Journal of Banking & Finance 50 (2015) 383–399

Contents lists available at ScienceDirect

Journal of Banking & Finance

journal homepage: www.elsevier.com/locate/jbf

How likely is contagion in financial networks?

Paul Glasserman [a,*], H. Peyton Young [b,c,d]

[a] Columbia Business School, Columbia University, United States
[b] Department of Economics, University of Oxford, United Kingdom
[c] Institute for New Economic Thinking, Oxford Martin School, United Kingdom
[d] Office of Financial Research, U.S. Treasury, United States

ARTICLE INFO

Article history:
Received 4 September 2013
Accepted 7 February 2014
Available online 25 February 2014

JEL classification:
D85
G21

Keywords:
Systemic risk
Contagion
Financial network

ABSTRACT

Interconnections among financial institutions create potential channels for contagion and amplification of shocks to the financial system. We estimate the extent to which interconnections increase expected losses and defaults under a wide range of shock distributions. In contrast to most work on financial networks, we assume only minimal information about network structure and rely instead on information about the individual institutions that are the nodes of the network. The key node-level quantities are asset size, leverage, and a financial connectivity measure given by the fraction of a financial institution's liabilities held by other financial institutions. We combine these measures to derive explicit bounds on the potential magnitude of network effects on contagion and loss amplification. Spillover effects are most significant when node sizes are heterogeneous and the originating node is highly leveraged and has high financial connectivity. Our results also highlight the importance of mechanisms that go beyond simple spillover effects to magnify shocks; these include bankruptcy costs, and mark-to-market losses resulting from credit quality deterioration or a loss of confidence. We illustrate the results with data on the European banking system.

1. Introduction

The interconnectedness of the modern financial system is widely viewed as having been a key contributing factor to the recent financial crisis. Due to the complex web of links between institutions, stresses to one part of the system can spread to others, leading to a system-wide threat to financial stability. Specific instances include the knock-on effects of the Lehman bankruptcy, potential losses to counterparties that would have resulted from a failure of the insurance company AIG, and more recently the exposure of European banks to the risk of sovereign default by some European countries. These and other examples highlight concerns that interconnectedness could pose a significant threat to the stability of the financial system.[1] Moreover there is a growing body of research that shows how this can happen in a theoretical sense.[2]

Although it is intuitively clear that interconnectedness has some effect on the transmission of shocks, it is less clear whether it significantly increases the likelihood and magnitude of losses compared to a financial system that is not interconnected. The contribution of this paper is to provide a general framework for analyzing this question. In contrast to much of the prior literature we do not subject the network to ad hoc shocks of different sizes. Instead we assume a full-fledged shock distribution and analyze the probability of default cascades and consequent losses of value that are attributable to network connections. A second distinguishing feature of our analysis is that, instead of estimating the absolute magnitude of default probabilities and losses, we estimate how much larger these quantities are in a networked system as compared to a similar system in which all links between financial institutions have been severed. In other words we estimate the extent to which defaults and losses are magnified by the interbank network over and above the original shocks.

It turns out that one can derive general bounds on the impact of the network with very little information about the network topology: our bounds hold independently of the degree distribution, node centrality, average path length, and so forth. This topology-free property of our results is one of the main contributions of the paper. We also show that these bounds hold for a wide range of distributions, including beta, exponential, normal, and many

* Corresponding author. Tel.: +1 212 854 4102.
E-mail addresses: pg20@columbia.edu (P. Glasserman), peyton.young@economics.ox.ac.uk (H.P. Young).

[1] See, for example, Bank of England (2011), International Monetary Fund (2012), and Office of Financial Research (2012).

[2] See in particular Allen and Gale (2000), Upper and Worms (2002), Degryse and Nguyen (2004), Goodhart et al. (2004), Elsinger et al. (2006), Allen and Babus (2009), Gai and Kapadia (2010), Gai et al. (2011), Haldane and May (2011), Upper (2011), Georg (2013), Rogers and Veraart (2013), Acemoglu et al. (2013), and Elliott et al. (2013).

http://dx.doi.org/10.1016/j.jbankfin.2014.02.006
0378-4266/

others. This robustness is important because detailed information about interbank liabilities is often unavailable and the exact form of the shock distributions is subject to considerable uncertainty. We are not claiming that the network topology is inconsequential, but that one can derive useful bounds on the financial system's susceptibility to contagion without knowing the details of the topology.

The starting point for our analysis is the elegant framework of Eisenberg and Noe (2001). Their model specifies a set of nodes that represent financial institutions together with the obligations between them. Given an initial shock to the balance sheets of one or more nodes, one can compute a set of payments that clear the network; that is, it provides a consistent way of valuing all the nodes conditional on an arbitrary shock to the system. This framework is very useful for analyzing how losses propagate through the financial system. A concrete example would be delinquencies in mortgage payments: if some fraction of a bank's mortgages are delinquent and it has insufficient reserves to cover the shortfall, then it will be unable to pay its creditors in full, who may be unable to pay their creditors in full, and so forth. The original shortfall in payments can cascade through the system, causing more and more banks to default through a domino effect. The Eisenberg–Noe framework shows how to compute a set of payments that clear the network, and it identifies which nodes default as a result of an initial shock to the system. The number and magnitude of such defaults depend on the network topology, and there is now a substantial literature characterizing those structures that tend to propagate default or alternatively that tend to dampen it (Gai and Kapadia, 2010; Gai et al., 2011; Haldane and May, 2011; Acemoglu et al., 2013; Elliott et al., 2013).

One limitation of the Eisenberg–Noe framework, as with most models of financial networks, is that it does not provide an account of link formation – that is, it does not model the dynamic process by which financial institutions enter into obligations to one another in the first place. This underscores the importance of having an estimation framework that does not rely too heavily on the specific features of the network, which is the subject of the present paper. We take the balance sheets of individual financial institutions as given and estimate how much they contribute to systemic effects over and above the impact of the initial shocks to asset values. In particular we shall examine the following two questions: How likely is it that a given set of nodes will default due to contagion from another node, as compared to the likelihood that they default from direct shocks to their own assets from sources outside the financial system, such as households and nonfinancial firms? And how much does the network increase the probability and magnitude of losses compared to a situation where there are no connections?

To compare systems with and without interconnections, we proceed as follows. First, we define the nodes to be financial institutions that borrow and lend on a significant scale, which together with their obligations to one another constitute the *financial network*. In addition, such institutions borrow and lend to the nonfinancial sector, which is composed of investors, households, and nonfinancial firms. We compare this system to one without connections that is constructed as follows. We remove all of the obligations between the financial nodes while keeping their links with the nonfinancial sector unchanged. We also keep node equity values as before by creating, for each node, a fictitious outside asset (or liability) whose value equals the net value of the connections at that node that were removed. We then apply the same shock distributions to both systems, with the shocks to real assets originating in the external sector and the fictitious assets (if any) assumed to be impervious to shocks. We can ascertain how much the network connections contribute to increased defaults and losses by comparing the outcomes in the two systems.

One might suppose that such a comparison is sensitive to the choice of shock distribution, but this turns out not to be the case: we show how to compute general bounds on the increased losses attributable to network connections that hold under a wide variety of distributions, including the beta, exponential, normal and many others. The bounds also hold whether the shocks are independent or positively associated and thus capture the possibility that institutions have portfolios that are exposed to common shocks (see for example Caccioli et al., 2012).

Two key findings emerge from this analysis, one concerning the probability of default cascades and the other concerning the expected losses from such cascades. We begin by computing the probability that default at a given node causes defaults at other nodes (via network spillovers), and compare this with the probability that all of these nodes default by direct shocks to their outside assets with no network transmission. We then derive a general formula that shows when the latter probability is larger than the former, in which case we say that contagion is *weak*. This characterization shows explicitly that substantial heterogeneity in node sizes makes a network more vulnerable to contagion through pure spillover effects. The network is particularly vulnerable to contagion when the originating node is large, highly leveraged, and, crucially, has a relatively high proportion of its obligations held by other financial institutions as opposed to the nonfinancial sector, what we will call high *financial connectivity*. These three factors – size, leverage, and financial connectivity – determine a *contagion index* for each institution that measures the potential impact of its failure on the rest of the financial sector.

Second, we apply our framework to estimate the expected system-wide loss in asset values that results from shocks that originate outside the financial sector. We derive a simple formula that compares the additional expected losses generated by the network with the expected losses that occur when the network links are severed. It turns out that the losses attributable to the network are typically quite modest under a wide range of shock distributions. Here, again, the network effect is highly dependent on the level of financial connectivity.

We emphasize that these results do not imply that *all* forms of network contagion are unimportant; rather they demonstrate that simple spillover or "domino" effects have only limited impact. These findings are consistent with the empirical and simulation literature on network stress testing, which finds that contagion is quite difficult to generate through the interbank spillover of losses (Degryse and Nguyen, 2004; Elsinger et al., 2006; Furfine, 2003; Georg, 2013; Nier et al., 2007). Put differently, our results show that contagion through spillover effects becomes most significant under the conditions described in Yellen (2013), when financial institutions inflate their balance sheets by increasing leverage and expanding interbank claims backed by a fixed set of real assets.

Indeed our results suggest that additional channels, aside from pure spillover effects, are needed to generate substantial losses from contagion. One such channel involves fire sales, in which firms dump assets on the market in order to cover their losses.[3] Another channel is the drying up of liquidity, which results when the default of one institution heightens uncertainty about the health of others, leading to a general tightening of credit. More generally, financial institutions may respond to changing market conditions in a variety of ways that exacerbate the impact of an initial negative shock and result in contagion. Although some of these dynamic effects are difficult to capture within an essentially static network

[3] See Shleifer and Vishny (2011) for a survey and Cifuentes et al. (2005) for an extension of the Eisenberg–Noe framework with fire sales.