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 (Charles Dickens, David Copperfield, Chap. 12, p. 185 (1950). First published 1849–1850.).

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

Keywords Systemic risk definition · Financial stability · Balance sheet · Contagion channel · Macroprudential regulation

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 always lead him and his loved ones to the unmitigated misery of the poorhouse. On the other hand, it seems that an aristocrat would usually escape from his debtors to the comforts of his dining club.

For people, firms, and in particular banks, bankruptcy nowadays is more complicated yet still retains some of the flavour of the olden days. When a bank fails, it often seems that the rich financiers responsible for its collapse and the collateral damage it inflicts walk away from the wreckage with intact compensation packages and bonuses. When a particularly egregious case arises and a scapegoat is needed, 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 remain 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.

T.R. Hurd, *Contagion! Systemic Risk in Financial Networks*,
SpringerBriefs in Quantitative Finance, DOI 10.1007/978-3-319-33930-6_1

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 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, often do just fine. Banks that are too big to fail are propped up, while failed banks are bailed out by governments, their debts taken over and paid by the taxpayers.

If anything is different since the crisis of 2007–2008, perhaps it is the widespread recognition that society needs to find ways and means 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. When banking incentives and regulation are better aligned with the needs of society, we might hope bank failures will be better anticipated, prepared for and managed to reduce their most catastrophic social consequences.

1.1 The Nature of This Book

The title "Contagion! Systemic Risk in Financial Networks" is intended to suggest that financial contagion is analogous to the spread of disease, and that damaging financial crises may be better understood by bringing to bear ideas gained from studying the breakdown of other complex systems in our world. It also suggests that the aim of systemic risk management is similar to the primary aim of epidemiology, namely to identify situations when contagion danger is high, and then to make targeted interventions to damp out the risk.[1]

The primary goal of this book is to present a unified mathematical framework for the transmission channels for damaging shocks that can lead to instability in financial systems. Models in science and engineering can usually be described as either explanatory or predictive. In the early stages of research in a field, explanatory models may make dramatic oversimplifications or counterfactual assumptions that are only justifiable to the extent they highlight and explain the most critical mechanisms underlying the phenomenon of interest. Later, when guided by such improvements in understanding, predictive models become feasible. Certainly, predictive models will be more complex, and must be carefully calibrated to the details of the observed system in question. Since financial systemic risk is a rather new field, this book focusses on certain explanatory models developed by economists that aim to explore how disruptions can arise in large financial systems. We will therefore make certain

[1] 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.

dramatic oversimplifications in the hope of gaining mathematical clarity and analytic tractability that can improve understanding of the different ways financial instability can arise.

This introductory chapter will develop the concepts and setting for systemic risk in financial networks. It provides a brief survey of how people have viewed and defined financial crises and systemic risk. It looks at how banks' balance sheets reflect the type of business they deal with, and the ways adverse shocks between banks can be transmitted and amplified. Finally, we review the key aspects of the new international regulatory regime for banks that is designed to safeguard global financial stability.

From Chap. 2 onwards, we delve more deeply into the mechanics of the interactions between banking counterparties. Chapter 2 puts a sharp focus on the type of bank behaviour that can negatively impact the functioning of the entire system, by surveying, dissecting and classifying a number of economic models for financial contagion that have been proposed in recent years. We will make the important discovery that a common mathematical structure unlies a variety of financial cascade mechanisms, namely such crises proceed through cascade mappings that approach a cascade equilibrium. To address the intrinsic opacity of financial institutions and their interconnections, we identify a particular point of view developed by Gai and Kapadia [44], Amini et al. [7] and others that argues for the usefulness of *random financial networks*, a statistical representation of networks of banks, their interconnections and their balance sheets. The design of this concept reflects the type of models that network science, reviewed in the book [73], has already developed in other domains.

The remainder of the book is devoted to studying cascade models on large random financial networks. Chapter 3 provides the mathematical underpinning we need by developing and adapting the theory of random graphs which describes the *skeleton* structure at the heart of the random financial network. Two distinct classes of random graphs, the Assortative Configuration Graph model and the Inhomogeneous Random Graph model, are characterized in detail by their stochastic construction algorithms. The first class, which will form the framework underlying the cascade channels studied in the remaining chapters, is an extensive generalization of the well-known configuration graph model that incorporates *assortative wiring* between nodes that represent banks, which means wiring probabilities depend on banks' degree. It has not been well studied before so we spend time to develop its key mathematical properties, the most important of which we call the *locally tree-like property*. The second class of random graph will in principle enable the meaning of nodes to represent *types* of financial institutions other than banks, such as asset classes or hedge funds. Chapter 4 is devoted to understanding the relation between the Watts 2002 model of information cascades [85] and the concept of bootstrap percolation in random networks, studied recently in [10]. The Watts model will be fully analyzed from first principles, providing us with a template for results on more specific cascade mechanisms on financial networks. We shall learn that its properties can be determined using the mathematics of *percolation*, the theory of the size distribution of connected network components. Chapter 5 returns to focus on the zero recovery default cascade mechanism introduced by Gai and Kapadia [44]. It develops a purely analytical method for computing the large network asymptotics of cascade equilibria, based on the locally treelike property of assortative configuration graphs. The main theorem on the asymptotic form of the default cascade extends the work of Amini, Cont and Minca in certain respects, and requires new proof techniques not previously developed. 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. Finally in Chap. 6 we indicate some of the ways this theory can be extended to encompass more complex contagion channels.

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 resilience and stability in all important dimensions? Are the mathematical conclusions robust and relevant to the real world of financial crisis regulation? We hope this book will be viewed as providing an emphatic "YES" in answer to these questions.

1.2 What Is Systemic Risk?

First it is helpful to identify what systemic risk is not. Duffie and Singleton [34] 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. To the extent that the first four risk categories are focussed on individual institutions, they are not deemed to be systemic risk. However, each of the four also has market wide implications: such market wide implications are wrapped up into the fifth category, systemic risk.

Kaufman and Scott [61], John B. Taylor [81] 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, such 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 [3] 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. It is therefore important to have an unambiguous definition of systemic risk supported by a widespread consensus.

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 no consensus definition over the years, consider the following examples of definitions proposed in the past.

1. Mishkin 1995 [68]: "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 [60] "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 [41] "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 [75] "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 2007–2008 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 [77] of systemic risk explicitly includes this important aspect:

Systemic risk: a definition 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, it has received support from a rather broad range of practitioners. We will therefore accept it as the closest thing we have to a concise definition of the spirit of systemic risk.

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 Adrian and Brunnermeier [4]; *systemic expected shortfall* introduced by Acharya et al. [2]; and *marginal expected shortfall* introduced by Acharya et al. [83]. For a recent and comprehensive review of these and many other systemic risk measures, please see [12].

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" [49]. 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 free thinking overview, we cannot 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*, meaning they may be resistant to average or typical shocks, yet highly susceptible to an attack that targets a highly connected or dominant node.

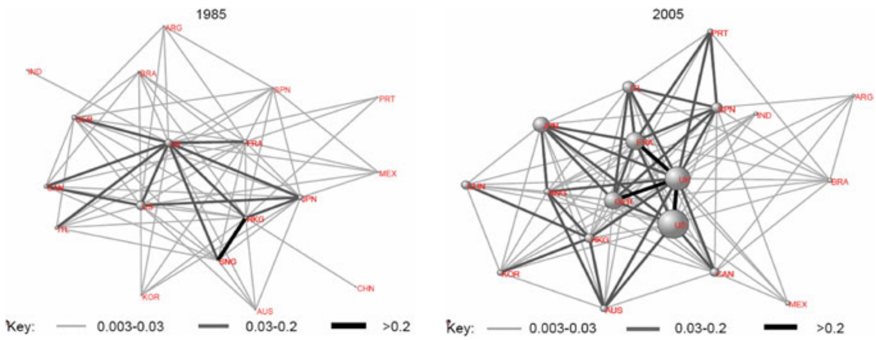

Fig. 1.1 The global financial network in 1985 (*left*) and 2005 (*right*). Here *line thickness* denotes link strength as fraction of total GDP (figure taken from Haldane [49])

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.

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, counterparty risk is known to be small, and thus "Knightian" uncertainty[2] 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 [78]). 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:

---

[2] 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.

> Haldane 2009 [49]: "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, in which all banks became vulnerable to infection by the same *virus*.

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 adaptive systems? Consider the following thought experiment, first proposed by Bak et al. [9]. 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, 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 is the nature of the likelihood of n grains falling off, for each single grain added. BTW's assertion, vindicated since by experiments, is that the frequency for between N and $2N$ grains to fall is roughly the same as for $2N$–$4N$ 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 occasionally be triggered by a very small event. They coined the term *self-organized criticality*, or "SOC", for this type of phenomenon, and boldly asserted that 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 (see [27]). Network scientists are thus not surprised to see results of [28, 80] and others that show evidence of Pareto tails in the size and connectivity of large financial networks, with highly connected *hub* banks that form a core within a periphery of intermediate and small banks.

It might sound naive to assert that something analogous to sand piling is happening in financial systems. However, as Minsky wrote in [67], "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 [5] provide a possible explanation. They demonstrate that in the 10 year period leading up to the 2007–2008 crisis, financial institutions exhibited strongly procyclical 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 a small trigger was needed to create the inevitable collapse.

1.3 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.

Table 1.1 shows the main classes of assets (what the bank owns) and liabilities (what the bank owes). All entries must be non-negative, and *equity*, which is the value of the firm to the share owners, is defined to be the difference $E = A - L$ between asset value A and liability value L. The most fundamental characteristic of a bank's balance sheet is its accounting ratio, *leverage* A/E. For banks, it has often

Table 1.1 The main components of a bank's balance sheet

Assets	Liabilities and equity
Loan portfolio: mortgages commercial loans credit cards	Deposits: retail deposits wholesale deposits certificates of deposit other banks' deposits
OTC securities: bonds OTC derivatives	OTC securities: bond issues OTC derivatives
Market securities exchange traded derivatives	Market securities exchange traded derivatives
Reverse repos	Repos
Cash cash equivalents	Hybrid capital preferred shares COCOs
Other sets	Other liabilities
	Equity

exceeded 50 in the past. Its reciprocal, E/A, measures the bank's safety buffer to absorb adverse balance sheet shocks.

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. 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.

1.3.1 Bank Asset Classes

Assets, what the firm owns, are entered on the left side of the balance sheet. The available classes of assets relevant in banking is extremely diverse, and this section gives only a schematic overview of some of the most important basic types. From a systemic risk perspective, assets' most relevant characteristics are: duration or maturity, credit quality (or collateral), interest rate and liquidity.

Loan portfolio: Banks, like any firm, invest heavily in endeavours for which they have a competitive advantage. Such *irreversible projects* are by their nature illiquid, and fail to recoup their full value when sold prematurely. For banks, this business line is called the *bank book*, and consists of a heterogeneous portfolio of loans and mortgages of all maturities, to counterparties ranging across the retail sector, small and medium enterprises (SMEs) and major corporates. Far from receding in importance since the crisis, [59] shows that real estate lending in the form of 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 as in the past, real estate bubbles will continue to be the dominant factor in triggering future systemic risk events.

From an accounting perspective, loans are regarded as to be "held to maturity", and are therefore assigned *book value*, which is typically the value at the time of origination. The book value may be marked down infrequently, if the asset is regarded as sufficiently stressed.

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, large 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. Counterparty risk management, reviewed for example in [29], 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 of the positive part of their exposure to the given counterparty on a given date in the near future. In the event that one bank suddenly defaults, PFE is a pessimistic estimate of losses to which its counterparties are exposed.

A systemically 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. Such *cash equivalents* are used to pay depositors on demand. Examples include money-market lending that pays the over-night rate, stocks, T-bills, and exchange traded derivatives. Since cash equivalent assets are regarded as "available for sale", instead of book value they are assigned their *mark-to-market* value. A new aspect of the

Basel III regulatory framework requires banks to exercise 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 debt obligations in a timely manner.

Reverse repo assets: As described in the next section on repos, a repo-lending bank 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: real estate, accounts receivable, goodwill and the like.

1.3.2 Debt and Liabilities

Debt and liability are regarded as the same thing, namely all obligations that the firm owes to others. Equity, on the other hand, even if entered on the right side of the balance sheet, is the value of firm ownership, and is *not* regarded as debt or liability.

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*, *repo* is the name of an important class of collateralized short term debt issued between banks and other institutional investors. 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* that 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. Preferred shares can be considered as hybrid capital: 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.

1.3.3 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 A/E is large. 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.

1.4 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 [30].

Asset Correlation: Different banks tend to hold common assets in their portfolios. Haldane [49] 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 dramatically increased most banks' leverage and hence the vulnerability of their asset portfolios. Such systemic events undermine the health of the system, in the same way that famine impairs the health of a community. They make it vulnerable to other types of contagion, but do not exhibit the amplification effect that characterizes contagion.

Default Contagion: Bank deposits held in other banks can be considered as a form of interbank lending, but banking practise in modern times has dramatically extended the range of interbank exposures. There are now 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 work hard to identify their potential future exposure (PFE) over different future time horizons. An insolvent bank, if it is not bailed out by a government agency, will be forced into bankruptcy thereby disrupting promised contractual payments. 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* and can in principle chain together like dominos to create a *default cascade*. Default cascades are most dangerous when interbank exposures are a high fraction of lending banks' equity, and [82] 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. Despite being the first type of contagion most economists consider, default cascades are rare in practise because defaulting banks are often bailed out by government.

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. The amplification characteristic of contagion occurs 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 [5] 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 [5] 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 "rollover" short term debt. Reference [71] 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 roll over their debt, will be more likely not to roll over 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*). Reference [8] extend this picture to a network setting by considering a multitude of simultaneous coordination games, leading to runs in the interbank network. In [46], 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.

1.5 Regulatory Capital and Constraints

The 1988 Basel Accord, also called Basel I, was a set of minimal capital requirements for banks that arose from the deliberations of international central bankers who formed the Basel Committee on Banking Supervision. Largely as a result of lessons hard learned during the 07–08 crisis, the world is now moving quickly beyond the Basel II regulatory regime, put in place in the early 2000s, that can be characterized as *microprudential* in emphasis, with regulations that were 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 now 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 counter-cyclical. During normal periods, capital requirements have a surcharge which can be removed to provide banks with more flexibility as the system begins to contract. Another example of macroprudential regulation is that Basel III now recognizes the existence of *SIFI*s (for *systemically important financial institutions*), also called *G-SIB*s (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 of 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 A/E. 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 [64] (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 [50], 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.