

Contagion in Financial Networks

Paul Glasserman

Office of Financial Research and Columbia University

paul.glasserman@treasury.gov

H. Peyton Young

Office of Financial Research, University of Oxford, and the London School of Economics

peyton.young@treasury.gov

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

Contagion in Financial Networks*

Paul Glasserman and H. Peyton Young

October 20, 2015

Among the many factors contributing to the financial crisis of 2007-2009, the increasing *interconnectedness* of the global financial system is perhaps the least well understood. The crisis exposed the fact that regulators and market participants had very limited information about the network of obligations between financial institutions. The crisis also revealed that there was little theoretical understanding of the relationship between interconnectedness and financial stability. Since that time, there has been a surge in empirical and theoretical work on these issues. This paper surveys this literature, focusing on what it tells us about the relationship between network structure and the vulnerability of the financial system to contagion.

The key challenges facing this research area include the following. What are the reasons for the growing interconnectedness of the financial system? Do more connections tend to amplify or dampen systemic shocks? Does the structure of the network matter? If so, what structural features are relevant for setting policy? A particularly important question is how network structure *interacts* with other potential sources of contagion. In the run-up to the crisis leverage levels increased and capital buffers at some banks were extremely thin. Large financial institutions had increased their reliance on short-term funding, which is quickly withdrawn when a crisis hits. There was an increasing concentration of assets in a few global institutions. The proliferation of derivatives and securitization also increased the complexity of banks' balance sheets, creating greater uncertainty about their financial condition during the crisis.

All of these factors affect the stability of the financial system to varying degrees. The key issue is how the network of obligations relates to other potential sources of contagion, and whether it serves to amplify or dampen them. Network connections can have a positive effect

* This paper was produced while Paul Glasserman was a consultant to the Office of Financial Research. The authors thank Steven Durlauf, Prasanna Gai, Iman van Lelyveld, Alizerza Tahbaz-Salehi and several anonymous reviewers for constructive comments on an earlier draft.

by diversifying risk exposures for individual banks, but they can also have a negative effect by creating channels through which shocks can spread. Much of the literature is devoted to examining the tension between these two forces. One branch focuses on the incentives that banks have to form links with one another. These include access to investment opportunities, risk diversification, liquidity management, and the provision of specialized financial products. Another branch of the literature investigates how shocks are transmitted through the web of obligations once they are in place. Here the key issue is the extent to which the network contributes to systemic risk by amplifying shocks to individual institutions or assets. The primary focus of this paper will be on the second branch of the literature, although along the way we shall also provide pointers to the first.

The question of whether interconnections amplify or dissipate shocks depends on many factors in addition to the network structure. To what extent are banks leveraged? Do interbank obligations have priority over obligations to the nonfinancial sector in case of default? When a default occurs what proportion of the nominal obligations are recovered? Another key question is what triggers a financial crisis in the first place. Is it a sudden demand for liquidity? An economic shock to the value of real assets? A loss of confidence in the credit-worthiness of particular institutions? Third, how should one measure the ‘size’ of a systemic event? Is it the number of banks that default? The shortfall in payments from the financial sector to households? The loss of bank equity? The overall contraction of credit? How do these indicators relate to the overall loss of welfare in the economy?

We will survey a variety of models that address these questions.¹ Not surprisingly the choice of assumptions leads to different conclusions about the importance of interconnectedness. However we also find that there is a common methodology underpinning many of these models, including the mechanisms that describe how initial shocks spread through the network and the use of fixed point arguments to describe the cumulative effect of these cascading shocks.

¹ Other surveys, each with its own focus, include Allen and Babus (2009), Summer (2013), Cabrales, Gale, and Gottardi (2015), and Hueser (2015).

The structure of the paper is as follows. In Section 2 we provide some context by recalling certain key events in the financial crisis of 2007-09, as well as the earlier savings and loan crisis of the late 1980s. The point of this discussion is to emphasize the multi-faceted nature of the modern financial network and the variety of ways that it can generate contagion in practice. This forms the backdrop for our subsequent discussion of the theoretical literature. In particular, the failure of the insurance giant AIG and the collapse of Lehman Brothers highlight two types of network effects that are the main focus of the rest of this article: AIG was ostensibly rescued to prevent a default cascade, whereas Lehman's failure led to (and was partly caused by) a funding run.

Section 3 lays the groundwork for network models of interconnected balance sheets. We introduce the basic elements that make up the balance sheet of a financial institution, including such concepts as assets, liabilities, leverage, and net worth. We then consider some of the reasons why institutions create links with one another, including the management of liquidity and diversification of risk. We illustrate the complexities that result from these interconnections by considering a hypothetical example, and trace how a shock that originates at one institution can spread through the entire network. The same framework can sometimes be used to model default cascades that result from actual shortfalls in payments, and also to model funding runs triggered by liquidity shocks or crises of confidence but not necessarily by outright defaults.

Section 4 considers network models of liquidity shocks. We begin by discussing the influential framework of Allen and Gale (2000) and related models. This line of research treats the formation of links in some detail, and then examines the relative vulnerability of simple, stylized network structures such as a completely connected network and a ring. We then consider a different approach that is exemplified by the framework of Gai, Haldane, and Kapadia (2011). This strand of the literature investigates the transmission of liquidity shocks in a richer variety of network structures, seeking to go beyond qualitative insights to provide quantitative comparisons through numerical simulations. The focus is on the trade-off mentioned earlier between the stabilizing effect of interconnections due to diversification, and the amplifying effect that may result from additional channels through which shocks can spread.

We turn in Section 5 to a discussion of cascading defaults, building on the pioneering model of Eisenberg and Noe (2001) and its extensions. The Eisenberg-Noe framework describes

how payment shortfalls that originate at one or more nodes spread through the network, causing a widening series of defaults. The original model treats all payment obligations as if they had equal priority; it also omits bankruptcy costs, which can significantly amplify the impact of a default. We consider extensions of the basic framework that accommodate these and other features of the real-world financial system. We also show how the framework can be adapted to model crises of confidence. Under this interpretation a change in market perceptions about the credit-worthiness of a given institution can have knock-on effects on the credit-worthiness of other institutions, thus creating a form of information contagion.

In Section 6 we examine the question of how to measure the amount of “damage” caused by a systemic event. Various measures have been proposed in the literature: the number of banks that fail; the total loss of bank capital; the cost of liquidating long-term assets to cover short-term liabilities; the deadweight costs of bankruptcy proceedings, and so forth. These modeling choices have different implications for the welfare effects of financial contagion and for setting policies that limit the contagiousness of the financial system. They also have implications for assessing the robustness of financial systems under the stress tests conducted by central banks.

In Section 7 we consider various indicators of contagiousness and vulnerability that can be computed without knowing the details of the entire network structure. We show that even with no detailed knowledge of the network structure one can say quite a lot about the inherent contagiousness and vulnerability of a node (an individual financial institution) using node-level information. These indicators include the node’s level of leverage, its relative size, its immediate exposure to other institutions, and its degree of connectivity with the rest of the financial system.

In Section 8 we consider a model of the joint distribution of shocks. Much of the financial network literature considers ad hoc shocks, like the failure of a specific institution. Stress testing exercises by central banks typically involve the construction of alternative scenarios rather than a probability distribution. Without a full distribution, however it is difficult to assess the inherent vulnerability of different network structures to financial contagion. Here we discuss an approach suggested by the present authors that is both flexible and tractable (Glasserman and Young, 2015). In particular we show how to bound the expected total loss inflicted by simultaneous financial shocks, and how the network topology enters into this calculation. Surprisingly, one can obtain useful bounds without assuming independence

among shocks, which is an important consideration given that banks often hold similar portfolios. A particularly important indicator is the “financial connectivity” of the system, which determines the extent to which shocks reverberate within the system before dissipating.

In Section 9 we discuss a variety of network concepts that have been proposed for assessing the inherent vulnerability of financial systems to contagion. One widely studied measure is the degree distribution, where the degree of a node (e.g., a bank) is the number of links that it has with other nodes. To keep the analysis tractable it is usually assumed that the nodes are the same in other respects, such as size, leverage, and asset quality, and the object is to study the vulnerability of such systems as the degree distribution is varied. Another strand of this literature focuses on measures of node centrality. Here the basic thesis is that more centrally placed nodes are potentially the most contagious. We discuss alternative notions of centrality in the network theory literature, and then consider empirical studies that investigate the relationship between centrality measures and the inherent vulnerability of the network to contagion. We argue that thus far the empirical work has not produced a compelling link between traditional network measures and financial stability. Much more remains to be done at both a theoretical and empirical level to identify the measures that are appropriate for predicting contagion in actual financial systems.

In Section 10 we return to a theme that has animated much of the literature, namely, the trade-off between connectivity as a channel of contagion and connectivity as a means of diversifying risk. A common strategy in this literature is to investigate the trade-off by varying the average number of links per node while assuming that the nodes are similar in other respects. One shortcoming of this approach is that it suppresses other factors that interact with connectivity in complex ways, such as heterogeneity in banks’ size and leverage. We argue that one should also distinguish between the level of connectivity *within* the financial system (as measured by the density of interbank links), and the level of connectivity *between* the financial and nonfinancial sectors. The latter is extremely important in amplifying (or dissipating) shocks to the system, as discussed in Section 8.

Section 11 examines contagion through common exposures, as opposed to contractual obligations. When financial firms have overlapping asset holdings, a shock to one firm can spread to others through a drop in the price of common assets, particularly if the initial shock leads to a fire sale. Optimal portfolio selection diversifies risk for individual firms but can

make the system as a whole more vulnerable if it leads many firms to hold similar portfolios of assets, increasing the likelihood that many firms will be in financial distress at the same time. We provide a brief survey of the literature on these issues.

In Section 12 we discuss the lack of information that bedevils much of the research on (and supervision of) financial networks. Even regulators who have access to confidential information do not have detailed knowledge of many parts of the system. There are various ways of dealing with this problem. One branch of the literature dispenses with any attempt to measure the links directly, and instead studies the co-movements between the market values of different financial institutions using statistical methods. Another strand of the literature fills in missing data on network connections based on the data available. Yet a third way of dealing with limited information is to formulate measures of contagiousness and vulnerability that do not rely on detailed knowledge of the network, as discussed in Section 7. We argue that the lack of information is not merely a problem for regulators and analysts; it also creates uncertainty for market participants that can become particularly acute in times of crisis. In other words, the opacity of the network due to lack of information is itself a contributor to contagion, and may lead to cascades and funding runs that would not occur if the network of obligations were known with greater certainty.

2. Financial Interconnections and Financial Stability in Practice

Let us begin by considering several recent episodes that illuminate the role of interconnections between financial institutions in periods of financial stress. This material will provide context and perspective for the discussion of theoretical models in subsequent sections. A theme of this discussion will be the variety of interconnections in the events described. In particular, the examples illustrate the following phenomena, sometimes operating simultaneously: direct loss spillovers through defaults, mark-to-market losses, funding runs, information contagion, contagion through correlation, and common exposures. The examples also highlight the heterogeneity and opacity that characterize financial networks.

2.1 AIG and Its Counterparties

Of all the dramatic events of 2008, the one that most nearly fits the image of defaults cascading through a network is the failure of the insurance giant AIG. The company's financial products unit had sold guarantees, in the form of credit default swaps, to several

large banks. The swaps guaranteed payments from certain pools of assets, including subprime mortgages. As market conditions deteriorated over the summer of 2008, AIG could not meet its payment obligations. The government stepped in, providing funds that allowed AIG to pay its bank counterparties on the swaps, sparing those banks losses that might otherwise have continued to spread through the financial system. The *Financial Crisis Inquiry Report* concludes (p.352), “Without the bailout, AIG’s default and collapse could have brought down its counterparties, causing cascading losses and collapses throughout the financial system.” The report also makes clear that the Federal Reserve and Treasury Department had an incomplete picture of AIG’s network of payment obligations.

The AIG story illustrates a further point to which we return in our discussion of models: the distinction between realized losses and “mark-to-market” losses. The credit default swaps AIG had sold required AIG to cover losses on pools of assets in excess of a threshold. Losses in a pool exceeding the threshold can be seen as analogous to a default, against which AIG was providing protection. In many if not most cases, actual losses did not exceed the required threshold so there was no default, and, in a narrow sense, no payments were due from AIG on its guarantees. However, as the condition of the underlying assets deteriorated, the *value* of the guarantees increased because it became more likely that they would be called on. AIG’s counterparties demanded collateral payments from AIG to reflect the increased mark-to-market value of its liabilities. AIG was brought down by these collateral calls, not by a direct failure to fulfill its guarantees. (See Chapter 19 of the *Financial Crisis Inquiry Report* for a fuller discussion of these events.) The point this illustrates is that losses can spread from one node in the network (a pool of assets) to another node (AIG) through changes in the market value of a guarantee, without an actual default.

2.2 *The Lehman Brothers Bankruptcy*

One of the puzzles of September 2008 is whether the bankruptcy of Lehman Brothers precipitated the events that followed it or whether the latter would have occurred anyway. The most significant direct consequence of the Lehman failure was the liquidation of the Reserve Primary Fund, a money market fund that held Lehman debt. This triggered a run on money market funds, including funds with no direct exposure to Lehman, a major source of short-term funding for financial and nonfinancial companies. Short-term funding, on which the financial system had become increasingly reliant, became scarce. Creditors became concerned about the viability of borrowers and they hoarded liquidity as a precaution against

a further deepening of the crisis. In short, the days following the Lehman bankruptcy look more like a funding run than a cascade of defaults. By the end of the week, the Treasury department stepped in to guarantee money market fund assets in an effort to revive short-term lending. Without a government guarantee, money market funds might have been forced to sell illiquid assets to meet redemptions. A fire sale of illiquid assets would have driven down their prices, decreasing the market value of similar assets held by other financial institutions.

The aftermath of the Lehman bankruptcy illustrates multiple avenues for contagion: a direct loss imposed on the Reserve Primary Fund, “information contagion” that spread fears to other money market funds, a funding run as creditors pulled back lending, and potential fire sales. Network opacity heightened uncertainty in the lead-up to the Lehman bankruptcy. According to the FCIC report (p.329), “there was no way to know who would be owed how much and when payments would have to be made—information that would be critically important to analyze the possible impact of a Lehman bankruptcy on derivatives counterparties and the financial markets.”

2.3 Mortgage Funding Chains

Subprime lending and changes in the mortgage market featured prominently in the unfolding of the financial crisis. In traditional mortgage lending, a bank evaluates and monitors the creditworthiness of the borrower and the value of the mortgaged property. With securitization, the ultimate lender and borrower are often separated by a chain of intermediaries.

Figure 1 (adapted from Shin 2010) illustrates a possible chain. Households take on mortgage debt, mortgages are pooled into securities, the securities are sold to dealers, the dealers pledge the securities as collateral to borrow from commercial banks, commercial banks fund themselves by issuing short-term debt to money market funds, and money market funds take deposits from households. The arrows in the figure indicate the direction of payment obligations.

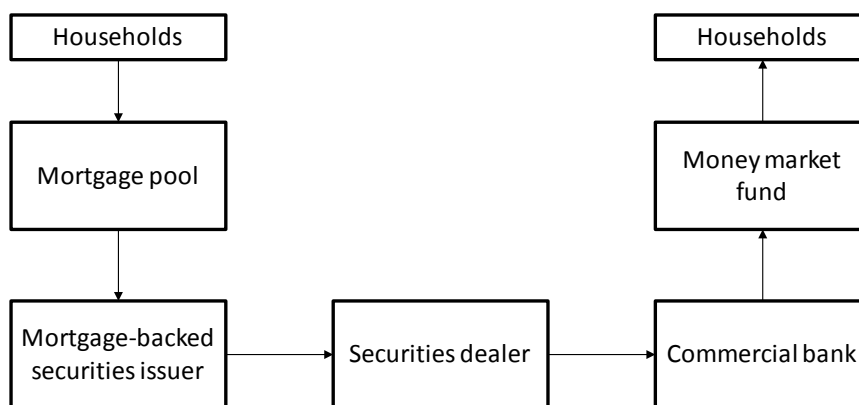


Figure 1: A funding chain. The arrows indicate the direction of payments due. Funding flows clockwise through the chain. Adapted from Shin (2010)

The chain illustrates several points: The nodes in the chain can be highly heterogeneous in their structure and function – heterogeneity that is often overlooked in network models. At each step in the chain, there is a potential loss of information about the quality of the underlying debt. Leverage along each step in the chain amplifies shocks to the underlying real estate value. Shin (2010) argues that shocks are amplified because financial institutions increase their borrowing and lending to each other to increase their leverage while the supply of real assets remains relatively fixed.

2.4 Hedge Funds and the Quant Meltdown

During the first week of August 2007, many hedge funds following quantitative trading strategies suffered large losses, in some cases leading the funds to close. The abnormal market behavior lasted just a few days, but the damage was enough to discredit many “quant” investment strategies. In retrospect, the turmoil of that week was also a harbinger of the crisis to come.

For our purposes, the incident is a valuable example of contagion through correlated assets and strategies, rather than through contractual obligations. The events of that week can be summarized as follows (Rothman 2007; Khandani and Lo 2007). Due to an initial shock outside the stock market – perhaps losses on mortgage-backed securities – some funds were forced to sell assets. To raise cash quickly, they sold their most liquid assets, namely U.S. equities. Algorithms at other quant funds reacted by joining the sell-off, and the pattern continued for a few days. Khandani and Lo (2007) detail the cascade of losses from one type of strategy to the next each day. This was not a conventional fire sale – the assets sold were highly liquid – but rather a rush to the exit that triggered a rapid deleveraging. Losses spread

through price correlations and exposures to common factors, rather than through direct connections between the affected funds.

2.5 The Savings and Loan Crisis

Our final historical incident serves as an example of widespread failures through common exposures. Rising interest rates in the early 1980s challenged savings and loan institutions which had to pay market rates on deposits but earned fixed rates on the mortgages they had issued. Deregulation allowed these institutions to shift to riskier commercial real estate lending in an attempt to make up losses. By the end of the 1980s, hundreds of savings and loan institutions had failed at an estimated cost of \$160 billion (FDIC 1997). From a network perspective, these were isolated nodes with minimal, if any, direct exposures to each other. They failed together because of common exposure to interest rate risk and because many of them were located in the same regional real estate market (the southwestern United States).

One could point to many other historical examples, but these suffice to illustrate the mechanisms network models need to take into account. In the rest of this article, we will revisit the mechanisms introduced in the examples of this section, with particular emphasis on the two that are primarily network effects – loss spillovers and funding runs. Fire sales, information contagion, correlated risks, and common exposures do not require a network of payment obligations, although they can amplify network effects. At several points, we will also stress heterogeneity and opacity as essential aspects of financial networks, though these features are often overlooked in simple models.

3. Networks of Interlinked Balance Sheets

3.1 Balance Sheet Components

Figure 2 shows a stylized balance sheet of a financial institution, which for brevity we shall call a *bank*. The bank (indexed by i) has two categories of assets – outside assets and in-network assets. Outside assets c_i are claims on nonfinancial entities, such as mortgages and commercial loans. In-network assets are claims on other banks; these include interbank loans and exposures through derivatives. We denote by \bar{p}_{ki} the payment obligation of bank k to bank i . The bank's liabilities include obligations b_i to nonfinancial entities – for example, depositors – and obligations \bar{p}_{ij} to other banks j . The difference between the bank's assets and liabilities yields the bank's *net worth* w_i . The figure is purely schematic and is not intended to indicate the relative magnitudes of the various parts of the balance sheet.

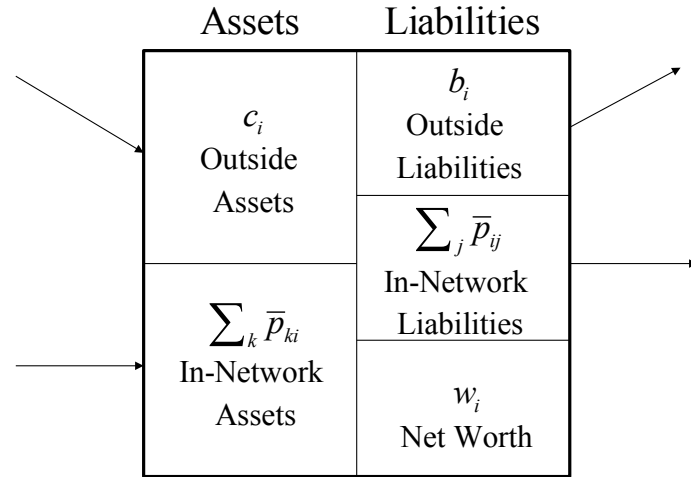


Figure 2: A stylized bank balance sheet for bank i .

The links between balance sheets define a network. Each node is a bank, and a directed edge runs from node i to node j if bank i has a payment obligation to node j . (All entities outside the network can be represented through a single node representing the “outside.”)

To put this network in context, we recall key functions of the financial system:

- Facilitating payments;
- Allocating capital by intermediating between lenders and borrowers, and also between investors and businesses;
- Managing liquidity and maturity transformation in intermediating between lenders who prefer to lend for short maturities and borrowers who prefer to borrow for longer maturities;
- Providing risk transfer from agents seeking to reduce risk to others willing to bear greater risk if compensated through higher returns.

The first of these roles is inherently a network operation, requiring links between buyers, sellers, and banks. Several studies have examined the network structure of large-value payment systems (such as Fedwire and CHIPS in the United States and TARGET2 in Europe), and we return to these in Section 9. Allocating capital deals primarily with investing the outside liabilities b_i into outside assets c_i . This process entails converting the illiquid, long-dated assets c_i into money-like claims b_i . The resulting maturity and liquidity transformation leaves banks vulnerable to a loss of funding, and interbank borrowing and

lending serves to help banks manage the risk of fluctuations in their outside funding. Banks also help corporations manage their exposures to exchange rates, interest rates, and commodity prices through derivatives and other contracts; the banks hedge this risk by trading with other banks.

Interbank claims thus arise as a mechanism for banks to share risks. But these same links become channels through which problems at one bank can spread to another. The tradeoff between the costs and benefits of interconnectedness is one of the main themes of the financial networks literature.

3.2 Channels of Contagion

We can use Figure 2 to revisit some of the channels for contagion touched on in Section 2. A shock to the value of bank i 's assets would cause a drop in c_i . This asset shock could be a drop in the value of real estate or a downturn in an industrial sector to which the bank has made loans. A drop in c_i is initially absorbed by the bank's net worth w_i . But if the shock is sufficiently large, the net worth is wiped out, the bank is unable to fully repay its liabilities, and it defaults. In particular, its actual payment p_{ij} to bank j will be less than its promised payment \bar{p}_{ij} . If the payment shortfall is sufficiently large, it can push bank j to default as well, and so on. An initial asset shock to one bank can spill over to other banks, creating a cascade of defaults.

As this discussion illustrates, we distinguish between a bank's promised payment and its actual payment to another bank. More generally, we may distinguish between the face value \bar{p}_{ij} of bank i 's obligation to bank j and the *market value* of this obligation. In the simplest case, the two remain the same until bank i defaults. More generally, the market value may decline as bank i 's net worth shrinks, even if the bank remains solvent (meaning $w_i > 0$). In this case, a shock to bank i can impose a loss on bank j through a "marking to market" of the value of i 's obligation to j , even without a default.

Another type of contagion arises when banks pull funding from one another. Such a funding run can be triggered by an unexpected liquidity shock as in Diamond and Dybvig (1983) and the literature builds on their framework. In the balance sheet of Figure 2, a liquidity shock arises as a reduction in b_i . With its funding reduced, the bank needs to reduce its assets as well. If it withdraws its lending to other banks – its in-network assets – those banks will in turn need to reduce their lending, creating a funding run.

Figure 2 suggests other ways in which bank balance sheets may interact. Banks may have common exposures through their outside assets – real estate, for example – and may thus be subject to correlated shocks. A disclosure by one bank about its assets may lead creditors to make inferences about the assets held by other banks, producing “information contagion.” If one bank is forced to sell illiquid assets and in so doing drives down the price of these assets, then other banks holding similar assets incur a fire-sale externality through the price drop. Although these dynamics can arise even in the absence of a network, that is, in a collection of isolated banks, they are further amplified when banks have obligations to one another. In what follows, however we shall focus mainly on those sources of contagion, such as asset and funding shocks, whose operation *depends* on interconnections between banks.

3.3 A Simple Network Example

To illustrate the spread of shocks through interlinked balance sheets consider the example in Figure 3. The number on each directed edge represents a payment obligation, and each node’s net worth is shown in bold. For example, consider bank C. It is owed 160 by outside entities – for instance, households with mortgage obligations – and it owes 50 to a possibly different set of households or depositors. Additionally C is owed 100 by bank B and it owes 100 to each of banks A and D; these obligations might take the form of unsecured overnight loans for example. The difference between bank C’s assets (160+100) and its liabilities (50+100+100) leave it with a net worth of 10.




Figure 3. A hypothetical financial network showing payments due. The net worth of each node is shown in bold. The outside sector consists of households and nonfinancial firms.

Suppose that the economy is hit by a shock that causes some households to default on their payments to bank C: instead of the promised 160 they pay only 40. Then C defaults because its assets total $100 + 40 = 140$ whereas it owes 50 to the outside sector and 200 to other banks. In this case, we assume that C's remaining assets are paid pro rata to C's creditors. As we shall see, C's assets may turn out to be worth even less than 140, because its default may trigger a chain of defaults that lead back to C.

To work through these spillover effects, we proceed by computing "interim" payoffs as follows. If we take the interim value of C's assets to be 140, the pro rata rule implies that C pays $(100/250) \times 140 = 56$ to D, 56 to A, and 28 to the outside depositors. Now D has assets worth $204 + 56 = 260$ and debts totalling 300, so D is in default. The pro rata rule implies that D pays 130 to A and 130 to its outside depositors. At this stage A's assets have interim value $120 + 130 + 56 = 306$, whereas its nominal obligations come to 360. Thus A defaults and the pro rata rule implies that it pays one-half of its assets to B, namely 153, and an equal amount to outside depositors. At this juncture B's assets are worth $153 + 30 = 183$, whereas its obligations total 200. Therefore B defaults and the pro rata rule implies that it pays 91.5 to C and 91.5 to outside depositors.

At this point, we discover that the value of 140 we used for C's assets was incorrect. That value reflected the initial outside shock of 40, but it assumed full repayment of 100 from bank B. In fact, B is able to pay at most 91.5, so C's assets are worth at most 131.5, and the cycle must be repeated. Because of this cascade of defaults, determining the consequences of the initial shock is a fixed-point problem, as analyzed by Eisenberg and Noe (2001), which we take up in greater detail in Section 5.

We now use the same example to illustrate a funding run. Reverse the direction of all arrows in Figure 3. In this interpretation, bank C has raised 260 in funding but extended only 250 in loans. The difference of 10 is a cash buffer – an asset now, not a liability. The initial shock reduces the amount bank C has raised from households from 160 to 40. This is now a funding shock rather than an asset shock; households have experienced an unexpected demand for liquidity and withdrawn 120 of their deposits.

Bank C can manage a small funding shock through its cash buffer of 10. But to meet the withdrawal of 120, it needs to reduce its lending. For simplicity, suppose it does so pro rata. Following the same arithmetic as before, it reduces its lending to banks D and A to 56 each, and it reduces its outside lending to 28. Eventually, we find that the funding run causes bank

B to pull some of its lending to bank C, so that C's funding loss ends up being larger than the initial shock.

The mechanics of this funding run are similar² to the mechanics of the default cascade we described first, but the interpretation is different. Both mechanisms are important in practice, and the two may operate simultaneously. We will discuss models based on liquidity shocks in Section 4 and then return to default cascades in Section 5.

We have kept the example of Figure 3 simple for illustration. Additional features will be discussed in later sections. In particular, shocks can be amplified through bankruptcy costs in the case of a default cascade and liquidation costs and liquidity hoarding in a funding run. One can replace the assumption of a pro rata allocation with different levels of priority. Balance sheets can be enriched to include more detail on assets and liabilities.

4. Liquidity Risk Sharing and Funding Runs

In this section, we examine a key trade-off in financial networks: interbank lending helps banks manage liquidity risk, but it also creates channels through which shocks can spread from one bank to another. The relevant literature is large, and we will not try to be comprehensive in our coverage. Instead, we will anchor our discussion around two models that highlight important features of the broader literature. The first model, due to Allen and Gale (2000), provides detailed foundations for the actions of banks and depositors and focuses on a network of four banks. A second line of work, exemplified by the model of Gai, Haldane, and Kapadia (2011), relies on reduced-form descriptions of the interactions between banks and strives to describe aggregate behavior in large networks.

4.1 The Allen and Gale (2000) Model

In the Diamond and Dybvig (1983) model of bank runs, a bank is exposed to the risk of a liquidity preference shock when a fraction of depositors withdraw their deposits early for reasons exogenous to the model. Early withdrawals may force the bank to liquidate illiquid long-term assets and incur costs in doing so.

In the model of Allen and Gale (2000), interbank lending networks allow banks in different regions or sectors to share liquidity risk because the liquidity shocks they face are negatively

² A formal duality between the two mechanisms is developed in Chapter 6 of Hurd (2015).

correlated. Lending by banks with excess liquidity to banks with a liquidity shortage can prevent costly early liquidation of long-term assets.

The model has the following elements within each region:

- Time evolves over three dates, 0, 1, 2.
- At time 0, consumers deposit their endowments in banks. At time 1, a fraction of consumers experience a liquidity shock and seek to withdraw their deposits. At time 2, the bank's remaining assets are paid to remaining depositors.
- Banks can invest in a liquid short asset and an illiquid long asset. The long asset earns a higher return if held to time 2 but a lower return if liquidated at time 1.

In the baseline case, bank holdings of the short asset are sufficient to meet the withdrawals of consumers who receive the liquidity shock at time 1, and the remaining “patient” consumers earn a higher return by waiting until time 2 to withdraw. But with a large liquidity shock, a bank is forced to incur the cost of early liquidation of the long asset, and “patient” consumers may optimally decide to withdraw their deposits early.³

Interbank lending across regions can help mitigate this risk. Allen and Gale (2000) first consider a case where initially identical regions differ in the severity of the liquidity shock they receive, holding the overall severity constant. When one region has a higher fraction of impatient consumers, another region has a lower fraction of impatient consumers. Because they are ex ante identical, banks in the different regions hold the same level of short assets. At time 1, the liquidity shortfall in regions with a high fraction of impatient consumers is offset by the excess supply of liquidity in regions with a low fraction of impatient consumers. Anticipating this outcome, but not knowing which regions will receive which shocks, banks lend to each other by exchanging a fraction of their deposits across regions, eliminating the risk that results from the disparity in the shocks.

Allen and Gale (2000) then show that interbank lending, while allowing risk sharing, can also create financial fragility. They introduce a state in which the fraction of early depositors takes its average value in all regions except for one in which the fraction is strictly higher. This state is unanticipated by banks and consumers at time 0, so the total supply of short

³ In particular, a bank run in Allen and Gale (2000) is rational and due to bank fundamentals. In Diamond and Dybvig (1983), a bank results from a coordination failure among depositors rather than from the condition of the bank.

assets held by banks is insufficient to meet withdrawals at time 1. Interbank lending cannot solve the problem of meeting the aggregate demand of early consumers. Banks in the region with the highest liquidity shock are forced to liquidate some of their long assets at time 1, incurring liquidation costs. If the shock is sufficiently large, depositors cannot be fully repaid, including banks from other regions holding deposits in the affected region. Interbank deposits, which enable risk sharing in the original model, thus become a channel for the spread of losses after an unexpectedly large liquidity shock.

The transmission mechanism combines the two mechanisms illustrated through Figure 3 in Section 3.3. The initial shock is a liquidity shock. But when a bank is forced to liquidate the long asset early, the cost it incurs can lower the value of its liabilities, including the deposits made by other banks. The other banks withdraw their deposits, as in a funding run, but they recover less than the promised amount, as in a default cascade. The two effects are determined simultaneously because the lending bank withdraws its deposits precisely when the deposits are worth less than the promised amount.

Allen and Gale (2000) compare the fragility of alternative network topologies. A ring network, in which each bank borrows from exactly one other bank, is particularly fragile. If the initial shock and liquidation costs are sufficiently severe to bankrupt one bank, then the loss of value in that bank's liabilities is concentrated on the one bank from which it borrowed, potentially pushing that bank into default. This leads to a default cascade around the ring, if the shock is sufficiently large. In contrast, in a completely connected network, where every bank lends to every other bank, the impact of the first default is diluted among other banks; a larger initial shock is required for defaults to spread. Banks may also be separated into groups, with interbank lending within groups but not across groups. This allows liquidity risk sharing within groups but creates a fire wall that prevents losses from spreading across groups. One of the themes of Allen and Gale (2000) is the interplay between network topology and the correlation in liquidity shocks across the network.

4.2. Some Closely Related Models

The model of Freixas, Parigi, and Rochet (2000) is driven by the movement of depositors from one region to another, rather than by exogenous liquidity shocks. The model builds on the Diamond-Dybvig (1983) model but sets parameters so that no depositors would need to withdraw their deposits from a solvent bank before the final date. In other words, the model has no impatient consumers.

Migration prompts consumers to consider moving deposits from one bank to another at the intermediate date. Interbank credit allows banks to avoid costly liquidation in meeting early withdrawals. But, as in Allen and Gale (2000), these links can also become a channel for contagion. A bank may become insolvent because of poor returns on its investments, and the response of depositors and the structure of interbank links determine whether the failure of one bank forces the early liquidation of others. In particular, Freixas, Parigi, and Rochet (2000) contrast two topologies: a ring network and a completely connected network. They show that depositors are less likely to run in the ring network because in this case an insolvent bank can pass a greater fraction of its losses to other banks. In the terminology of Freixas, Parigi, and Rochet (2000), depositors impose greater market discipline on banks in the completely connected network, making this configuration less resilient.

Because interbank lending allows for the sharing of liquidity risk, it creates moral hazard: after a bank reduces its vulnerability to a liquidity shock, it has an incentive to take on greater risk in its investments. Brusco and Castiglionesi (2007) modify the Allen and Gale (2000) model to incorporate this effect. (See Zawadowski 2013 for a different type of network model with moral hazard.) The main source of risk in Brusco and Castiglionesi (2007) comes from banks' investments in risky assets (through moral hazard), and more links between banks create more opportunities for bad investment outcomes at one bank to spill over to other banks. In their model, a completely connected network is less stable than a ring network.

4.3 The Gai, Haldane, and Kapadia (2011) model

Like the model of Allen and Gale (2000), the model of Gai, Haldane, and Kapadia (2011) examines the interaction between liquidity shocks and interconnectedness. The two papers are representative of two broader strands of literature. The Allen-Gale framework seeks to endogenize the actions of banks and outside depositors; doing so generally limits the analysis to simple networks. Gai, Haldane, and Kapadia (2011) (briefly, GHK) consider more general networks and more varied bank balance sheets. To accommodate this generality, their analysis relies on simple rules for bank behavior and on numerical experiments for conclusions. Their simulation framework offers the flexibility to compare policy alternatives. It builds on an analytic framework in Gai and Kapadia (2010), but the earlier model describes default cascades rather than funding runs.

To describe their model, we return to the example of Figure 3 under the funding run interpretation in Section 3.3 – the second of the two interpretations given there. In Figure 4,

we zoom in on the balance sheets of nodes B and C. Bank B has extended a loan of 100 to bank C. It has 100 in fixed assets and 10 in liquid assets. This terminology follows GHK and introduces a distinction between two types of outside assets in Figure 2: liquid assets can be easily disposed of, but prohibitive liquidation costs preclude the selling of fixed assets. Bank B has borrowed 180 from other banks. It also has 30 in what we refer to in Figure 4 as stable funding – this term is not used in GHK. Stable funding includes both outside deposits and equity or net worth. We combine these two because they function equivalently in the model – neither gets withdrawn in a funding run.

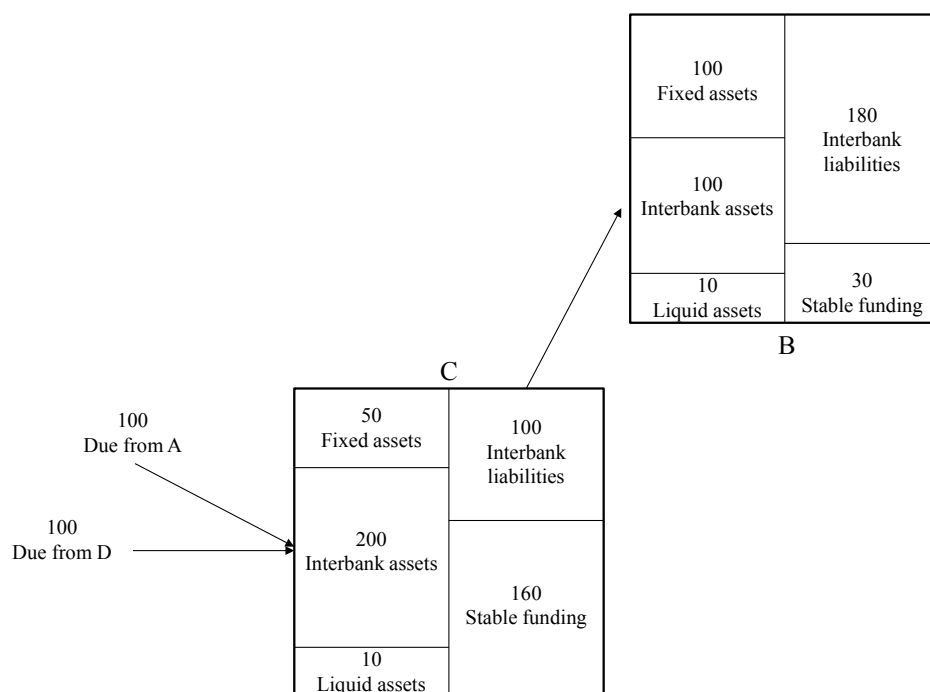


Figure 4 Interlinked balance sheets for a model of liquidity hoarding

Suppose now that bank B withdraws 50 of its funding from bank C. As discussed in GHK, bank C cannot quickly make up the loss of funding by increasing its outside deposits or raising new equity – that is, by increasing its stable funding. It may be unable or unwilling to find another bank to borrow from because of the negative signal that the need to borrow would send to the market. If it sold its fixed assets, it would incur liquidation costs. Bank C’s best option is to withdraw funding from banks A and D, thus propagating the initial shock through the network.

In our discussion of Figure 3, we assumed a minimal withdrawal by each bank. In Figure 4, the minimal response would have bank C pull 20 in funding from each of banks A and D, which combined with C's liquid assets would make up the shortfall of 50. But GHK observe that in a crisis – notably in the most recent one – banks respond to shocks by *hoarding liquidity*, meaning that they reduce their lending more than they need in order to cover their own funding shortfall. In particular, banks may view an initial shock as a warning of greater scarcity to come and, as a precaution, increase their buffer of liquid assets by reducing their lending. In the setting of Figure 4, bank C may respond to its initial loss of 50 in funding by pulling up to 200 in loans from banks A and D. In GHK, the extent of liquidity hoarding is controlled by a multiplier, and the larger this value, the greater the amplification of the initial shock. For most of their results, GHK assume a maximum multiplier, that is, a bank withdraws *all* of its lending to other banks if its buffer of liquid assets is depleted.

Figure 4 omits important features from the more detailed balance sheets in GHK. In particular, their interbank loans include repurchase agreements, or “repo” transactions. These are secured loans, backed by assets pledged as collateral. Repo transactions are very short term (typically one to seven days) and therefore easily withdrawn; indeed, the growth in short-term funding prior to the crisis is widely viewed as a major contributor to the turmoil of 2007-2009.

A repo transaction carries a “haircut,” which is the difference between the amount lent and the value of the collateral. For example, in Figure 4, banks A and D might each provide 105 in collateral to borrow 100 from bank C, and bank C may then pledge 105 of this collateral to borrow 100 from bank B. In GHK, liquidity shocks in the interbank lending market are expressed through wider haircuts: if bank B widens the haircut it applies to bank C, this reduces the amount C can borrow from B, holding fixed C's supply of collateral assets.⁴ Wider haircuts may be idiosyncratic, reflecting heightened concerns about an individual bank, or they may be systematic, reflecting general concerns about liquidity or deteriorating quality of collateral assets (such as mortgage bonds). GHK examine both types of shocks. They apply these shocks to randomly generated networks to evaluate the consequences. There are no defaults in GHK; in contrast to most of the financial networks literature, they focus on the likelihood and severity of what might be called a funding crisis. More specifically, they estimate the probability that at least 10 percent of banks will need to

⁴ See Gorton (2010) for a discussion of the role of widening repo haircuts in 2007.

withdraw funding from other banks and, conditional on reaching this threshold, the fraction of banks affected.

Their experiments consider many variations, with a particular focus on the effect of increasing connectivity among banks, which they measure through the average number of banks from which each bank borrows. They find that increasing connectivity first increases and then decreases the probability of a funding crisis; however, the severity of a crisis, conditional on the occurrence of a crisis, consistently increases with connectivity.

Despite major differences in models, this finding is broadly consistent with Allen and Gale (2000) and several other papers. It reflects the trade-off we highlighted earlier: Greater interbank links serve both to dilute shocks and to propagate shocks. The GHK model allows a systematic examination of this effect under various assumptions. In contrast, the configurations in Allen and Gale (2000) are much simpler, but they illustrate the interplay between network topology and the correlation in initial funding shocks across banks.

A few features of the GHK model are worth mentioning because they raise issues that will recur in our discussion. In their numerical experiments they hold banks' total level of interbank borrowing fixed as they increase the number of links, and each bank's total interbank borrowing is evenly divided among the banks from which it borrows. Increasing the number of links between banks thus reduces the amount borrowed over each link. An alternative notion of connectivity, to which we return in Section 7.2, counts dollars borrowed from banks rather than the number of banks from which they are borrowed, and this can lead to rather different conclusions. The random networks in GHK generate banks that are ex ante identical; we will argue that heterogeneity is important – both empirically and for model implications.

5. Basic Network Models

In this section we describe a model due to Eisenberg and Noe (2001) that forms the basis of much subsequent work on contagion in financial networks. We start by assuming that the network of obligations is given; the object is to study how shocks to particular institutions or assets propagate through the network. The model has four key ingredients (see Figure 2): i) a set of n nodes $N = \{1, 2, \dots, n\}$ representing different financial entities such as banks, broker-dealers, insurance companies, and the like; ii) an $n \times n$ liabilities matrix $\bar{P} = (\bar{p}_{ij})$ where

$\bar{p}_{ij} \geq 0$ represents the *payment due* from node i to node j at the end of the current period; iii) a vector $c = (c_1, c_2, \dots, c_n) \in R_+^n$ where $c_i \geq 0$ represents the total payments due from nonfinancial entities to node i ; iv) a vector $b = (b_1, b_2, \dots, b_n) \in R_+^n$ where $b_i \geq 0$ represents the total payments due from node i to nonfinancial entities. The numbers c_i and b_i will be called i 's *outside assets* and *outside liabilities* respectively.⁵

The *asset* side of node i 's balance sheet is given by $c_i + \sum_{j \neq i} \bar{p}_{ji}$, and the *liability* side by $\bar{p}_i = b_i + \sum_{j \neq i} \bar{p}_{ij}$. The node's *value* is the total value of its assets, namely, $c_i + \sum_{j \neq i} \bar{p}_{ji}$. The node's *net worth* is

$$w_i = c_i + \sum_{j \neq i} \bar{p}_{ji} - \bar{p}_i. \quad (1)$$

We shall assume that initially the net worth of every node is strictly positive. We then examine what happens when the outside assets suffer a negative shock, possibly causing the net worth of one or more nodes to become negative. A shock realization will be represented by an n -vector $x = (x_1, x_2, \dots, x_n)$ where $0 \leq x_i \leq c_i$ for $1 \leq i \leq n$. The *direct effect* of the shock x is to reduce the net worth of each node i to the value

$$w_i(x) = c_i - x_i + \sum_{j \neq i} \bar{p}_{ji} - \bar{p}_i. \quad (2)$$

If $w_i(x)$ is negative, node i *defaults*. We shall assume that all debt obligations have equal priority, and that in case of default the assets are distributed to the creditors in proportion to the nominal amounts they are owed. (The equity holders are wiped out, since their claim is on the firm's net worth provided the latter is positive.) The problem is to determine a consistent set of payments conditional on the initial shock. We have already encountered this problem in our illustration of a default cascade using Figure 3 in Section 3.3.

⁵ The notion of outside assets and outside liabilities is a variant of the Eisenberg-Noe framework that is common in the literature (see among others Elsinger 2009 and Glasserman and Young 2015). It can be seen as a special case of Eisenberg-Noe by positing a fictitious node that holds all of the outside liabilities as assets and all of the outside assets as liabilities.

To this end let us define the *relative liabilities matrix* $A = (a_{ij})$ to be the $n \times n$ matrix with entries

$$\begin{aligned} a_{ij} &= \bar{p}_{ij}/\bar{p}_i \text{ if } \bar{p}_i > 0 \\ a_{ij} &= 0 \quad \text{if } \bar{p}_i = 0 \end{aligned} \quad (3)$$

The term a_{ij} represents the proportion that i 's obligations to node j represent of its total liabilities to all other nodes and to the external sector. In what follows we shall assume that from every node i there exists a chain of positive obligations to *some* node j such that $b_j > 0$. This seems like a reasonable assumption for real-world financial networks; it also guarantees that the spectral radius (the absolute value of the largest eigenvalue) of the relative liabilities matrix A has is less than one. This is useful in characterizing the equilibria of the system, a subject that we turn to next.

5.1 Clearing Payments

Suppose that the outside assets suffer a shock x . We shall say that node i suffers a *direct default* if $x_i > w_i = c_i + \sum_{j \neq i} \bar{p}_{ji} - \bar{p}_i$. The pro rata allocation rule implies that i 's payments are proportional to the various claims against i 's assets. The complication is that the value of i 's assets depend on the payments made by others *to* i . Thus i 's payment to j (conditional on x) is given by the expression

$$p_{ij}(x) = a_{ij}(c_i - x_i + \sum_{k \neq i} p_{ki}(x)). \quad (4)$$

where $\sum_{k \neq i} p_{ki}(x)$ is the sum of payments to i from the other nodes in the system. In particular, the payment $p_{ki}(x)$ will be less than \bar{p}_{ki} if node k is also in default. We shall say that the payments $p_{ij}(x)$ are *consistent* if

$$\forall i, j, \quad p_{ij}(x) = \bar{p}_{ij} \wedge a_{ij}(c_i - x_i + \sum_{k \neq i} p_{ki}(x)). \quad (5)$$

This condition can be expressed in a more compact form as follows. Let $p_i(x)$ denote the *total* payment from i to all other nodes in the financial system plus its payments to the outside sector. Let $p(x) = (p_1(x), \dots, p_n(x))$ be the corresponding payments vector. These payments are consistent if

$$\forall i, \quad p_i(x) = \bar{p}_i \wedge (c_i - x_i + \sum_j a_{ji} p_j(x)) . \quad (6)$$

Expression (5) follows from (6) by applying the pro rata distribution rule in case of default. Any vector $p(x) \in R_+^n$ satisfying (6) is called a *clearing vector*. As shown by Eisenberg and Noe (2001), a clearing vector always exists; moreover under our assumption on the matrix A the clearing vector is unique for any shock realization x .

Clearing vectors can be computed via the following recursive procedure known as the *fictitious default algorithm* (Eisenberg and Noe, 2001).⁶ Fix a shock realization x and define the mapping $\Phi : R_+^n \rightarrow R_+^n$ as follows:

$$\forall i, \quad \Phi_i(p) = \bar{p}_i \wedge (c_i - x_i + \sum_j p_j a_{ji}) . \quad (7)$$

Starting with $p^0 = \bar{p}$ let

$$p^1 = \Phi(p^0), p^2 = \Phi(p^1), \dots \quad (8)$$

This iteration yields a monotone decreasing sequence $p^0 \geq p^1 \geq p^2 \dots$. Since the sequence is bounded below by the zero vector, it has a limit, say $p' = p'(x)$. Since Φ is continuous, $p'(x)$ satisfies (6), hence it is a clearing vector.

In Section 3 we considered the problem of finding a clearing vector for the example in Figure 3. The calculation given there represents one iteration of this algorithm. The reader may verify that successive application of the algorithm leads to the solution shown in Figure 5. Here the incoming payments equal the outgoing payments at every node, and the payments from each node are distributed in proportion to the nominal amounts owed. Thus we obtain a mutually consistent (equilibrium) set of payments; moreover under our assumptions it is unique.

⁶ This procedure is analogous to solving a Leontief input-output system via iteration.




Figure 5. Clearing payments when the outside sector pays 40 instead of 160 to bank C.

The existence of a clearing vector is a consequence of a general fixed-point theorem on lattices due to Tarski (1955). Consider the complete lattice $L = \{x \in R^n : 0 \leq x \leq \bar{p}\}$ with the partial order $x \leq y \Leftrightarrow x_1 \leq y_1, \dots, x_n \leq y_n$. Tarski's theorem states that *every monotone function $F : L \rightarrow L$ has at least one fixed point; moreover it has a greatest fixed point and a least fixed point*. This result immediately implies the existence of a vector satisfying (7), because the mapping $\Phi(p)$ is monotone nondecreasing in p . This result underpins various extensions of the Eisenberg-Noe model, as we shall see in subsequent sections.

5.2 Extensions of the Basic Model

The preceding model provides a simple and transparent way of representing the interlocking assets and liabilities in the financial system. In particular the model illustrates how defaults can be transmitted from one node to another, potentially amplifying the impact of an initial shock. Nevertheless the model is oversimplified in several respects. One limitation is that when default occurs, it is assumed that the available assets are distributed pro rata to the creditors without further impairment. In practice, defaults can seriously disrupt the distribution of assets, and entail significant costs that further reduce the value of the assets to the downstream claimants. These costs amplify the initial shock as they course through the financial network. In the next section we shall show how to incorporate bankruptcy costs into the model in a general way; we shall revisit this topic in more detail in Section 8.7.

A second limitation is that the model puts all claims on the same footing in case of default. In reality banks have a great variety of claims against one another, including secured and

unsecured loans, derivatives, equity stakes, and other financial instruments that have different degrees of priority when default occurs. In Section 5.4 we shall show how to extend the model to incorporate interbank equity claims as well as interbank debt obligations; other types of claims and levels of seniority can be accommodated in similar fashion.

A third limitation of the model is that it treats financial crises as if they originated solely from a reduction in payments. In fact there are many other forms of financial contagion where interconnections play a key role but the source of contagion does not arise from payment shortfalls. One example is the funding run discussed in Section 4.3, which can trigger a widespread contraction of credit in the financial system. Another example is a fire sale. Suppose that one bank must sell some of its illiquid assets to raise cash or meet regulatory capital requirements. The price of these assets will then be driven down, which has a negative impact on the balance sheets of other banks, triggering further sales and a downward price spiral. The basic mechanics of this model are discussed in Section 5.5 and revisited in Section 11.

Yet another source of financial contagion is a loss of confidence in the creditworthiness of particular institutions. Doubts about the quality of one bank's assets will have a negative impact on the perceived value of its liabilities to other banks, which will have a negative impact on *their* balance sheets, and so forth. Thus, what began as a loss of confidence in one bank (but not its outright default) can have knock-on effects that cascade through the financial system. In our view this loss-of-confidence effect, which we discuss further in Section 5.6, is one of the most important channels through which the financial network amplifies systemic risk in practice.

5.3 Bankruptcy Costs and Asset Recovery Rates

When a bank defaults there will typically be delays in paying its creditors in addition to legal and administrative costs. Thus only a fraction of the firm's assets will be available for distribution to the creditors. These costs increase both the magnitude and the likelihood of default cascades. Rogers and Veraart (2013) show how such costs can be incorporated into the Eisenberg-Noe framework using a recovery function that drops discontinuously at the default boundary and then decreases linearly with the amount of assets available. More generally, one can posit a recovery function $r(\alpha, \bar{p})$ that represents the amount paid to

creditors as a function of the bank's assets α and its nominal obligations \bar{p} . We assume that r is monotone non-decreasing in both of its arguments, and that

$$0 \leq r(\alpha, \bar{p}) \leq \alpha \text{ if } \alpha < \bar{p} \quad (9)$$

$$r(\alpha, \bar{p}) = \bar{p} \text{ if } \alpha \geq \bar{p}$$

Let $[\bar{p}_{ij}]_{1 \leq i \neq j \leq n}$ be the matrix of nominal payment obligations between financial nodes, and let $[p_{ij}(x)]_{1 \leq i \neq j \leq n}$ be the matrix of realized payments between nodes following a shock x . As before let $p_i(x)$ denote the total realized payments from i to all other entities including the outside sector. The realized value of i 's assets equals $c_i - x_i + \sum_j a_{ji} p_j(x)$, hence the clearing condition takes the form

$$p_i(x) = r(c_i - x_i + \sum_j a_{ji} p_j(x), \bar{p}_i) . \quad (10)$$

Under our assumptions the mapping implied by (10) is monotone non-decreasing in p (although it need not be continuous). Hence Tarski's theorem implies that at least one solution to (10) exists. Specific examples of recovery functions and empirical estimates of the amounts recovered in practice will be discussed in Section 8.7.

5.4 Claims of Different Seniority

The Eisenberg-Noe model assumes a pro rata allocation of payments in case of default. In practice, different claims may have different seniority. This holds, for example, when financial institutions hold equity stakes in other financial institutions. These differ from debt obligations because they constitute residual claims on assets after creditors have been paid, and there is no nominal limit to how much a claim is worth. Cross-holdings have network spillover effects that can be modelled by an extension of the Eisenberg-Noe framework that involves solving a nested pair of fixed point problems (as in Elsinger 2009 and related work by Gournieroux, Heam, and Monfort 2013). We briefly outline the argument here. Let $\theta_{ij} \in (0,1)$ be the fraction of bank i 's equity owned by bank j . Thus j has a claim on the net worth of i provided the latter is nonnegative. This constitutes an asset on j 's balance sheet that can be used to discharge j 's current debt obligations if needed. To avoid degeneracy let us assume that a positive fraction of each bank is owned by the outside sector, that is, the

matrix $\Theta = [\theta_{ij}]$ is strictly row substochastic.⁷ Given a shock x , let $p \equiv p(x)$ be the resulting payments vector. We may define the “interim” net worth of each node i as follows:

$$\forall i, \quad w_i(p) = [c_i + \sum_j a_{ji} p_j + \sum_j \theta_{ji} (w_j(p) \vee 0)] - \bar{p}_i. \quad (11)$$

The first three terms in (11) represent i 's assets while the last term represents i 's nominal liabilities. (If $w_j(p)$ is negative for some j , then i 's equity claim on j is worthless.) It can be shown that for every $p \in [0, \bar{p}]$ there exists a unique vector of net worths $w(p)$ satisfying (11); moreover $w(p)$ is monotone increasing in p (Elsinger 2009).

To close the model, we need to find a vector of payments $p \in [0, \bar{p}]$ that is a fixed point of the mapping

$$p = ([w(p) + \bar{p}] \vee 0) \wedge \bar{p}. \quad (12)$$

Note that (12) involves two fixed points, one nested inside the other. Because the fixed point $w(p)$ is monotone nondecreasing in p , a solution of (12) always exists, although it need not be unique (Elsinger 2009, Theorem 1).

5.5 Fire Sales

Network models can be applied to situations where contagion arises not only through payment shortfalls per se, but also from spillover effects that arise from common exposures. Suppose for example that to meet capital requirements a given bank has to sell illiquid assets to shore up its balance sheet by increasing cash reserves. Suppose further that other banks are exposed to the same asset. Selling by the first bank puts downward pressure on its price, which has a negative impact on the balance sheets of the other banks, forcing them to raise cash as well. The resulting contagion is known as a fire sale, and arises when banks are exposed to the same asset classes, and value their assets at current market prices instead of

⁷ In other words, each row sum is strictly less than one. In fact it suffices that there is no subgroup of banks all of which are completely owned by other banks in the group (Elsinger 2009).

historical costs.⁸ Duarte and Eisenbach (2014) estimate the impact of these effects empirically.

Here we outline a model of this process due to Cifuentes et al. (2005). Assume that the assets of bank i consist of three parts: cash reserves e_i , a quantity of illiquid assets q_i whose current price is θ , and payments from other banks $\sum_{k \neq i} p_{ki}$. Its liabilities consist of interbank obligations and obligations to depositors. The asset side of the balance sheet can be written as follows:

$$\theta q_i + e_i + \sum_{k \neq i} p_{ki} . \quad (13)$$

Suppose that, due to a liquidity shock, bank i must increase its cash reserves by selling some of its holdings of the illiquid asset. For simplicity let us assume that all banks hold the same illiquid asset. (The situation where banks have overlapping exposures to multiple assets is discussed in Section 11). Let $q'_i \leq q_i$ be the amount that bank i is forced to sell, and assume that its price $\theta(q')$ decreases as q' increases.⁹ The assets of bank i can now be expressed as follows:

$$\theta(q')(q_i - q'_i) + (e_i + \theta(q')q'_i) + \sum_{k \neq i} p_{ki}(q') . \quad (14)$$

The first term is the value of the remaining illiquid assets, the second term is the new and higher amount of cash reserves, and the third term represents the current payments from other banks (which may be affected by the forced sales).

Cifuentes et al. (2005) posit a capital requirement that forces banks to raise more cash the lower the price of the illiquid asset. These forced sales further depress the price, leading to a downward price spiral and possibly to outright default by some banks. Under appropriate regularity conditions this process has an equilibrium set of sales q' and payments p ; the argument is analogous to the preceding cases and depends on the fact that the price impact function $\theta(q')$ is monotone decreasing in q' .

⁸The potentially destabilizing effects of mark-to-market accounting are discussed by Sapra, Plantin and Shin (2008) and Allen and Carletti (2008)..

⁹ Cifuentes et al. (2005) assume that $\theta(q')$ is a negative exponential, but a similar analysis holds for many other functional forms, as demonstrated in Chen, Liu, and Yao (2014).

5.6 Mark-to-market Valuations and Crises of Confidence

The preceding examples illustrate how the basic logic of the Eisenberg–Noe framework can be adapted to model many different types and sources of network contagion. In this section we show that a similar framework can be used to model, in reduced form, changes in market perceptions about the quality of banks’ balance sheets. This framework shows how a decline of confidence in the creditworthiness of a particular institution can spread through the system and develop into a general crisis of confidence.

Consider a directed network on n nodes, where each node represents a financial or nonfinancial (“outside”) institution, and a directed edge $i \rightarrow j$ corresponds to a claim that node j has on node i . Let \bar{v}_{ij} denote the nominal value of such a claim and let v_{ij} denote its current mark-to-market value, where $0 \leq v_{ij} \leq \bar{v}_{ij}$. Let $V = [v_{ij}]_{1 \leq i, j \leq n}$ be the value matrix. The balance sheet of node i consists of the column vector of i ’s assets $v_{\cdot i}$ and the row vector of i ’s liabilities v_i . Let us posit a function $\phi_i : R_+^n \rightarrow R_+^n$ that maps the value of i ’s assets to the value of its liabilities. We assume that ϕ_i is nondecreasing and that $0 \leq \phi_i(v_{\cdot i}) \leq \bar{v}_i$ whenever $0 \leq v_{\cdot i} \leq \bar{v}_i$. The matrix V of values is *consistent* if $0 \leq V \leq \bar{V}$ and $\phi_i(v_i) = v_i$ for every i . (The existence of such a matrix follows directly from Tarski’s theorem.)

This framework has direct application to situations where contagion is triggered by changes in market perceptions about the creditworthiness of particular institutions, and the values of their assets and liabilities are marked to market. Suppose, for example, that there is a loss of confidence in the balance sheet of a given institution i (as happened in the case of Lehman Brothers). This initial loss of confidence causes a decline in the market value of i ’s liabilities (as modelled by the function ϕ_i), and thus a decline in the asset values of i ’s creditors, and hence to a general decline in asset values throughout the network. Furthermore these declines can lead to the outright default of some institutions, even though no one defaulted to begin with. This application illustrates how the basic logic of the Eisenberg–Noe model can be extended to more general sources of contagion than those arising from simple payment shortfalls.

6. Systemic Losses and Systemic Risk

The preceding sections have identified various ways in which the financial network can amplify shocks that result from payment shortfalls or declines in asset values. What does this

framework tell us about the contribution of the network to systemic risk? There appears to be no accepted definition of this term in the literature. Roughly speaking it describes the possibility of widespread losses in the financial system due to shocks that originate in particular parts of it.

The preceding framework suggests several ways of measuring the losses triggered by such shocks. One is the total loss of bank equity, which is an indicator of the ability of the financial sector to extend credit (see for example Cont, Santos, and Moussa, 2012). Another measure is the aggregate loss inflicted on the nonfinancial sector, that is, the total shortfall in payments from the financial sector to households and nonfinancial firms. Yet a third and more comprehensive measure is the total loss in asset values summed over all entities in the system. We shall call this the *systemic loss in value* (SLV). Using the notation introduced in Section 5.6 this is the total amount by which the value of all claims -- including interbank claims and claims by households and nonfinancial institutions on the financial sector -- are reduced relative to their nominal values:¹⁰

$$SLV = \sum_{0 \leq i, j \leq n} (\bar{v}_{ij} - v_{ij}). \quad (15)$$

In the setting where values are identified with clearing payments, SLV is the total amount by which payments are reduced as the result of an exogenous shock. Specifically, suppose that $x = (x_1, \dots, x_n)$ represents a shock to the outside assets of the n financial nodes and $\bar{p} = (\bar{p}_1, \dots, \bar{p}_n)$ represents the nominal payments due from these nodes to each other and to the external sector. Let $p(x)$ be the resulting clearing vector. Then

$$SLV = \sum_{1 \leq i \leq n} x_i + \sum_{1 \leq i \leq n} (\bar{p}_i - p_i(x)). \quad (16)$$

The first term in (15) is the direct loss in asset values by the financial sector, while the second term is the indirect loss in values due to reduced payments to both the financial and nonfinancial sectors. (Recall that $p_i(x)$ represents i 's total payments to all entities.) This measure treats losses on interbank obligations in the same way that it treats losses on obligations between financial and nonfinancial actors. As the discussion in Section 3 shows, contractual obligations between financial firms play a critical role in managing liquidity

¹⁰ A variant of this idea was proposed in Glasserman and Young (2015).

needs, intermediating between lenders and borrowers, and diversifying risk exposures. These obligations can be thought of as intermediate goods that have economic value to the contracting parties, and their impairment corresponds to genuine economic losses.

In addition, financial losses have other consequences for the real economy. Credit is less available for funding new investment projects; moreover, existing investments may have to be liquidated early in order to meet short-term obligations, thus leading to inefficiencies. Second, bankruptcy carries significant administrative and legal costs (a topic we shall return to in Section 8), Third, financial losses have an impact on household balance sheets with consequent reductions in consumption and underutilization of productive capacity in the economy at large. Although there has been some research attempting to quantify these effects,¹¹ much remains to be done in pinning down the relationship between financial losses and the resulting reduction in economic welfare.

Finally, we should emphasize that uncertainty in shock sizes creates uncertainty in losses. In other words, we should think of a systemic loss measure as having a probability distribution, and it is important to understand how the network of financial obligations transforms a distribution of initial shocks into a distribution of system-wide losses. Such a framework will be discussed in the next two sections.

7. Measures of Vulnerability and Contagion

We turn first to the problem of assessing the potential contagiousness of different banks given their individual characteristics – such as size, leverage, and asset quality – as well as their position in the financial network. Although there is a growing literature on the susceptibility of different network topologies to contagion, much of this literature does not separate the impact of the topology from differences in balance sheet characteristics that make some banks inherently more contagious than others. In this section we shall focus on how to measure the inherent contagiousness of a bank, and the inherent vulnerability of a bank to contagion, using just individual balance sheet information about them. It turns out that some useful conclusions can be drawn from individual-level data without knowing the details of the banks' position in the network or the probability distribution that governs shocks to their assets. These topics will be taken up in subsequent sections.

¹¹ See among others Gertler and Kyotaki (2010), Adrian and Boyarchenko (2012), Chen, Iyengar, and Moallemi (2013), Bassett et al. (2014), and Brunnermeier and Sannikov (2014).

7.1 Pairwise Measures

To fix ideas let us begin by focusing on a particular pair of banks, i and j , one of which has obligations to the other. They are embedded in a larger network, the details of which are left unspecified (see Figure 6).




Figure 6. Node i owes \bar{p}_{ij} to node j . The outside assets of node i are hit by a shock $x_i \leq c_i$

Let $X_i \in [0, c_i]$ be a random shock with realization $X_i = x_i$ that reduces the value of i 's outside assets to $c_i - x_i$ and its net worth to $w_i - x_i$. (Note that the latter may be negative.) The amount of the loss that spills over onto j is $a_{ij}(x_i - w_i)_+$, which reduces j 's net worth. In particular i 's default triggers j 's default if

$$a_{ij}(x_i - w_i)_+ > w_j. \quad (17)$$

Recalling that x_i is at most c_i , we can measure j 's *vulnerability to i* by the ratio

$$a_{ij}(c_i - w_i) / w_j = a_{ij}(\lambda_i - 1)w_i / w_j, \quad (18)$$

where $\lambda_i = c_i/w_i$ is i 's *outside leverage*. If the ratio in (18) is less than unity then j is *relatively immune* to a shock from i . Note that this condition does not guarantee *complete* immunity, because simultaneously j might be weakened by shocks to its own outside assets or by shocks transmitted from nodes other than i . Furthermore a shock to node i could ripple through the financial network, triggering a cascade of defaults that *eventually* take out node j . In this sense the criterion in (18) is only a first-order measure of vulnerability that does not take into account the transmission of one or more shocks in the full network. This situation will be considered in more detail in Section 8.

We can rewrite the relative immunity condition as follows. Recall that $a_{ij} = \bar{p}_{ij} / \bar{p}_i$. Hence j is relatively immune to a shock from i if $\bar{p}_{ij}(\lambda_i - 1) / w_j < \bar{p}_i / w_i$. The *total leverage* of i (its assets divided by its net worth) equals $\lambda_i^* = (\bar{p}_i + w_i) / w_i$. Clearly λ_i^* is at least as great as i 's outside leverage λ_i . It follows from (18) that j is relatively immune to a shock from i if

$$\bar{p}_{ij} / w_j < (\lambda_i^* - 1) / (\lambda_i - 1). \quad (19)$$

This measure is related to a test of j 's vulnerability that has been put forth by the Basel Committee on Banking Supervision (2014) through its "large exposure limit." Consider j 's maximum exposure to any given counterparty divided by j 's net worth

$$\max_i \bar{p}_{ij} / w_j. \quad (20)$$

If this ratio is less than unity, the only way that j can default is if more than one of its creditor banks fails completely and/or j suffers substantial shocks to its own assets, thus reducing its net worth. In fact the Basel Committee's large exposure limit caps a bank's exposure to any single counterparty at 25 percent of the bank's Tier 1 capital; for global systemically important banks (G-SIBs) a more stringent cap of 15 percent has been proposed (Basel Committee 2014).

7.2 General Measures: Financial Connectivity and the Contagion Index

This approach can be extended to measure the potential vulnerability of *subsets* of nodes to the default of a single node. Define the *financial connectivity of node i* , β_i , to be the fraction of i 's obligations that are owed to the financial sector, that is, $\beta_i = (\bar{p}_i - b_i) / \bar{p}_i$. This number is of fundamental importance in assessing the extent to which i contributes to systemic risk as we shall see in subsequent sections. For the moment let us consider how it can be used to assess vulnerability and contagion in a partial setting without specifying the details of the network or the distribution of shocks. Fix a node i and a subset of nodes S that does not include i . When can a shock to i 's outside assets cause all the nodes in S to default purely through network spillover effects (including spillovers from intermediate nodes) but without any further amplification from bankruptcy costs? Evidently this is *impossible* if

$$\sum_{j \in S} w_j > w_i \beta_i (\lambda_i - 1). \quad (21)$$

The right-hand side of (21) is called *i*'s *contagion index* (Glasserman and Young 2015). A bank with a high contagion index poses a systemic risk in the sense that it is more susceptible to failure (due to high leverage), its failure has large consequences (due to its size), and its failure has a potentially large impact on the rest of the financial system (due to its high financial connectivity).

8. Shock Distributions

The preceding analysis of default cascades is based on worst-case scenarios and does not assume an underlying shock distribution. In this section we introduce shock distributions explicitly in order to study the likelihood of default cascades and the losses that they generate. Let $F(x_1, x_2, \dots, x_n)$ be a joint c.d.f. of shocks to the outside assets of the various nodes in the financial system. In practice these shocks may be positively correlated due to common exposures, and thus the probability of multiple bank failures will generally be higher than if the shocks were independently distributed. Note however that systemic risk from common exposures and correlated shocks is present *whether or not* the banks are interconnected. To address the question of how much the financial network amplifies systemic losses, we will first consider the case of independent shocks. It turns out that many of these estimates are also valid for the case of correlated shocks.

8.1 Beta distributions

To be specific, let us assume that the shocks to the various nodes are *independent* and the size of the shock x_i is proportional to the value of the outside assets c_i . In this case we can write

$$F(x_1, x_2, \dots, x_n) = \prod_i H_i(x_i / c_i), \quad (22)$$

where $H_i(y)$ is a c.d.f. with domain $y \in [0, 1]$. A particularly convenient and flexible family of distributions are the monotone beta distributions of form

$$H_i(y) = 1 - (1 - y)^{q_i}, \quad q_i \geq 1. \quad (23)$$

The parameter q_i can be interpreted as an asset *quality index*: the larger the value of q_i , the smaller the expected loss and the variance of the losses.¹² Beta distributions can be used to approximate more complex distributions, such as the Gaussian copula model used for setting Basel capital requirements. The Gaussian copula model describes the distribution of losses in a loan portfolio using two parameters – an average probability of default PD , and a parameter ρ describing the correlation between loans,

$$F(x) = \Phi \left(\frac{\sqrt{1-\rho}\Phi^{-1}(x) - \Phi^{-1}(PD)}{\sqrt{\rho}} \right) \quad (24)$$

Figure 7 compares the Gaussian copula with $PD = 5\%$ and $\rho=0.13$ with the monotone beta distribution when $q_i = 19$.




Figure 7. Densities for beta with c.d.f. $F(x) = 1 - (1-x)^{19}$ and the Gaussian copula distribution with probability of default $PD = 0.05$ and $\rho = 0.13$.

The *probability* that i defaults due to losses on its outside assets is given by the expression

$$\delta_i = P(X_i > w_i) = 1 - H_i(1/\lambda_i). \quad (25)$$

In the case of monotone beta distributions we therefore obtain the following simple formula relating leverage λ_i , asset quality q_i , and probability of default δ_i :

¹² The expected loss is $1/(1+q_i)$ and the variance is $1/(2+q_i)(1+q_i)^2$, both of which are monotone decreasing in q_i .

$$\delta_i = (1 - 1/\lambda_i)^{q_i}. \quad (26)$$

8.2 Estimating the Probability of Default Cascades

We now employ this framework to estimate the probability of default cascades. To keep matters simple we shall continue to assume (for the time being) that there is no amplification due to other mechanisms such as bankruptcy costs or common exposures. These will be discussed in subsequent sections.

Consider the simple chain of obligations shown in Figure 8. To be concrete we can think of A as a mortgage company whose assets consist of outstanding mortgages to households worth 100, and whose liabilities consist of short-term loans from investment bank B worth 95. Thus A's net worth is 5. Bank B has liabilities of 60 to outside depositors and 30 in obligations to bank C, so its net worth is also 5. Bank C owes 49 to depositors and is owed 20 by nonfinancial firms so its net worth is 1.




Figure 8. A chain of obligations. The net worth of each node is in bold.

Suppose that node A suffers a loss that is large enough to knock out B. Then A must lose its net worth of 5 plus B's net worth of 5. By assumption the probability that A loses at least 5 is $\delta_A = (1 - 5/100)^q$, where $q \geq 1$ depends on the quality of A's assets. Although we do not know the value of q , we do know that whenever $a > 1$ and $0 < b < 1$ the following inequality holds: $(1 - b)^a > 1 - ab$. If $ab < 1$ we can conclude that $(1 - ab)^q < (1 - b)^{aq}$ for all $q > 0$. Therefore the probability that A loses enough to topple B is $(1 - 10/100)^q < \delta_A^2$. Under normal conditions we would expect δ_A to be small, say less than 1 percent. Thus, even though we do not know the value of q , we can deduce that it is highly improbable that A's losses will be enough to topple B unless there is loss amplification from other sources.

Next let us consider the probability that A loses enough to knock out both B and C. Clearly it suffices that A lose at least 13: it takes a loss of 10 to topple A and B, plus another 3 to knock out C. (Note that only one-third of B's losses spill over onto C's balance sheet, because only one-third of B's obligations are to C.) The probability of this event is $(1-13/100)^q < \delta_A^{2.6}$ which again is very small under normal conditions.

Of course there may be times, such as the run-up to the crisis in 2008-09, when risk-taking behaviour leads to elevated probabilities of default at the individual bank level, in which case δ_A may not be particularly small. However, the point of the preceding analysis is not to estimate the probability of network spillovers in an absolute sense, but to illustrate how they compare to the probability of direct default by individual banks. We should also emphasize that this is a partial analysis that does not account for simultaneous shocks to the assets of different nodes in the network, nor does it account for cascades that might emerge from other parts of the network. In the next few sections we examine how to estimate the systemic impact from all of these effects taken together.

8.3 Estimating Systemic Losses

As we have argued earlier, a comprehensive measure of systemic impact is the decline in asset values across the entire system, that is, the systemic loss in value defined in Section 6. To illustrate this concept consider again the example in Figure 8. Suppose that node A suffers a loss of 25 on payments from its mortgage-holders. Then A's payment to B is reduced from 95 to 75, and B's payment to C is reduced from 30 to 25. Furthermore B's depositors lose 10 and C's depositors lose 4. Thus in this example the systemic loss in value caused by A's default is $25 + 20 + 5 + 10 + 4 = 64$. The network amplifies the initial shock by partially destroying the value of a whole series of obligations between financial actors. In this section we shall show how to bound the expected loss in value when relatively little information about the network topology is available.

For simplicity, we shall employ the payments framework of Eisenberg and Noe. As in Figure 2, each financial node i has nominal outside assets c_i , nominal outside obligations b_i , and nominal obligations to other nodes $\bar{p}_{ij}, j \neq i$. Suppose that the outside assets of the various nodes are simultaneously hit by nonnegative shocks $x = (x_1, \dots, x_n)$ with cumulative distribution function $F(x)$. Unlike in previous sections we shall not require that the shocks be independent.

Let $D(x)$ denote the set of nodes that default given x . Our standing assumption is that from every node there exists a chain of obligations that leads to the outside sector, so the clearing vector $p(x)$ is unique (see Section 5). The resulting *shortfall* in payments by node i is

$$s_i(x) = \bar{p}_i - p_i(x). \quad (27)$$

In particular $s_i(x) > 0$ for all $i \in D(x)$ and $s_i(x) = 0$ for all $i \notin D(x)$. In this setting the *systemic loss in value* as defined in (16) can be written as follows:

$$L(x) = \sum_i x_i + \sum_i s_i(x). \quad (28)$$

The first term $\sum_i x_i$ is the direct loss in value from reductions in payments by the outside sector to the financial sector. The second term $\sum_i s_i(x)$ is the indirect loss in value due to reductions in payments by financial entities to one another and also to the outside sector.

Note that the reduction in payments to the outside sector can be expressed in the form $\sum (1 - \beta_i) s_i(x)$, where $\beta_i \in [0, 1]$ is i 's financial connectivity.¹³

Let A_D be the $|D| \times |D|$ matrix obtained by restricting the relative liabilities matrix A to $D(x)$. Let I_D be the $|D| \times |D|$ identity matrix. Similarly let w_D be the vector of initial net worths of the nodes in $D(x)$ and let s_D be the vector of shortfalls. The following shortfall equation follows directly from (27) and the fact that p is a clearing vector:

$$s_D A_D - (w_D - x_D) = s_D. \quad (29)$$

The logic of this equation can be seen as follows. For each defaulting node $i \in D$, $[s_D A_D]_i$ represents the incoming shortfall in payments to i from other nodes. These shortfalls are reduced by i 's initial net worth w_i and they are augmented by the direct shock x_i .

¹³ Recall that the financial connectivity β_i of node i is the fraction of i 's liabilities owed to other entities in the financial sector, and $1 - \beta_i$ is the fraction owed to nonfinancial entities.

Under our hypotheses there is a positive chain of obligations from every financial node to some node that has positive obligations to the external sector. Hence the spectral radius of A_D is less than unity so $\lim_{k \rightarrow \infty} A_D^k \rightarrow 0$. It follows that $I_D - A_D$ is invertible and

$$[I_D - A_D]^{-1} = I_D + A_D + A_D^2 + \dots \quad (30)$$

8.4 Node Depth

We now introduce a topological measure called *node depth* that plays a key role in the analysis of systemic risk. Intuitively the node depth measures the extent to which each additional dollar of loss at a given node becomes amplified as it cascades through the network, causing additional losses at other nodes. Consider a Markov chain on the default set D with probability transition matrix A_D . Given $i \in D$ compute the expected number of periods u_i that it takes to exit D starting from i . Let $u_i = 0$ for all $i \notin D$. We shall call u_i the *depth of node i in D* . Recalling that D and u depend on the shock vector x , we can write

$$\begin{aligned} \forall i \in D(x), u_i(x) &= [I_{D(x)} + A_{D(x)} + A_{D(x)}^2 + \dots] \cdot \mathbf{1}_{D(x)} \\ \forall i \notin D(x), u_i(x) &= 0 \end{aligned} \quad (31)$$

From the definition of $L(x)$ and (29)-(31) it follows that the systemic loss given a shock x is given by the expression

$$L(x) = \sum_i x_i + \sum_i (x_i - w_i) u_i(x). \quad (32)$$

Thus we see that the node depths $u_D = [I_D + A_D + A_D^2 + \dots] \cdot \mathbf{1}_D$ measure the *amplification* of losses due to interconnections among financial entities. This concept is related to the notion of eigenvector centrality in the networks literature, as we shall see in Section 9.

For the present let us illustrate the concept of node depth with a couple of examples. Suppose that all three nodes in Figure 8 default. The depth of the last node in the chain is $u_C = 1$, because the probability of exiting to the external sector equals one. The number of periods to exit from node B is 1 with probability $\frac{1}{2}$ and $1 + u_C = 2$ with probability $\frac{1}{2}$, hence the depth of node B is $u_B = 1.5$. Finally the depth of A is one more than the depth of B, hence $u_A = 1 + u_B = 2.5$.

In similar fashion one can compute the node depths for the example in Figure 3. Let us assume that all nodes default. Then the transition probabilities together with the node depths are as shown in Figure 9. Since the node depths are monotonically increasing with the default set D , these are upper bounds on the node depths with respect to *any* subset of defaulting nodes.




Figure 9. Transition probabilities and node depths for the example in Figure 3.

8.5 Bounding the Node Depths

We can derive a useful lower bound on the amount of amplification by adapting a concept from the social networks literature called “cohesiveness” (Morris 2000). A set of nodes D is

α -cohesive if $\sum_{j \in D} a_{ij} \geq \alpha$ for every $i \in D$, that is, every node in D has at least α of its

obligations to other nodes in D . The *cohesiveness* of D is the maximum such α , which we denote by α_D . Evidently the probability of exiting the set D in any given period is at most $1 - \alpha_D$. This yields the lower bound

$$\forall i \in D, u_i \geq (1 - \alpha_D)^{-1} . \quad (33)$$

It follows that the more cohesive the default set, the greater the depth of the nodes in the default set and the greater the potential amplification of an initial shock.

Using (32), we can bound the systemic loss $L(x)$ by bounding the node depths from above. Recall that node i 's financial connectivity β_i is the proportion of i 's obligations to other nodes in the financial system. Let $\beta_D = \max\{\beta_i : i \in D\}$. The probability of exiting D in any period is at least $1 - \beta_D$, hence

$$\forall i \in D, u_i \leq (1 - \beta_D)^{-1} . \quad (34)$$

In practice the default set D will depend on the shocks and on the topology of the network that transmits the shocks, neither of which may be known. Nevertheless one can derive a useful upper bound on the node depths with virtually no information about the network or the shocks. Namely let $\beta^+ = \max_{1 \leq i \leq n} \beta_i$ be the *maximum financial connectivity* among all nodes. Then the node depths satisfy the uniform upper bound

$$\forall i, u_i \leq (1 - \beta^+)^{-1} . \quad (35)$$

The number $1 - \beta^+$ is a lower bound on the rate at which losses in the financial system are *dissipated* to the outside sector. For some banking systems this bound can be estimated from publicly available data. Glasserman and Young (2015) estimate the financial connectivities of 90 large European banks using European Banking Authority stress test data, under the assumption that the fraction of in-network liabilities is the same as the fraction of in-network assets. The estimated connectivities range from 0.021 to 0.43 with a median value of .119. In this case $\beta^+ = 0.43$, so none of the node depths can exceed $1/(1 - .43) = 1.76$. Allahrakha, Glasserman, and Young (2015) estimate somewhat lower connectivities for the largest U.S. bank holding companies.

8.6 Bounding Systemic Loss

Although the systemic loss in value $L(x)$ clearly depends on the topology of the network and on the specific shock vector x , we claim that the *expected systemic loss* $\int L(x) dF(x)$ can be bounded for a wide range of shock distributions with minimal information about the network topology. The idea is to compare the expected systemic losses in the presence of the network, with the expected systemic losses when all network connections are severed and the balance sheets of the various nodes are held fixed. We can create a similar financial system with no interconnections as follows: as before each node i has outside assets c_i and outside liabilities b_i . To keep their net worths fixed as we sever connections, we introduce fictitious assets and liabilities. In particular, if $w_i > c_i - b_i$ we give i a new class of outside assets in the amount $c'_i = w_i - (c_i - b_i)$; if $w_i < c_i - b_i$ we give i a new class of outside liabilities in the

amount $b'_i = c_i - b_i - w_i$. It is assumed that these fictitious assets are impervious to shocks, while the original outside assets are subject to the same shock distribution as before.

We illustrate the idea in Figure 10 using the example in Figure 3. Consider node A. It still has its original outside assets (worth 120) and its original outside liabilities (worth 180). Its original in-network assets were 250 in payments due from nodes C and D: these have been aggregated into a fictitious asset worth 250. Similarly its original in-network liability (a payment of 180 due to node D) has become a fictitious asset worth 180. The entries for the other nodes are derived in similar fashion.




Figure 10. Example of Figure 3 with interbank connections severed and balance sheets held fixed by the introduction of fictitious assets and liabilities (in bold).

Given a shock vector x let $L(x) = \sum_i x_i + \sum_i s_i(x)$ denote the systemic loss in the original interconnected system. In the corresponding severed system the systemic loss is simply the sum of the initial shocks plus the shortfall in payments from defaulting nodes to the outside sector:

$$L^o(x) = \sum_i x_i + \sum_i (x_i - w_i)_+ . \quad (36)$$

The *systemic impact of the network* can be defined as the ratio $\bar{L} / \bar{L}^o - 1$, where $\bar{L} = E[L(x)]$ and $\bar{L}^o = E[L^o(x)]$.

One can bound the systemic impact of the network for a wide range of shock distributions as follows. Let $F(x_1, x_2, \dots, x_n)$ be the joint distribution of outside shocks. We assume that the shocks are identically distributed and *homogeneous* in the underlying assets, that is, $F(x_1, x_2, \dots, x_n) = G(x_1 / c_1, \dots, x_n / c_n)$ for a symmetric c.d.f. G with domain $[0, 1]^n$.

We do not need to assume that the shocks are independent. Let $G_i(y_i)$ be the marginal distribution of G with respect to i , where $y_i \in [0,1]$. We say that G has an *increasing failure rate* (IFR) if $g_i(y_i)/(1-G_i(y_i))$ is an increasing function for every i . Examples of IFR distributions include the normal, exponential, monotone beta with $q \geq 1$, and all log-concave distributions.

The probability that i defaults directly can be expressed as follows:

$$\delta_i = P(x_i > w_i) = 1 - G_i(w_i / c_i) . \quad (37)$$

Proposition 1.¹⁴ *If the joint shock distribution is homogenous in assets and IFR, then the systemic loss in value is bounded as follows:*

$$\bar{L} / \bar{L}^o - 1 \leq \frac{\sum \delta_i c_i}{(1 - \beta^+) \sum c_i} . \quad (38)$$

The right-hand side of (38) will typically be small if the direct default probabilities δ_i are small. Suppose for example that all financial institutions are alike in the sense that they have the same quantity and quality of outside assets, and the same probability of direct default δ . Then (39) implies that the relative increase in expected systemic losses due to network effects is at most $\delta / (1 - \beta^+)$. If the financial institutions are not alike, the expression highlights which of them contribute most to systemic risk: namely, those that have a large base of outside assets and/or a high probability of default. Another way of thinking about expression (38) is that it captures the interaction between key microprudential variables (the δ_i) and the macroprudential variable β^+ , which measures financial connectivity at the systemic level.

8.7 The Amplification Effect of Bankruptcy Costs

Up to this point we have made the simplifying assumption that when a bank defaults, its remaining assets are distributed pro rata without further impairment by the resolution process itself. This assumption is clearly unrealistic. Unfortunately there is relatively little empirical work on how large the impairments are in practice, nor is there much agreement on how to

¹⁴ This result is proved in Glasserman and Young (2015) for a slightly different loss function, namely,

$\tilde{L}(x) = \sum_i (x_i \wedge w_i) + \sum_i s_i(x)$. In fact Proposition 1 holds under either definition of systemic loss.

model the amount of impairment as a function of fundamentals. We see this as a crucial topic where more research is needed. Here we shall highlight the main components of bankruptcy costs, and suggest alternative ways of modelling them.

Broadly speaking the costs associated with default may be divided into three categories: administrative and legal costs, the costs of delay in making payments, and mark-downs in the valuation of assets in order to sell them quickly and reduce delays in making payments. In practice it is difficult to distinguish between liquidation costs and unrealized losses that occurred prior to failure.¹⁵ From a modelling perspective, the distinction is between a loss due to an initial shock and an amplification that results from bankruptcy costs.

Bennett and Unal (2014) estimate the legal and administrative costs of bank resolutions using FDIC data over the period 1986-2007. For the largest 25 banks in the sample these costs ranged from 0.33 percent to 13.19 percent of the book value of the bank's assets at the time of failure, and the median value was 5.69 percent.¹⁶ The losses could be even greater (as a percentage of assets) when a large and complex institution fails in a disorderly manner, as occurred in the case of Lehman.¹⁷

We now turn to the question of how to model bankruptcy costs theoretically. In other words, when a bank's assets are worth less than its obligations what portion of the assets' value is eventually recovered by creditors and what portion is lost through the resolution process? Recall from Section 5.3 that the *recovery function* $r(\alpha, \bar{p})$ represents the amount paid to creditors as a function of the bank's assets α and its nominal obligations \bar{p} . As before we assume that $r(\alpha, \bar{p})$ is monotone nondecreasing in both arguments and that $0 \leq r(\alpha, \bar{p}) \leq \alpha$ when $\alpha < \bar{p}$ and $r(\alpha, \bar{p}) = \bar{p}$ when $\alpha \geq \bar{p}$. The *deadweight loss* is $\alpha - r(\alpha, \bar{p})$, which is the difference between the assets available and the amount distributed to creditors.

¹⁵ An asset may lose value prior to the failure of a bank but the loss is realized only when the asset is sold or acquired. James (1991) provides estimates of these effects.

¹⁶ Bennett and Unal (2014), Table 2, direct expenses discounted plus receivership expenses discounted.

¹⁷ Fleming and Sarkar (2014) estimate that Lehman's creditors recovered 28 percent of what they were owed. This loss rate reflects both unrealized losses prior to bankruptcy and additional losses due to the bankruptcy itself, but they note that the abruptness of the bankruptcy may have substantially reduced the value of the Lehman estate.

Rogers and Veraart (2013) propose a discontinuous recovery function such that for some $\theta < 1$,

$$r(\alpha, \bar{p}) = \theta\alpha \text{ whenever } 0 \leq \alpha < \bar{p}, \quad (39)$$

$$r(\alpha, \bar{p}) = \bar{p} \text{ whenever } \alpha \geq \bar{p}. \quad (40)$$

Elliott et al. (2014) suggest that when α is sufficiently close to \bar{p} there is full recovery of the available assets, but when α falls below a critical threshold the recovery rate $r(\alpha, \bar{p})$ drops discontinuously. Glasserman and Young (2015) propose a model in which the amount of unrecovered assets increases linearly in the gap between the assets and the liabilities, that is, $\alpha - r(\alpha, \bar{p}) = \gamma(\bar{p} - \alpha)$ for some $\gamma > 0$. Taking into account the fact that $r(\alpha, \bar{p})$ must be nonnegative this leads to a linear recovery function of form

$$r(\alpha, \bar{p}) = ((1 + \gamma)\alpha - \gamma\bar{p}) \wedge 0. \quad (41)$$

Various functional forms are illustrated in Figure 11. We remark that even if the recovery function $r(\alpha, \bar{p})$ is nonlinear, the formulation in (41) provides a useful way of *bounding* the expected losses from above. In particular it suffices to find a value of γ such that the function $(1 + \gamma)\alpha - \gamma\bar{p}$ bounds the rate of recovery from below (see the dashed lines in Figure 11). The extreme case of zero recovery at default corresponds to an infinite γ and can be approximated by a large finite γ .




Figure 11. Three models of the recovery function. Left panel: discontinuous and piecewise linear. Middle panel: continuous and nonlinear. Right panel: continuous and linear. In the left and middle panels the dashed lines bound the rate of recovery from below.

We now show how to incorporate bankruptcy costs into the previous framework. Consider a recovery function that is bounded below by a linear function with slope $1 + \gamma$.

$$r(\alpha, \bar{p}) \geq ((1 + \gamma)\alpha - \gamma\bar{p}) \wedge 0. \quad (42)$$

Let x be a shock vector and let $D = D(x)$ be the resulting default set. A derivation similar to (29) shows that the shortfalls at the nodes in D satisfy the inequality

$$(1 + \gamma)[s_D A_D - w_D + x_D] \geq s_D. \quad (43)$$

Suppose in addition that $(1 + \gamma)A_D$ has spectral radius less than one, in which case $I_D - (1 + \gamma)A_D$ is invertible. Define the modified node depth vector as follows

$$u_D(\gamma) = [I_D + (1 + \gamma)A_D + (1 + \gamma)^2 A_D^2 + \dots] \mathbf{1}_D. \quad (44)$$

Then the system-wide shortfall satisfies the upper bound

$$\sum_i s_i(x) \leq s_D u_D(\gamma) = (1 + \gamma)(x_D - w_D) u_D(\gamma), \quad (45)$$

This in turn provides an upper bound on the systemic loss in value $L(x) = \sum_i x_i + \sum_i s_i(x)$.

These bounds depend on the matrix A_D , which depends on the structure of the default set D , which depends in turn on the shock vector x . In practice it is difficult to estimate these quantities with any precision. We do know, however, that the node depths are bounded above by the node depths that result when D is the set of all nodes. As before let $\beta^+ = \max_i \beta_i$ and suppose that $(1 + \gamma)\beta^+ < 1$. Then the node depths are uniformly bounded above as follows:

$$\forall i, u_i \leq (1 - (1 + \gamma)\beta^+)^{-1}. \quad (46)$$

In analogy with Proposition 1 one can then bound the expected systemic losses attributable to the network as follows.

Proposition 2. *Suppose that the joint shock distribution $F(x)$ is homogenous in assets and IFR. Suppose further that the amount recovered in default is bounded below by the linear recovery function with slope $1 + \gamma$. If $(1 + \gamma)\beta^+ < 1$, the expected systemic loss in value relative to the expected loss with no network connections is bounded as follows:*

$$\bar{L} / \bar{L}^0 - 1 \leq \frac{\sum \delta_i c_i}{(1 - (1 + \gamma)\beta^+) \sum c_i}. \quad (47)$$

We can interpret the product $(1 + \gamma)\beta^+$ as a *critical threshold* that differentiates two regimes: when it is less than unity the amplification of losses due to spillover effects is more than offset by the dissipation of these losses to the outside sector; when $(1 + \gamma)\beta^+$ is greater than unity the losses within the financial system can escalate dramatically. As we have emphasized throughout this section, the bound in (47) uses information about individual nodes in the network but it does not assume detailed knowledge of the network topology.

9. Network Measures

Whereas the economics literature on financial networks has often focused on simple configurations – ring networks and fully connected networks – a line of research drawing on network science has studied the features of actual networks, which are much more complex. One goal of this line of work is to identify simple network features that describe the vulnerability of a network or part of a network. We will review work applying these ideas to interbank networks defined through payment obligations.

To discuss network measures, we need to introduce some notation and terminology. A network has a set of nodes $\{1, \dots, N\}$ and edges that connect pairs of nodes. The network's $N \times N$ *adjacency matrix* B is defined by setting $B_{ij} = 1$ if an edge connects nodes i and j and $B_{ij} = 0$ otherwise. In a *directed* network, each edge has a direction. In this case, $B_{ij} = 1$ indicates the presence of an edge running from node i to node j . In most financial networks, each edge has a weight as well as a direction. In particular, in previous sections we defined financial networks through a matrix \bar{p} of liabilities, where $\bar{p}_{ij} \geq 0$ denotes the amount node i owes to node j , as in Figure 2. We interpret these as gross amounts so both \bar{p}_{ij} and \bar{p}_{ji} may be positive. Both the liabilities matrix \bar{p} and the adjacency matrix B have zeros on their diagonals.

9.1 Degree Distributions

Empirical descriptions of networks frequently include information about degree distributions. In a directed network, a node's in-degree is the number of edges directed to that node, and the out-degree is the number of edges directed from that node. These are given by the column sums and row sums of the adjacency matrix B . In an undirected network, a node's degree is the number of edges connected to that node. The degree distribution is just the empirical distribution of degrees across the nodes in a network.

Across many types of networks (including social networks and the Internet, for example), one often finds that a power law (a distribution with a Pareto-like tail) describes the degree distribution. In their analysis of the Austrian interbank lending network, Boss et al. (2004) estimate power laws for the distributions of in-degree, out-degree, and total degree. Their estimates are based on partial network information and a method for imputing unobserved links, to which we return in Section 12. They also find a power law in the distribution of imputed interbank liabilities. Similar conclusions can be found in Bastos and Cont (2010) and Cont, Santos, and Moussa (2013) for the Brazilian interbank network, in Martinez-Jaramillo et al. (2014) for the Mexican banking system. Power law degree distributions are also documented in the Japanese and U.S. payment systems by Inaoka et al. (2004) and Soramaki et al. (2007), respectively. Fricke and Lux (2013) challenge findings of power laws in interbank networks using data on overnight interbank lending in Europe. They find that parameter estimates for fitted power laws are unstable over time, and they conclude that distributions with an exponential tail often provide a better fit.

What do these findings tell us about interbank networks? Leaving aside the question of the precise decay rates of degree distributions, at least two observations seem robust. First, degree distributions in interbank networks are highly skewed: most banks have few links, and a few banks have a large number of links. Second, there is substantial asymmetry between in-degree and out-degree distributions (as observed, in particular, by Bech and Atalay 2010 and Pühr, Seliger, and Sigmund 2012): when a bank is highly connected, it is usually because it borrows from many other banks and not because it lends to many other banks. Pritsker (2014) offers a supporting theory for these patterns, to which we return in Section 12. The dynamic model of Blasques, Bräuning, and van Lelyveld (2015) generates networks of interbank relationship lending with these features, which the authors attribute to credit risk uncertainty and peer monitoring costs.

9.2 Core-Periphery Structure

The skewed distribution of interbank links arises in part from a core-periphery structure typical of interbank networks, with a small number of highly interconnected banks at the core and all others in the periphery. To formalize this idea, Craig and von Peter (2014) propose a precise, idealized definition of this structure. In their definition, every core bank lends to

every other core bank; periphery banks do not lend to each other; every core bank lends to and borrows from at least one periphery bank.¹⁸

In practice, the separation between the two tiers of banks is not that sharp: Craig and von Peter (2014) propose a measure of the distance from an ideal core-periphery structure and use it to test the fit of the model to data. Using data on interbank lending among 1,802 German banks, they identify a core of 45 banks. Not surprisingly, the core banks are large banks that provide wholesale services as well as traditional banking operations.

Van Lelyveld and in't Veld (2012) apply the method of Craig and von Peter (2014) to the Dutch banking system and identify 10 to 20 percent of banks as core. They also find that the average capital buffer at core banks is only about one-half to one-third of the average capital buffer for periphery banks, meaning that the most interconnected banks are the least well capitalized. Anand et al. (2015) estimate core-periphery structures across 25 markets as part of their comparison of methods for reconstructing networks from partial information, the topic of Section 12.

9.3 Centrality and Node Depth

Centrality measures seek to identify the importance of a node in a network, and the networks literature has generated many such measures; Chapter 7 of Newman (2010) provides an overview. A node's degree is a simple measure of its importance; other measures try to capture more information about a node's overall position in the network. In this section, we digress from our discussion of empirical work on networks to discuss centrality measures and a connection with the idea of node depth used in Section 8.

The eigenvector centrality v_j of node j (also called Bonacich centrality after Bonacich 1972) satisfies

$$\lambda v_j = \sum_{i=1}^N v_i B_{ij} , \quad (48)$$

where B is the adjacency matrix and λ is an eigenvalue of B . In this expression, a node has high centrality if it is connected to other nodes with high centrality. In the directed case, the

¹⁸ Capponi, Chen, and Yao (2015) develop a different notion of concentration using the majorization partial ordering of payment vectors. In Castiglionesi and Navarro (2008), a core-periphery structure arises endogenously through a network formation game.

expression says that the centrality of j depends on the centrality of nodes with edges pointing to j . Defining centrality through a right eigenvector instead flips the direction of influence and makes the centrality of a node the weighted average of the centrality of the nodes to which it points.

Suppose that the network is *irreducible* in the sense that every node can be reached from every other node. In other words, for any pair of nodes i and j , there is some positive integer n for which the ij -entry of B^n is strictly positive. In this case, the Perron-Frobenius theorem implies that B has an eigenvector with strictly positive entries, and this eigenvector is unique up to multiplication by a constant. The associated eigenvalue is real and positive and equals the spectral radius of B . This is the eigenvalue-eigenvector pair used to define eigenvector centrality, often with the normalization that the centrality values sum to one.

If we divide each entry of B by the sum of entries in its row, we obtain the matrix P with entries

$$P_{ij} = B_{ij} / \sum_{k=1}^N B_{ik} . \quad (49)$$

This is a stochastic matrix describing the motion of a dollar moving randomly through the network, following each of the edges out of a node with equal probability. The Perron-Frobenius eigenvalue is 1, and the corresponding left eigenvector is the equilibrium distribution for the Markov chain defined by P . In this case, a node's centrality is the fraction of time the dollar spends at that node as it moves randomly through the network¹⁹.

Centrality measures can also be defined from the liabilities matrix \bar{p} and suggest different interpretations depending on whether we take a left or right eigenvector

$$\lambda v_j = \sum_{i=1}^N v_i \bar{p}_{ij}, \quad \lambda u_j = \sum_{i=1}^N \bar{p}_{ij} u_i . \quad (50)$$

The left eigenvector assigns greater centrality to nodes that have claims on nodes with greater centrality. We may interpret this as *funding* centrality. The right eigenvector assigns greater centrality to nodes that have obligations to nodes with greater centrality, which suggests an

¹⁹ Other, closely related measures of centrality (such as Katz centrality) are available when the irreducibility condition is not satisfied.

interpretation as *borrowing* centrality.²⁰ The failure of a node with high borrowing centrality would result in defaults on large obligations and could set off a default cascade; in contrast, the failure of a node with high funding centrality could create a liquidity shock at other nodes through the withdrawal of funding.

The notion of node depth, which we introduced in Section 8 to bound the expected loss in a network, is in a precise sense dual to eigenvector centrality, as we now explain. As in Section 8, let A denote the relative liabilities matrix for a set of nodes. The set could be the nodes in default or it could be the entire network; we simply assume that A is substochastic and that its largest eigenvalue is strictly less than 1. The associated node depth vector u for the network represented by A is given by

$$u = [I + A + A^2 + \dots]\mathbf{1}, \quad (51)$$

where $\mathbf{1}$ is the vector of ones. If we interpret A as the transition matrix of an absorbing Markov chain, then u_i is the expected number of steps until absorption, starting from node i , and absorption corresponds to leaving the network.²¹

The contrast with centrality is easiest to see when the matrix A^T is also substochastic and thus also defines an absorbing Markov chain. Consider a modified chain that restarts at a node picked uniformly at random following absorption; between restarts, the motion of the chain is governed by A^T .²² The ergodic distribution of the modified chain²³ is proportional to

$$\mathbf{1}^T [I + A^T + (A^T)^2 + \dots] = u^T. \quad (52)$$

In this sense, node depth with respect to A corresponds to centrality with respect to A^T .

²⁰ These ideas parallel upstream and downstream effects in network models of the macroeconomy; see Acemoglu, Akcigit, and Kerr (2015). A general treatment of centrality measures is given in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015b).

²¹ In more detail, the states of the Markov chain are the nodes of the network. When the chain is in state i , it leaves the network and gets absorbed upon its next transition with probability equal to 1 minus the sum of the entries in the i th row of A .

²² This random restart mechanism is similar to the PageRank version of eigenvector centrality associated with the Google search algorithm, but PageRank uses a constant probability of random restart from every node. See Newman (2010), Section 7.4.

²³ The modified chain has transition matrix $P = A^T + [I - A^T]\mathbf{1}\mathbf{1}^T/n$. This matrix satisfies $u^T P = u^T$ even if A^T is not substochastic, but P may have some negative entries in that case.

Figure 12 illustrates the contrast between depth and centrality. Node C has only outside obligations whereas all other nodes have only in-network obligations. Node C is central because a dollar flowing through the network inevitably passes through C, but node C is shallow because a dollar leaving C immediately leaves the network. In contrast, the node at the start of each chain is peripheral and deep. Switching between A and A^T reverses the directions of the arrows and flips between centrality and depth.




Figure 12. A network in which node C is central but not deep, and the peripheral nodes are deep but not central.

Staum and Liu (2012) note a related Markov chain interpretation in their calculation of sensitivities in the Eisenberg-Noe (2001) model. Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015a) use the Markov chain analogy in defining a measure they call the “harmonic distance” between nodes in a network. Their measure counts the expected number of steps to reach one node from another node through a Markov chain defined by a matrix of relative liabilities. In their setting, both the relative liabilities matrix and its transpose are stochastic so there is no notion of leaving the network.

9.4 Empirical Network Measures and Measures of Stability

Network measures help describe complex interbank networks. Whether they also help identify vulnerabilities is less clear. We review efforts to establish these connections.

Using data from a German credit register, Craig, Koetter, and Krüger (2014) use bank centralities (measured through a variant of eigenvector centrality) as explanatory variables in regressions of individual bank risk and conclude that higher centrality predicts a lower probability of default. In their study of the Mexican banking system, Martinez-Jaramillo et al. (2012) combine many different centrality measures into a composite measure using principal components analysis. They contrast the structure of payment systems and interbank lending networks and find that centrality is not necessarily determined by size. In their study of the

Fed funds market, Bech and Atalay (2010) find that centrality measures are useful predictors of the interest rate banks charge each other on overnight loans, but they do not directly investigate implications for financial stability. Gabrieli (2011) undertakes a similar analysis using European data. Battiston et al. (2012b) propose a modified centrality measure and apply it to data on emergency loans made by the Federal Reserve in the period 2008-2010.

Puhr, Seliger, and Sigmund (2012) find Katz centrality (which is similar to eigenvector centrality with discounting of longer paths) to be a useful indicator in measuring systemic risk. They start with data on the Austrian banking system and use this as input to a simulation of bank failure. Puhr et al. (2012) distinguish contagiousness (a propensity to topple others) and vulnerability (a propensity to be toppled); this is similar to the contrast between left and right eigenvectors discussed in Section 9.3. They find that contagiousness increases with centrality, but vulnerability depends on several network measures, including a clustering coefficient. However, it is unclear to what extent the importance of the centrality measures is a feature of the simulation model or of the original banking system.

Alter, Craig, and Raupach (2014) propose using a bank's centrality in setting its capital requirements. They combine lending data from the German credit register with a network simulation in which random shocks to outside loans propagate through the network according to the Rogers-Veraart (2013) extension of the Eisenberg-Noe (2001) algorithm. They compare alternative ways to reallocate capital requirements from banks with low centrality to banks with high centrality using various measures of centrality. They find that eigenvector centrality based on the directed adjacency matrix is most effective in minimizing bankruptcy losses. Using centrality or almost any network measure in setting capital requirements poses a practical challenge because each bank's requirement becomes dependent on the actions of other banks.

Clustering is a focus of the study in Boss et al. (2004) of the Austrian banking system. As explained there, the Austrian banks are organized in three tiers, and this is captured by the clustering methods used in the paper. Santos and Cont (2010) examine clustering and a related assortativity property in the Brazilian interbank network.

Overall, empirical work has not yet produced a compelling link between traditional network measures and financial stability. Empirical work is limited by the confidentiality of interbank transactions and the low frequency of financial crises. But it may well be that standard network measures developed to highlight features in other applications are not well suited to

identifying vulnerabilities in the financial system, in which case it may be more productive to focus on new measures that reflect features specific to interbank networks.

10. The Impact of Connectivity

A recurring theme of the financial networks literature is the double-edged nature of connectivity: interbank links allow risk sharing, but they also provide channels for the spread of shocks. We encountered this trade-off in our discussion of funding runs in Section 4. Several studies have investigated this trade-off numerically or through simplified tractable models.

Nier et al. (2007) simulate random networks in which each pair of banks has a fixed probability of connection through an interbank loan, independent of all other pairs. Shocks to a bank's assets spill over to other banks through cascading defaults.²⁴ As Nier et al. (2007) increase the connection probability, they find a nonmonotonic (in fact, M-shaped) effect on the total number of defaults. Increasing connectivity increases shock transmission and shock absorption, with the first effect dominating at low connectivity and the second effect dominating at higher connectivity. Gai and Kapadia (2010) and Haldane and May (2011) draw similar conclusions but also observe that higher connectivity produces more severe, if less frequent, crises. Elliott, Golub, and Jackson (2014) make a similar point.

The effect of connectivity depends on what is held constant as connectivity varies. Several studies hold each bank's total interbank liabilities fixed as the number of links grows; this is the case in the nonmonotonic (M-shaped) finding of Nier et al. (2007). But connectivity can also be measured through the *size* of a bank's interbank liabilities, rather than through the number of banks from which it borrows. This is arguably a more important measure of vulnerability, as discussed in Section 8. When they increase the proportion of interbank assets to total assets, Nier et al. (2007) find a monotonic effect on the number of defaults. The effect plateaus because they increase net worth at the same time, but it does not decline.

Measuring connectivity through the probability of random connections is convenient, but it misses an important feature of the Allen and Gale (2000) model, where the structure and not just the level of connectivity matters. In particular, contagion in Allen and Gale (2000) is

²⁴ They also consider an extension with fire sales, as in Cifuentes, Ferrucci, and Shin (2005).

driven by the relationship between network structure and the correlation in funding shocks across banks.

Gofman (2014) introduces important structural features to his model and takes a different approach to studying the consequences of connectivity. Using a preferential attachment model of link formation, he calibrates features of his model to properties of the U.S. Fed funds market, as reported in Bech and Atalay (2010), which leads to a core-periphery structure. He holds fixed the average number of lending counterparties per bank and instead varies the maximum number of counterparties allowed per bank. In this sense, his comparison is orthogonal to most other comparisons of connectivity. Increasing the maximum number of counterparties while holding fixed the average increases the concentration of the core – an important dimension of connectivity in practice. He shocks the system with the failure of the most interconnected bank and finds a nonmonotonic relationship between the total number of failures and the cap on the number of counterparties. This pattern results from two competing effects: lowering the cap increases the size of the core while decreasing the exposures between core banks. Gofman (2013) also develops a measure of surplus loss that increases with the length of chains through the network because of inefficiencies in intermediation. His overall conclusion is that a cap on the number of lending counterparties can decrease both stability and efficiency.

Relevant to this discussion is the Basel Committee’s (2014) standard for limiting large exposures between banks, which was developed specifically to mitigate contagion risk. As discussed in Section 7.1, the rule limits the exposure of a bank to any other bank to no more than 25 percent of the lending bank’s capital. For exposures between global systemically important banks (roughly the core banks in the international core-periphery structure), the limit is 15 percent. Using the notation of Figure 2, we may formulate this rule through constraints of the form

$$\bar{p}_{ij} / w_i \leq 15\% \text{ or } 25\% \quad . \quad (53)$$

In particular, this limits the size of interbank exposures but not the number of exposures. This constraint applies on the asset side. A recent U.S. proposal (Federal Reserve 2014) would set higher capital requirements for banks with high levels of “wholesale” funding. The rule is complex but may be roughly approximated in a network model through a constraint of the form

$$w_i \geq k \beta_i \bar{p}_i, \quad (54)$$

for some multiplier k . In other words, a higher level of in-network (wholesale) funding requires a higher level of capital. These two rules substantially limit the interconnectedness of the network. But, appropriately, they do so by limiting exposure magnitudes, not the in-degree or out-degree allowed at a node.

Some of the main results in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015a) examine the effect of connectivity holding the network topology fixed and instead varying the size of an initial shock. They show that a ring network always produces the greatest number of defaults because, as in Section 4.1, it concentrates the spillover from one node to another. A completely connected network flips from producing the least number of defaults to the greatest number of defaults at a sufficiently large shock.

We illustrate the simplest version of this result through Figure 13. All nodes are ex ante identical. Each has c in outside asset and b in outside liabilities; each has an identical payment obligation of \bar{p} to (and from) every other bank. Outside liabilities are senior to interbank liabilities. Payments are determined through a fixed-point argument similar to Eisenberg and Noe (2001).

Node 1 suffers a shock X to its outside asset. All other nodes are identical, so we may focus on node 2. By symmetry, banks $2, \dots, n$ make identical payments to each other regardless of the shock, so those payments cancel each other. Banks $2, \dots, n$ will either all fail or all survive.




Figure 13: A symmetric completely connected network. Connections between nodes 2, ..., n are not shown.

The interesting case has node 1 default, which requires $c - X < b$. Whatever payment p node 1 receives from each of the other nodes it will be able to return $(c - X - b + (n-1)p) / (n-1)$ to each node after paying its senior outside obligations (as shown in the figure) provided this quantity is positive; otherwise, it returns zero. Now consider node 2's ability to pay. Node 2 defaults when its assets are worth less than its liabilities; that is,

$$c + \bar{p} + \frac{c - X - b}{n-1} < b + \bar{p}, \quad (55)$$

which is equivalent to $X > n(c - b)$. This is the threshold proved in Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015a) for what they call a phase transition. Notice however that $c - b$ is each node's net worth, so $n(c - b)$ is the total net worth of the banking system. For any asset shock short of that high level, the fully connected network produces the smallest number of defaults.

Interestingly, Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015a) show that there can be other configurations under which some nodes survive even a catastrophic shock of this size. The alternative configurations partition the network into loosely connected subsets, creating a firewall that limits the spread of losses but without the complete separation of the containment example in Allen and Gale (2000). Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015a) give an elegant characterization of the degree of separation through a bottleneck parameter.

The spread of losses is contained by the sizes of interbank liabilities – a bank cannot fail to pay more than 100 percent of its obligations. Indeed, in Figure 13, for node 1 to impose the indicated loss on each of the other nodes requires

$$\bar{p} > \frac{c - X - b}{n-1}. \quad (56)$$

Thus even the catastrophic shock $X \geq c(n - b)$ will not bring down the network unless node 1's obligation to each other bank exceeds its net worth $c - b = w$. This would violate the Basel large exposure limit discussed above; the large exposure limit is therefore consistent with the containment strategy suggested by a network perspective.

11. Common Exposure Networks and Fire Sales.

Our primary focus has been on networks defined by payment obligations from borrowers to lenders, but networks can also be used to model the situation where financial institutions have overlapping exposures to different asset classes. In Section 5.5 we discussed the simple case where they all hold the same illiquid asset and contagion results when some of them are forced to sell to meet liquidity requirements.




Figure 14. A bipartite network of firms and assets. Edges indicate ownership of assets by firms.

A similar dynamic can arise when institutions hold multiple illiquid assets and there is overlap in their exposures. Consider the bipartite graph of exposures in Figure 14. A shock to Asset 1 (for example, Florida real estate) creates a cash shortfall for Firm 1 (a bank or a hedge fund), forcing it to sell its holdings in Asset 2 (for example, emerging market bonds). The forced sale could result from a leverage constraint in the case of a hedge fund or a capital requirement in the case of a bank. If the selloff is large enough to drive down the price of Asset 2, it imposes losses on Firm 2 and Firm N. Thus, these firms can suffer indirect effects from the initial shock to Asset 1, even though they have no direct exposure to that asset.

The Office of Financial Research (2013), p.15, uses this type of network representation to illustrate the overlap – and potential spillover – in holdings among the ten largest U.S. prime money market funds. Caccioli et al. (2014) propose a branching process model of contagion through this type of common exposure network.

Chen, Iyengar, and Moallemi (2014) develop a model of this type in which investment firms optimize their portfolio holdings and then optimally rebalance in response to an asset shock; this endogenizes the contagion effect. They distinguish low and high leverage regimes and

identify a stark contrast between the two. In the first case, network amplification of shocks is minimized when all firms hold an identical mix of assets, and in the second case it is minimized by splitting the banks into two groups with nonoverlapping exposures. These effects do not depend on payment obligations between banks.

As discussed in Chen, Iyengar, and Moallemi (2014), their results bear on a related question of the trade-off between diversification and diversity: investment diversification reduces risk for each firm, but when all firms hold similar diversified positions, the system as a whole may be more vulnerable. Models of this phenomenon include Acharya (2009), Allen, Babus, and Carletti (2012), Battiston et al. (2012a), Bimpikis and Tahbaz-Salehi (2014), Cabrales, Gottardi, and Vega-Redondo (2014), Ibragimov, Jaffee, and Walden (2011), and Wagner (2011).

12. Assessing Systemic Risk with Limited Network Information

Network models of the financial system usually assume complete information about interbank exposures. In practice, this level of detail is never available to the public or to the banks themselves, and even regulators rarely have a comprehensive view of all payment obligations between financial institutions – obligations which in any case are constantly changing.²⁵ In this section, we discuss various approaches to the problem of limited network information.

12.1 Networks Defined Through Comovements

One way to get around the problem of limited information on interbank links is to define a network through comovements in stock prices. This approach taps an abundance of data to uncover potential links between firms, with the caveat that the links may be difficult to interpret and need not correspond to a structural network defined through specific contracts or payments.

²⁵ The most detailed analyses of financial networks have used confidential data on short-term (mostly overnight) lending – for example, Soramäki et al. (2007), Bech and Atalay (2010), Afonso, Kovner, and, Schoar (2011), Gabrieli and Georg (2014), and Blasques, Bräuning, and van Lelyveld (2015). Even these studies have relied on incomplete data: they use funds transfers through centralized payment systems to infer payment obligations. This is often done through the Furfine (1999) algorithm, which interprets nearly offsetting payments between banks as a loan and a repayment with interest.

Billio et al. (2011) construct networks of financial institutions through Granger causality relationships in stock returns. In more detail, they estimate models of stock returns R of the form

$$\begin{aligned} R_{t+1}^i &= a^i R_t^i + b^{ij} R_t^j + e_{t+1}^i \\ R_{t+1}^j &= a^j R_t^j + b^{ji} R_t^i + e_{t+1}^j \end{aligned} \quad (57)$$

where i and j index financial firms and t indexes time. The error terms are assumed uncorrelated across firms and across times. A nonzero value of b^{ij} indicates Granger causality from firm j to firm i ; it measures the strength of the spillover from a shock to j in one period to a shock to i in the next period. Including firm i 's own lagged stock return in the regression helps ensure that b^{ij} captures a genuine spillover effect and not simply the impact of a common shock in the prior period.

Billio et al. (2011) use the statistically significant b^{ij} to construct directed networks between financial firms and to define several associated measures of connectivity. They compare estimates from different time periods and find that connectivity increases in advance of financial crises, suggesting the possibility of an early warning signal.

Bonaldi, Hortaçsu, and Kastl (2015) estimate networks of banks based on changes in funding costs: a directed edge from one bank to another indicates that an increase in the funding cost in one bank is associated with an increase in the funding cost of the other bank in the next period. Their analysis is based on confidential bids in the European Central Bank's short-term refinancing auctions. They find that using public data through the method of Billio et al. (2011) leads to lower estimates of interconnectedness between banks. They attribute the difference to market efficiency – in an informationally efficient market, none of the b^{ij} in (57) should be significant.

Diebold and Yilmaz (2014) estimate networks from market data using variance decompositions for the forecast errors in multivariate time series. Off-diagonal entries in the variance decomposition reflect the influence of shocks to one variable on the forecast errors in another variable, and these cross-variable influences can be used to construct directed networks. In their main application, Diebold and Yilmaz (2014) apply their method to stock return *volatilities*, rather than the returns themselves. They find high connectivity between Fannie Mae and Freddie Mac and among some of the largest commercial banks. In 2007-

2008, they find spikes in net connectedness from AIG, Bear Stearns, Lehman Brothers, and Wachovia at critical dates.

These methods exploit the abundance of market data to identify links between institutions. In contrast, network models based on payment obligations are constrained by the confidentiality and relatively low frequency of the data they require. Network measures based on correlations may be able to detect important patterns of comovements in market data, but they are generally unable to identify the underlying mechanism generating those comovements.

12.2 Imputing Missing Data

An alternative approach, which goes back to some of the earliest research on interbank networks, imputes values for unavailable information to complete the specification of the network. Let \bar{p}_{ij} denote the (unknown) liability of bank i to bank j , and, as before, think of the \bar{p}_{ij} as entries of an $N \times N$ matrix. Suppose we know only the row sums $r_i, i=1, \dots, N$, and the column sums $c_j, j=1, \dots, N$, of this matrix: r_i represents bank i 's total obligations to other banks, and c_j represents bank j 's total claims on other banks.

This was the situation faced by Sheldon and Maurer (1998) in their analysis of the Swiss interbank network. Of the infinitely many sets of interbank liabilities \bar{p}_{ij} consistent with the marginal assets c_j and liabilities r_i , they proposed using the values from a maximum entropy solution, interpreting this as the solution that minimizes the amount of external information imposed on the solution.

In more detail, they proposed fixing the interbank liabilities by solving the constrained convex optimization problem

$$\begin{aligned} & \min_{\{\bar{p}_{ij}\}} \sum_{i,j} \bar{p}_{ij} \ln \bar{p}_{ij} \\ & \text{subject to } \sum_j \bar{p}_{ij} = r_i, i = 1, \dots, N \\ & \sum_i \bar{p}_{ij} = c_j, j = 1, \dots, N \end{aligned} \tag{58}$$

with the convention that $0 \cdot \ln 0 = 0$. Without loss of generality, we may normalize the total amount of lending (the sum over all \bar{p}_{ij}) to 1, which makes $\sum_i r_i = \sum_j c_j = 1$. The optimal

solution is then

$$\bar{p}_{ij} = r_i c_j. \quad (59)$$

The simplicity of this solution is appealing, but other features make it implausible. As noted by Sheldon and Maurer (1998), it suggests stochastic independence across banks in their decisions to borrow and lend to other banks. It also yields nonzero diagonal entries although banks cannot meaningfully lend to themselves.

Upper and Worms (2004) address the second of these concerns. They change the optimization objective to

$$\min_{\{\bar{p}_{ij}\}} \sum_{i,j} \bar{p}_{ij} \ln(\bar{p}_{ij} / x_{ij}), \quad (60)$$

where $x_{ij} = r_i c_j$ for $i \neq j$ and $x_{ii} = 0$. This modification ensures that the optimal solution will have zero diagonal entries (taking $\ln(0/0) = 0$). It produces the solution with this property that is closest to the Sheldon-Maurer solution in the sense of relative entropy, and it is computationally convenient.

This type of approach has been widely used in studies of national banking networks, including Degryse and Nguyen (2007) and Wells (2004). Elsinger, Lehar, and Summer (2006) have data on most, but not all, interbank liabilities in the Austrian banking system; they use entropy optimization to fill in the missing values. Upper (2011) compares the results of these and other studies.

The maximum entropy and minimum relative entropy methods produce interbank liabilities that are as uniform as possible, given other constraints imposed on the problem. In particular, without additional constraints they produce fully connected networks. These features make these solutions unrealistic, given the prevalence of power law degree distributions and core-periphery structure reviewed in Section 3.

Mistrulli (2011) compares actual interbank exposures in the Italian banking system with exposures imputed using maximum entropy and finds, indeed, that the imputed networks are too dense, given the presence of large, core banks in the Italian system. Through simulation experiments, he concludes that maximum entropy networks understate the severity of contagion unless the assumed loss given default is quite high. This is to be expected – spreading liabilities too uniformly dilutes the effect of a shock.

Van Lelyveld and Liedorp (2006) address these shortcomings using additional information. For their analysis of the Dutch banking system, they have detailed bilateral information reported by banks on their largest exposures. For all other interbank exposures, they have only partial (row sum and column sum) information. They hypothesize that the network of large exposures is representative of the network of all exposures and can therefore be used as a guide to filling in missing data. They select the matrix of interbank assets that is closest to the matrix of large exposures, in the sense of relative entropy, subject to the marginal constraints.

An alternative to imputing a fixed set of values for interbank liabilities is to generate random networks consistent with partial network information and then to analyze contagion across a distribution of networks. Examples include Anand, Craig, and von Peter (2014), Drehmann and Tarashev (2013), Gandy and Veraart (2015), Halaj and Kok (2013), and Sachs (2014). These types of methods offer the potential for a richer analysis consistent with limited partial information. Their sampling methods are somewhat heuristic and their properties therefore not yet well understood. The comparison of methods in Anand et al. (2015) is an important step in understanding which methods for network reconstruction work well with which types of networks.

12.3 Bounds Based on Partial Information

As discussed in Section 8, Glasserman and Young (2015) address the problem of incomplete information by using node-level data to bound the amount of loss amplification due to the (unknown) network. Their analysis assumes four pieces of information about each node: its total claims inside and outside the network, and its total liabilities inside and outside the networks. As discussed in Glasserman and Young (2015) and Allahrakha, Glasserman, and Young (2015), this information is well approximated through public sources. Their analysis does not assume knowledge of the network topology – it does not assume information about claims and obligations between pairs of nodes.

12.4 Within-Network Uncertainty

In his discussion of challenges facing the measurement of systemic risk, Hansen (2013) highlights the importance of uncertainty and distinguishes two types: uncertainty on the part of modelers observing the financial system, and uncertainty on the part of participants within

the financial system. The discussion in this section has focused on the first issue; we now touch on models of the second.

Uncertainty amplifies contagion in the interbank network model of Caballero and Simsek (2013). They consider a simple ring network in which each bank lends to exactly one other bank; the banks also hold risky outside assets and cash. One bank is hit with an unexpected payment obligation, and all banks decide whether to sell their risky assets or use their cash to buy more risky assets. The unexpected shock causes a domino effect of failures at some banks and a flight to quality at banks that liquidate their risky assets. Caballero and Simsek (2013) contrast two settings: one in which all banks have full information about the network, and one in which banks do not know their distance from the bank that suffers the initial shock. In the uncertain case, banks choose their actions to optimize against the worst-case configuration. With a sufficiently large shock, uncertainty about the network leads to more bank failures, a universal flight to quality, and a fire sale price for the risky asset. Alvarez and Barlevy (2014) extend the model and investigate when mandatory disclosure of which banks have suffered losses is welfare improving.

Whereas the network topology is exogenous in Caballero and Simsek (2013), Pritsker's (2013) model seeks to explain the structure of the interbank lending market through uncertainty aversion. The Fed funds market has a tiered structure in which smaller banks are net lenders to larger banks, and large banks lend to each other. The market features both relationship-based lending and, in the top tier, anonymously brokered lending. In Pritsker's (2013) model, lending banks are uncertain about the asset quality (hence creditworthiness) of borrowing banks and demand an uncertainty premium. Uncertainty has a greater effect on the worst-case default probability for small banks, and this leads to a tiered structure in equilibrium. Pritsker (2013) also discusses the role of the government in improving the functioning of the interbank market by producing information to reduce uncertainty.

Acharya and Yorulmazer (2008) develop a model in which "information contagion" contributes to systemic risk. The failure of one bank leads depositors to update their beliefs about a surviving bank, leading them to demand a higher interest rate or withdraw their deposits. Anticipating this response, banks choose correlated investments so that they are more likely to fail together than separately. A line of research on over-the-counter markets studies the spread of information in trading networks; see Duffie (2012), Duffie, Malamud, and Manso (2009), and Gofman (2011).

Outside the network setting, Dang et al. (2014) argue that opacity is intrinsic to the function of banks as providers of safe assets like deposits. In their framework, a financial crisis is associated with safe, *information insensitive* assets becoming *information sensitive*, undermining their function as money-like securities. In Caballero and Krishnamurthy (2008) a financial crisis unfolds as a shock to one asset amplifies uncertainty about other assets, prompting market participants to hoard liquidity.

13. Conclusion and Open Problems

The focus of this paper has been on what the financial networks literature has to say about the relationship between interconnectedness and financial stability. Do more interbank connections promote stability through risk sharing, or do they lead to greater fragility by creating channels for contagion? One of the key lessons of the literature is that this question cannot be answered without considering other factors that contribute to contagion, including leverage levels and heterogeneity in size. As we have seen the interaction between these factors and the network topology is quite complex and not fully understood even at a theoretical level. From a practical standpoint one would like to know what the literature teaches us about the contagiousness of particular institutions. How can one identify the institutions that are most likely to generate contagion and which institutions are the most vulnerable? Various concepts from the social networks literature been proposed to tackle this problem. Nevertheless more theoretical and empirical research needs to be done before we can be confident that these measures are valid predictors of financial contagion in practice.

Another issue that needs more systematic study is how to measure the “size” of a systemic event. Simply counting the number of bank failures is clearly not sufficient. A more comprehensive measure is the total loss in value of financial and nonfinancial assets. Of course from a policy standpoint one would like to be able to quantify the loss in economic welfare associated with a given loss in asset values. Although there has been some work on this question, much remains to be done -- both empirically and theoretically -- to identify the relationship between financial losses, the supply of credit, and their impact on the real economy. As we mentioned at the outset, the focus of this paper has been on the dynamics of contagion in the financial system when the network of obligations is given. There is also a growing literature on the dynamics of link formation (see among others Babus (2007), Castiglionesi and Navarro (2008), Georg (2013), Erol and Vohra (2014), Farboodi (2014), In ‘t Veld, van der Leij, and Hommes (2014), and Vuillemeij and Breton (2014)). The

conventional view is that institutions establish links with one another as a way of diversifying risk and facilitating intermediation. While this is certainly true, we would argue that a realistic model of network formation must include other factors. In particular, one must acknowledge that links between financial institutions are often created in a decentralized fashion within particular lines of business, such as commercial lending, foreign exchange, derivatives trading, repo desks, and the like. Building realistic models of the resulting dynamics will require a high degree of institutional knowledge and a clear understanding of the incentives faced by the individuals who are forming (and severing) these links. We believe that this is one of the most important challenges for future research in this area.

Another challenge is to develop a more comprehensive and integrated account of the various mechanisms that cause contagion within the financial system. The most often-studied mechanism is the transmission of defaults through payment shortfalls, as in the Eisenberg-Noe model and its extensions. A quite different transmission mechanism is the contraction in short-term funding that can occur when institutions pull funding in response to liquidity shocks. To date there has been relatively little research on how institutions (and individual traders within these institutions) actually behave when determining how to ration short-term funding in times of financial stress.

Other mechanisms for contagion can operate without network linkages while reinforcing network effects. For example, the downward spiral in asset prices that results from the dumping of illiquid assets (fire sales) in response to liquidity shocks can spread losses through markets, amplifying a cascade of defaults. Yet another mechanism is “information contagion.” This is triggered by changes in market perceptions about the creditworthiness of particular institutions and the value of their assets, which can feed through the system and create a general crisis of confidence. More empirical research needs to be done to disentangle these various mechanisms and to estimate their relative magnitudes.

Much of the literature on financial network models presupposes complete information about the network. We have argued that network opacity is a first-order concern for agents within the network, for regulators monitoring the network, and therefore for researchers developing models. More work is needed on inference from partial observations of network data and on understanding how opacity itself may contribute to contagion.

References

Acemoglu, D., Akcigit, U., and Kerr, W. (2015). Networks and the macroeconomy: an empirical exploration, NBER Macroa Annual 2015, forthcoming.

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015a). Systemic risk and stability in financial networks. *American Economic Review* 105(2), 564-608.

Acemoglu, D., Ozdaglar, A., and Tahbaz-Salehi, A. (2015b). Networks, shocks, and systemic risk. *Oxford Handbook of Economics*, to appear.

Acharya, V. V. (2009). A theory of systemic risk and design of prudential bank regulation. *Journal of Financial Stability*, 5(3), 224-255.

Acharya, V. V., and Yorulmazer, T. (2008). Information contagion and bank herding. *Journal of Money, Credit and Banking*, 40(1), 215-231.

Adrian, T., and Boyarchenko, N. (2012). Intermediary leverage cycles and financial stability. *Becker Friedman Institute for Research in Economics Working Paper*, (2012-010).

Afonso, G., Kovner, A. and Schoar, A. (2011). Stressed, Not Frozen: The Federal Funds Market in the Financial Crisis. *The Journal of Finance*, 66: 1109–1139

Allahrakha, M., Glasserman, P., and Young, H.P. (2015). Systemic importance indicators for 33 U.S. bank holding companies: an overview of recent data. Office of Financial Research Research Brief 15-01, Washington, D.C.

Allen, F., and Babus, A. (2009). Networks in finance. In *The Network Challenge: Strategy, Profit and Risk in an Interlinked World*, P. R. Kleindorfer and Y. Wind (eds.), pp. 367-382, Wharton School Publishing.

Allen, F., Babus, A., and Carletti, E. (2012). Asset commonality, debt maturity and systemic risk. *Journal of Financial Economics*, 104 (3), 519-534.

Allen, F., and Carletti, E. (2008). Mark-to-Market Accounting and Liquidity Pricing, *Journal of Accounting and Economics*, 45, 358-378.

Allen, F. and D. Gale (2000) Financial Contagion. *Journal of Political Economy* 108: 1-33

Alter, A., Craig, B. R., and Raupach, P. (2014). Centrality-Based Capital Allocations and Bailout Funds. Available at SSRN 2450788.

Alvarez, F., and Barlevy, G. (2014). Mandatory disclosure and financial contagion. Working paper 2014-04, Federal Reserve Bank of Chicago.

Amini, H., Cont, R., and Minca, A. (2013). Resilience to contagion in financial networks. *Mathematical Finance*.

Anand, K., Craig, B., and von Peter, G. (2014). Filling in the blanks: network structure and interbank contagion. Deutsche Bundesbank Working Paper No. 02/2014.

Anand, K., van Lelyveld, I., Banai, A., Friedrich, S., Garratt, R., Halaj, G., Howell, B., Hansen, I., Jaramillo, S., Lee, H., Molina-Borboa, J., Nobili, S., Rajan, S., Salakhova, D., Silva, T., Silvestri, L., Stancato de Souza, S. (2015). The missing links: a global study on uncovering financial network structure from partial data. Working paper, forthcoming.

Babus, A. (2007) "The Formation of Financial Networks". *Fondazione Eni Enrico Mattei Working Papers*. Paper 129. <http://services.bepress.com/feem/paper129>

Basel Committee on Banking Supervision (2014). *Supervisory framework for measuring and controlling large exposures – final standard*. Bank for International Settlements, Basel, Switzerland.

Bassett, W.F., M.B. Chosak, J.C. Driscoll, E. Zakrajsek 2014. Changes in bank lending standards and the macroeconomy. *Journal of Monetary Economics* 62, 23-40.

Battiston, S., Delli Gatti, D., Gallegati, M., Greenwald, B., and Stiglitz, J. E. (2012a). Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8), 1121-1141.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., and Caldarelli, G. (2012b). Debtrank: Too central to fail? financial networks, the Fed and systemic risk. *Scientific Reports* 2.

Bech, M. L., and Atalay, E. (2010). The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22), 5223-5246.

Bennett, R. L., and Unal, H. (2014) Understanding the components of bank resolution costs. Center for Financial Research working paper 2014-04, Federal Deposit Insurance Corporation, Washington, D.C.

Billio, M., Getmansky, M., Lo, A. W., and Pelizzon, L. (2012). Econometric measures of connectedness and systemic risk in the finance and insurance sectors. *Journal of Financial Economics* 104(3), 535-559.

Bimpikis, K., and Tahbaz-Salehi, A. (2014). Inefficient diversification. SSRN working paper 2198053.

Blasques, F., Bräuning, F., and van Lelyveld, I. (2015). A dynamic network model of the unsecured interbank lending market. Working paper, VU University Amsterdam, Netherlands.

- Bonacich, P. (1972). Factoring and weighting approaches to status scores and clique identification. *Journal of Mathematical Sociology*, 2(1), 113-120.
- Bonaldi, P., Hortaçsu, A., and Kastl, J. (2015). An empirical analysis of funding costs spillovers in the euro-zone with application to systemic risk. Working paper 21462, National Bureau of Economic Research, Cambridge, Massachusetts.
- Boss, M., Elsinger, H., Summer, M., and Thurner, S. (2004). Network topology of the interbank market. *Quantitative Finance*, 4(6), 677-684.
- Brunnermeier, M.K., and Y. Sannikov (2014). A macroeconomic model with a financial sector. *American Economic Review* 104, 379-421.
- Brusco, S., and Castiglionesi, F. (2007). Liquidity coinsurance, moral hazard, and financial contagion. *Journal of Finance* 62(5), 2275-2302.
- Caballero, R. J., and Krishnamurthy, A. (2008). Collective risk management in a flight to quality episode. *Journal of Finance* 63(5), 2195-2230.
- Caballero, R. J., and Simsek, A. (2013). Fire sales in a model of complexity. *Journal of Finance* 68(6), 2549-2587.
- Cabrales, A., Gale, D., and Gottardi, P. (2015). Financial contagion in network. In *Oxford Handbook on the Economics of Networks*, Y. Bramouille, A. Galleotti, and B. Rogers (eds.), Oxford University Press.
- Cabrales, A., Gottardi, P., and Vega-Redondo, F. (2014). Risk-sharing and contagion in networks. SSRN working paper 2425558.
- Caccioli, F., Shrestha, M., Moore, C., and Farmer, J.D. (2014). Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking and Finance*.
- Capponi, A., Chen, P.C., and Yao, D.D. (2014). Liability concentration and losses in financial networks. SSRN working paper 2517755.
- Castiglionesi, F., and Navarro, N. (2008). Optimal fragile financial networks. SSRN working paper 1089357.
- Chen, C., Iyengar, G., and Moallemi, C. C. (2013). An axiomatic approach to systemic risk. *Management Science* 56(6), 1373-1388.
- Chen, C., Iyengar, G., and Moallemi, C. C. (2014). Asset-based Contagion Models for Systemic Risk. Working paper, Columbia Business School, New York.
- Chen, N., Liu, X., and Yao, D.D. (2014). An optimization view of systemic risk modeling: the network effect and the market liquidity effect. SSRN Working paper 2463545.
- Cifuentes, R., Ferrucci, G., and Shin, H. S. (2005). Liquidity risk and contagion. *Journal of the European Economic Association* 3(2-3), 556-566.

- Cont, R., Santos, E.B., and Moussa, A. (2013) Network structure and systemic risk in banking systems.in: JP Fouque and J Langsam (eds.): *Handbook of Systemic Risk*, Cambridge University Press, pp. 327-368.
- Craig, B. R., Koetter, M., and Krüger, U. (2014). Interbank lending and distress: Observables, unobservables, and network structure. Deutsche Bundesbank Discussion Paper No. 18/2014.
- Craig, B., and Von Peter, G. (2014). Interbank tiering and money center banks. *Journal of Financial Intermediation*.
- Dang, T. V., Gorton, G., Holmström, B., and Ordonez, G. (2014). Banks as secret keepers. Working paper w20255, National Bureau of Economic Research.
- Degryse, H., and Nguyen, G. (2007). Interbank exposures: An empirical examination of contagion risk in the Belgian banking system. *International Journal of Central Banking* 3(2), 123-171.
- Diamond, D. W., and Dybvig, P. H. (1983). Bank runs, deposit insurance, and liquidity. *The Journal of Political Economy*, 401-419.
- Diebold, F. X., and Yılmaz, K. (2014). On the network topology of variance decompositions: Measuring the connectedness of financial firms. *Journal of Econometrics* 182(1), 119-134.
- Drehmann, M., and Tarashev, N. (2013). Measuring the systemic importance of interconnected banks. *Journal of Financial Intermediation* 22(4), 586-607.
- Duarte, F., and Eisenbach, T.M. (2014). Fire-sale spillovers and systemic risk. Staff Report 645, Federal Reserve Bank of New York.
- Duffie, D. (2012). *Dark markets: Asset pricing and information transmission in over-the-counter markets*. Princeton University Press.
- Duffie, D., Malamud, S., and Manso, G. (2009). Information percolation with equilibrium search dynamics. *Econometrica*, 77(5), 1513-1574.
- Eisenberg, L., and Noe, T. H. (2001). Systemic risk in financial systems. *Management Science* 47(2), 236-249.
- Elliott, M., Golub, B., and Jackson, M. O. (2014). Financial Networks and Contagion. *American Economic Review* 104(10), 3115-53.
- Elsinger, H., Lehar, A., and Summer, M. (2006). Risk Assessment for Banking Systems, *Management Science* 52, 1301–1314.
- Erol, S., and Vohra, R. (2014). Network formation and systemic risk. SSRN working paper 2500846.

Farboodi, M. (2014). Intermediation and Voluntary Exposure to Counterparty Risk. Working paper, University of Chicago.

Federal Deposit Insurance Corporate (1997). *History of the 80s, Volume 1: An examination of the banking crises of the 1980s and early 1990s*. <https://www.fdic.gov/bank/historical/history/vol1.html>

Financial Crisis Inquiry Commission. (2011). *The Financial Crisis Inquiry Report*. US Government Printing Office.

Fleming, M., and Liu, W. (2013). Near failure of Long Term Capital Management, in *100 Years: The Federal Reserve System*. <http://www.federalreservehistory.org/Events/DetailView/52>

Freixas, X., Parigi, B. M., and Rochet, J. C. (2000). Systemic risk, interbank relations, and liquidity provision by the central bank. *Journal of money, credit and banking*, 611-638.

Fricke, D., and Lux, T. (2013). On the distribution of links in the interbank network: Evidence from the e-mid overnight money market (No. 1819). Kiel Working Paper.

Fricke, D., and Lux, T. (2014). Core-periphery structure in the overnight money market: evidence from the e-mid trading platform. *Computational Economics*, 1-37.

Furfine C., (1999) The microstructure of the federal funds market. *Financial Markets, Institutions, and Instruments* 8(5), 24-44.

Gabrieli, S. (2011). Too-interconnected versus too-big-to-fail: Banks' network centrality and overnight interest rates. Available at SSRN 1801390.

Gabrieli, S., and Georg, C. P. (2014). A network view on interbank market freezes. SSRN working paper 2544603.

Gai, P., and Kapadia, S. (2010). Contagion in financial networks. *Proceedings of the Royal Society Series A*, 466(2120), 2401-2423.

Gai, P., Haldane, A., and Kapadia, S. (2011). Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5), 453-470.

Gandy, A., and Veraart, L. (2015). A Bayesian methodology for systemic risk assessment in financial networks. SSRN working paper 2580869.

Georg, C.-P. (2013). The effect of the interbank network structure on contagion and common shocks. *Journal of Banking and Finance* 37, 2216-2228.

Gertler, M. and Kyotaki, N. (2010). Financial Intermediation and Credit Policy in Business Cycle Analysis," in B. M. Friedman and M. Woodford, *Handbook of Monetary Economics*, 3, 547-599.

- Glasserman, P., and Young, H. P. (2015). How likely is contagion in financial networks?. *Journal of Banking and Finance* 50, 383-399.
- Gofman, M. (2011, July). A network-based analysis of over-the-counter markets. In *AFA 2012 Chicago Meetings Paper*.
- Gofman, M. (2014). Efficiency and stability of a financial architecture with too interconnected to fail institutions. Available at SSRN.
- Gorton, G. B. (2010). *Slapped in the Face by the Invisible Hand: The Panic of 2007*, Oxford University Press.
- Gourieroux, C., Heam, J.C., and Monfort, A. (2014). Liquidation equilibrium with seniority and hidden CDO, *Journal of Banking and Finance* 37(12), 5261-5274.
- Hałaj, G., and Kok, C. (2013). Assessing interbank contagion using simulated networks. *Computational Management Science* 10(2-3), 157-186.
- Hansen, L.P. (2013) Challenges in Identifying and Measuring Systemic Risk. Working paper, University of Chicago.
- Haldane, A. G., and May, R. M. (2011). Systemic risk in banking ecosystems. *Nature* 469(7330), 351-355.
- Hurd, T. (2015). *Contagion! The Spread of Systemic Risk in Financial Networks*, Springer (forthcoming).
- Hüser, A.-C. (2015), Too interconnected to fail: a survey of the interbank networks literature. SAFE Working Paper 91, Goethe University, Frankfurt, Germany.
- Ibragimov, R., Jaffee, D., and Walden, J. (2011). Diversification disasters. *Journal of financial economics*, 99(2), 333-348.
- Inaoka, H., Ninomiya, T., Taniguchi, K., Shimizu, T., and Takayasu, H. (2004). Fractal Network derived from banking transaction—An analysis of network structures formed by financial institutions. *Bank of Japan Working Papers*, 4.
- In't Veld, D. L., Van der Leij, M., and Hommes, C. H. (2014). The formation of a core periphery structure in heterogeneous financial networks. SSRN working paper 2473131.
- James, C. (1991). The Losses Realized in Bank Failures. *Journal of Finance* 46, 1223-1242.
- Khandani, A. E., and Lo, A. W. (2011). What happened to the quants in August 2007? Evidence from factors and transactions data. *Journal of Financial Markets*, 14(1), 1-46.
- Martinez-Jaramillo, S., Alexandrova-Kabadjova, B., Bravo-Benitez, B., and Solórzano-Margain, J. P. (2014). An empirical study of the Mexican banking system's network and its implications for systemic risk. *Journal of Economic Dynamics and Control*, 40, 242-265.

- Mistrulli, P. E. (2011). Assessing financial contagion in the interbank market: Maximum entropy versus observed interbank lending patterns. *Journal of Banking and Finance* 35(5), 1114-1127.
- Newman, M. (2010). *Networks: An Introduction*. Oxford University Press, USA.
- Nier, E., Yang, J., Yorulmazer, T., and Alentorn, A. (2007). Network models and financial stability. *Journal of Economic Dynamics and Control* 31(6), 2033-2060.
- Office of Financial Research (2013). *Annual Report*. Available at www.treasury.gov.
- Pritsker, M. (2013). Knightian uncertainty and interbank lending. *Journal of Financial Intermediation* 22(1), 85-105.
- Puhr, C., Seliger, R., and Sigmund, M. (2012). Contagiousness and vulnerability in the Austrian interbank market. Oesterreichische Nationalbank Financial Stability Report, 24.
- Rogers, L. C. G., and Veraart, L. A. (2013). Failure and rescue in an interbank network. *Management Science* 59(4), 882-898.
- Sachs, A. (2014). Completeness, interconnectedness and distribution of interbank exposures—a parameterized analysis of the stability of financial networks. *Quantitative Finance*, in press.
- Santos, E. B., and Cont, R. (2010). *The Brazilian interbank network structure and systemic risk* (Working Paper No. 219), Banco Central do Brasil.
- Sapra, H., Plantin, G., and Shin, H. S. (2008). Marking to Market: Panacea or Pandora's Box? *Journal of Accounting Research*, 46, 435-460.
- Sheldon, G., and Maurer, M. (1998). Interbank Lending and Systemic Risk: An Empirical Analysis for Switzerland. *Swiss Journal of Economics and Statistics* 134(4), 685-704.
- Shin, H. S. (2010). Financial intermediation and the post-crisis financial system. Working paper 304, Bank for International Settlements, Basel, Switzerland.
- Soramäki, K., Bech, M. L., Arnold, J., Glass, R. J., and Beyeler, W. E. (2007). The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications*, 379(1), 317-333.
- Staum, J., and Liu, M. (2012). Systemic Risk Components in a Network Model of Contagion. SSRN working paper 1726107.
- Summer, M. (2013). Financial contagion and network analysis. *Annual Review of Financial Economics*, 5(1), 277-297.

Tarski, A. (1955). A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics* 5(2), 285-309.

van Lelyveld, I., and in 't Veld, D. (2014). Finding the core: Network structure in interbank markets. *Journal of Banking and Finance* 49, 27-40.

van Lelyveld, I., and Liedorp, F. (2006). Interbank contagion in the Dutch banking sector: A sensitivity analysis. *International Journal of Central Banking* 2(2), 99-133.

Upper, C., (2011). Simulation Methods to Assess the Danger of Contagion in Interbank Markets, *Journal of Financial Stability* 7, 111-125.

Upper, C., and Worms, A. (2004). Estimating bilateral exposures in the German interbank market: Is there a danger of contagion? *European Economic Review* 48(4), 827-849.

Vuillemeys, G., and Breton, R. (2014). Endogenous derivative networks. Banque de France.

Wagner, W. (2011). Systemic liquidation risk and the diversity–diversification trade-off. *Journal of Finance* 66, 1141–1175.

Wells, S. (2004). Financial interlinkages in the United Kingdom's interbank market and the risk of contagion. Working paper 230, Bank of England.

Zawadowski, A. (2013). Entangled financial systems. *Review of Financial Studies*.