%. In Figure (b), the simulation involved $N_{\text{sim}} = 2500$ realizations of the $N = 10000$ node graph and a single random initial seed node.

Taken together, the two figures 4.2 verify the vulnerable percolation picture that global cascades can arise whenever \mathcal{C}_V is a significant fraction of the entire network (say any positive fraction if $N = \infty$, or 0.5% if $N = 10000$). When this happens, a single isolated seed that happens to have at least one neighbour in \mathcal{C}_V will trigger all

of \mathcal{C}_V and more, a possibility that happens with probability f given by the frequency formula (4.36).

4.7 Extensions of the Watts Model

The basic construction and main results of the previous section has been worked out in many different guises. We outline here some of the main lines that have been developed.

1. **General degree distributions:** The Poisson random graph model fails to capture most of the structural features observed in real world social networks. In particular, it has a thin tailed degree distribution that is incompatible with the type of fat-tails and power laws observed in real world networks. It turns out to be completely straightforward to analyze the Watts model on general configuration graphs, obtaining the same result as given by Theorem 11, for an arbitrary degree distribution P_k .
2. **Mixtures of directed and undirected edges:** In [12], percolation results are proved in random networks with both directed and undirected edges and arbitrary two-point correlations. Analysis of the Watts model in this network seems to be a straightforward next step.
3. **Assortative graphs:** In Section 3.2, we observed that the usual configuration structure can be generalized to allow for general assortativity, that is, non-independent edge-type distributions Q_{kj} . This generalization is important because observed financial networks are typically disassortative, and this fact is expected to strongly influence a network's resilience to default. It was demonstrated in [43] that for the Gai-Kapadia 2010 default cascade model on assortative directed configuration graphs, similar arguments lead to a characterization of the limiting default probabilities in terms of the fixed points of a vector valued cascade mapping $G : \mathbb{R}^{\mathbb{Z}^+} \rightarrow \mathbb{R}^{\mathbb{Z}^+}$.
4. **Random edge weights:** The Watts model allows only equal link weights, whereas in reality, we might expect the strength of friendship links to be stochastic. We might also expect that the statistical properties of friendship strength will be intimately related to the degree of connectivity of social contacts, i.e to the edge-type. As first shown in [44], even capturing the adoption cascade with purely deterministic dependence between link-strength and edge-degree in the Watts model turns out to require an analysis of random link weights. In this more general setting, that paper finds that limiting adoption probabilities are characterized in terms of the fixed points of another vector-valued cascade mapping $G : \mathbb{R}^{\mathbb{Z}^+} \rightarrow \mathbb{R}^{\mathbb{Z}^+}$.

In the next chapter, we will explore such extensions, not in the Watts model but in the context of systemic risk, by studying variations of the Gai-Kapadia default cascade model. First however, we pay a bit more attention to the mathematical structure at its core that leads to such elegant large graph asymptotic results.

4.8 Dependence in Random Financial Networks

A random financial network, as defined in Section 2.6, consists of a skeleton graph, decorated by additional random structures, namely balance sheets for each node and exposures for each directed link. Up to this point however, we haven't really considered how these high-dimensional collections of random variables might be provided with multivariate distributions that reflect financial reality about the nature of banks and their balance sheets, as well as the often high uncertainty in our observations and knowledge of the system at an arbitrary moment in time. Let us consider what reasonable constraints should be made on the possible dependence relating different balance sheet entries and exposures.

The first reasonable hypothesis is that banks control their balance sheets while subject to constraints imposed by the financial regulator. Important constraints are usually expressed in the form of capital ratios, such as the minimum core tier one capital or "capital adequacy ratio" (CAR), the liquidity coverage ratio (LCR) and net stable funding ratio (NSFR). However, banks are monitored only intermittently by the regulators, and therefore constraints are not strictly binding. Banks have non-negligible probabilities of being non-compliant at a random moment in time. Indeed, a financial crisis, almost by definition, typically occurs at a moment when one or more banks have been hit by shocks that have made them non-compliant.

A second reasonable hypothesis is that individual banks control their balance sheets independently of any knowledge of the precise balance sheets of other banks in the system. Similarly, interbank exposure sizes are the result of a bargaining game between two counterparties whose outcome depends only on their connectivity and balance sheets. Notice the intrinsic conceptual difficulty in imagining the results of the myriad of overlapping games being played within the network: to move forward will require overarching assumptions that capture some salient features of the observed interbank exposure sizes. Furthermore, each interbank link leads to some dependence between the interbank components of balance sheets of counterparty banks, through the consistency relations (2.34).

Pushing harder, we might make a third hypothesis that banks' balance sheets (or at least some components) can depend on the "type" of the bank, but when conditioned on the type of bank, can be treated as independent of other banks. In the simplest RFN settings where we equate bank type to node-degree, imagining that bank size (measured by total assets) is strongly correlated with bank connectivity, implies that balance sheet dependency arises only through the skeleton graph.

As a final consideration, we recognize that real financial network data is often highly aggregated across counterparties and the problem is to specify bank specific random variables consistent with these observations. Treating such aggregated measurements as constraints leads to system-wide dependences amongst banks' balance sheets that are impractical to model. Instead, we will usually adopt the strategy to enforce such aggregate observations in expectation, providing linear constraints on the means of the more basic (and independent) bank-specific random variables. This type of inconsistency is reduced by the law of large numbers effect that usually as-

tures that such aggregated random variables become equal to their means in the large network limit.

We are about to commit to making a conceptual leap that will tame balance sheet and exposure dependence by tying it directly to the dependence built into the skeleton. Before doing so, one should imagine the alternative in a large, partly measured financial network. Such an object is extremely complex. Even the marginal distributions of essential quantities necessarily have high uncertainty. Specifying dependence in such a setting is orders of complexity beyond specifying marginals. It is the difference between specifying univariate functions and specifying functions on a space approaching infinite dimension. Our aim in this book is to be reasonable: we will make a strong assumption about dependence that allows both Monte Carlo and analytical experiments to run side by side, and will then use these computational tools to explore the consequences of this assumption. We do not claim that our dependence assumptions rest on fundamental economic principles.

We consider a RFN adapted to a given cascade mechanism, parametrized by the quadruple $(\mathcal{N}, \mathcal{E}, \bar{\Delta}, \bar{\Omega})$. The key to taming dependence is to allow the skeleton $(\mathcal{N}, \mathcal{E})$ to determine the overall dependence in the network. Conceptually, we first draw a random skeleton, finite or infinite. Then, conditioned on the realization of the skeleton, we draw remaining random variables independently. Moreover, the dependence of these random variables on the skeleton will be *local*. That is, the marginal distributions of node variables (buffers) depend only on the type of the node itself and not on its neighbourhood. Similarly, the marginal distribution of each edge variable (exposure) depends only on the type of the edge.

It is helpful to express this definition in terms of general probability theory as developed for example in Chapters One and Two in Volume Two of the popular book on stochastic calculus by Shreve [67]. We express the probability framework in terms of the triple $(\Omega', \mathcal{F}, \mathbb{P})$, where Ω' is the sample space, \mathbb{P} is the probability measure, and \mathcal{F} is the information set, or sigma-algebra. \mathcal{F} can be reduced to a union of sub-sigma-algebras

$$\mathcal{F} = \mathcal{G} \vee \mathcal{F}_\Delta \vee \mathcal{F}_\Omega .$$

where: \mathcal{G} denotes the information contained in the skeleton \mathcal{N}, \mathcal{E} ; \mathcal{F}_Δ denotes the information of the collection of buffers $\bar{\Delta}_v$; \mathcal{F}_Ω denotes the information of the collection of exposures $\bar{\Omega}_\ell$. Then the dependence of $\bar{\Delta}_v$ only on the type of the node v can be expressed in terms of expectations conditioned on knowing all information apart from the information of $\bar{\Delta}_v$ itself (which is measured by the sigma algebra $\sigma(\bar{\Delta}_v)$ it generates): it holds that for any bounded Borel function f and a randomly selected node $v \in \mathcal{N}$, there is a bounded Borel function $g : \mathbb{Z}_+^2 \rightarrow \mathbb{R}$ such that

$$\mathbb{E}[f(\bar{\Delta}_v) | \mathcal{F} \setminus \sigma(\bar{\Delta}_v)] = g(j_v, k_v) .$$

This implies in particular that the conditional CDF of $\bar{\Delta}_v$ is a function that depends only on its degree: for all $x \geq 0$, there are Borel functions D_{jk} such that

$$\mathbb{E}[\bar{\Delta}_v \leq x | \mathcal{F} \setminus \sigma(\bar{\Delta}_v), v \in \mathcal{N}_{jk}] = D_{jk}(x) . \quad (4.37)$$

In exactly the same way, for exposures it holds that there are Borel functions W_{kj} such that

$$\mathbb{E}[\bar{\Omega}_\ell \leq x | \mathcal{F} \setminus \sigma(\bar{\Omega}_\ell), \ell \in \mathcal{E}_{kj}] = W_{kj}(x). \quad (4.38)$$

While the dynamics of cascades is determined by the specification of the reduced RFN, questions of interpretation, and in particular, the impact of the cascade on the financial system and the larger economy require also specifying the remaining balance sheet random variables. Of course, the interbank assets and liabilities Z_v, X_v are determined endogenously by the exposures $\bar{\Omega}$, and the buffer variables $\bar{\Delta}$ typically amount to additional linear constraints on external balance sheet entries. In absence of more detailed information than this, it makes sense to model all remaining random variables as linear combinations of a minimal collection of additional independent random variables.

Finally, we should reassure ourselves that the proposal to build RFNs this way is common in the literature on random financial networks. For example, Nier et al in [62], start with a random directed Poisson graph with 25 nodes, and conditioned on the model for the skeleton, followed by specifying random variables for buffers and exposures that depend on node and edge types, but are otherwise independent.

4.8.1 The LTI Property

In the subsequent sections of the book, our analysis will focus on random skeletons of configuration type, which in the large N limit have the locally tree-like property. In the configuration graph setting, our dependence hypothesis becomes the following definition.

Definition 9. A random financial network (RFN) $(\mathcal{N}, \mathcal{E}, \bar{\Delta}, \bar{\Omega})$ is said to have the *locally tree-like independence (LTI) property* when it satisfies the following conditions:

1. The skeleton graph is an infinite (directed, undirected or mixed) configuration graph $(\mathcal{N}, \mathcal{E})$, with arbitrary node and edge type distributions $\{P, Q\}$. In this general setting, each node v is assigned its type τ_v , that is jk or k in the case of directed and undirected graphs and more general for mixed graphs. Similarly, each edge ℓ is assigned its type τ_ℓ , which may be kj or kk' or more general. Here, the indices j, k, k' run over the collection of possible integer degrees.
2. Conditioned on the realization of the skeleton graph $(\mathcal{N}, \mathcal{E})$, the buffer random variables $\bar{\Delta}_v$, $v \in \mathcal{N}$ and exposure random variables $\bar{\Omega}_\ell$, $\ell \in \mathcal{E}$ form a mutually independent collection. Moreover, the buffer distribution of $\bar{\Delta}_v$ depends only on the type τ_v of v and the exposure distribution of $\bar{\Omega}_\ell$ depends only on the type τ_ℓ of ℓ .

4.8.2 *Ramifications of the LTI Property*

As we shall see in the remainder of this book, the LTI property is designed to enable a cascade analysis in complex RFNs that follows the computational template laid out in our simple version of the Watts cascade model. This works essentially because under LTI, the shocks transmitted through different edges entering a node never have a chance to become dependent. The mathematical advantages that stem from LTI will become clear as we proceed to other models. To reiterate a second point, an essentially arbitrary dependence structure over the balance sheet random variables $\bar{\Delta}_v, \bar{\Omega}_\ell$ is dramatically reduced under LTI to the dependence encoded into the skeleton graph. Thus LTI reins in the amount of information needed to completely specify an RFN model. If one feels a given dependence structure induced by the skeleton is overly restrictive, yet wish to retain the LTI property, one can explore the further option to add complexity to the RFN at the skeleton level, for example, by extending the concept of node and edge types.

The question to ask of our proposed conceptual framework is thus pragmatic: why or when can we hope that computations based on the LTI assumption will have some predictive power when transferred to a real world network model where the assumption is apparently far from true. Pushing further, is it possible to weaken the LTI assumption, and to extend its exact consequences to a broader setting? Some measure of optimism that LTI computations are predictive even when the skeleton graph has high clustering or is otherwise far from locally-treelike can be gained by extensive investigations of the question in the random network literature. For example, [55] make a strong case for the “unreasonable effectiveness” of tree-based theory for networks with clustering, and give some guidelines to predict the circumstances under which tree-based theory can provide useful approximations.

Chapter 5

Zero Recovery Default Cascades

Abstract This chapter realizes the central aim of the book, which is to understand a simple class of cascades on financial networks as a generalization of percolation theory. The main theorems assume random financial networks with locally tree-like independence and characterize zero-recovery default cascade equilibria as fixed points of certain cascade mappings. The proofs of the main results follow a new and distinctive template presented here for the first time, that has the important virtue that its logic extends to LTI financial networks of arbitrary complexity. Numerical computations, both large network analytics and finite Monte Carlo simulations, verify that essential characteristics such as cascade extent and cascade frequency can be derived from the properties of the cascade fixed points.

Keywords: Directed configuration graph, random buffers, random edge weights, cascade mapping theorem, numerical implementation.

In Chapter 3 we have understood something about possible models for the skeleton of a financial network, and in Chapter 4 we showed how two concepts, bootstrap percolation and site percolation, reveal the nature of the cascades observed in the Watts model. It is time to return our attention to financial networks and to investigate how well the various systemic cascade mechanisms described in Chapter 2 can be adapted to the setting of random financial networks, using the Watts model as a template. For clarity in this chapter, we confine the discussion to simple default cascades, under the condition of zero recovery as assumed by [38]: banks recover none of their assets loaned to a defaulted bank. Of course, the common mathematical features of the different financial cascade mechanisms identified earlier in Chapter 2 mean the narrow formulation presented in this chapter will extend to a rich diversity of cascade phenomena.

At a theoretical level, perhaps the most important new contribution of this chapter is the proof of Theorem 12 which provides the explicit large N asymptotic limit of the cascade mapping for the extended GK default cascade model. This proof is different in spirit, detail and scope from the well-known proof given in a similar setting by Amini, Cont and Minca [5].

We begin again with a network of N banks whose schematic balance sheets, as shown in Table 2.1, consist of the collection of *book values* $(\bar{Y}, \bar{Z}, \bar{D}, \bar{X}, \bar{\Omega})$ with interbank exposures $\bar{\Omega}_\ell = \bar{\Omega}_{wv}$ (which we recall denotes the amount w owes v) that are consistent with the interbank assets and liabilities

$$\bar{Z}_v = \sum_w \bar{\Omega}_{wv}, \quad \bar{X}_v = \sum_w \bar{\Omega}_{vw}.$$

The initial capital buffer of bank v is defined by $\bar{\Delta}_v = \bar{Y}_v + \bar{Z}_v - \bar{D}_v - \bar{X}_v$. At the onset of our schematic financial crisis, certain banks are found to be initially insolvent, defined by the condition $\bar{\Delta}_v \leq 0$. By the law of limited liability, insolvency is assumed to trigger immediate default. Under an assumption about the loss given default, losses will be transmitted to both external creditors, and, importantly for systemic risk, to creditor banks. This implies the possibility of a default cascade crisis.

We may suppose that counterparty relations, i.e. links, are expensive to maintain, and so in a large network, the matrix of counterparties is sparse. These relationships are also changing in time, and not easily observable, and can therefore be considered to be random: the skeleton graph is a large, sparse, directed random graph $(\mathcal{N}, \mathcal{E})$. The balance sheets likewise change rapidly, are not well observed, and can be thought of as random variables. Under such conditions, we have argued that it is appropriate to adopt the notion of random financial network (RFN) with the *locally tree-like independent* (LTI) assumption of Section 4.8.1 as the underlying dependence hypothesis.

In the present chapter we focus our attention on the simplest of default mechanisms, namely the zero recovery mechanism of Gai and Kapadia [38], governed by the cascade equation (2.12) we derived in Chapter 2 :

$$\Delta_v^{(n)} = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} (1 - \tilde{h}(\Delta_w^{(n-1)})) = \bar{\Delta}_v - \sum_w \bar{\Omega}_{wv} \mathbf{1}(\Delta_w^{(n-1)} \leq 0) \quad (5.1)$$

with initial values $\Delta_v^{(0)} = \bar{\Delta}_v$. As we observed earlier, these equations, and therefore the entire cascade dynamics, depend only on the skeleton graph $(\mathcal{N}, \mathcal{E})$ and the reduced balance sheet variables $\bar{\Delta}, \bar{\Omega}$. Our central aim is to apply this cascade mechanism to a RFN that satisfies the LTI hypothesis, and to determine the large scale properties of the behaviour that results.

5.1 The Gai-Kapadia Model

The specification of the Gai-Kapadia model we first present generalizes the default cascade introduced by [38] in several respects, but retains a restrictive condition on exposures $\bar{\Omega}$ that will be removed later in the chapter. It also satisfies the LTI property of Section 4.8.1.

- Assumption 7.** 1. As in the description of the model in Section 2.3, banks have limited liability and receive a zero recovery¹ of interbank liabilities from any defaulted debtor bank.
2. The skeleton graph is a directed assortative configuration random graph with consistent probability matrices (P_{jk}, Q_{kj}) that satisfy Assumption 5 with mean degree z .
3. Conditionally on the skeleton, banks' capital buffers $\bar{\Delta}_v$ are a collection of independent non-negative random variables whose cumulative probability distributions depend on the node type (j_v, k_v) and are given by

$$\mathbb{P}[\bar{\Delta}_v \leq x | v \in \mathcal{N}_{jk}] := D_{jk}(x), \quad x \geq 0. \quad (5.2)$$

4. Interbank exposures $\bar{\Omega}_{wv}$ are deterministic constants that are equal across the debtors $w \in \mathcal{N}_v^-$ of each bank v , and depend on the degree type (j_v, k_v) of v . This implies $\bar{\Omega}_{wv} = \bar{\Omega}_{jk}$ when $v \in \mathcal{N}_{jk}$ where $\bar{\Omega}_{jk} = \bar{Z}_{jk}/j$ for a collection of interbank asset parameters \bar{Z}_{jk} ².
5. The remaining balance sheet quantities are freely specified.

Apart from the restrictive condition on exposures $\bar{\Omega}$, this seems to be a natural setting for the model. The GK cascade mechanism has a symmetry that allows the rescaling of every bank and its exposures, while leaving the sequence of defaults unchanged. As one checks easily, if $\lambda_v, v \in \mathcal{N}$ is any collection of positive numbers, specifications $\bar{\Delta}_v, \bar{\Omega}_{wv}$ and $\lambda_v \bar{\Delta}_v, \lambda_v \bar{\Omega}_{wv}$ lead to identical cascades. Therefore, our restriction on exposures is equivalent to saying that they can simultaneously be set to 1 under this symmetry, by taking $\lambda_v = \lambda_{jk} = j/\bar{Z}_{jk}$ for each $v \in \mathcal{N}_{jk}$. For example, taking $\lambda_{jk} = 5j$ changes the benchmark exposures chosen by Gai and Kapadia in their paper to constant exposures $\bar{\Omega} = 1$.

5.1.1 Shocks and the Solvency Condition

Like the early adopters that initiate Watts' adoption cascades, our schematic financial crises are triggered by a set \mathcal{D}_0 of initially defaulted banks. Perhaps they default as the result of a collective shock to the system, leaving the remaining banks with depleted capital buffers. Or perhaps, one bank defaults for idiosyncratic reasons. All such cases can be modelled by supposing these initially defaulted banks have $\bar{\Delta}_v = 0$ while the remaining banks have positive buffers. The set of initially defaulted banks $\mathcal{D}_0 = \{v : \bar{\Delta}_v = 0\}$ thus has conditional probability

$$p_{jk}^{(0)} := \mathbb{P}[v \in \mathcal{D}_0 | v \in \mathcal{N}_{jk}] = D_{jk}(0). \quad (5.3)$$

¹ It is a trivial matter to extend the model slightly to allow for a constant fractional recovery with $R < 1$.

² Temporarily, we allow $\bar{\Omega}_{wv}$ to depend only on the node type of v rather than the edge type of wv . In the next section we will revert to the usual convention.

If we define $D_v^{-1} = 0$ for all v , then the indicator functions for the set \mathcal{D}_n of defaulted banks at step $n \geq 0$ are defined recursively by the formula

$$D_v^n := \mathbf{1}(v \in \mathcal{D}_n) = \mathbf{1}\left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} \bar{\Omega}_{w'v} D_{w'}^{n-1}\right). \quad (5.4)$$

Note that the directed graph condition means that $\bar{\Omega}_{w'v} \neq 0$ only if $w' \in \mathcal{N}_v^-$, which implies the sum in (5.4) is actually over \mathcal{N}_v^- .

We now present a formal derivation of a recursive formula for the conditional default probabilities $p_{jk}^{(n)} := \mathbb{P}[v \in \mathcal{D}_n | v \in \mathcal{N}_{jk}]$ after $n \geq 0$ cascade steps. Consider first (5.4) for bank v at the $n = 1$ step, conditioned on the locally tree-like skeleton $(\mathcal{N}, \mathcal{E})$. By the LTI property, when $v \in \mathcal{N}_{jk}$, the debtor nodes $w_i \in \mathcal{N}_v^-$, $i \in \{1, \dots, j\} := [j]$ are distinct and $\{\bar{\Delta}_v, D_{w_i}^0\}$ are independent random variables. Therefore,

$$\begin{aligned} \mathbb{P}[v \in \mathcal{D}_1 | \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}] &= \mathbb{P}\left[\bar{\Delta}_v \leq \bar{\Omega}_{jk} \cdot \sum_{i \in [j]} D_{w_i}^0 \mid \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}\right] \\ &= \sum_{m=0}^j \mathbb{P}\left[\bar{\Delta}_v \leq \bar{\Omega}_{jk} m \mid v \in \mathcal{N}_{jk}\right] \mathbb{P}\left[\sum_{i \in [j]} D_{w_i}^0 = m \mid \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}\right] \end{aligned} \quad (5.5)$$

where we sum over all possible values m for the number of defaulted neighbours. Knowing the skeleton determines the node types (j_i, k_i) of each w_i , and so summing over σ , denoting the possible size m subsets of the index set $[j]$, leads to

$$\begin{aligned} &\mathbb{P}\left[\sum_{i \in [j]} D_{w_i}^0 = m \mid \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}\right] \\ &= \sum_{\sigma \subset [j], |\sigma|=m} \left(\prod_{i \in \sigma} p_{j_i k_i}^{(0)}\right) \left(\prod_{i \notin \sigma} (1 - p_{j_i k_i}^{(0)})\right). \end{aligned} \quad (5.6)$$

Now, to take the expectation over the skeleton $(\mathcal{N}, \mathcal{E})$ we refer back to the details of the assortative configuration graph construction of Section 3.5.1, in particular Theorem 5, which we apply to the tree graph g which consists of v joined to its debtor nodes $w_i, i \in [j]$. Using a conditional version of (3.67) for this tree graph leads to the asymptotic large N formula

$$\begin{aligned} &\mathbb{E}\left[\left(\prod_{i \in \sigma} p_{j_i k_i}^{(0)}\right) \left(\prod_{i \notin \sigma} (1 - p_{j_i k_i}^{(0)})\right) \mid v \in \mathcal{N}_{jk}\right] \\ &= \sum_{j_1, k_1} \cdots \sum_{j_j, k_j} \prod_{i \in \sigma} p_{j_i k_i}^{(0)} \prod_{i \notin \sigma} (1 - p_{j_i k_i}^{(0)}) \prod_{i \in [j]} \left(\frac{k_i P_{j_i k_i} Q_{k_i j}}{z Q_{k_i}^+ Q_j^-}\right). \end{aligned}$$

By introducing

$$\tilde{\pi}_j^{(0)} := \sum_{j',k'} p_{j'k'}^{(0)} \left(\frac{k' P_{j'k'} Q_{k'|j}}{z Q_{k'}^+ Q_j^-} \right) = \sum_{j',k'} p_{j'k'}^{(0)} P_{j'|k'} Q_{k'|j}$$

where $P_{j'|k'} := P_{j'k'}/P_{k'}^+$ and $Q_{k'|j} = Q_{k'j}/Q_j^-$, this expectation can be written

$$(\tilde{\pi}_j^{(0)})^m (1 - \tilde{\pi}_j^{(0)})^{j-m}. \quad (5.7)$$

Finally, we put (5.6), (5.6), (5.7) together with the cumulative distribution function for $\bar{\Delta}_v$ to obtain the formula for the $n = 1$ conditional default probability:

$$p_{jk}^{(1)} = \mathbb{P}[v \in \mathcal{D}_1 | v \in \mathcal{N}_{jk}] = \sum_{m=0}^j D_{jk}(m \bar{Q}_{jk}) \text{Bin}(j, \tilde{\pi}_{jk}^{(0)}, m) \quad (5.8)$$

where $\bar{Q}_{jk} = \bar{Z}_{jk}/j$, $\text{Bin}(\cdot, \cdot, \cdot)$ denotes a binomial probability just as in (4.28), and

$$\tilde{\pi}_{jk}^{(0)} = \mathbb{P}[w \in \mathcal{D}_0 | w \in \mathcal{N}_v^-, v \in \mathcal{N}_{jk}] = \sum_{j',k'} p_{j'k'}^{(0)} P_{j'|k'} Q_{k'|j}. \quad (5.9)$$

Note that $\tilde{\pi}_{jk}^{(0)} = \tilde{\pi}_j^{(0)}$ is independent of k .

The extension of this argument to all $n \geq 1$ rests on the special relation between the GK mechanism and the LTI dependence structure of the model. We observe that conditioned on $(\mathcal{N}, \mathcal{E})$, the event $v \in \mathcal{D}_1$ depends on $\bar{\Delta}_v, \bar{\Delta}_{w_i}, i \in [j_v]$ where the $w_i \in \mathcal{N}_v^-$ are nodes “up-stream” from v . We see that this upstream dependence extends to all n when $(\mathcal{N}, \mathcal{E})$ is a tree graph, and will be asymptotically true as $N \rightarrow \infty$ on configuration graphs. Thus, under the LTI assumption, by induction on n , the event $v \in \mathcal{D}_n$ depends on the events $w_i \in \mathcal{D}_{n-1}$ for $w_i \in \mathcal{N}_v^-, i \in [j_v]$ which are conditionally independent and have identical default probabilities

$$\tilde{\pi}_j^{(n-1)} := \mathbb{P}[w \in \mathcal{D}_{n-1} | w \in \mathcal{N}_v^-, v \in \mathcal{N}_{jk}] \quad (5.10)$$

that do not depend on k . For reasons that will become clear in the next section, it is natural to reexpress (5.10) in terms of a further collection of conditional probabilities:

$$\tilde{\pi}_j^{(n-1)} = \sum_{k'} \pi_{k'}^{(n-1)} Q_{k'|j} \quad (5.11)$$

$$\pi_{k'}^{(n-1)} := \mathbb{P}[w \in \mathcal{D}_{n-1} | k_w = k'] = \sum_{j'} p_{j'k'}^{(n-1)} P_{j'|k'}. \quad (5.12)$$

Moreover, the conditional default probabilities $p_{jk}^{(n)} = \mathbb{P}[v \in \mathcal{D}_n | v \in \mathcal{N}_{jk}]$ will be given by a formula just like (5.8). This completes the formal justification for cascade mapping formulas for the basic GK model:

Proposition 6. [*Gai-Kapadia Cascade*] Consider the LTI sequence of GK financial networks $(N, P, Q, \bar{\Delta}, \bar{\Omega})$ satisfying Assumption 7. Let $p_{jk}^{(0)} = D_{jk}(0)$ and $\pi_k^{(0)} = \mathbb{P}[w \in \mathcal{D}_0 | k_w = k]$. Then the following formulas hold in the limit as $N \rightarrow \infty$:

1. $\tilde{\pi}_j^{(0)} = \mathbb{P}[w \in \mathcal{D}_0 | w \in \mathcal{N}_v^-, v \in \mathcal{N}_{jk}] = \sum_{k'} \pi_{k'}^{(0)} Q_{k'|j}$, which is independent of k .
2. For any $n = 1, 2, \dots$, the quantities $\tilde{\pi}_j^{(n-1)}, p_{jk}^{(n)}, \pi_k^{(n)}$ satisfy the recursive formulas

$$\tilde{\pi}_j^{(n-1)} = \sum_{k'} \pi_k^{(n-1)} Q_{k'|j}, \quad (5.13)$$

$$p_{jk}^{(n)} = \sum_{m=0}^j D_{jk}(m \bar{\Omega}_{jk}) \text{Bin}(j, \tilde{\pi}_j^{(n-1)}, m), \quad (5.14)$$

$$\pi_k^{(n)} = \sum_{j'} p_{j'k}^{(n)} P_{j'|k} \quad (5.15)$$

3. The new probabilities $\pi^{(n)} = (\pi_k^{(n)})$ are a vector valued function $G(\pi^{(n-1)})$ which is explicit in terms of the specification $(N, P, Q, \bar{\Delta}, \bar{\Omega})$.
4. The cascade mapping G maps $[0, 1]^{\mathbb{Z}^+}$ onto itself, and is monotonic, i.e. $G(\mathbf{a}) \leq G(\mathbf{b})$ whenever $\mathbf{a} \leq \mathbf{b}$, under the partial ordering relation defined by $\mathbf{a} \leq \mathbf{b}$ if and only if $a_k \leq b_k$ for all k . Since $\pi^{(0)} = G(0)$, the sequence $\pi^{(n)}$ converges to the least fixed point $\pi^* \in [0, 1]^{\mathbb{Z}^+}$, that is

$$\pi^* = G(\pi^*). \quad (5.16)$$

Remark 6. When the skeleton is non-assortative, meaning $Q_{k'|j} = Q_{k'}^+$ is independent of j , then each $\tilde{\pi}_j^{(n)} = \tilde{\pi}^{(n)}$ is independent of j . Under this condition, the cascade mapping can be reduced to a scalar function $\tilde{G} : \tilde{\pi} \in [0, 1] \rightarrow [0, 1]$ and our theorem is subsumed in the scalar cascade mapping of Theorem 3.6 in [5].

As compelling as Proposition 6 is, its power is weakened by the restrictive assumption on the form of the interbank exposures. Let us therefore repair this deficiency in the next section, before exploring the broader implications of our cascade mapping theorem.

5.2 The GK Model with Random Link Weights

The primary motivation for allowing the link weights to become random variables is to correct the asymmetry in the way interbank exposures are specified in the previous section. According to Assumption 7.4, “interbank exposures are deterministic constants that are equal across the debtors of each bank”, which means essentially that the creditor bank is choosing to lend equally to its counterparties. One can easily argue the opposite case that the debtor bank might be the counterparty that chooses

to borrow equally. And that the reality is likely closer to a compromise, wherein exposure sizes depend on attributes of both the borrower and the lender.

A look at the argument leading to Proposition 6 will convince one that if the exposures $\bar{\Omega}_\ell$ for $\ell \in \mathcal{E}_v^-$ depend deterministically on k_ℓ rather than j_ℓ , then they will be effectively random variables when conditioned on the degree $j_\ell = j_v$. And the main Proposition 6 cannot handle this fact. As we will now find, dealing with the simple variation where the $\bar{\Omega}_\ell$ depend deterministically on k_ℓ rather than j_ℓ is just as difficult as to deal with the more general case where the exposures $\bar{\Omega}_\ell$ are arbitrary random variables depending on both k_ℓ, j_ℓ . It is in this more general specification that we find what is arguably the most natural setting for LTI default models with zero recovery.

Thus, in this section we analyze the final version of the model assumptions, namely Assumption 7 with 7.4 replaced by 8.4

Assumption 8. 4. Each interbank exposure $\bar{\Omega}_\ell$ depends randomly on its edge type (k_ℓ, j_ℓ) . Conditionally on the skeleton, they form a collection of independent non-negative random variables, independent as well from the default buffers $\bar{\Delta}_v$. Their cumulative distribution functions (CDFs) and probability distribution functions (PDFs) are given for $x \geq 0$ and any k, j by

$$\begin{aligned} W_{kj}(x) &= \mathbb{P}[\bar{\Omega}_\ell \leq x | \ell \in \mathcal{E}_{kj}], \\ w_{kj}(x) &= dW_{kj}(x)/dx. \end{aligned} \quad (5.17)$$

Note that these generalized assumptions, particularly the assumed independence are still consistent with the basic LTI structure identified in Section 4.8. But now that exposures $\bar{\Omega}_\ell$ are random variables, we must reconsider the previous analysis of the default cascade.

First we reconsider (5.5) which now takes the form

$$\mathbb{P}[v \in \mathcal{D}_1 | \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}] = \mathbb{P}\left[\bar{\Delta}_v \leq \sum_{i \in [j]} \bar{\Omega}_{\ell_i} D_{w_i}^0 \mid \mathcal{N}, \mathcal{E}, v \in \mathcal{N}_{jk}\right]$$

where $\ell_i = (w_i v)$ for each i . Now, under this condition, $\bar{\Omega}_{\ell_i}, D_{w_i}^0, i \in [j]$ is a collection of independent random variables, and we can calculate that

$$\mathbb{P}[\bar{\Omega}_{\ell_i} D_{w_i}^0 \leq x | \mathcal{N}, \mathcal{E}, \ell_i \in \mathcal{E}_v^-, v \in \mathcal{N}_{jk}] = \tilde{W}_j^{(0)}(x)$$

is independent of k and given by

$$\tilde{W}_j^{(0)}(x) = \sum_{k'} Q_{k'|j} \left((1 - \pi_{k'}^{(0)}) \mathbf{1}(x \geq 0) + \pi_{k'}^{(0)} W_{k'j}(x) \right). \quad (5.18)$$

As before, $\pi_{k'}^{(0)} = \sum_{j'} p_{j'k'}^{(0)} P_{j'|k'}$. By the LTI property, this conditional independence of edge weights $\bar{\Omega}_{\ell_i}$ and indicators $D_{w_i}^{n-1}$ continues for higher n , and we find that (5.18) extends to

$$\tilde{W}_j^{(n-1)}(x) = \sum_{k'} \mathcal{Q}_{k'|j} \left((1 - \pi_{k'}^{(n-1)}) \mathbf{1}(x \geq 0) + \pi_{k'}^{(n-1)} W_{k'j}(x) \right) \quad (5.19)$$

for all $n \geq 1$. For this, we need to make the inductive verification using the LTI property that for $n \geq 0$ the collection of defaulted node events $w_i \in \mathcal{D}_{n-1}$ is mutually independent, and independent of all buffer and exposure random variables downstream to w_i . Then it follows that, conditioned on $(\mathcal{N}, \mathcal{E})$ and $v \in \mathcal{N}_{jk}$, the collection of random variables $\mathbf{1}(w_i \in \mathcal{D}_{n-1})$, $\bar{\Omega}_{\ell_i}$ and $\bar{\Delta}_v$ are mutually independent.

Bringing these pieces together, we find that for $v \in \mathcal{N}_{jk}$, the conditional default event $v \in \mathcal{D}_n = \{\Delta_v \leq \sum_{i \in [j]} \bar{\Omega}_{\ell_i} D_{w_i}^{n-1}\}$ has the form

$$X \leq Z, \quad Z := \sum_{i=1}^j Y_i$$

where X, Y_1, \dots, Y_j are independent non-negative random variables. Moreover, their CDFs are $F_X(x) = D_{jk}(x)$ and $F_{Y_i}(x) = \tilde{W}_j^{n-1}(x)$ with $Y_i \sim Y$, respectively. The probability in question is thus

$$\mathbb{P}[X \leq Z] = \int_0^\infty \int_0^\infty \mathbf{1}(x \leq z) f_X(x) f_Z(z) dx dz = \int_{\mathbb{R}} F_X(z) f_Z(z) dz = \langle F_X, f_Z \rangle \quad (5.20)$$

where $\langle f, g \rangle := \int_{\mathbb{R}} \bar{f}(x) g(x) dx$ is the usual (Hermitian) inner product over \mathbb{R} . To determine the PDF $f_Z(z) = F_Z'(z)$ of Z , let us stop for a moment to review the convolution of probability density functions. The sum of independent random variables $Z = \sum_{i=1}^j Y_i$ with known PDFs $f_i(x)$ has a PDF $f_Z(x)$ given by convolution:

$$f_Z(x) = [f_1 \otimes f_2 \otimes \dots \otimes f_j](x)$$

where the convolution product of functions is defined by

$$[f \otimes g](x) = \int_{\mathbb{R}} f(x-y) g(y) dy \quad (5.21)$$

In the present context, this means f_Z is a convolution power, $f_Z = (f_Y)^{\otimes j}$ with

$$f_Y(x) = \tilde{w}_j^{(n-1)}(x) = \frac{d\tilde{W}_j^{(n-1)}(x)}{dx} = \sum_{k'} \mathcal{Q}_{k'|j} \left((1 - \pi_{k'}^{(n-1)}) \delta_0(x) + \pi_{k'}^{(n-1)} w_{k'j}(x) \right). \quad (5.22)$$

Note that here we need to represent the point masses at $Y_i = 0$ and other potential point masses using delta functions $\delta_a(x)$.

We can now put together what is perhaps the main result of this book, which is the generalization of Proposition 6 to account for random link weights. Given its importance to the overall theory of financial cascades, we will also provide in the subsequent section a rigorous justification for the formal $N = \infty$ arguments we have been contented with until now.

Theorem 12. [Gai-Kapadia Cascade, RLW version] Consider the LTI sequence of GK financial networks $(N, P, Q, \bar{\Delta}, \bar{\Omega})$ satisfying Assumption 8 with random link weights. Let $p_{jk}^{(0)} = D_{jk}(0)$ and $\pi_k^{(0)} = \mathbb{P}[w \in \mathcal{D}_0 | k_w = k]$ be initial default probabilities. Then the following formulas hold as $N \rightarrow \infty$:

1. For any $n = 1, 2, \dots$, the quantities $p_{jk}^{(n)}, \pi_k^{(n)}$ satisfy the recursive formulas

$$p_{jk}^{(n)} = \langle D_{jk}, (\tilde{w}_j^{(n-1)})^{\otimes j} \rangle, \quad (5.23)$$

$$\pi_k^{(n)} = \sum_{j'} p_{j'k}^{(n)} P_{j'|k} \quad (5.24)$$

where the PDFs $\tilde{w}_j^{(n-1)}(x)$ are given by (5.22).

2. The new probabilities $\pi^{(n)}$ are a vector valued function $G(\pi^{(n-1)})$ which is explicit in terms of the specification $(N, P, Q, \bar{\Delta}, \bar{\Omega})$.
3. The cascade mapping G maps $[0, 1]^{\mathbb{Z}_+}$ onto itself, and is monotonic. Since $\pi^{(0)} = G(0)$, the sequence $\pi^{(n)}$ converges to the least fixed point $\pi^* \in [0, 1]^{\mathbb{Z}_+}$, that is

$$\pi^* = G(\pi^*). \quad (5.25)$$

Now that we have a final form of our main Theorem for the GK model, without unnatural and unnecessary restrictions, we can proceed to draw some interesting conclusions.

- Remark 7.*
1. We observe that the main conclusion of Proposition 6, namely the existence of a vector-valued monotonic cascade mapping $G : [0, 1]^{\mathbb{Z}_+} \rightarrow [0, 1]^{\mathbb{Z}_+}$ remains true with general random link weights.
 2. The vector-valued fixed point equation seems to be the natural feature of this type of model. In order to get a scalar fixed point condition such as the one obtained in [5], we need to assume both non-assortativity, i.e. $Q_{kj} = Q_k^+ Q_j^-$, and that the probability distributions of $\bar{\Omega}_{wv}$ depend on the degree type of v but not w .
 3. The numerical computation of this cascade mapping now makes intensive use of equation 5.23, which involves repeated high-dimensional integrations. Fortunately, as we shall see, we can make use of Fourier Transform methods to make this type of computation extremely efficient.

5.3 Proof of Theorem 12

The aim is now to fill the gap in the heuristic proof, based on exact LTI properties that hold only formally for $N = \infty$, that leads to the statement of Theorem 12. Amini, Cont and Minca, in [5], have proved a similar theorem in a somewhat narrower setting by a beautiful and lengthy argument that extends random graph methods developed by [59] and many others. However, it seems difficult to extend their method and ideas to more general settings. The method we now present provides an

alternative approach with a completely different logic that is applicable to a broad range of LTI model specifications. To keep the argument as transparent as possible, we focus on computing the probability of events $v \in \mathcal{D}_n$: By extension however, the same argument applies to more refined events such as $v \in \mathcal{D}_n \cap \mathcal{N}_{kj}$.

The aim is to show

$$\lim_{N \rightarrow \infty} \frac{1}{N} \mathbb{E}^N \left[\sum_{v \in \mathcal{N}} \mathbf{1}(v \in \mathcal{D}_n) \right] = p^{(n)} := \sum_{jk} p_{jk}^{(n)} P_{jk} \quad (5.26)$$

where $p_{jk}^{(n)}$ is given by (5.23). The strategy is as follows. First, regularize the infinite sums occurring in the problem by cutting all edge weights on high degree edges, effectively truncating the node and edge degrees to be less than a finite K . Second, compute $\mathbb{E}^{N,K}[\mathbf{1}(v \in \mathcal{D}_n) | \mathcal{N}, \mathcal{E}]$, which is an expectation of a function of the collection of *independent* random variables $(\bar{\Delta}, \bar{\Omega}^K)$ conditional on the skeleton \mathcal{N}, \mathcal{E} . Thirdly, compute the expectation over skeletons \mathcal{N}, \mathcal{E} as a sum over configurations of the type analyzed in detail in Section 3.5.2. The limit $N \rightarrow \infty$ can then be taken for fixed K , leading to a truncated version of (5.26). Finally, the limit $K \rightarrow \infty$ exists by monotonicity, and inspection reveals that the limit satisfies the recursive properties consistent with the statement of the theorem.

For each $K < \infty$, the degree truncation cuts off all edges that have a k or j degree greater than K :

$$\Omega_\ell^K = \Omega_\ell \mathbf{1}(k_\ell \leq K, j_\ell \leq K). \quad (5.27)$$

Note that for each K , the truncated model simply modifies the edge weight distributions so it satisfies the hypotheses of the Theorem and possesses the requisite $N = \infty$ properties. For any positive integers n, K, N , the regularized default probabilities at cascade step n are defined to be

$$p^{(n,K,N)} := \mathbb{E}^{N,K} \left[\frac{1}{N} \sum_{v \in \mathcal{N}} \mathbf{1}(v \in \mathcal{D}_n) \right]. \quad (5.28)$$

Now consider one term $\mathbb{E}^{N,K}[\mathbf{1}(v \in \mathcal{D}_n) | (\mathcal{N}, \mathcal{E})]$ conditional on a fixed (P, Q) skeleton graph $(\mathcal{N}, \mathcal{E})$ of size $|\mathcal{N}| = N$. By the local nature of the cascade dynamics, the event $\mathbf{1}(v \in \mathcal{D}_n)$ clearly depends only on the random variables $(\bar{\Delta}, \bar{\Omega})$ restricted to the n -th order in-neighbourhood $g := \mathcal{N}_v^{n,-}$: this is a collection of conditionally independent random variables whose distributions are determined by the degree types of the nodes of g . Such a connected subgraph g , labelled by degrees, is an example of a *configuration* that was the object of study in Section 3.5.2. Performing the conditional expectation then leads to an N -independent quantity $X(g, n, K)$ that is a functional only of the labels on g and n, K .

Next, the unconditional expectation will be

$$\mathbb{E}^{N,K}[\mathbf{1}(v \in \mathcal{D}_n)] = \mathbb{E}^{N,K} \left[\frac{1}{N} \sum_{v \in \mathcal{N}} \mathbf{1}(v \in \mathcal{D}_n) \right] = \sum_g P(g, K, N) X(g, n, K) \quad (5.29)$$

where the sum is over all possible n -th order in-neighbourhoods g , labelled by degrees, and $P(g, K, N)$ denotes the probability of this labelled g . For K, n finite, this is a finite sum since degree truncation renders finite both the sum over possible topologies of g as well as the sums over possible degrees. Finally, Assumption 8 includes the relevant conditions on the collection (P_{jk}, Q_{kj}) so that Theorem 5 controls the limit $N \rightarrow \infty$ of $P(g, K, N)$ for n, K fixed. It implies that $\lim_{N \rightarrow \infty} P(g, K, N) = P(g, K, \infty)$ is finite if g is a tree and zero if g has any cycles at all. Therefore,

$$\lim_{N \rightarrow \infty} p^{(n, K, N)} = \sum_{g \text{ tree}} P(g, K, \infty) X(n, g, K) := p^{(n, K, \infty)} \quad (5.30)$$

where the sum is restricted to tree graphs g , labelled by their bounded degrees.

It is trivial to verify two assertions about the collection of $N = \infty$ probabilities $p^{(n, K, \infty)}$: they are jointly monotonically increasing in both K and n and they are bounded by 1. Therefore $p^{(n, \infty, \infty)} = \lim_{K \rightarrow \infty} p^{(n, K, \infty)}$ exists and is a monotone sequence in n . Next recall the consequence of Theorem 5 that configuration probabilities for tree graphs with $N = \infty$ encode the locally tree-like (LT) property that is at the core of formal verification of the recursive cascade dynamics. This implies that the collection of probabilities $p^{(n, K, \infty)}$, when combined with the collection of more refined probabilities $p_{jk}^{(n, K, \infty)}$ that can be analyzed in the same way, is a realization of a system of quantities that are functionally related by the recursion relations specified by our Theorem. To reiterate, because the conditions of the theorem apply to the cascade model with truncated edge weights $\bar{\Omega}^K$, all terms in the finite sum over trees in (5.30) have the LT property that provides the formal justification of the recursive cascade mapping (5.23), (5.24).

From these observations, one concludes that the Theorem is true for the set of probabilities $p^{(n, K, \infty)}$ for K finite. By boundedness and monotonicity, it is also true of the limiting collection $p^{(n, \infty, \infty)}$. \square

Remark 8. 1. The structure of this proof is indeed not sensitive to the specific nature of the cascade dynamics and random financial network, but is highly dependent on the LTI property. Therefore, we have confidence that large N asymptotics will continue to hold for a wide range of more complex network models as long as they have a generalized LTI property. Note also that this strategy clearly works to yield rigorous proofs for the Watts model of Section 4.5.2 and the basic GK model.

2. Under the conditions on the collection (P_{jk}, Q_{kj}) , these results can be extended to show convergence of the empirical probability of default: by the use of the Markov Inequality, one can show that with high probability

$$\lim_{N \rightarrow \infty} \frac{1}{N} \sum_{v \in \mathcal{N}} \mathbf{1}(v \in \mathcal{D}_n) = p^{(n)}. \quad (5.31)$$

5.4 Measures of Cascade Impact

In cascade modelling, it is always important to have a number of economic measures of the impact of such hypothetical crises. In this section, we explore some simple statistical measures that can be computed analytically using the cascade mapping theorem. The first global measure of cascade impact is of course the unconditional probability of eventual default, computed by

$$p^* := \mathbb{P}[v \in \mathcal{D}_\infty] = \sum_{jk} P_{jk} p_{jk}^{(\infty)}. \quad (5.32)$$

Also of interest are two variants of default probability that condition on a bank being either a debtor or creditor of another bank. That is, if $w \in \mathcal{N}_v^-$,

$$p^{*-} = \mathbb{P}[w \in \mathcal{D}_\infty | w \in \mathcal{N}_v^-] = \sum_{jk} P_{jk} Q_k^+ p_{jk}^{(\infty)} \quad (5.33)$$

$$p^{*+} = \mathbb{P}[v \in \mathcal{D}_\infty | w \in \mathcal{N}_v^-] = \sum_{jk} P_{k|j} Q_j^- p_{jk}^{(\infty)}. \quad (5.34)$$

Recall from Section 2.3 that the zero-recovery assumption implies there is a large cost to the system at the time any bank defaults, and we are certainly interested in computing measures of this impact. The first measure was defined by (2.13):

$$\text{Cascade impact: } \text{CI} = N^{-1} \sum_v \bar{X}_v \mathbf{1}(v \in \mathcal{D}_\infty).$$

The expected cascade impact per bank can easily be computed in terms of the fixed point π^* and the eventual default probabilities $p^* = F(\pi^*)$:

$$\mathbb{E}[\text{CI}] = \sum_{jk} P_{jk} \mathbb{E} \left[\sum_{\ell \in \mathcal{E}_v^+} \bar{\Omega}_\ell \mathbf{1}(v \in \mathcal{D}_\infty) \middle| v \in \mathcal{N}_{jk} \right].$$

In an LTI model, the collection $\bar{\Omega}_\ell, \ell \in \mathcal{E}_v^+$ and $\mathbf{1}(v \in \mathcal{D}_\infty)$ is mutually independent conditioned on $v \in \mathcal{N}_{jk}$, and for large N we have the asymptotic formula

$$\mathbb{E}[\text{CI}] = \sum_{jk} k P_{jk} p_{jk}^{(\infty)} \sum_{j'} Q_{j'|k} \mathbb{E} \bar{\Omega}_{kj'} \quad (5.35)$$

where $\mathbb{E} \bar{\Omega}_{kj} := \mathbb{E}[\bar{\Omega}_\ell | \ell \in \mathcal{E}_{kj}] = \int_{\mathbb{R}^+} x dW_{kj}(x)$.

By comparing cascade impact to the average buffer size $\mathbb{E}[\bar{\Delta}_v] = \sum_{jk} P_{jk} \int_{\mathbb{R}^+} x dD_{jk}(x)$ we also get a measure of the average eventual buffer size:

$$\mathbb{E}[\Delta_v^{(\infty)}] = N^{-1} \mathbb{E} \left[\sum_v \left(\bar{\Delta}_v - \sum_{w \in \mathcal{N}_v^-} \bar{\Omega}_{vw} \mathbf{1}(w \in \mathcal{D}_\infty) \right) \right] = E[\bar{\Delta}_v - \text{CI}]. \quad (5.36)$$

A slightly different measure of the cascade impact is the *average shortfall* of banks' buffers compared to the default shocks impacting them:

$$\text{Average Shortfall: } AS = N^{-1} \sum_v \left(\sum_{w \in \mathcal{N}_v^-} \tilde{\Omega}_{wv} \mathbf{1}(w \in \mathcal{D}_\infty) - \tilde{\Delta}_v \right)^+.$$

The expected average shortfall also has a nice formula which we can compute using a formula related to (5.20): for two independent random variables X, Z ,

$$\mathbb{E}[(Z-X)^+] = \int_{\mathbb{R}^2} (z-x)^+ f_X(x) f_Z(z) dx dz = \int_{\mathbb{R}} \tilde{F}_X(z) f_Z(z) dz = \langle \tilde{F}_X, f_Z \rangle \quad (5.37)$$

where $\tilde{F}_X(z) = \int_{-\infty}^z (z-x) f_X(x) dx$ is the integrated CDF of X . Thus we find the expected average shortfall is

$$EAS = \sum_{jk} P_{jk} \langle \tilde{D}_{jk}, (\tilde{w}_j^{(\infty)})^{\otimes j} \rangle \quad (5.38)$$

where $\tilde{D}'_{jk}(x) = D_{jk}(x)$ and $\tilde{w}_j^{(\infty)}$ is the limit of (5.22).

Default Correlation: We are also interested in *default correlation* and more refined measures of joint default such as *CoVaR* introduced by [3]. The most basic network measure of dependence is the joint probability of eventual default for counterparty pairs (w, v) :

$$p_{joint}^{(\infty)} := \mathbb{P}[w \in \mathcal{D}_\infty, v \in \mathcal{D}_\infty | w \in \mathcal{N}_v^-].$$

Equivalently, one can compute the default probabilities of a bank conditioned on default of one of its counterparties:

$$\mathbb{P}[v \in \mathcal{D}_\infty | w \in \mathcal{D}_\infty, w \in \mathcal{N}_v^-] = \frac{p_{joint}^{(\infty)}}{p^{*-}}, \quad \mathbb{P}[w \in \mathcal{D}_\infty | v \in \mathcal{D}_\infty, w \in \mathcal{N}_v^-] = \frac{p_{joint}^{(\infty)}}{p^{*+}} \quad (5.39)$$

where the denominators are the two natural conditional marginal default probabilities from (5.33), (5.34). The easiest way to compute such quantities is first to disaggregate the joint probability of *nondefault* over the types of w and v . We then have

$$\begin{aligned} & \mathbb{P}[w \notin \mathcal{D}_\infty, v \notin \mathcal{D}_\infty | w \in \mathcal{N}_v^- \cap \mathcal{N}_{j'k'}, v \in \mathcal{N}_{jk}] \\ &= \mathbb{P} \left[\Delta_w \leq \sum_{w' \in \mathcal{N}_w^-} \Omega_{w'w} D_{w'}^\infty, \Delta_v \leq \sum_{w' \in \mathcal{N}_v^-} \Omega_{w'v} D_{w'}^\infty \mathbf{1}(w' \neq w) \mid w \in \mathcal{N}_v^- \cap \mathcal{N}_{j'k'}, v \in \mathcal{N}_{jk} \right] \\ &= \left(1 - \langle D_{j'k'}, (\tilde{w}_{j'}^{(\infty)})^{\otimes j'} \rangle \right) \left(1 - \langle D_{jk}, (\tilde{w}_j^{(\infty)})^{\otimes j-1} \rangle \right). \end{aligned}$$

Note that the convolution power in the second factor is reduced by 1 since it is known that one counterparty of v has definitely not defaulted. Therefore, the uncon-

ditional joint probability of *non-default* is

$$\begin{aligned} & \mathbb{P}[w \notin \mathcal{D}_\infty, v \notin \mathcal{D}_\infty | v \in \mathcal{N}_w^+] \\ &= \sum_{j'k',jk} P_{j'|k'} Q_{k'j} P_{k|j} \left(1 - \langle D_{j'k'}, (\tilde{w}_{j'}^{(\infty)})^{\otimes j'} \rangle \right) \left(1 - \langle D_{jk}, (\tilde{w}_j^{(\infty)})^{j-1} \rangle \right) \\ &= \sum_{jk} P_{k|j} (1 - \tilde{\pi}_j^*) \left(1 - \langle D_{jk}, (\tilde{w}_j^{(\infty)})^{\otimes j-1} \rangle \right) \end{aligned}$$

where $\tilde{\pi}_j^* = \sum_{k'} \pi_{k'}^* Q_{k'j}$. Since the marginal probabilities of default are p^{*-} and p^{*+} one can see that

$$p_{joint}^{(\infty)} = \mathbb{P}[w \notin \mathcal{D}_\infty, v \notin \mathcal{D}_\infty | v \in \mathcal{N}_w^+] + p^{*-} + p^{*+} - 1. \quad (5.40)$$

5.5 The Cascade Condition

Just as in Section 4.5.3 we derived a cascade condition that characterizes the growth of small cascades in the Watts model, we now consider the possibilities for cascades on the Gai-Kapadia RFN that start with small initial default probabilities. That is, we let $D_{jk}(0) = p_{jk}^{(0)} := \varepsilon_{jk} \geq 0$ and suppose these are uniformly bounded by a small positive number $\bar{\varepsilon} > 0$:

$$|\varepsilon| := \max_{jk} |\varepsilon_{jk}| \leq \bar{\varepsilon}.$$

Then the default buffer CDF (5.41) becomes

$$D_{jk}(x) = \varepsilon_{jk} + (1 - \varepsilon_{jk}) \tilde{D}_{jk}(x), \quad x \geq 0 \quad (5.41)$$

with \tilde{D} interpreted as the CDF of $\tilde{\Delta}_v$ conditioned on v not initially defaulted.

As we have remarked already, a very low density of initially defaulted banks means they are likely to be far apart in the network and the only probable way for a large cascade to develop is that there should be a positive probability for any single bank to trigger an increasing sequence of defaults, without regard to other initially defaulted banks. In Section 4.5.3 we found for the Watts model that this statement can be related to the existence of a giant connected “vulnerable” cluster.

We again write $G(\pi) = G(\pi; \varepsilon)$ to highlight the dependence on the parameters ε and suppress the dependence on \tilde{D} . Now the sequence $\pi^{(n)}$ starting from the initial values $\pi^{(0)} = G(0; \varepsilon)$ must converge to the least fixed point $\pi^*(\varepsilon)$. The question now is: is $\pi^*(\varepsilon)$ of order ε or of order 1 as $\varepsilon \rightarrow 0$? In other words, is there a *cascade condition* that determines if an infinitesimally small initial “seed” fraction will grow to a large-scale cascade? In view of the vector valued nature of our present problem, it turns out that the answer depends on the spectral radius of the derivative matrix ∇G with $\nabla G_{k,k'} = \partial G_k / \partial \pi_{k'}$ evaluated at $\pi = 0; \varepsilon = 0$. Recall that the *spectral radius* of ∇G , defined by $\|\nabla G\| := \max_{a:|a|=1} |\nabla G \cdot a|$, is the largest eigenvalue of ∇G in absolute value.

In our framework, the derivative ∇G is easy to calculate:

$$\nabla G_{k,k'} = \sum_j j (\langle \tilde{D}_{jk}, w_{k'j} \rangle - \tilde{D}_{jk}(0)) Q_{k'|j} P_{j|k}. \quad (5.42)$$

Note each component of ∇G is non-negative: To enable an elementary proof of the following result, we assume each component is strictly positive and the degrees are bounded.

Proposition 7. (*Cascade Condition*) *Suppose the extended GK financial network $(N = \infty, P, Q, \tilde{A}, \tilde{\Omega})$ is such that ∇G defined by (5.42) is a component-wise positive finite dimensional matrix.*

1. *If $\|\nabla G\| > 1$, then there is $\bar{\varepsilon} > 0$ such that for all ε with $0 < |\varepsilon| < \bar{\varepsilon}$, $|\pi^*(\varepsilon)| > \bar{\varepsilon}$. That is, in this network, any uniform seed with a positive fraction will trigger a cascade with default fraction bigger than $\bar{\varepsilon}$ almost surely.*
2. *If $\|\nabla G\| < 1$, then there is $\bar{\varepsilon} > 0$ and C such that for all ε with $0 < |\varepsilon| < \bar{\varepsilon}$, $|\pi^*(\varepsilon)| \leq C\varepsilon$. That is, this network will almost surely not exhibit large scale cascades for any infinitesimal seed.*

Proof: Part 1: We write $\nabla G = M_\varepsilon$ where $M_\varepsilon = \partial G / \partial \pi|_{\pi=0, \varepsilon}$. By continuous dependence in ε , there are values $\bar{\varepsilon}_1 > 0$ and $\lambda > 1$ such that the matrix M_ε is positive and has spectral radius $\|M_\varepsilon\| \geq \lambda$ for all ε with $0 \leq \varepsilon_{jk} < \bar{\varepsilon}_1$. Let us fix any such ε .

By the Perron-Frobenius Theorem for positive matrices, there is a unique normalized eigenvector v such that $M_\varepsilon \cdot v = \|M_\varepsilon\|v$: it has all positive entries and normalization $|v| = 1$. Taylor's Theorem with a second order remainder implies that for $\bar{\varepsilon}_1$ small enough there is $C' > 0$ such that

$$G(a; \varepsilon) = G(0; \varepsilon) + M_\varepsilon \cdot a + R(a), \quad |R(a)| \leq C'|a|^2$$

for all $a \in [0, 1]^{\mathbb{Z}^+}$ with $|a| \leq \bar{\varepsilon}_1$ (note we drop the \cdot notation in the following).

Now we show that the sequence $a^{(1)} = G(0; \varepsilon)$, $a^{(n+1)} = G(a^{(n)}; \varepsilon)$ leaves the set $|a| \leq \bar{\varepsilon}$ provided $\bar{\varepsilon}$ is chosen small enough (independently of ε). For this, since $\bar{\varepsilon}_1 > 0$ there is $\beta_1 > 0$ and a non-negative vector y_1 such that $a^{(1)} = \beta_1 v + y_1$. Assuming inductively that $a^{(n)} = \beta_n v + y_n$ for some $\beta_n > 0$ and a non-negative vector y_n and that $|a^{(n)}| \leq \bar{\varepsilon}_1$, the monotonic property of G combined with Taylor's Theorem implies

$$\begin{aligned} a^{(n+1)} &= G(a^{(n)}; \varepsilon) \geq G(\beta_n v; \varepsilon) \\ &= G(0; \varepsilon) + \beta_n M_\varepsilon \cdot v + R(\beta_n v) \\ &\geq \beta_1 v + y_1 + \frac{1}{2}(1 + \lambda)\beta_n v + \left(\frac{1}{2}(\lambda - 1)\beta_n v + R(\beta_n v) \right) \end{aligned}$$

Let $\beta_{n+1} = \beta_1 + \frac{1}{2}(1 + \lambda)\beta_n$ and note that $y_{n+1} = a^{(n+1)} - \beta_{n+1}v \geq 0$ provided $\bar{\varepsilon} \leq \min(\bar{\varepsilon}_1, \frac{1}{2C'}(\lambda - 1) \min_j v_j)$. Since the sequence β_n increases without bound, we can iterate the inductive argument only a finite number of steps before $|a^{(n+1)}| > \bar{\varepsilon}$.

Part 2: By continuous dependence in both a and ε , there are now values $\bar{\varepsilon} > 0$ and $\lambda = \frac{1}{2}(1 + \|\nabla G\|) < 1$ such that the matrix $M_{a; \varepsilon} = \partial G / \partial a|_{a; \varepsilon}$ has spectral radius

$\|M_{a;\varepsilon}\| \leq \lambda$ for all $0 \leq \varepsilon < \bar{\varepsilon}$ and $|a| \leq \bar{\varepsilon}$. Fix any such ε . Now we note that for vectors a, b with $|a|, |b| \leq \bar{\varepsilon}$ we can use Taylor's Theorem again to write

$$G(a; \varepsilon) - G(b; \varepsilon) = M_\varepsilon \cdot (a - b) + R(a, b)$$

where the remainder has bound $C''|a - b|^2$ for some $C'' > 0$. Then provided $|a^{(n+1)}|, |a^{(n)}| \leq \bar{\varepsilon}$ and $\bar{\varepsilon} \leq \frac{1 - \|\nabla G\|}{4C''}$

$$\begin{aligned} |a^{(n+1)} - a^{(n)}| &= |G(a^{(n)}; \varepsilon) - G(a^{(n-1)}; \varepsilon)| \\ &\leq \frac{1}{2}(\lambda + 1)|a^{(n)} - a^{(n-1)}| + \left(\frac{1}{2}(\lambda - 1)|a^{(n)} - a^{(n-1)}| + |R(a^{(n)}, a^{(n-1)})| \right) \\ &\leq \frac{1}{2}(\lambda + 1)|a^{(n)} - a^{(n-1)}| \end{aligned}$$

for all $n \geq 1$. Since $|G(0; \varepsilon)| \leq C'\varepsilon$ for some $C' > 0$, we can iterate this inequality to show $|a^{(\infty)}| \leq C\varepsilon$ with $C = \frac{4C'}{1 - \|\nabla G\|}$. \square

We can understand the cascade condition more clearly by introducing the notion of *vulnerable edge* which means a directed edge $\ell = (wv)$ whose weight $\bar{\Omega}_\ell$ exceeds the default buffer of its downstream node $v = \mathcal{N}_\ell^+$. For the following argument, we suppose the network has only solvent banks, i.e. $D_{jk}(0) = 0$ for all j, k . An edge $\ell = (wv)$ is thus vulnerable if and only if $\bar{\Delta}_v \leq \bar{\Omega}_\ell$. The matrix element $\nabla G_{kk'}$ has a simple explanation that gives more intuition about the nature of the cascade condition: it is the expected number of vulnerable edges ℓ with $k_\ell = k'$ that enter a node v with $k_v = k$. Then for small values of π , one has a linear approximation for the change in π in a single cascade step:

$$\pi_k^{m+1} - \pi_k^m = \sum_{k'} G_{k,k'} (\pi_{k'}^m - \pi_{k'}^{m-1}) + O(|\pi|^2). \quad (5.43)$$

The condition for a global cascade starting from an infinitesimal seed is that the matrix ∇G must have an expanding direction, i.e. an eigenvalue with magnitude bigger than 1.

It turns out that the cascade condition is indeed a strong measure of systemic risk in simulated networks. One can check that in the setting of independent edge probabilities $Q_{kj} = Q_k^+ Q_j^-$ and deterministic edge weights $\bar{\Omega}_{wv} = \bar{\Omega}_{jk}$ when $v \in \mathcal{N}_{jk}$, the spectral radius becomes

$$\|\nabla G\| = \sum_{jk} \frac{jk}{z} P_{jk} \mathbb{P}[\bar{\Delta}_{j'k} \leq \bar{\Omega}_{j'k}],$$

a result that has been derived in a rather different fashion in [38] and [5]. [38] also extends the [72] percolation theory approach from undirected networks to the case of directed nonassortative networks and we will see in the next section that the percolation approach to the cascade condition also extends to our general setting of directed assortative networks with random edge weights.

5.5.1 Frequency and Size of Global Cascades

We learned in Section 4.5.3 that the possibility of a large scale cascade in the Watts model depends on the connectivity of the directed subnetwork of vulnerable edges and nodes, a problem related to “site percolation”. The previous section addressed the potential for a small seed to grow into a global GK cascade, and now we wish to understand how the frequency of global cascades in large random networks is related to the so-called extended in-component associated to the giant vulnerable cluster. In the present context, a vulnerable cluster has the meaning of a connected subgraph of the network consisting of *vulnerable* directed edges, where a vulnerable directed edge is an edge whose weight is sufficient to exceed the default buffer of its downstream node. We define:

- $\mathcal{E}_V \subset \mathcal{E}$, the set of vulnerable directed edges;
- \mathcal{E}_s , the largest strongly connected set of vulnerable edges (the *giant vulnerable cluster* of \mathcal{E}_V);
- \mathcal{E}_i and \mathcal{E}_o , the *in-component* and *out-component* of the giant vulnerable cluster, i.e. the set of vulnerable edges that are connected to or from \mathcal{E}_s by a directed path of vulnerable edges;
- $1 - b_k := \mathbb{P}[\ell \in \mathcal{E}_i | k_\ell = k]$, a conditional probability of an edge being in \mathcal{E}_i ;
- $a_{k,jk'} = \mathbb{P}[\bar{\Delta}_v \leq \bar{\Omega}_{wv} | \ell \in \mathcal{E}_v^-, k_\ell = k, v \in \mathcal{N}_{jk'}]$, the conditional probability of an edge being vulnerable.

Now note that $\ell = (w, v) \in \mathcal{E}_i^c$ (i.e. the complement of \mathcal{E}_i) means either $\bar{\Delta}_v > \bar{\Omega}_{wv}$ or $\bar{\Delta}_v \leq \bar{\Omega}_{wv}$ and all the k_v “downstream” directed edges $\ell' \in \mathcal{E}_v^+$ are in the set \mathcal{E}_i^c . Thus, invoking the LTI, one determines that for all $k \in \mathbb{Z}_+$

$$b_k = \sum_j Q_{j|k} \sum_{k'} P_{k'|j} \left(1 - a_{k,jk'} + a_{k,jk'} b_{k'} \right) := H_k(b) \quad (5.44)$$

where $a_{k,jk'} = \langle \bar{D}_{jk'}, w_{kj} \rangle$. In other words, the vector $b = (b_k)$ satisfies the fixed point equation $b = H(b)$ where

$$H_k(b) = \sum_{jk'} Q_{j|k} P_{k'|j} \left(1 - \langle \bar{D}_{jk'}, w_{kj} \rangle + \langle \bar{D}_{jk'}, w_{kj} \rangle b_{k'} \right), \quad k \in \mathbb{Z}_+. \quad (5.45)$$

The equation $b = H(b)$ has a trivial fixed point $e = (1, 1, \dots)$. In case $e = (1, 1, \dots)$ is a stable fixed point, we expect that the set \mathcal{E}_i will have probability zero. We now verify that the cascade condition $\|\nabla G\| > 1$ of Proposition 7 is equivalent to the condition that e is an unstable fixed point, in which case there will be a nontrivial fixed point $0 \leq b_\infty < e$ that corresponds to the set \mathcal{E}_i having positive probability. A sufficient condition for e to be an unstable fixed point is that $\|\nabla H\| > 1$ where the derivative $\nabla H_{kk'} = (\partial H_k / \partial b_{k'})|_{b=e}$ is given by

$$\nabla H_{kk'} = \sum_j k' Q_{j|k} P_{k'|j} \langle \bar{D}_{jk'}, w_{kj} \rangle. \quad (5.46)$$

One can verify directly that

$$\nabla H = \Lambda^{-1} \cdot (\nabla G)' \cdot \Lambda$$

for the diagonal matrix $\Lambda_{kk'} = \delta_{kk'} k Q_k^+$ and from this it follows that the spectrum, and hence the spectral radii and spectral norms, of ∇H and ∇G are equal. Hence $\|\nabla H\| > 1$ if and only if $\|\nabla G\| > 1$.

As long as the cascade condition $\|\nabla H\| > 1$ is satisfied, a global cascade will arise from a random single seed v if it triggers at least one edge $(v, v') \in \mathcal{E}_i$. The *cascade frequency* f is at least as large as the probability that this occurs, and is therefore bounded from below:

$$f \geq \sum_k (1 - b_k^k) P_k^+ . \quad (5.47)$$

Given that the single seed triggers an edge in the giant in-cluster \mathcal{E}_i , how large will the resultant global cascade be? Well, certainly, the cascade will grow to the strongly-connected giant cluster \mathcal{E}_s , and continue to include all of the extended out-component \mathcal{E}_o of the giant cluster. From this point, higher order defaults become likely, so the cascade may grow much further. But, without restriction, we can say that when the cascade condition holds, whenever the giant vulnerable cluster is triggered, the resultant cascade will include all of \mathcal{E}_o . To compute the fractional size of this set, it is convenient to introduce the conditional probability

$$c_k = \mathbb{P}[v \notin \mathcal{E}_o | k_v = k] \quad (5.48)$$

where the event $v \notin \mathcal{E}_o$ is defined to mean that it has no in-edges ℓ that are in \mathcal{E}_o . For this calculation, we recognize that the event $\ell = (w, v) \notin \mathcal{E}_o$ means either $w \notin \mathcal{E}_o$ or $w \in \mathcal{E}_o$ and $\bar{\Omega}_{wv} < \bar{\Delta}_v$. The events $\ell = (w, v) \notin \mathcal{E}_o$ for all $w \in \mathcal{N}_v^-$ are mutually independent only when conditioned on the state of v and $\bar{\Delta}_v$, and conditioning carefully leads to the fixed point equation for the vector of probabilities $c = (c_k)$:

$$c_k = \sum_j P_{j|k} \int d\tilde{D}_{jk}(x) \left[\sum_{k'} Q_{k'|j} (W_{k'j}(x) + (1 - W_{k'j}(x)) c_{k'}) \right]^j . \quad (5.49)$$

Again, the cascade condition $\|\nabla G\| > 1$ is sufficient for the trivial fixed point $c = e$ to be unstable, meaning the size of the vulnerable out-component \mathcal{E}_o is a positive fraction of the network, computable by the formula

$$\mathbb{P}[v \in \mathcal{E}_o] = 1 - \sum_{jk} P_{jk} \int d\tilde{D}_{jk}(x) \left[\sum_{k'} Q_{k'|j} (W_{k'j}(x) + (1 - W_{k'j}(x)) c_{k'}) \right]^j \quad (5.50)$$

This is a lower bound on the size of the default cascade that results when the initial seed triggers a global cascade. It is interesting that the form of the expectation in (5.49) involves the point-wise power of the W distribution rather the convolution power that appears in (5.23).

5.6 Testing and Using the Cascade Mapping

The cascade mapping framework, both mathematical and conceptual, was developed for a larger purpose, namely to help understand the nature of real cascades in real world networks. So, how well does it work after all this effort?

While we have considered in this chapter only the simplest cascade mechanism, namely the zero-recovery default cascade of Gai-Kapadia, we have placed this mechanism on RFNs that have a rich complexity. Determining what our cascade mapping has to say about actual cascades is still a question of experimental computation. Since we have made a number of uncontrolled approximations, the analytical method should be validated by comparing with the results of Monte Carlo simulation experiments under a broad range of model specifications. Before continuing, it is helpful to consider the types of questions we wish to address.

Robustness of the formalism: The LTI formalism on RFNs, and the resultant cascade mapping results, rest on bold assumptions that somehow need to be checked. First, we can gain from the experience of others who have studied the Locally Tree-like ansatz for a great number of models on random networks. As a general rule, when benchmarked against Monte Carlo simulation experiments, the analytical formulas have been found to work effectively under a broad range of conditions. This is hard, slow work that must continue to push back the modelling frontiers as new models are introduced.

The framework developed in this chapter extends the original Gai-Kapadia paper in many respects. We have extended the skeleton construction to allow configuration graphs with general node and edge degree distributions. We have allowed buffers and exposures to be random with arbitrary marginal distributions. We have new formulas for the cascade condition and the economic impact of the cascade. The corresponding Monte Carlo validation experiments have not yet been completed, and will take considerably more time and effort.

Usefulness of the formalism: While a systematic Monte Carlo survey to validate the LTI analytic framework is pending, that will hopefully build confidence in the reliability of the framework, it is worthwhile to press forward using the analytical framework as a freestanding tool. Without the need for Monte Carlo, there are many purely analytical experiments on simple networks that are easy and instructive to implement on a computer. The scope of our default cascade formalism now includes flexibility in new dimensions never before explored. The skeletons are now assortative. Our balance sheets and exposures are now stochastic, and their variances are key parameters that represent our uncertainty in the network. These new features have unknown implications on the cascades that can result in simple models.

Learning about real cascades: As network data for financial systems grows and becomes available, network models will grow in complexity to incorporate new, more realistic features. Moreover, the problem of calibrating network models to

data will become increasingly challenging. Our framework has been designed to scale up in complexity to adapt to such needs.

Economic and financial implications: Analytic models can provide insight into purely economic problems. One important example is capital adequacy: Regulators want to know how a fixed amount of regulatory capital can be spread optimally across different banks to reduce systemic risk. The answer to such a question can be used to decide how much excess capital should be held by systemically important financial institutions (SIFIs). The behaviour of bankers is complex: they use evolving strategies and game theory continuously in time to navigate their firm through tortuous waters. Analytic models can be used by policy makers, regulators and market participants to test such strategies under stressful scenarios.

5.7 Default Cascades with Asymmetric Shocks

It turns out that the Watts model of Chapter 4 can be unified with the Gai-Kapadia zero recovery mechanism of the present chapter to give an economically natural default cascade model on a undirected network with bivariate link exposures that represent the unnetted positive exposures between counterparties.

The links in the Gai-Kapadia model are directed, and represent an idealized debtor-creditor relationship that is usually described in terms of unsecured overnight lending. The reality of banking is that counterparty relations are arbitrarily complex, and certainly cannot be boiled down to a single value at a moment in time. Counterparty banks will likely have overnight lending relations, will likely share a portfolio of interest rate and FX swap exposures, will owe each other certificates of deposit and the like, will trade repos, and so on. Determining the value of such exposures at any moment is exceedingly complicated: banks exert major efforts to compute their exposures to other banks at all times, following regulatory guidelines on counterparty credit risk. If at some moment one bank were to default, these exposures will often have opposite signs, and it is the positive part of the unnetted exposure that will impact the counterparty. To reduce counterparty risk, banks enter into master netting agreements (MNAs) with all of their important counterparties. This standardized agreement aims to specify exactly how the positive and negative exposures in the portfolio of different contracts between them can be offset in the event one bank defaults. An ideal fully netted position would leave only a single exposure in one direction at any time. It follows that the pervasiveness of MNAs in the interbank network provides a partial justification for taking edges to point in only one direction and for neglecting reflexive edges (those with links of both directions), as we have done in the Gai-Kapadia model. However, despite the existence of MNAs, allowing counterparty exposures to be unnetted or partially netted, and therefore bi-directional, is clearly a very natural modeling generalization.

In this section, we introduce a default cascade model that combines features of the Watts and Gai-Kapadia models while allowing edges to take on a more nuanced

meaning. It views the financial system as a network of banks connected by *undirected edges*, with edges are placed where there are deemed to be strong counterparty relations. For example, these edges should certainly include all pairs of banks that share a master netting agreement. It is known that building and maintaining counterparty relations is expensive for banks, particularly when the relationship is governed by the MNA. Thus it is reasonable to expect the network to be sparse, and that the existence of edges may be slowly varying while the exposures they carry might change quickly.

Given an edge $\ell = (w, v)$ between two banks v and w , the exposure \bar{Q}_ℓ carried by the edge will now be assumed to be multi-dimensional to represent the aggregated exposures across different types of securities. In the simplest variant we now consider, the multidimensional vector \bar{Q}_ℓ can be reduced to a pair of non-negative exposures $(\bar{Q}_{w,v}, \bar{Q}_{v,w})$. We interpret $\bar{Q}_{w,v}$ as the loss to v given the default of w . The model is then akin to the Watts model, but with asymmetric shocks that may be transmitted in either direction across edges. If $\min(\bar{Q}_{w,v}, \bar{Q}_{v,w}) = 0$ (i.e. only one direction is ever non-zero), this new setting reduces to a slightly non-standard specification of the Gai-Kapadia model.

With an aim to develop large N asymptotics, we now provide an LTI-compatible RFN specification for this model:

- Assumption 9.** 1. Banks have limited liability and receive zero recovery of inter-bank liabilities from any defaulted bank.
2. The skeleton consists of a random graph $(\mathcal{N}, \mathcal{E})$ on N banks which is an undirected assortative configuration model with node and edge type distributions $(P_k, Q_{kk'})$ with mean degree $z = \sum_k k P_k$ and satisfying the consistency conditions:

$$\begin{aligned} Q_{kk'} &= Q_{k'k}, \quad \sum_{k \leq k'} Q_{kk'} = 1 \\ Q_k &:= kP_k/z = \frac{1}{2} \sum_{k' \neq k} Q_{k'k} + Q_{kk} \\ Q_{k'|k} &:= (1 + \delta_{kk'}) \frac{Q_{kk'}}{2Q_k}. \end{aligned} \tag{5.51}$$

3. Conditionally on the skeleton, the default buffers $\bar{\Delta}_v$ are a collection of independent non-negative random variables whose distributions depend only on the degree k_v :

$$\mathbb{P}[\bar{\Delta}_v \leq x | v \in \mathcal{N}_k] := D_k(x), \quad k, x \geq 0$$

for cumulative probability distributions $D_k(\cdot)$ parametrized by k .

4. For each undirected link $\ell = (w, v) \in \mathcal{E}$, the exposure is a bivariate random variable $(\bar{Q}_{w,v}, \bar{Q}_{v,w})$ on \mathbb{R}_+^2 . Conditioned on the skeleton, the collection of edge exposures is independent with a bivariate distribution function that depends only on the bi-degree (k_w, k_v)

$$\mathbb{P}[\bar{Q}_{w,v} \leq x, \bar{Q}_{v,w} \leq y | w \in \mathcal{N}_k \cap \mathcal{N}_v, v \in \mathcal{N}_{k'}] := W_{kk'}(x, y) = W_{k'k}(y, x), \quad x, y \geq 0.$$

The conditional marginals are

$$W_{kk'}(x) := W_{kk'}(x, \infty) = W_{k'k}(\infty, x).$$

5. The remaining balance sheet quantities are freely specified.

Since $\bar{\Omega}_{v,w}$ represents the shock that will be transmitted from v to w at the moment v defaults, and $\bar{\Delta}_w \geq 0$ represents the threshold for the default of w , the set of defaulted banks \mathcal{D}_n and its indicator $D_v^n = \mathbf{1}(v \in \mathcal{D}_n)$ after n steps of the cascade again follows a recursion for all $n \geq 0$ and $v \in \mathcal{N}$:

$$D_v^n = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w \in \mathcal{N}_v} \bar{\Omega}_{w,v} D_w^{n-1} \right) \quad (5.52)$$

starting with $D_v^{-1} = 0$. Like the Watts model, we can show that this model has the WOR threshold property. First we define WOR default events for directed edges (v, w) by the recursion

$$D_{v,w}^n := \mathbf{1}(v \in \mathcal{D}_n \text{ WOR } w) = \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w' \in \mathcal{N}_v} \bar{\Omega}_{w',v} D_{w',v}^{n-1} \mathbf{1}(w' \neq w) \right) \quad (5.53)$$

starting with $D_{v,w}^{-1} = 0$ for all directed edges. Then, one finds that the default cascade has the WOR form because for all $n \geq 0$

$$D_v^n = \tilde{D}_v^n \quad (5.54)$$

where

$$\tilde{D}_v^{(n)} := \mathbf{1} \left(\bar{\Delta}_v \leq \sum_{w \in \mathcal{N}_v} \bar{\Omega}_{w,v} D_{w,v}^{n-1} \right). \quad (5.55)$$

The following proposition is analogous to Theorem 12 and can be justified by similar arguments:

Proposition 8. *Consider the LTI sequence of Watts-GK financial networks $(N, P, Q, \bar{\Delta}, \bar{\Omega})$ satisfying Assumptions 9. Let $\hat{p}_k^{(0)} := \mathbb{P}[w \in \mathcal{D}_0 \text{ WOR } v | w \in \mathcal{N}_v \cap \mathcal{N}_k] = D_k(0)$ denote initial WOR default probabilities. Then the following formulas hold in the limit as $N \rightarrow \infty$:*

1. *For any $n = 1, 2, \dots$, the quantities $\hat{p}_k^{(n)} = \mathbb{P}[w \in \mathcal{D}_n \text{ WOR } v | w \in \mathcal{N}_v \cap \mathcal{N}_k]$ satisfy the recursive formulas*

$$\hat{p}_k^{(n)} = \langle D_k, (\tilde{w}_k^{(n-1)})^{\otimes k-1} \rangle, \quad (5.56)$$

and the full default probabilities $p_k^{(n)} = \mathbb{P}[w \in \mathcal{D}_n | w \in \mathcal{N}_k]$ are given by

$$p_k^{(n)} = \langle D_k, (\tilde{w}_k^{(n-1)})^{\otimes k} \rangle. \quad (5.57)$$

Here the marginal exposure PDFs $\tilde{w}_k^{(n-1)}(x)$ are given by

$$\tilde{w}_k^{(n-1)}(x) = \sum_{k'} Q_{k'|k} \left((1 - \hat{p}_{k'}^{(n-1)}) \delta_0(x) + \hat{p}_k^{(n-1)} W_{k'k}^+(x) \right) \quad (5.58)$$

with $Q_{k'|k}$ defined by (5.51).

2. The new probabilities $\hat{p}^{(n)} = (\hat{p}_k^{(n)})$ are given recursively by $\hat{p}^{(n)} = G(\hat{p}^{(n-1)})$ for a vector valued function which is explicit in terms of the specification $(N, P, Q, \bar{\Delta}, \bar{\Omega})$. The cascade mapping G maps $[0, 1]^{\mathbb{Z}^+}$ onto itself, and is monotonic. Since $\hat{p}^{(0)} = G(0)$, the sequence $\hat{p}^{(n)}$ converges to the least fixed point $\hat{p}^* \in [0, 1]^{\mathbb{Z}^+}$, that is

$$\hat{p}^* = G(\hat{p}^*) . \quad (5.59)$$

Note that a consequence of the WOR property is that the asymptotic cascade mapping depends only on the collection of marginal distributions $W_{kk'}^+(x)$, and not the full bivariate distribution for $\bar{\Omega}_{wv}, \bar{\Omega}_{vw}$. In other words, without affecting the cascade probabilities, $\bar{\Omega}_{wv}, \bar{\Omega}_{vw}$ can be taken to be conditionally independent, so that

$$W_{kk'}(x, y) = W_{kk'}^+(x) W_{k'k}^+(y)$$

for an arbitrary collection $W_{kk'}^+(x)$ of univariate CDFs. Alternatively, we may assume $\bar{\Omega}_{wv} = \bar{\Omega}_{vw}$ for all $(w, v) \in \mathcal{E}$. Then the theorem reduces to the cascade mapping theorem proved in [44] for the Watts model on an assortative skeleton with random edge weight CDFs

$$\mathbb{P}[\bar{\Omega}_{w,v} \leq x | w \in \mathcal{N}_k, v \in \mathcal{N}_{k'}] = W_{kk'}(x), \quad x \geq 0 .$$

It is interesting that the more realistic meaning attached to exposures is consistent with a skeleton that is undirected instead of directed. One nice modeling feature is therefore that the node-degree k_v can be unambiguously correlated with the size of v 's balance sheet. In contrast, by focussing on the directionality of edges, the Gai-Kapadia model on RFNs is forced to live on a directed skeleton, whose bi-degrees (j_v, k_v) need to be specified, creating an additional complexity that now seems of less importance.

We have seen already in the Watts and Gai-Kapadia models that the cascade mapping function G determines essential features beyond the eventual default probabilities \hat{p}^* . We will not be surprised that it is the spectral norm of $\nabla G_{kk'} = \partial_{p_{k'}} G_k |_{p=0, \varepsilon=0}$ that determines whether the model admits global cascades or not. Or that estimates of the frequency and size of global cascades can be computed using percolation ideas. We leave such exercises to the interested reader.

5.8 Cascade Computations

While the version of the Gai-Kapadia cascade mapping given in Proposition 6 is straightforward to implement on a computer, the structure of the convolution power in (5.23) and similar equations at the heart of the generalized cascade mapping of Theorem 12 is problematic from the point of view of numerical approximations. Numerical evaluation of the implied integrals leads to truncation errors and discretization errors, both of which will be difficult to handle in our setting. In this section, we analyze the case where the random variables $\{\tilde{\Delta}_v, \tilde{\Omega}_\ell\}$ all take values in the finite discrete set

$$\mathcal{M} = \{0, \delta x, \dots, (M-1)\delta x\} \quad (5.60)$$

with a large value M and a common grid spacing δx . We can think of this as specifying both the truncation and discretization of non-negative continuous random variables. In such a situation, the convolutions in (5.23) can be performed exactly and efficiently by use of the discrete Fast Fourier Transform (FFT), whose detailed properties are summarized in Appendix A.2. For the moment, let us take $\delta x = 1$ for simplicity.

Let X, Y be two independent random variables with probability mass functions (PMF) p_X, p_Y taking values on the non-negative integers $\mathbb{Z}_+ := \{0, 1, 2, \dots\}$. Then the random variable $X + Y$ also takes values on this set and has the probability mass function (PMF) $p_{X+Y} = p_X \otimes p_Y$ where the convolution of two functions f, g is defined to be

$$(f \otimes g)(n) = \sum_{m=0}^n f(m)g(n-m), \quad n \in \mathbb{Z}_+. \quad (5.61)$$

Note that p_{X+Y} will not necessarily have support on the finite set \mathcal{M} if p_X, p_Y have support on \mathcal{M} , a fact that can lead to so-called *aliasing problems*.

We now consider a probability given as in (5.23)

$$P = \mathbb{P}[X \leq \sum_{i=1}^j Y_i]$$

for independent random variables X, Y_1, Y_2, \dots, Y_j where each Y_i is distributed with PMF $g_i(m)$ and X has PMF $f(m)$. We suppose that f has support on \mathcal{M} while each g_i has support on $\{0, 1, \dots, \lfloor M-1/j \rfloor\}$. We also define the CDF $F(n) = \sum_{m=0}^n f(m) = \mathbb{P}(X \leq n)$ for $n \in \mathcal{M}$. Then these M -vectors have M -dimensional FFTs $\hat{f} = \mathcal{F}(f)$, $\hat{g}_i = \mathcal{F}(g_i)$, $\hat{F} = \mathcal{F}(F)$. Using the FFT identities (A.1) to (A.4), we are led to a means to compute P for any j involving matrix operations and the FFT:

$$\begin{aligned} \mathbb{P}[X \leq \sum_{i=1}^j Y_i] &= \sum_{m \in \mathcal{M}} (g_1 \otimes \dots \otimes g_j)(m) \sum_{n=0}^m f(n) = \sum_{m \in \mathcal{M}} (g_1 \otimes \dots \otimes g_j)(m) F(m) \\ &= \langle F, g_1 \otimes \dots \otimes g_j \rangle = \frac{1}{M} \langle \hat{F}, \widehat{(g_1 \otimes \dots \otimes g_j)} \rangle = \frac{1}{M} (\hat{F})' \cdot (\hat{g}_1 \cdot * \dots * \hat{g}_j). \end{aligned} \quad (5.62)$$

Here in the final expression, A' denotes the conjugate transpose of a matrix A , \cdot denotes matrix multiplication between a size $[1, M]$ matrix and a size $[M, 1]$ matrix, and $\cdot*$ denotes component-wise (“Hadamard”) multiplication of vectors and matrices.

Remark 9. There is no aliasing problem and the identity (5.62) is true if $g_1 \otimes \cdots \otimes g_j$ has support on \mathcal{M} . We formalise this requirement as the Aliasing Assumption.

To summarize, as long as no aliasing errors arise, we can compute equation (5.23) efficiently using the FFT identity (5.62). Having realized this fact, it becomes sufficient to store as initial data only the Fourier transformed probability data for the random variables $\bar{\Delta}_v, \bar{Q}_\ell$, that is

$$\hat{D}_{jk} := \mathcal{F}(D_{jk}) \in \mathbb{C}^M, \quad \hat{w}_{kj} := \mathcal{F}(w_{kj}) \in \mathbb{C}^M, \quad j, k \in \mathbb{Z}_+. \quad (5.63)$$

Then the node update step becomes

$$p_{jk} = \frac{1}{M} \left\langle \hat{D}_{jk}, \widehat{(\tilde{w}_j)}^{\otimes j} \right\rangle \quad (5.64)$$

where

$$\widehat{(\tilde{w}_j)} = \sum_{k'} [1 - \pi_{k'} + \pi_{k'} \hat{w}_{k'j}] Q_{k'|j}.$$

Remark 10. The overall computation time for any numerical implementation of the cascade formulas in Theorem 12 will be dominated by computing convolutions such as those in (5.23). In particular, one can check that negligible total time will be taken in precomputing FFTs (each of which take of the order of $M \log_2 M$ additions and multiplications). One can compare the efficiency of our recommended FFT approach to the direct approach by considering a single evaluation of the convolution: to compute $f \otimes g$ using (5.61) for M -vectors f, g requires M^2 additions and multiplications, whereas to compute $\hat{f} \cdot * \hat{g}$ requires only M multiplications. All else being equal, we can expect a speedup by a factor of $O(M)$ when the FFT method is used instead of the direct method. Since in our implementations we often have $M \gtrsim 2^{10}$, this is a huge improvement. This speedup factor is too optimistic in practice when one takes care of the aliasing problem by taking a large, conservative value for M .

5.9 Numerical Experiments

The formalism developed in this chapter already raises a great variety of questions and unanswered issues concerning the GK model, which is of course the simplest of all default cascade models. Most such questions can only be studied with computer experiments. In this section, we report briefly on some of the simplest experiments that test the usefulness and validity of the LTI analytics developed so far in this chapter.

5.9.1 Experiment 1: Benchmark Gai-Kapadia Model

We choose as our benchmark the model specification given in the paper of [38], and develop variations on the theme of Gai and Kapadia that explore certain interesting dimensions away from this benchmark. Here is their original model specification:

1. The skeleton graph comprises $N = 1000$ banks (which we change to $N = 10000$) taken from the Poisson random directed graph model with mean in and out degree z , and thus $P = \text{Bin}(N, z/(N-1)) \times \text{Bin}(N, z/(N-1))$ and $Q = Q^+ Q^-$.
2. Capital buffers and assets are identical across banks, with $\bar{\Delta}_v = 4\%$ and $\bar{Z}_v = 20\%$.
3. Exposures are equal across the debtors of each bank, and so $\bar{\Omega}_{wv} = \frac{20}{j_v}$.
4. The number of Monte Carlo simulations performed was $N_{\text{sim}} = 1000$.

It is also interesting to use the invariance property of the cascade mapping to rescale exposures $\bar{\Omega}_{wv}$ and buffers $\bar{\Delta}_v$ by $j_v/20$, leading to (i) $\bar{\Delta}_v = j_v/5$; (ii) $\bar{Z}_v = j_v$; (iii) $\bar{\Omega}_{wv} = 1$. This rescaled specification is very similar to the benchmark parametrization used in the Watts model.

Figure 5.1(a) shows the dependence of the mean cascade size as a function of the mean degree z computed both analytically using Proposition 6 and by Monte Carlo simulation. The analytical results were obtained with seed probabilities $d_k(0) = 10^{-3}$ for all k and for the Monte Carlo simulation we took the initial seed to be 1% of the network (100 banks).

Figure 5.1(b) shows how the analytical formula (5.47) for the frequency of global cascades compares to the frequency of global Monte Carlo cascades that started from a single random seed, where in the Monte Carlo simulations, a global cascade was deemed to be one that exceeds 50 banks (0.5% of the network).

Comparison of these figures with the Watts model experiments in Chapter 4 exhibits striking similarities. Clearly this GK model specification is very similar to the benchmark Watts 2002 model implemented in [72]. In fact, in both models, nodes with in-degree j are vulnerable if and only if $j \leq 5$. Apart from differences stemming from the directed nature of the GK model, the cascade results shown in Figure 5.1 are almost identical to those found in Chapter 4. When z is smaller than 2 or 3, almost all nodes are vulnerable, and thus the contagion condition is essentially the condition for the giant cluster to exist, which suggests a phase transition at the value $z = 1$. On the other hand, when z is larger than about 7, most nodes are not vulnerable, and there is no giant vulnerable cluster. As z increases into this range it becomes harder for a single seed default to trigger a cascade. Occasionally, however a seed may have a large degree, opening the possibility for finite size effects causing higher order defaults that can lead to a large scale cascade. Although infrequent, usually such cascades appear to trigger almost 100% of the network to default.




Fig. 5.1 These two graphs show (a) the mean fractional cascade size, and (b), the fractional extended vulnerable cluster size (analytic) and global cascade frequency (Monte Carlo), in the benchmark Gai-Kapadia model as a function of mean degree z , as computed using the large N analytics of Theorem 11 (blue curve) and by Monte Carlo simulation (red crosses). In Figure (a), the Monte Carlo computation involved $N_{\text{sim}} = 50$ realizations of the $N = 10000$ node graph, each with an initial seed generated by selecting nodes independently with probability 0.5%. In Figure (b), the simulation involved $N_{\text{sim}} = 2500$ realizations of the $N = 10000$ node graph and a single random initial seed node.

5.9.2 Experiment 2: Assortative Networks

For various reasons, configuration graphs with general edge assortativity parametrized by the Q matrix have been little studied in network science. On the other hand, numerous statistical studies of financial networks in the real world, notably [68], [10] and [27], have pointed out their strongly negative assortativity and speculated that this property is likely to strongly influence the strength of cascades that the network will exhibit. In this experiment, we explore the GK model close to the benchmark specification but with different parametrizations of the edge type matrix Q . Addressing this question provides us with a good starting point to test a number of points. How effective is the simulation algorithm for assortative configuration graphs given in Section 3.5? Does edge assortativity have a strong effect on the possibility of cascades?

We twist the benchmark GK model described above by replacing the independent edge probability matrix $Q_{kj} = Q_k^+ Q_j^-$ by an interpolated matrix that exhibits positive or negative assortativity

$$Q(\alpha)_{kj} = \begin{cases} \alpha Q_{kj}^+ + (1 - \alpha) Q_k^+ Q_j^- & \alpha \geq 0 \\ |\alpha| Q_{kj}^- + (1 - |\alpha|) Q_k^+ Q_j^- & \alpha < 0 \end{cases}$$

Here $Q_{kj}^+ = \delta_{kj}Q_k^+$ has 100% correlation and is the maximally assortative matrix with the prescribed marginals. On the other hand, the maximally negative assortative matrix turns out to be given by

$$Q_{kj}^- = \mathbb{P}\left[U \in (b_{k-1}, b_k], 1 - U \in (b_{j-1}, b_j]\right]$$

where U is uniform $[0, 1]$ and $b_k = \sum_{k' \leq k} Q_{k'}^+$ for $k = 0, 1, \dots, K_{max}$.

With some surprise we noticed that with P_{jk} given as in Experiment 1, the mean cascade size had no variation with α ! After a moment's thought, however, we realize that because this model has an independent bivariate distribution of node degrees, any possible dependence on α is wiped out. We need some dependence between in and out node degrees to allow for the cascade to depend on assortativity. Figure 5.2 shows the dependence of the mean cascade size on the parameter α for four values of z , when we replace the independent matrix P by the fully correlated matrix

$$P_{jk} = P_k^+ \delta_{jk}.$$

As we hoped, we see a strong nonmonotone dependence of cascade frequency on the assortativity parametrized by α , with a maximum effect for a positive assortativity value $\alpha \sim 0.4$.

Throughout this experiment, we need to pay attention to discrepancies between the Monte Carlo results and the analytics. While for the most part the match is quite good, surprisingly it breaks down completely for $z = 7.5$ and negative assortativity where the Monte Carlo results are particularly erratic. One important “rule of thumb” that can partly account for such anomalies is that they seem likely to occur in the neighbourhood of discontinuous phase transitions, in this case the one near $z \sim 7$. It seems likely we will see even larger discrepancies when the finite Monte Carlo samples are taken from a less heterogeneous model.

This experiment also shows that Monte Carlo simulation can become quite challenging as the random graph model becomes more complex, and provides motivation to explore widely using large N analytics which are relatively straightforward to implement.

5.9.3 Experiment 3: Random Buffers and Link Weights

What is the effect of uncertainty in the network? In reality, even with excellent databases we might never expect to know actual exposures and buffers with any precision. We can model this uncertainty by taking these to be random variables and testing how the variance parameters affect the resultant cascade sizes and frequencies. Furthermore, we can check whether the analytic approximations get much worse or not.

To test the impact of this idea, we can introduce collections of log normally distributed buffer and exposure random variables with four additional parameters




Fig. 5.2 These two graphs show (a) the mean cascade size, and (b) the fractional extended vulnerable cluster size (analytic) and global cascade frequency (Monte Carlo), in the GK model, as a function of positive and negative assortativity when P is maximally correlated. The analytical results are shown by the solid curves, with four values of z : 1.5 (red), 3.5 (blue), 5.5 (green), 7.5 (black). In Figure (a), the Monte Carlo computation involved $N_{\text{sim}} = 50$ realizations of the $N = 10000$ node graph, each with an initial seed generated by selecting nodes independently with probability 0.5%. In Figure (b), the analytical formula for the extended giant vulnerable cluster is compared to the result of an $N_{\text{sim}} \times N_{\text{freq}}$ Monte Carlo simulation of global cascades consisting of $N_{\text{sim}} = 25$ realizations of the $N = 10000$ node graph, and for each graph realization $N_{\text{freq}} = 2000$ random initial seed nodes.

$\sigma_{\Delta}, \sigma_{\Omega} \geq 0$ and $\rho_{\Delta}, \rho_{\Omega} \in (-1, 1)$. That is,

$$\Delta_v = \frac{j_v}{5} \exp[\sigma_{\Delta}(\rho_{\Delta}Z_{\Delta} + \sqrt{1 - \rho_{\Delta}^2}X_v) + \sigma_{\Delta}^2/2] \quad (5.65)$$

$$\Omega_{\ell} = \exp[\sigma_{\Omega}(\rho_{\Omega}Z_{\Omega} + \sqrt{1 - \rho_{\Omega}^2}Y_{vw}) + \sigma_{\Omega}^2/2] \quad (5.66)$$

with a collection $\{Z_{\Delta}, X_v, Z_{\Omega}, Y_{\ell}\}$, $v \in \mathcal{N}$, $\ell \in \mathcal{E}$ of independent standard normals. Note that the benchmark Gai-Kapadia model of Experiment 1 arises by taking $\sigma_{\Delta} = \sigma_{\Omega} = 0$, and for all parameters, the new specification agrees with the benchmark in expectation. The case of $\sigma_{\Omega} = 0$ and varying $\sigma_{\Delta}, \rho_{\Delta}$ has been studied extensively elsewhere and so is of lesser interest. Also, by the rescaling property, the cascade mapping when $\rho_{\Omega} \neq 0$ can be transformed into the case $\rho_{\Omega} = 0$. However, the case of random exposures has not been studied previously, and therefore we take $\sigma_{\Delta} = 0, \rho_{\Omega} = 0$ and focus on the dependence of the default cascade on the shape parameter σ_{Ω} .

In our implementation of equations (5.23) and (5.24) of Theorem 12 we truncated and discretized the log normal exposure random variables by placing their values on a grid of $M = 2^{12}$ points with spacing $\delta x = 0.1$. Then we applied the FFT algorithm for computing the cascade equilibrium of Theorem 12. Prior to plotting our results

we confirmed that they were insensitive to this approximation by comparing to the results when the grid was coarsened to $M = 2^{10}$ and $\delta x = 0.2$ (results not shown).

Figure 5.3 plots the dependence of the mean cascade size and cascade frequency as a function of the shape parameter of the exposure distribution, σ_Ω , for four values of the Poisson mean degree z : 1.5 (red), 3.5 (blue), 5.5 (green), 7.5 (black).




Fig. 5.3 These two graphs show (a) the mean cascade size, and (b) the fractional extended vulnerable cluster size (analytic) and global cascade frequency (Monte Carlo), in the GK model, as a function of σ_Ω for $\sigma_\Delta = 0$. The analytical results are shown by the solid curves, with four values of z : 1.5 (red), 3.5 (blue), 5.5 (green), 7.5 (black). In Figure (a), the Monte Carlo computation involved $N_{\text{sim}} = 10$ realizations of the $N = 10000$ node graph, each with an initial seed generated by selecting nodes independently with probability 0.5%. In Figure (b), the analytical formula for the extended giant vulnerable cluster is compared to the result of Monte Carlo simulation of global cascades consisting of $N_{\text{sim}} = 10$ realizations of the $N = 10000$ node graph, and for each graph realization averaging over $N_{\text{freq}} = 1000$ random initial seed nodes.

Again, we see a general agreement between the Monte Carlo and analytical results, but with some significant departures at $z = 1.5$. Further exploration shows that the lower phase transition at $z = 1$ is pushed to higher values of z as σ_Ω grows: it appears that the anomalous behaviour we observed is connected with this fact.

These three experiments illustrate only a very limited set of possible analytical experiments on the simplest of all default cascade models. Already we observe that while the analytical formulas agree with the Monte Carlo simulation results for the most part, there are always circumstances that undermine the accuracy of the agreement. It is important to try to develop “rules of thumb” that supply the intuition of the circumstances when the agreement is not acceptable.

Chapter 6

Future Directions for Cascade Models

Abstract The prospects are considered for extending the mathematical framework of cascade mechanisms on locally tree-like random financial networks to address problems of real financial importance.

Keywords: Random financial networks, local tree-like independence, cascade mechanism, balance sheet models, bank behaviour.

This book has set out to show that systemic risk and contagion in financial networks have a rich underlying mathematical structure: contagion shocks spread out in a way that can be understood in terms of percolation theory. Along the way, we have explored a variety of combinations of the three acronyms, a RFN (random financial network) with the LTI property (locally tree-like independence) on which we place a CM (cascade mechanism), to create contagion models of insolvency, illiquidity and fire sales. Towards the end of the book, we developed analytical methods in the context of default contagion, but which extend to other types of contagion as well. These models exhibit key stylized features of market contagion that economists have identified in financial crises from the past. In some circumstances, these ingredients fit together in a natural way that retains mathematical properties of percolation analysis as it has been developed in probability theory, leading to explicit cascade mappings and a host of associated concepts.

The concept of RFN, or random financial network, was introduced as a general stochastic setting appropriate either for hypothetical banking systems, or for simplifying the description of real-life banking systems that are sufficiently large and homogeneous. It was taken to comprise a skeleton of nodes representing banks connected by edges representing counterparty links, together with a stochastic specification of bank balance sheets and interbank exposures. Typical cascade models on RFNs can be computed using Monte Carlo simulation. What we have shown is that these models sometimes also admit an explicit analytical formulation from which we can understand much more. Ultimately, RFNs can be created with ever more complex combinations of random skeletons, balance sheets and cascade mechanisms.

LTI combines an assumption about the stochastic properties of the underlying skeleton related to sparsity of the edges together with an assumption about the multivariate dependence of balance sheet random variables. Roughly, it means that the dependence structure of the entire network is coded into the random skeleton: the multivariate distributions of remaining balance random variables are built with independence conditionally on the skeleton. Since there are a variety of approaches to modelling skeletons available in the random graph literature, such a framework provides a useful compromise between tractability, parsimony and flexibility.

A number of cascade mechanisms have been proposed in the literature, notably the model of Eisenberg and Noe [32]. Typically they amount to assumptions or behavioural rules that determine how banks adjust their balance sheets in response to the evolution of the financial crisis. We have reviewed in Chapter 2 different CMs describing a variety of effects such as liquidity shocks, bank defaults, and forced asset sales. The key observation was that these CMs tend to have a common threshold structure. Noticing this commonality is helpful in developing intuition about how different cascades will behave. Of course, such CMs provide only a caricature of the complex decision making of banks and must certainly develop in sophistication in future research, perhaps by following the theory of global games developed for example in [61].

What then are the advantages of having systemic risk models with such ingredients in which crises can be viewed as cascade mappings that lead to fixed points that describe new network equilibria? Typically, this structure leads to a number of advantageous features. Analytics in such models can often be easier to program and faster to compute than Monte Carlo simulations, facilitating exploration of sensitivities to key structural parameters. Sometimes simple specifications of such models can be directly understood by comparison with previously studied models. As well we have seen that certain systemic risk measures are computable directly from the fixed point.

Simple static cascade models are often criticized because they exclude important institutional elements, most notably *central clearing parties (CCP)*, *central banks* and regulators. However, principles of model building in science suggest that it is important to understand the behaviour of uncontrolled isolated systems before trying to understand how to control them. Cascade mappings provide a level of understanding about the uncontrolled network that will guide the actions of regulators and governments. For a concrete example, it is now fully recognized in Basel III that systemically important financial institutions (SIFIs) must be subjected to regulatory capital surcharges. How such punitive measures are implemented should be defensible by economic and scientific principles. With a reliable cascade model, one can in principle provide such a rationale by finding the optimal distribution of regulatory capital amounts across the network, subject to a fixed aggregated budget.

The three elements, RFNs with an LTI specification upon which a CM is proposed, are intended to be a template for further explorations of systemic risk in ever more complex settings. So what are some of the most promising avenues to follow?

First, RFNs can be far from as simple as those investigated here. Nodes might take on richer characteristics: in our language, the *type* of the node will no longer

mean the node degree, but may come to include information about the geographical location, the style of banking undertaken, the size of the institution and so on. For this, skeletons modelled by inhomogeneous random graphs generalizing the construction of Section 3.4 will be useful. Similarly, the network itself will expand to include insurance firms, hedge funds, mutual funds, pension funds, asset classes and even corporates. See [17] for an agent-based model of such an extended network. Describing *community structure* is an active topic of network science that will directly impact systemic risk research. Similarly, the meaning of edges will become more nuanced, and depend in subtle ways on the types of the adjoined banks. Skeletons will evolve into hypergraphs to account for important classes of tri-party contracts such as credit default swaps, repos and asset return swaps.

The balance sheets of banks in this book all have a simple stylistic structure. [40] has shown how more complex debt seniority structures can be included in the cascade picture. Similar ideas applied to the asset side will allow rich structures that account for features of different asset classes such as credit rating, maturity and liquidity properties.

The structure of interbank contracts described here is similarly stylistic: for the most part, these have been described as short term unsecured lending. In reality, such contracts are complex, and have structure and magnitude that are the result of bilateral strategic games played continuously between pairs of counterparty banks who seek to control risk and maximize returns.

This book has focussed on systemic risk, that is to say, the risk of large scale or even catastrophic disruptions of the financial markets. But good models of the financial system must account for the behaviour of markets in normal times as well. This is already recognized in macroprudential models used by central banks that link together modules describing parts of the economy, such as the RAMSI model used by the Bank of England and the MFRAF model used by the Bank of Canada. One important module in such economic models is always the financial network submodel that should behave realistically in normal times, and should generate endogenously an adequate spectrum of crisis scenarios.

Systemic risk modelling has advanced enormously on a range of fronts since the great financial crisis of 2007-08. It is my hope that this book will be viewed as a timely contribution to the field that provides researchers with two different things: first, a flexible set of mathematical techniques for analyzing cascades in general categories of networks, and second, a scientific modelling framework of the financial economy that can scale up to account for all the important dimensions of a type of risk that has the potential to critically impact the entire global economy.

Appendix A

Background Material

A.1 Special Notation

Matrix and Vector Notation: For vectors $x = [x_v]_{v=1,\dots,N}, y = [y_v]_{v=1,\dots,N} \in \mathbb{R}^N$ define relations

$$\begin{aligned}
 x \leq y & \text{ means } \forall v, x_v \leq y_v, \\
 x < y & \text{ means } x \leq y, \exists v : x_v < y_v, \\
 \min(x, y) & = [\min(x_v, y_v)]_{v=1,\dots,N} \\
 \max(x, y) & = [\max(x_v, y_v)]_{v=1,\dots,N} \\
 (x)^+ & = \max(x, 0), \\
 (x)^- & = \max(-x, 0).
 \end{aligned}$$

Whenever $x \leq y$ we can define the hyperinterval $[x, y] = \{z : x \leq z \leq y\}$. Column N -vectors are often treated as $[1, N]$ matrices, $X \cdot Y$ denotes the matrix product, and X' denotes the Hermitian (complex conjugate) transpose.

Function Notation: For functions f, g , $f \circ g$ denotes composition and $f \otimes g$ denotes convolution.

Graph Notation: For a complete description of the graphical notation used in this book, please see Section 3.1.

Set Notation and Indicators: For any set $A \subset \Omega$ in the sample space Ω , $A^c := \Omega \setminus A$ is its complement. For elements $x \in \Omega$ and Boolean propositions P :

$$\begin{aligned}
 \text{Indicator } \mathbf{1}(P) & = \begin{cases} 1 & \text{if } P \text{ is true} \\ 0 & \text{if } P \text{ is false} \end{cases} \\
 \text{Indicator function } \mathbf{1}_A(x) & \text{ means } \mathbf{1}(x \in A).
 \end{aligned}$$

A.2 The Discrete Fourier Transform

We consider the space \mathbb{C}^M of \mathbb{C} -valued functions on $\mathcal{M} = \{0, 1, \dots, M-1\}$. The discrete Fourier transform, or *fast Fourier transform (FFT)*, is the linear mapping $\mathcal{F} : a = [a_0, \dots, a_{M-1}] \in \mathbb{C}^M \rightarrow \hat{a} = \mathcal{F}(a) \in \mathbb{C}^M$ defined by

$$\hat{a}_k = \sum_{l \in \mathcal{M}} \zeta_{kl} a_l, k \in \mathcal{M}.$$

where the coefficient matrix $Z = (\zeta_{kl})$ has entries $\zeta_{kl} = e^{-2\pi i kl/M}$.

It is easy to prove that its inverse, the *inverse FFT (IFFT)*, is given by the map $a \rightarrow \tilde{a} = \mathcal{G}(a)$ where

$$\tilde{a}_k = \frac{1}{M} \sum_{l \in \mathcal{M}} \bar{\zeta}_{kl} a_l, k \in \mathcal{M}.$$

If we let \bar{a} denote the complex conjugate of a , we can define the Hermitian inner product $\langle a, b \rangle := \sum_{m \in \mathcal{M}} \bar{a}_m b_m$. We also define the convolution product of two vectors:

$$(a \otimes b)(n) = \sum_{m \in \mathcal{M}} a(m) b(n-m \text{ modulo } M), \quad n \in \mathcal{M}.$$

Now we note the following easy-to-prove identities which hold for all $a, b \in \mathbb{C}^M$:

1. Inverse mappings:

$$a = \mathcal{G}(\mathcal{F}(a)) = \mathcal{F}(\mathcal{G}(a)); \quad (\text{A.1})$$

2. Conjugation:

$$\overline{\mathcal{G}(a)} = \frac{1}{M} \mathcal{F}(\bar{a}); \quad (\text{A.2})$$

3. Parseval Identity:

$$\langle a, b \rangle = M \langle \tilde{a}, \tilde{b} \rangle = \frac{1}{M} \langle \hat{a}, \hat{b} \rangle; \quad (\text{A.3})$$

4. Convolution Identities:

$$\tilde{a} \cdot \tilde{b} = \widetilde{(a \otimes b)}; \quad \hat{a} \cdot \hat{b} = \widehat{(a \otimes b)} \quad (\text{A.4})$$

where \cdot denotes the component-wise product.

The primary application of the FFT in this book is its use in accelerating the computation of convolutions of probability mass functions supported on the set CM . If the support of the sum of two \mathcal{M} -valued random variables X, Y is itself in \mathcal{M} , that is, if $\text{supp}(X+Y) = \{n | \exists m \in \text{supp}(X) \text{ s.t. } n-m \in \text{supp}(Y)\} \in \mathcal{M}$, then p_{X+Y} , the PMF of $X+Y$, is given by $p_X \otimes p_Y$. On the other hand, if there are $m \in \text{supp}(X), n \in \text{supp}(Y)$ such that $m+n \geq M$, then $p_{X+Y} - p_X \otimes p_Y$ is not zero. Such a difference is called an “aliasing error”, and implies that $p_{X+Y} \neq \mathcal{G}(\mathcal{F}(p_X)\mathcal{F}(p_Y))$. In any application, we must take care to keep all such aliasing errors sufficiently small by choosing M sufficiently large but finite.

References

- [1] Viral V. Acharya, Lasse H. Pedersen, Thomas Philippon, and Matthew P. Richardson. Measuring Systemic Risk. *SSRN eLibrary*, 2010.
- [2] Anat Admati and Martin Hellwig. *The Bankers' New Clothes*. Princeton University Press, 2014.
- [3] Tobias Adrian and Markus K. Brunnermeier. Covar. Working Paper 17454, National Bureau of Economic Research, October 2011.
- [4] Tobias Adrian and Hyun S. Shin. Liquidity and Leverage. *Journal of Financial Intermediation*, 19(3):418–437, July 2010.
- [5] Hamed Amini, Rama Cont, and Andreea Minca. Resilience to contagion in financial networks. *Mathematical Finance*, pages 1–37, 2013.
- [6] Kartik Anand, Prasanna Gai, and Matteo Marsili. Financial crises and the evaporation of trust. *arXiv preprint arXiv:0911.3099*, 2009.
- [7] Per Bak, Chao Tang, and Kurt Wiesenfeld. Self-organized criticality: An explanation of the $1/f$ noise. *Phys. Rev. Lett.*, 59:381–384, Jul 1987.
- [8] József Balogh and Boris G. Pittel. Bootstrap percolation on the random regular graph. *Random Structures & Algorithms*, 30(1-2):257–286, 2007.
- [9] Albert-Laszlo Barabasi and Reka Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, 1999.
- [10] M. Bech and E. Atalay. The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22):5223–5246, 2010.
- [11] Dimitrios Bisias, Mark D. Flood, Andrew W. Lo, and Stavros Valavanis. A survey of systemic risk analytics. Technical Report 1, U.S. Department of Treasury, Office of Financial Research, 01 2012.
- [12] Marián Boguñá and M. Ángeles Serrano. Generalized percolation in random directed networks. *Phys. Rev. E*, 72:016106, 2005.
- [13] B. Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *Eur. J. Comb.*, 1:311, 1980.
- [14] B. Bollobás. *Random Graphs*. Cambridge studies in advanced mathematics. Cambridge University Press, 2 edition, 2001.
- [15] Béla Bollobás, Christian Borgs, Jennifer Chayes, and Oliver Riordan. Directed scale-free graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Sym-*

- posium on Discrete Algorithms*, SODA '03, pages 132–139, Philadelphia, PA, USA, 2003. Society for Industrial and Applied Mathematics.
- [16] Béla Bollobás, Svante Janson, and Oliver Riordan. The phase transition in inhomogeneous random graphs. *Random Struct. Algorithms*, 31(1):3–122, August 2007.
- [17] Richard Bookstaber. Using agent-based models for analyzing threats to financial stability. Working Paper Series 3, Office of Financial Research, December 2012.
- [18] Tom Britton, Maria Deijfen, and Anders Martin-Löf. Generating simple random graphs with prescribed degree distribution. *Journal of Statistical Physics*, 124(6):1377–1397, 2006.
- [19] Andrei Broder, Ravi Kumar, Farzin Maghoul, Prabhakar Raghavan, Sridhar Rajagopalan, Raymie Stata, Andrew Tomkins, and Janet Wiener. Graph structure in the web. *Comput. Netw.*, 33(1-6):309–320, June 2000.
- [20] Markus K. Brunnermeier and Lasse Heje Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- [21] Fabio Caccioli, Munik Shrestha, Cristopher Moore, and J. Doyne Farmer. Stability analysis of financial contagion due to overlapping portfolios. arXiv:1210.5987v1 [q-fin.GN], October 2012.
- [22] J. Chalupa, P. L. Leath, and G. R. Reich. Bootstrap percolation on a Bethe lattice. *Journal of Physics C: Solid State Physics*, 12(1):L31+, 1979.
- [23] Ningyuan Chen and Mariana Olvera-Cravioto. Directed random graphs with given degree distributions. *Stochastic Systems*, 3(1):147–186, 2013.
- [24] Fan Chung and Linyuan Lu. Connected components in random graphs with given expected degree sequences. *Annals of Combinatorics*, 6(2):125–145, 2002.
- [25] R. Cifuentes, G. Ferrucci, and H. S. Shin. Liquidity risk and contagion. *Journal of the European Economic Association*, 5:556–566, 2005.
- [26] A. Clauset, C. Shalizi, and M. Newman. Power-law distributions in empirical data. *SIAM Review*, 51(4):661–703, 2009.
- [27] Rama Cont, Amal Moussa, and Edson B. Santos. Network Structure and Systemic Risk in Banking Systems. In Jean-Pierre Fouque and Joseph A. Langsam, editors, *Handbook on Systemic Risk*, chapter 13, pages 327–368. Cambridge University Press, 2013.
- [28] Stephane Crepey, T. R. Bielecki, and D. Brigo. *Counterparty Risk and Funding: A Tale of Two Puzzles*. Chapman & Hall/CRC, Boca Raton, FL, 2014.
- [29] O. De Bandt, P. Hartmann, and J.L. Peydró. Systemic risk in banking: an update. In A.N. Berger, P. Molyneux, and et al. Wilson, J., editors, *Oxford Handbook of Banking*. Oxford University Press, Oxford, 2009.
- [30] Gabrielle Demange. Contagion in financial networks: a threat index. PSE Working Papers n2012-02 ;halshs-00662513_{v1}, 2012.
- [31] D. Duffie and K. Singleton. *Credit risk: pricing, measurement and management*. Princeton University Press, 2003.
- [32] L. Eisenberg and T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.

- [33] P. Erdős and A. Rényi. On random graphs. *I. Publ. Math. Debrecen*, 6:290–297, 1959.
- [34] Bank for International Settlements. *64th Annual Report*. Bank for International Settlements, Basel, Switzerland, 1994.
- [35] Craig H. Furfine. Interbank exposures: Quantifying the risk of contagion. *Journal of Money, Credit and Banking*, 35(1):111–128, 2003.
- [36] P. Gai, A. Haldane, and S. Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58:453–470, 2011.
- [37] P. Gai and S. Kapadia. Liquidity hoarding, network externalities, and interbank market collapse. *Proc. R. Soc. A*, 466:2401–2423, 2010.
- [38] Prasanna Gai and Sujit Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A*, 466(2120):2401–2423, 2010.
- [39] Gary Gorton and Andrew Metrick. Haircuts. *Review*, November:507–520, 2010.
- [40] C. Gourieroux, J. C. Heam, and A. Monfort. Bilateral exposures and systemic solvency risk. *Canadian Journal of Economics*, 45:1273–1309, 2012.
- [41] A. G. Haldane. Rethinking the financial network, April 2009. Speech delivered at the Financial Student Association, Amsterdam.
- [42] A. G. Haldane. The dog and the frisbee, August 2012. Speech Given at the Federal Reserve Bank of Kansas City’s 36th economic policy symposium, “The Changing Policy Landscape”, Jackson Hole, Wyoming.
- [43] T. R. Hurd and James P. Gleeson. A framework for analyzing contagion in banking networks. arXiv:1110.4312 [q-fin.GN], October 2011.
- [44] T. R. Hurd and James P. Gleeson. On Watts cascade model with random link weights. *Journal of Complex Networks*, 1(1):25–43, 2013.
- [45] Svante Janson. On percolation in random graphs with given vertex degrees. arXiv:0804.1656, 2008.
- [46] Svante Janson. The probability that a random multigraph is simple. *Combinatorics, Probability and Computing*, 18:205–225, 3 2009.
- [47] Svante Janson and Malwina J. Luczak. A new approach to the giant component problem. *Random Struct. Algorithms*, 34(2):197–216, March 2009.
- [48] Óscar Jordá, Moritz Schularick, and Alan M. Taylor. Betting the house. Working Paper Series 28, Federal Reserve Bank of San Francisco, December 2014.
- [49] George G. Kaufman. Comment on systemic risk. In George Kaufman, editor, *Research in Financial Services: Banking, Financial Markets, and Systemic Risk*, volume 7, pages 47–52. Greenwich, CT: JAI Press, 1995.
- [50] George G. Kaufman and Kenneth E. Scott. What is systemic risk, and do bank regulators retard or contribute to it? *Independent Review*, 7(3):371–91, 2003.
- [51] Seung Hwan Lee. Systemic liquidity shortages and interbank network structures. *Journal of Financial Stability*, 9(1):1–12, 2013.
- [52] Nelly Litvak and Remco van der Hofstad. Uncovering disassortativity in large scale-free networks. *Phys. Rev. E*, 87:022801, Feb 2013.
- [53] Robert Jr Lucas. Econometric policy evaluation: A critique. *Carnegie-Rochester Conference Series on Public Policy*, 1(1):19–46, January 1976.

- [54] Robert M. May and Nimalan Arinaminpathy. Systemic risk: the dynamics of model banking systems. *Journal of The Royal Society Interface*, 7(46):823–838, 2010.
- [55] Sergey Melnik, Adam Hackett, Mason A. Porter, Peter J. Mucha, and James P. Gleeson. The unreasonable effectiveness of tree-based theory for networks with clustering. *Phys. Rev. E*, 83:036112–036123, 2011.
- [56] H. P. Minsky. *John Maynard Keynes*. New York, NY: Columbia University Press, 1975.
- [57] Frederic Mishkin. Comment on systemic risk. In George Kaufman, editor, *Research in Financial Services: Banking, Financial Markets, and Systemic Risk*, volume 7, pages 31–45. Greenwich, CT: JAI Press, 1995.
- [58] Michael Molloy and Bruce Reed. A critical point for random graphs with a given degree sequence. *Random Structures & Algorithms*, 6(2-3):161–180, 1995.
- [59] Michael Molloy and Bruce Reed. The size of the largest component of a random graph on a fixed degree sequence. *Combin. Probab. Comput.*, 7:295–306, 2008.
- [60] S. Morris and H. S. Shin. Illiquidity component of credit risk. Working paper, 2009.
- [61] Stephen Morris and Hyun Song Shin. Global games: Theory and applications. In L. Hansen M. Dewatripont and S. Turnovsky, editors, *Advances in Economics and Econometrics*. Cambridge University Press, 2003.
- [62] Erlend Nier, Jing Yang, Tanju Yorulmazer, and Amadeo Alentorn. Network Models and Financial Stability. *J. Econ. Dyn. Control*, 31:2033–2060, 2007.
- [63] Board of Governors of the Federal Reserve System. Policy statement on payments system risk. Docket No. R-1107, 1–13. Washington, D.C., May 2001.
- [64] Mohammad Farzad Pourbabae. Assortative configuration random graphs. Mathematics m.sc. thesis, McMaster University, May 2015.
- [65] Steven L. Schwarcz. Systemic risk. *Georgetown Law Journal*, 97(1), 2008. Duke Law School Legal Studies Paper No. 163, Available at SSRN: <http://ssrn.com/abstract=1008326>.
- [66] H. S. Shin. Securitisation and financial stability. *Economic Journal*, 119:309–332, 2009.
- [67] Steven Shreve. *Stochastic Calculus for Finance II: Continuous-Time Models*. Springer Verlag, Berlin Heidelberg New York, 2004.
- [68] Kimmo Soramäki, M. Bech, J. Arnold, R. Glass, and W. Beyeler. The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications*, 379(1):317–333, 2007.
- [69] John B. Taylor. Defining systemic risk operationally. In Kenneth E. Scott, George P. Shultz, and John B. Taylor, editors, *Ending Government Bailouts As We Know Them*, chapter 4. Hoover Institution, Stanford University, 2010.
- [70] Christian Upper. Simulation methods to assess the danger of contagion in interbank markets. *J. Financial Stability*, 2011.
- [71] R. van der Hofstad. Random graphs and complex networks. Book, to be published, 2014.

- [72] Duncan J. Watts. A simple model of global cascades on random networks. *PNAS*, 99(9):5766–5771, 2002.
- [73] N. C. Wormald. The differential equation method for random network processes and greedy algorithms. In *Lectures on Approximation and Randomized Algorithms*, pages 73–155. Citeseer, 1999.
- [74] Huqiu Zhang and Aad Moorsel. Fast generation of scale free networks with directed arcs. In *Proceedings of the 6th European Performance Engineering Workshop on Computer Performance Engineering*, EPEW '09, pages 131–148, Berlin, Heidelberg, 2009. Springer-Verlag.

Index

- adjacency matrix, 52
- asset fire sale, 18
- asset fire sales, 42
- assortative configuration graph model, 68
- assortativity, 57
- average shortfall, 119

- balance sheets, 49
- bank book, 19
- betweenness centrality, 81
- book values, 27
- bootstrap percolation, 91
- branching process, 86

- capital adequacy ratio, 22, 42
- cascade condition, 98, 121
- cascade equilibrium, 26
- cascade frequency, 124
- cascade impact, 37
- cascade mechanism, 26
- central bank, 138
- central clearing party, 19
- central clearing party (CCP), 138
- centrality measure, 80
- clipping, 56
- clustering coefficient, 81
- community structure, 139
- complex adapted system, 51
- configuration, 73
- connectivity, 79
- contingent convertible bonds, 21
- convergence in probability, 55
- CoVaR, 119
- credit default swap, 20
- critical, 87

- default buffer, 34

- default cascade, 17, 25
- default contagion, 17
- default correlation, 37, 119
- degree centrality, 80
- disassortative, 68

- early adopters, 93
- edge, 52
- edge type, 52
- edge-type distribution, 54
- eigenvalue centrality, 80
- equity, 22
- exposures, 49
- external assets, 19
- extinction probability, 87

- face values, 27
- fast Fourier transform, 142
- fixed-mix strategy, 46
- funding illiquidity, 17

- G-SIB, 23
- Galton-Watson process, 86
- generalized random graph, 62
- giant cluster, 85
- giant component, 85
- giant strongly connected component, 82
- global cascade, 91
- global systemically important bank, 23
- graph
 - directed graph, 52
 - undirected graph, 52

- haircut, 21
- hard default, 36
- hub, 80
- hybrid capital, 21

- in-component, 82
- in-degree, 52
- in-neighbourhood, 53
- in-stub, 56, 69
- in-subgraph, 32
- independent edge condition, 57
- inhomogeneous random graph, 62
- interbank assets, 19
- inverse FFT, 142
- irreducible in-subgraph, 32
- irreversible project, 19

- Katz centrality, 81

- limited liability, 22
- link, 52
- liquidity cascade, 26
- liquidity contagion, 17
- liquidity coverage ratio, 23
- liquidity hoarding, 38
- locally tree-like, 77
- locally tree-like independence, 105
- locally tree-like independence property, 77
- LTI, 105

- macroprudential regulation, 22
- master netting agreement, 20
- microprudential regulation, 22

- net stable funding ratio, 23
- network topology, 79
- node, 52
- node type, 52, 67
- node-type distribution, 54
- nominal value, 27

- out-component, 82
- out-degree, 52
- out-neighbourhood, 53
- out-stub, 56, 69
- over-the-counter securities, 19

- Pareto distribution, 16
- Pareto tail, 80
- percolation
 - bootstrap percolation, 91
 - site percolation, 90

- potential future exposure, 20
- preferential attachment model, 57
- procyclicality, 23

- random financial network, 49
- random graph, 54, 55
 - Poisson graph, 55
 - random regular graph, 56
- reflexive edge, 81
- reflexivity, 81
- rehypothecation, 21
- repo, 21
- repo haircut, 18
- repurchase agreement, 21
- risk weight, 23
- rollover risk, 18

- scale free graph, 57
- seed, 26
- seed node, 93
- self-organized criticality, 16
- SIFI, 23
- skeleton, 49
- spectral radius, 120
- static default cascade, 26
- stress buffer, 39
- stress shocks, 39
- stressed bank, 39
- strongly connected component, 82
- subcritical, 87
- supercritical, 87
- surplus set, 32
- systemically important financial institution, 23, 126

- tail exponent, 80
- tendril, 82
- threshold function, 92
- total return swaps, 20
- tube, 82

- vulnerable edge, 122, 123
- vulnerable node, 98

- w.h.p., 87
- weakly connected component, 82
- with high probability, 90