Systemic Risk and Stability in Financial Networks*

Daron Acemoglu† Asuman Ozdaglar‡ Alireza Tahbaz-Salehi§

This version: December 2013
First version: December 2011

Abstract

We provide a framework for studying the relationship between the financial network architecture and the likelihood of systemic failures due to contagion of counterparty risk. We show that financial contagion exhibits a form of *phase transition* as the extent of interbank interconnectivity increases: as long as the magnitude and the number of negative shocks affecting financial institutions are sufficiently small, a more equal distribution of interbank obligations enhances the stability of the system. However, beyond a certain point, such dense interconnections start to serve as a mechanism for the propagation of shocks and lead to a more fragile financial system. Our results thus highlight the "robust-yet-fragile" nature of financial networks: the same features that make the system more resilient under certain conditions may function as significant sources of systemic risk and instability under another.

Keywords: Contagion, financial network, systemic risk, counterparty risk.
JEL Classification: G01, D85.

*We thank the editor and five anonymous referees for very helpful remarks and suggestions. We are grateful to Ali Jadbabaie for extensive comments and conversations. We also thank Fernando Alvarez, Gadi Barlevy, Ozan Candogan, Andrew Clausen, Paul Glasserman, Gary Gorton (discussant), Jennifer La'O (discussant), John Moore, Francesco Nava (discussant), Jean-Charles Rochet, Alp Simsek, Ali Shourideh, Larry Wall, and seminar participants at the Atlanta Fed, Becker-Friedman Institute, Chicago Booth, Columbia, Federal Reserve Board, The House of Finance, LBS, MIT, Penn State, Stanford GSB, Economic Networks and Banking Conference at LSE, Macroeconomic Financial Modeling Conference, and the Budapest-Essex Networks Meeting. Acemoglu and Ozdaglar gratefully acknowledge financial support from the Army Research Office, Grant MURI W911NF-12-1-0509.

†Department of Economics, Massachusetts Institute of Technology.
‡Laboratory for Information and Decision Systems, Massachusetts Institute of Technology.
§Columbia Business School, Columbia University.

1 Introduction

Since the global financial crisis of 2008, the view that the architecture of the financial system plays a central role in shaping systemic risk has become conventional wisdom. The intertwined nature of the financial markets has not only been proffered as an explanation for the spread of risk throughout the system,[1] but also motivated much of the policy actions both during and in the aftermath of the crisis.[2] Such views have even been incorporated into the new regulatory frameworks developed since.[3] Yet, the exact role played by the financial system's architecture in creating systemic risk remains, at best, imperfectly understood.

The current state of uncertainty about the nature and causes of systemic risk is reflected in the potentially conflicting views on the relationship between the structure of the financial network and the extent of financial contagion. Pioneering works by Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000) suggested that a more interconnected architecture enhances the resilience of the system to the insolvency of any individual bank. Allen and Gale, for example, argue that in a more densely interconnected financial network, the losses of a distressed bank are divided among more creditors, reducing the impact of negative shocks to individual institutions on the rest of the system. In contrast to this view, however, others have suggested that dense interconnections may function as a destabilizing force, paving the way to systemic failures. For example, Vivier-Lirimont (2006) argues that as the number of a bank's counterparties grows, the likelihood of a systemic collapse increases. This perspective is also shared by Blume et al. (2011) who model interbank contagion as an epidemic.

In view of the conflicting perspectives noted above, this paper provides a framework for studying the network's role as a shock propagation and amplification mechanism. Though stylized, our model is motivated by a financial system in which different institutions are linked to one another via unsecured debt contracts and hence are susceptible to counterparty risk. Our setup enables us to provide a number of theoretical results that highlight the implications of the network's structure on the extent of financial contagion and systemic risk.[4]

More concretely, we focus on an economy consisting of n financial institutions that lasts for three dates. In the initial date, banks borrow funds from one another to invest in projects that yield returns both in the intermediate and final dates. The liability structure that emerges from such interbank loans determines the *financial network*. In addition to its commitments to other financial institutions, each bank also has to make other payments with claims that are more senior to those of other banks. These claims may correspond to payments due to retail depositors or other types of commitments such as wages, taxes or claims by other senior creditors. We assume that the returns at the

[1] See, for example, Plosser (2009) and Yellen (2013).

[2] For an account of the policy actions during the crisis, see Sorkin (2009).

[3] An example of recent policies motivated by this perspective is the provision on "single counterparty exposure limits" in the Dodd-Frank Act, which attempts to prevent the distress at an institution from spreading to the rest of the system by limiting each firm's exposure to any single counterparty.

[4] In what follows, the stylized nature of our model notwithstanding, for ease of terminology, we refer to our network as a "financial network" and to its comprising entities as "financial institutions" or "banks".

final date are not pledgeable, so all debts have to be repaid at the intermediate date. Thus, a bank whose short-term returns are below a certain level may have to liquidate its project prematurely (i.e., before the final date returns are realized). If the proceeds from liquidations are also insufficient to pay all its debts, the bank defaults. Depending on the structure of the financial network, this may then trigger a cascade of failures: the default of a bank on its debt may cause the default of its creditor banks on their own counterparties, and so on.

The main focus of the paper is to study the extent of financial contagion as a function of the structure of interbank liabilities. By generalizing the results of Eisenberg and Noe (2001), we first show that, regardless of the structure of the financial network, *a payment equilibrium* — comprising of a mutually consistent collection of asset liquidations and repayments on interbank loans — always exists and is generically unique.

We then characterize the role of the structure of the financial network on the resilience of the system. We restrict our attention to *regular* financial networks in which the total claims and liabilities of all banks are equal. Such a normalization guarantees that any variation in the fragility of the system is due to the financial network's structure rather than size or leverage heterogeneity across banks.[5]

Our first set of results shows that when the magnitude of negative shocks is below a certain threshold, a result similar to those of Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000) holds: a more equal distribution of interbank liabilities leads to a less fragile financial system. In particular, the complete financial network, in which the liabilities of each institution are equally held by all other banks, is the configuration least prone to contagious defaults. At the opposite end of the spectrum, the ring network — a configuration in which all liabilities of a bank are held by a single counterparty — is the most fragile of all financial network structures. The intuition underlying these results is simple: a more equal distribution of interbank liabilities guarantees that the burden of any potential losses is shared among more counterparties and hence, in the presence of relatively small shocks, the excess liquidity of the non-distressed banks can be efficiently utilized in forestalling further defaults.

Our next set of results show that as the magnitude or the number of negative shocks cross certain thresholds, the types of financial networks that are most prone to contagious failures change dramatically. In particular, a more interconnected architecture is no longer a guarantee for stability. Rather, in the presence of large shocks, denser interbank liabilities facilitate financial contagion and create a more fragile system. On the other hand, "weakly connected" financial networks — in which different subsets of banks have minimal claims on one another — are significantly less fragile.[6] The intuition underlying such a sharp "phase transition" is that, with large negative shocks, the excess liquidity of the banking system may no longer be sufficient for absorbing the losses. Under such a scenario, a more sparse interbank network guarantees that the losses are shared with the senior

[5] The role of size differences across banks is discussed in Section 4.5.

[6] Such weakly connected financial networks are somewhat reminiscent of the old-style unit banking system, in which banks within a region are only weakly connected to the rest of the financial network, even though there might be strong intra-region ties.

creditors of the distressed banks, protecting the rest of the system.

Our results thus confirm a conjecture of Andrew Haldane (2009), the Executive Director for Financial Stability at the Bank of England, who suggested that highly interconnected financial networks may be "robust-yet-fragile" in the sense that "within a certain range, connections serve as shock-absorbers [and] connectivity engenders robustness." However, beyond a certain range, interconnections start to serve as a mechanism for the propagation of shocks, "the system [flips to] the wrong side of the knife-edge," and fragility prevails. More broadly, our results highlight that the same features that make a financial system more resilient under certain conditions may function as sources of systemic risk and instability under others.

In addition to illustrating the role of the network structure on the stability of the financial system, we introduce a new notion of distance over the financial network, the *harmonic distance*, which captures the susceptibility of each bank to the distress at any other. We show that, in the presence of large shocks, all banks whose harmonic distances to a distressed bank is below a certain threshold default. This characterization shows that, in contrast to what is often presumed in the empirical literature, various off-the-shelf (and popular) measures of network centrality — such as eigenvector or Bonacich centralities — may not be the right notions for identifying systemically important financial institutions. Rather, if the interbank interactions exhibit non-linearities similar to those induced by the presence of unsecured debt contracts, then it is the bank closest to all others according to our harmonic distance measure that may be "too-interconnected-to-fail".

Related Literature Our paper is part of a recent but growing literature that focuses on the role of the architecture of the financial system as an amplification mechanism. Kiyotaki and Moore (1997), Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000) provided some of the first formal models of contagion over networks. Using a multi-region version of Diamond and Dybvig (1983)'s model, Allen and Gale, for example, show that the interbank relations that emerge to pool region-specific shocks may at the same time create fragility in response to unanticipated shocks.[7] Dasgupta (2004) studies how the cross-holdings of deposits motivated by imperfectly correlated regional liquidity shocks can lead to contagious breakdowns. Shin (2008, 2009), on the other hand, constructs an accounting framework of the financial system as a network of interlinked balance sheets. He shows that securitization enables credit expansion through greater leverage of the financial system as a whole, drives down lending standards, and hence, increases fragility.

More recently, Allen, Babus, and Carletti (2012) have argued that the pattern of asset commonalities between different banks determines the extent of information contagion and hence, the likelihood of a systemic crisis. Also related is the work of Castiglionesi, Feriozzi, and Lorenzoni (2010), who show that a higher degree of financial integration leads to more stable interbank interest rates

[7] Allen and Gale also note that compared to a four-bank ring network, a pairwise-connected (and thus overall disconnected) network can be less prone to financial contagion originating from a single shock. Their paper, however, does not contain any of our results on the central role played by the size of the shocks in the fragility of the system (such as the phase transition of the complete network, flipping from being the most to the least stable and resilient financial network as the size of the shock grows).

in normal times, but to larger interest rate spikes during crises.[8] None of the above papers, however, provide a comprehensive analysis of the relationship between the structure of the financial network and the likelihood of systemic failures due to contagion of counterparty risk.

Our paper is also related to several recent, independent works, such as Elliott, Golub, and Jackson (2013) and Cabrales, Gottardi, and Vega-Redondo (2013), that study the broad question of propagation of shocks in a network of firms with financial interdependencies. These papers, however, focus on a contagion mechanism different from ours. In particular, they study whether and how cross-holdings of different organizations' shares or assets may lead to cascading failures. Elliott et al. (2013) consider a model with cross-ownership of equity shares and show that in the presence of bankruptcy costs, a firm's default may induce losses on all firms owning its equity; hence, triggering a chain reaction. On the other hand, Cabrales et al. (2013) study how securitization — modeled as exchange of assets among firms — may lead to the instability of the financial system as a whole. Our work, in contrast, focuses on the likelihood of systemic failures due to contagion of counterparty risk.

Focusing on financial contagion through direct contractual linkages, Alvarez and Barlevy (2013) use a model similar to ours to study the welfare implications of a policy of mandatory disclosure of information in the presence of counterparty risk. Glasserman and Young (2013) also rely on a similar model, but rather than characterizing the fragility of the system as a function of the financial network's structure, they provide a network-independent bound on the probability of financial contagion.[9] Our paper is also related to Eboli (2012), who studies the extent of contagion in some classes of networks. In contrast to our paper, his main focus is on the indeterminacy of interbank payments in the presence of cyclical entanglement of assets and liabilities.[10]

Gai, Haldane, and Kapadia (2011) also study a network model of interbank lending with unsecured claims. Using numerical simulations, they show how greater complexity and concentration in the financial network may amplify the fragility of the system.

The role of the networks as shock propagation and amplification mechanisms has also been studied in the context of production relations in the real economy. Focusing on the input-output linkages between different sectors, Acemoglu, Carvalho, Ozdaglar, and Tahbaz-Salehi (2012) and Acemoglu, Ozdaglar, and Tahbaz-Salehi (2013a) show that in the presence of linear (or log-linear) economic interactions, the volatility of aggregate output and the likelihood of large economic down-

[8]Other related contributions include Rochet and Tirole (1996), Cifuentes, Ferrucci, and Shin (2005), Leitner (2005), Rotemberg (2011), Zawadowski (2011), Gofman (2011, 2013), Caballero and Simsek (2013), Denbee, Julliard, Li, and Yuan (2013), Cohen-Cole, Patacchini, and Zenou (2013), and Georg (2013). For a more detailed discussion of the literature, see the survey by Allen and Babus (2009). A more recent and smaller literature focuses on the formation of financial networks. Examples include Babus (2013), Zawadowski (2013) and Farboodi (2013), as well as the working paper version of the current work (Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2013b).

[9]In addition to focusing on different questions, the liability structures of the financial institutions are also different in the two papers. In particular, due to the absence of the outside senior claims, the model of Glasserman and Young (2013) imposes an implicit upper bound on the size of the negative shocks, essentially limiting the extent of contagion.

[10]A different strand of literature studies the possibility of *indirect* spillovers in the financial markets. In particular, rather than taking place through direct contractual relations as in our paper, the amplification mechanisms studied in this literature work through the endogenous responses of various market participants. Examples include Holmström and Tirole (1998), Brunnermeier and Pedersen (2005), Lorenzoni (2008) and Krishnamurthy (2010). For a recent survey, see Brunnermeier and Oehmke (2013).

turns are independent of the sparseness or denseness of connections, but rather depend on the extent of asymmetry in different entities' interconnectivity. The contrast between the insights on prorogation of shocks in production economies with (log) linear interactions and those in the presence of default (due to debt-like financial instruments) presented in this paper highlights that the role of networks in contagion crucially depends on the nature of economic interactions between different entities that constitute the network.

Outline of the Paper The rest of the paper is organized as follows. Our model is presented in Section 2. In Section 3, we define our solution concept and show that a payment equilibrium always exists and is generically unique. Section 4 contains our results on the relationship between the extent of financial contagion and the network structure. Section 5 concludes. A discussion on the properties of the harmonic distance and the proofs are presented in the Appendix, while an Online Appendix contains several omitted proofs.

2 Model

2.1 Financial Institutions

Consider a single-good economy, consisting of n risk-neutral banks indexed by $N = \{1, 2, \dots, n\}$. The economy lasts for three dates, $t = 0, 1, 2$. At the initial date, each bank i is endowed with k_i units of capital that it can either hoard as cash (denoted by c_i), lend to other banks, or invest in a project that yields returns at the intermediate and final dates. More specifically, bank i's project yields a random return of z_i at $t = 1$, and if held to maturity, a fixed, non-pledgeable long-term return of A at $t = 2$. The bank can (partially) liquidate its project prematurely at $t = 1$, but can only recover a fraction $\zeta < 1$ of its full value.

Interbank lending takes place through *standard debt contracts* signed at $t = 0$. Let k_{ij} denote the amount of capital borrowed by bank j from bank i. The face value of j's debt to i is thus equal to $y_{ij} = R_{ij}k_{ij}$, where R_{ij} is the corresponding interest rate.[11] In addition to its liabilities to other banks, each bank must also meet an outside obligation of magnitude $v > 0$ at $t = 1$, which is assumed to have seniority relative to its other liabilities. These more senior commitments may be claims by the bank's retail depositors, senior debts (to other financial institutions outside the network), wages due to its workers or taxes due to the government. The total liabilities of bank i is thus equal to $y_i + v$, where $y_i = \sum_{j \neq i} y_{ji}$.[12]

Given our assumption that the long-term returns are not pledgeable, all debts have to be cleared at date $t = 1$. If bank j is unable to meet its $t = 1$ liabilities in full, it has to liquidate its project

[11] In this version of the paper, we take the interbank lending decisions and the corresponding interest rates as given, and thus do not formally treat the banks' actions at $t = 0$. This stage of the game is studied in detail in our working paper version (Acemoglu et al., 2013b).

[12] This formulation also allows for liabilities to outside financial institutions that have the same level of seniority as interbank loans by simply setting one of the banks, say bank n, to have claims but no inside-the-network liabilities, i.e., $y_{in} = 0$ for all $i \neq n$.

prematurely (in part or in full), where the proceeds are distributed among its creditors. We assume that all junior creditors — that is, the other banks — are of equal seniority. Hence, if bank j can meet its senior liabilities, v, but defaults on its debt to the junior creditors, they are repaid in proportion to the face value of the contracts. On the other hand, if j cannot meet its more senior outside liabilities v, its junior creditors receive nothing.[13]

2.2 The Financial Network

The lending decisions of the banks and the resulting counterparty relations can be represented by an interbank network. In particular, we define the *financial network* corresponding to the bilateral debt contracts in the economy as a weighted, directed graph on n vertices, where each vertex corresponds to a bank and a directed edge from vertex j to vertex i is present if bank i is a creditor of bank j. The weight assigned to this edge is equal to y_{ij}, the face value of the contract between the two banks. Throughout the paper, we denote a financial network with the collection of interbank liabilities $\{y_{ij}\}$.

We say a financial network is *symmetric*, if $y_{ij} = y_{ji}$ for all pairs of banks i and j. On the other hand, a financial network is said to be *regular* if all banks have identical interbank claims and liabilities; i.e., $\sum_{j \neq i} y_{ij} = \sum_{j \neq i} y_{ji} = y$ for some y and all banks i. Figures 1(a) and 1(b) illustrate two regular financial networks, known as the *ring* and the *complete* networks, respectively. The ring financial network represents a configurations in which bank $i > 1$ is the sole creditor of bank $i-1$ and bank 1 is the sole creditor of bank n; that is, $y_{i,i-1} = y_{1,n} = y$. Hence, for a given value of y, the ring network is the regular financial network with the sparsest connections. In contrast, in the complete network, the liabilities of each banks are held equally by all others; that is, $y_{ij} = y/(n-1)$ for all $i \neq j$, implying that the interbank connections in such a network are maximally dense.

Finally, we define,

Definition 1. The financial network $\{\tilde{y}_{ij}\}$ is a γ-*convex combination* of financial networks $\{y_{ij}\}$ and $\{\hat{y}_{ij}\}$ if there exists $\gamma \in [0, 1]$ such that $\tilde{y}_{ij} = (1-\gamma)y_{ij} + \gamma\hat{y}_{ij}$ for all banks i and j.

Thus, for example, a financial network that is a γ-convex combination of the ring and the complete financial networks exhibits an intermediate degree of density of connections: as γ increases, the financial network approaches the complete financial network.

3 Payment Equilibrium

The ability of a bank to fulfill its promises to its creditors depends on the resources it has available to meet those liabilities, which include not only the returns on its investment and the cash at hand, but also the realized value of repayments by the bank's debtors. In this section, we show that a mutually

[13] As we will discuss in detail in Section 4.2, the presence of these senior claims is required to guarantee the possibility of the systemic event in which all banks default. In particular, in the absence of such senior claims that divert some of the funds out of the system, at least one bank would always end up with a positive balance, and hence, avoid default.

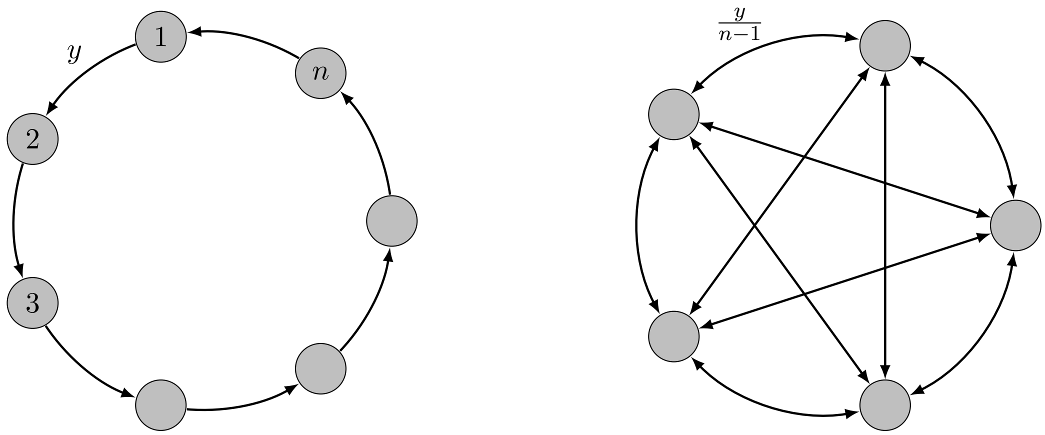

(a) The ring financial network (b) The complete financial network

Figure 1. The ring and the complete financial networks

consistent collection of repayments on interbank loans and asset liquidations always exists and is generically unique.

Let x_{js} denote the repayment by bank s on its debt to bank j at $t = 1$. By definition, $x_{js} \in [0, y_{js}]$. The total cash flow of bank j when it does not liquidate its project is thus equal to $h_j = c_j + z_j + \sum_{s \neq j} x_{js}$, where c_j is the cash held by the bank at the initial date. If h_j is larger than the bank's total liabilities, $v + y_j$, then the bank is capable of meeting its liabilities in full, and as a result, $x_{ij} = y_{ij}$ for all $i \neq j$. If, on the other hand, $h_j < v + y_j$, the bank needs to start liquidating its project in order to avoid default. Given that liquidation is costly, the bank liquidates its project up to the point where it can cover the shortfall $v + y_j - h_j$, or otherwise in its entirety to pay back its creditors as much as possible. Mathematically, the bank's liquidation decision, $\ell_j \in [0, A]$, is given by

$$\ell_j = \left[\min\left\{\frac{1}{\zeta}(v + y_j - h_j), A\right\}\right]^+, \tag{1}$$

where $[\cdot]^+$ stands for $\max\{\cdot, 0\}$ and guarantees that the bank does not liquidate its project if it can meet its liabilities with a combination of the cash it holds, the short-term return on its project, and the repayment by its debtor banks.

If the bank cannot pay its debts in full even with the full liquidation of its project, it defaults and its creditors are repaid according to their seniority. If $h_j + \zeta A$ is less than v, the bank defaults on its senior liabilities and its junior creditors receive nothing; that is, $x_{ij} = 0$. On the other hand, if $h_j + \zeta A \in (v, v + y_j)$, senior liabilities are paid in full and the junior creditors are repaid in proportion to the face value of their contracts. Thus, the $t = 1$ payment of bank j to a creditor bank i is equal to

$$x_{ij} = \frac{y_{ij}}{y_j}\left[\min\{y_j, h_j + \zeta\ell_j - v\}\right]^+, \tag{2}$$

where recall that $h_j = c_j + z_j + \sum_{s \neq j} x_{js}$ denotes the funds available to the bank in the absence of any liquidation and ℓ_j is its liquidation decision given by (1). Thus, equations (1) and (2) together determine the liquidation decision and the debt repayments of bank j as a function of its debtors' repayments on their own liabilities.

Definition 2. For a given realization of the projects' short-term returns and the cash available to the banks, the collection $(\{x_{ij}\}, \{\ell_i\})$ of interbank debt repayments and liquidation decisions is a *payment equilibrium* of the financial network if (1) and (2) are satisfied for all i and j simultaneously.

A payment equilibrium is thus a collection of mutually consistent interbank payments and liquidations at $t = 1$. The notion of payment equilibrium in our model is a generalization of the notion of a clearing vector introduced by Eisenberg and Noe (2001) and utilized by Shin (2008, 2009). In contrast to these papers, however, banks in our model not only have financial liabilities of different seniorities, but also can obtain extra proceeds by (partially or completely) liquidating their long-term projects.

Via equations (1) and (2), the payment equilibrium captures the possibility of financial contagion in the financial system. In particular, given the interdependence of interbank payments across the network, a (sufficiently large) negative shock to a bank not only leads to that bank's default, but may also initiate a cascade of failures, spreading to its creditors, its creditors' creditors, and so on. The next proposition shows that, regardless of the structure of the financial network, the payment equilibrium always exists and is uniquely determined over a generic set of parameter values and shock realizations.[14]

Proposition 1. *For any given financial network, cash holdings, and realization of the shocks, a payment equilibrium always exists and is generically unique.*

Finally, for any given financial network and the corresponding payment equilibrium, we define the (utilitarian) social surplus in the economy as the sum of the returns to all agents; that is,

$$u = \sum_{i=1}^{n} (\pi_i + T_i),$$

where $T_i \leq v$ is the transfer from bank i to its senior creditors and π_i is the bank's profit.

4 Financial Contagion

As discussed above, the interdependence of interbank payments over the network implies that distress at a single bank may induce a cascade of defaults throughout the financial system. In this section, we study how the structure of the financial network determines the extent of contagion.

[14] As we show in the proof of Proposition 1, in any connected financial network, the payment equilibrium is unique as long as $\sum_{j=1}^{n}(z_j + c_j) \neq nv - n\zeta A$. In the non-generic case in which $\sum_{j=1}^{n}(z_j + c_j) = nv - n\zeta A$, there may exist a continuum of payment equilibria, in almost all of which banks default due to "coordination failures". For example, if the economy consists of two banks with $c_1 = c_2 = v$, bilateral contracts of face values $y_{12} = y_{21}$, no shocks and no proceeds from liquidation (that is, $\zeta = 0$), then defaults can occur if banks do not pay one another, even though both are solvent. See Alvarez and Barlevy (2013) for a similar characterization in financial networks with some weak form of symmetry.

For most of our analysis, we focus on regular financial networks in which the total claims and liabilities of all banks are equal. Such a normalization guarantees that any variation in the fragility of the system is simply due to how interbank liabilities are distributed, while abstracting away from effects that are driven by other features of the financial network, such as size or leverage heterogeneity across banks.[15]

To simplify the analysis and the exposition of our results, we also assume that the short-term returns on the banks' investments are independent and can only take two values $z_i \in \{a, a-\epsilon\}$, where $a > v$ is the return in the "business as usual" regime and $\epsilon \in (a - v + \zeta A, a)$ corresponds to the magnitude of a negative shock. The upper bound on ϵ simply implies that the returns of the project is always positive, whereas the lower bound guarantees that absent any payments by other banks, a "distressed bank" — that is, a bank directly hit by the negative shock — would not be able to pay its senior creditors. Finally, in what follows we assume that all banks hold the same amount of cash, which we normalize to zero.

Proposition 2. *Conditional on the realization of p negative shocks, the social surplus in the economy is equal to*

$$u = n(a + A) - p\epsilon - (1 - \zeta)\sum_{i=1}^{n} \ell_i.$$

As expected, the social surplus is decreasing in the extent of liquidation in the corresponding payment equilibrium. In particular, in the case that proceeds from liquidation are "trivial", that is, $\zeta = 0$, the social surplus is simply determined by the number of bank failures, that is,[16]

$$u = na - p\epsilon + (n - \sharp\text{defaults})A.$$

Under this assumption, it is natural to measure the performance of a financial network in terms of the number of banks in default.

Definition 3. Consider two regular financial networks $\{y_{ij}\}$ and $\{\tilde{y}_{ij}\}$. Conditional on the realization of p negative shocks,

(i) $\{y_{ij}\}$ is more *stable* than $\{\tilde{y}_{ij}\}$ if $\mathbb{E}_p u \geq \mathbb{E}_p \tilde{u}$, where $\mathbb{E}_p$ is the expectation conditional on the realization of p negative shocks.

(ii) $\{y_{ij}\}$ is more *resilient* than $\{\tilde{y}_{ij}\}$ if $\min u \geq \min \tilde{u}$, where the minimum is taken over all possible realizations of the p negative shocks.

Thus, stability and resilience capture the expected and worst-case performances of the financial network in the presence of p negative shocks, respectively. Clearly, both measures of performance not only depend on the number (p) and the size (ϵ) of the realized shocks, but also on the structure

[15] For example, Acemoglu et al. (2012) show that asymmetry in the degree of interconnectivity of different industries as input suppliers in the real economy plays a crucial role in the propagation of shocks.

[16] Here $\zeta = 0$ stands for $\zeta \to 0$, since in the limit where $\zeta = 0$ there is no economic reason for liquidation and in fact, equation (1) is not well-defined.

of the financial network. To illustrate the relation between the extent of financial contagion and the network structure in the most transparent manner, we initially assume that exactly one bank is hit with a negative shock and that the proceeds from liquidations are trivial, i.e., $p = 1$ and $\zeta = 0$. We relax these assumptions in Sections 4.3 and 4.4.

4.1 Small Shock Regime

We first characterize the fragility of different financial networks when the size of the negative shock is less than a critical threshold.

Proposition 3. *Let $\epsilon^* = n(a - v)$ and suppose that $\epsilon < \epsilon^*$. Then, there exists y^* such that for $y > y^*$,*

(a) The ring network is the least resilient and least stable financial network.

(b) The complete network is the most resilient and most stable financial network.

(c) The γ-convex combination of the ring and complete networks becomes (weakly) more stable and resilient as γ increases.

The above proposition thus establishes that as long as the size of the negative shock is below be critical threshold ϵ^*, the ring is the financial network most prone to financial contagion, whereas the complete network is the least fragile. Furthermore, a more equal distribution of interbank liabilities leads to less fragility. Proposition 3 is thus in line with, and generalizes, the observations made by Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000). The underlying intuition is that a more equal distribution of interbank liabilities implies that the burden of any potential losses is shared among more banks, creating a more robust financial system. In particular, in the extreme case of the complete financial network, the losses of a distressed bank are divided among as many creditors as possible, guaranteeing that the excess liquidity in the financial system can fully absorb the transmitted losses. On the other hand, in the ring financial network, the losses of the distressed bank — rather than being divided up between multiple counterparties — are fully transferred to its immediate creditor, leading to the creditor's possible default.

The condition that $\epsilon < \epsilon^*$ means that the size of the negative shock is less than the total "excess liquidity" available to the financial network as a whole.[17] Proposition 3 also requires that interbank liabilities (and claims) are above a certain threshold y^*, which is natural given that for small values of y, no contagion would occur, regardless of the structure of financial network.

The extreme fragility of the ring financial network established by Proposition 3 is in contrast with the results of Acemoglu et al. (2012, 2013a), who show that if the interactions over the network are linear (or log linear), the ring is as stable as any other regular network structure. This contrast reflects the fact that, with linear interactions, negative and positive shocks cancel each other out in exactly the same way, independently of the structure of the network. However, the often non-linear

[17] Recall that in the absence of any shock, $a - v$ is the liquidity available to each bank after meeting its liabilities to the senior creditors outside the network.

nature of financial interactions (captured in our model by the presence of unsecured debt contracts) implies that the effects of negative and positive shocks are not necessarily symmetric. Stability and resilience are thus achieved by minimizing the impact of the distress at any given bank on the rest of the system. The ring financial network is highly fragile precisely because the adverse effects of a negative shock to any bank are fully transmitted to the bank's immediate creditor, triggering maximal financial contagion. In contrast, a more equal distribution of interbank liabilities reduces the impact of a bank's distress on any single counterparty.

Our next result shows that this intuition extends to a broad set of network structures.

Definition 4. For two given subset of banks S and T, the financial network $\{\tilde{y}_{ij}\}$ is a (S, T, γ)-*rewiring* of the regular financial network $\{y_{ij}\}$ if

$$\tilde{y}_{ij} = \begin{cases} (1-\gamma)y_{ij} + \dfrac{\gamma}{|S^c|}\displaystyle\sum_{k \in S^c} y_{kj} & \text{if } i \in S^c, j \in S. \\ y_{ij} & \text{if } i, j \in T. \end{cases}$$

In other words, as γ increases, the liabilities of banks in S to banks in S^c become more equally distributed, while the liabilities of banks in T to one another remain unchanged.[18] This notion of rewiring is distinct from a γ-convex combination with the complete network, according to which the liabilities of *all* banks become more equally distributed as γ increases.

Proposition 4. *Suppose that $\epsilon < \epsilon^*$ and $y > y^*$. For a given financial network, let D denote the set of banks in default and suppose that the distressed bank can meet its liabilities to its senior creditors. Then, any (D, D, γ)-rewiring of the financial network does not increase the number of defaults.*

The intuition behind this result is similar to that of Proposition 3: a rewiring of a financial network that spreads the financial liabilities of banks in default to the rest of the system guarantees that the excess liquidity available to the non-distressed banks are utilized more effectively. In the presence of small enough shocks, this can never lead to more defaults. An immediate corollary of this proposition extends Proposition 3 to the γ-convex combination of a given network with the complete network.

Corollary 1. *Suppose that $\epsilon < \epsilon^*$ and $y > y^*$. If there is no contagion in a financial network, then there is no contagion in any γ-convex combination of that network and the complete network.*

Our results thus far show that as long as $\epsilon < \epsilon^*$, a more equal distribution of interbank liabilities, formalized by the notions of convex combinations and rewirings, can never increase — but may reduce — the fragility of an already stable financial network. Our next example, however, illustrates that not all transformations that equalize interbank liabilities lead to a less fragile system.

[18] This definition does not put any extra restrictions on the liabilities of banks in S^c to those in S or on those of banks in T^c to one another. Thus, the rewiring of a financial network is not necessarily unique.

Example 1. Consider the financial network depicted in Figure 2, in which interbank liabilities are given by

$$y_{i,i+1} = y_{i+1,i} = \begin{cases} qy & \text{if } i \text{ odd} \\ (1-q)y & \text{if } i \text{ even} \end{cases}$$

where $1/2 \le q < 1$ and $y > y^*$; i.e, the financial network consists of pairs of interconnected banks located on a ring-like structure, with weaker inter-pair liabilities. The liabilities of a given bank i to banks $i-1$ and $i+1$ become more equalized as q approaches $1/2$.

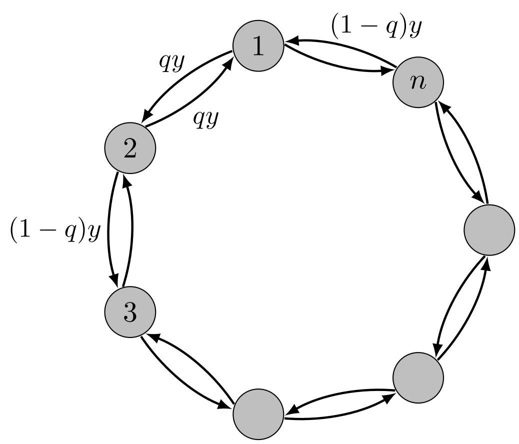

Figure 2. The liabilities of each bank to its two counterparties become more equalized as q approaches $1/2$.

Now suppose that bank 1 is hit with a negative shock of size $\epsilon = (3+\omega)(a-v)$ for some small $\omega > 0$. If $q = 1/2$, then, by symmetry, banks 1, 2 and n cannot meet their liabilities in full, and in particular, banks 2 and n default due to a small shortfall of size $\omega(a-v)$. However, given that n is just at the verge of solvency, increasing q slightly above $1/2$ guarantees that bank n no longer defaults, as a larger fraction of the losses would now be transferred to bank 2. More specifically, one can show that if $q = (1+\omega)/2$, then only banks 1 and 2 default, whereas all other banks can meet their liabilities in full.

To summarize, though Propositions 3(c) and 4 and Corollary 1 show that, in the presence of small shocks, γ-convex combinations with the complete network and various rewirings increase the stability and the resilience of the financial system, the same logic does not apply to all transformations that equalize interbank liabilities. For instance, in the preceding example, lower values of q which make the liabilities of a given bank to its two counterparties more equal, may nevertheless increase systemic risk by transferring resources away from the bank that relies on them for survival.

4.2 Large Shock Regime

Propositions 3 and 4, along with Corollary 1, show that as long as the magnitude of the negative shock is below the threshold ϵ^*, a more equal distribution of interbank liabilities leads to less fragility. In particular, the complete network is the most stable and resilient financial network: except for the bank that is directly hit with the negative shock, no other bank defaults. Our next set of results, however, shows that when the magnitude of the shock is above the critical threshold ϵ^*, this picture changes dramatically.

Definition 5. A regular financial network is δ*-connected* if there exists a collection of banks $S \subset N$ such that $\max\{y_{ij}, y_{ji}\} \leq \delta y$ for all $i \in S$ and $j \notin S$.

In other words, in a δ-connected financial network, the fraction of liabilities of banks inside and outside of S to one another is no more than $\delta \in [0,1]$. Hence, for small values of δ, the banks in S have weak ties — in terms of both claims and liabilities — to the rest of the financial network. We have the following result:

Proposition 5. *Suppose that* $\epsilon > \epsilon^*$ *and* $y > y^*$. *Then,*

(a) The complete and the ring networks are the least stable and least resilient financial networks.

(b) For small enough values of δ, *any* δ*-connected financial network is strictly more stable and resilient than the ring and complete financial networks.*

Thus, when the magnitude of the negative shock crosses the critical threshold ϵ^*, the complete network exhibits a form of *phase transition*: it flips from being the most to the least stable and resilient network, achieving the same level of fragility as the ring network. In particular, when $\epsilon > \epsilon^*$, all banks in the complete network default. The intuition behind this result is simple: since all banks in the complete network are creditors of the distressed bank, the adverse effects of the negative shock are directly transmitted to them. Thus, when the size of the negative shock is large enough, all banks — including those originally unaffected by the negative shock — default.

Not all financial systems, however, are as fragile in the presence of large shocks. In fact, as part (b) shows, for small enough values of δ, any δ-connected financial network is strictly more stable and resilient than both the complete and the ring networks. The presence of such "weakly connected" components in the network guarantees that the losses — rather than being transmitted to all other banks — are borne in part by the distressed bank's senior creditors.

Taken together, Propositions 3 and 5 illustrate the "robust-yet-fragile" property of highly interconnected financial networks conjectured by Haldane (2009). They show that more densely interconnected financial networks, epitomized by the complete network, are more stable and resilient in response to a range of shocks. However, once we move outside this range, these dense interconnections act as a channel through which shocks to a subset of the financial institutions transmit to the entire system, creating a vehicle for instability and systemic risk.

The intuition behind such a phase transition is related to the presence of two types of "shock absorbers" in our model, each of which is capable of reducing the extent of contagion in the network. The first absorber is the excess liquidity, $a - v > 0$, of the non-distressed banks at $t = 1$: the impact of a shock is attenuated once it reaches banks with excess liquidity. This mechanism is utilized more effectively when the financial network is more "complete", an observation in line with the results of

Allen and Gale (2000) and Freixas, Parigi, and Rochet (2000). However, the claim v of senior creditors of the distressed bank also function as a shock absorption mechanism. Rather than transmitting the shocks to other banks in the system, the senior creditors can be forced to bear (some of) the losses and hence, limit the extent of contagion. In contrast to the first mechanism, this shock absorption mechanism is best utilized in weakly connected financial networks and is the least effective in the complete network. Thus, when the shock is so large that it cannot be fully absorbed by the excess liquidity in the system — which is exactly when $\epsilon > \epsilon^*$ — financial networks that significantly utilize the second absorber are less fragile. This intuition extends to the rewirings of an arbitrary financial network:

Proposition 6. *Suppose that $\epsilon > \epsilon^*$ and $y > y^*$. Also suppose that bank j is hit with the negative shock and let D denote the set of banks other than j that default. Then, any $(\{j\}, D, \gamma)$-rewiring of the financial network does not decrease the number of defaults.*

The above result which is the large shock counterpart to Proposition 4 shows that, in contrast to the small shock regime, rewirings that lead to a more equal distribution of the the distressed bank's liabilities do not reduce — but may increase — the extent of contagion.

The remainder of this subsection provides a characterization of the set of banks that default in a general financial network and shows that the intuition on the role of interconnectivity in the fragility of the system remains valid for a broad set of network structures. We first define a new notion of distance over the financial network.

Definition 6. The *harmonic distance* from bank i to bank j is

$$m_{ij} = 1 + \sum_{k \neq j} \left(\frac{y_{ik}}{y} \right) m_{kj}, \tag{3}$$

with the convention that $m_{ii} = 0$ for all i.[19]

The harmonic distance from bank i to bank j depends not only on how far each of its immediate debtors from j are, but also on the intensity of their liabilities to i. Such a definition implies that the harmonic distance between any pair of banks can be considerably different from the shortest-path, geodesic distance defined over the financial network. In particular, the more direct or indirect liability chains exist between banks i and j, the closer the two banks are to one another.

Proposition 7. *Suppose that bank j, hit with the negative shock, defaults on its senior liabilities. Then, there exists m^* such that,*

(a) If $m_{ij} < m^*$, then bank i defaults.

(b) If all banks in the financial network default, then $m_{ij} < m^*$ for all i.

[19]Strictly speaking, the harmonic distance is a *quasimetric*, as it does not satisfy the symmetry axiom (that is, in general, $m_{ij} \neq m_{ji}$). Nevertheless, for ease of reference, we simply refer to m_{ij} as the distance from bank i to bank j. For a discussion on the properties of the harmonic distance, see Appendix A.

This result implies that banks that are closer to the distressed bank in the sense of the harmonic distance are more vulnerable to default. Consequently, financial networks in which the pairwise harmonic distances between any pairs of banks is smaller are less stable and resilient in the presence of large shocks.[20] Proposition 7 generalizes Proposition 5. In particular, one can verify that the harmonic distance between any pair of banks is minimized in the complete financial network as predicted by Proposition 5(a).[21] On the other hand, in a δ-connected network (for sufficiently small δ), there always exists a pair of banks whose pairwise harmonic distance is greater than m^*, ensuring that the network is strictly more stable and resilient than the complete financial network, and thus establishing Proposition 5(b) as a corollary.

Proposition 7 also highlights that in a given financial network, the bank that is closest to all others in the sense of harmonic distance is the most "systemically important" financial institution: a shock to such a bank would lead to the maximal number of defaults. This observation contrasts with much of the recent empirical literature that relies on off-the-shelf measures of network centrality — such as eigenvector or Bonacich centralities — for identifying systemically important financial institutions.[22] Such standard network centrality measures would be appropriate if interbank interactions are linear. In contrast, Proposition 7 shows that if interbank interactions exhibit non-linearities similar to those induced by the presence of debt contracts, it is the harmonic distances of other banks to a financial institution that determine its importance from a systemic perspective.[23]

Our last result in this subsection relates the interbank harmonic distances to an intuitive structural property of the financial network.

Definition 7. The *bottleneck parameter* of a financial network is

$$\phi = \min_{S \subseteq N} \sum_{i \in S} \sum_{j \notin S} \frac{(y_{ij}/y)}{|S||S^c|}.$$

Roughly speaking, ϕ quantifies how the financial network can be partitioned into two roughly equally-sized components, while minimizing the extent of interconnectivity between the two.[24] In particular, for a given partition of the financial network into two subsets of banks, S and S^c, the quantity $\sum_{i \in S, j \notin S} y_{ij}$ is equal to the total liabilities of banks in S^c to those in S. The bottleneck parameter thus measures the minimal extent of interconnectivity between the banks in any partition (S, S^c), while ensuring that the neither set is significantly smaller than the other (see Figure 3). Thus, a highly interconnected financial network, such as the complete network, exhibits a large bottleneck parameter, whereas $\phi = 0$ for any disconnected network. We have the following result:

[20] Note that $\epsilon > \epsilon^*$ implies that the bank hit by the negative shock defaults on its senior creditors. For a proof, see Lemma B.4 in the Appendix.

[21] See Appendix A.

[22] See, for example, Bech, Chapman, and Garratt (2010), Bech and Atalay (2010), Akram and Christophersen (2010), Bisias, Flood, Lo, and Valavanis (2012), and Craig, Fecht, and Tümer-Alkan (2013).

[23] In their numerical simulations, Soramäki and Cook (2013) make a similar observation and propose a measure of relative importance of banks that is related to our notion of harmonic distance.

[24] The bottleneck parameter is closely related to the notion of conductance in spectral graph theory. It can also be interpreted as the discrete counterpart of the isoperimetric constant of a compact Riemannian manifold in Riemannian geometry (Chung, 1997).

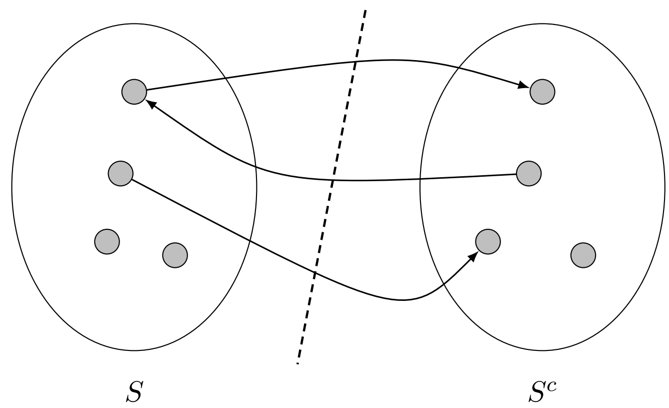

Figure 3. A partition of the financial network into subsets S and S^c. Removing the liabilities over the dashed line disconnects the network. The bottleneck parameter measures how easily one can disconnect the network into two subsets of roughly equal sizes, by removing as few interbank liabilities as possible.

Lemma 1. *For any symmetric financial network,*

$$\frac{1}{2n\phi} \leq \max_{i \neq j} m_{ij} \leq \frac{16}{n\phi^2}, \tag{4}$$

where ϕ is the corresponding bottleneck parameter.

The above lemma thus provides bounds on the maximum harmonic distance between any pairs of banks in the financial network in terms of the network's bottleneck parameter. More importantly, it shows that the the relationship between the extent of interbank connectivity and the financial network's fragility discussed after Proposition 5 holds for a broad set of network structures. In particular, the interbank harmonic distances are smaller when the financial network is more interconnected, guaranteeing more defaults in the presence of a large shock. The following corollary to Proposition 7 and Lemma 1 formalizes this observation:

Corollary 2. *Suppose that $\epsilon > \epsilon^*$. Then, there exist constants $\overline{\phi} > \underline{\phi}$ such that for any symmetric financial network,*

(a) if $\phi > \overline{\phi}$, then all banks default;

(b) if $\phi < \underline{\phi}$, then at least one bank does not default.

We end this discussion by demonstrating the implications of the above results by the means of a few examples. First, consider the complete financial network. It is clear that for any partition (S, S^c) of the set of banks, $\sum_{i \in S, j \notin S} y_{ij} = \frac{y}{n-1}|S||S^c|$, and as a result, $\phi^{\text{comp}} = 1/(n-1)$. On the other hand, choosing $S = \{i\}$ in any arbitrary financial network guarantees that $\phi \leq 1/(n-1)$. Therefore, the complete network has the largest bottleneck parameter among all regular financial networks. Corollary 2 thus implies that, if a large shock leads to the default of all banks in any financial network, it would also do so in the complete network, as predicted by Proposition 5.

At the other end of the spectrum, in a δ-connected financial network, there exists a partition (S, S^c) of the set of banks for which $\max\{y_{ij}, y_{ji}\} \leq \delta y$ for all $i \in S$ and $j \in S^c$. It is then immediate to verify that the bottleneck parameter of any such network satisfies $\phi \leq \delta$. Hence, for small enough values of δ and in the presence of large shocks, the financial network is strictly more stable and resilient than the complete network, again in line with the predictions of Proposition 5.

Finally, given a regular financial network with interbank liabilities $\{y_{ij}\}$ and bottleneck parameter ϕ, let

$$y_{ij}(\gamma) = (1-\gamma)y_{ij} + \gamma y_{ij}^{\text{comp}}$$

denote the interbank liabilities in the γ-convex combination of the former with the complete network. One can show that the corresponding bottleneck parameter satisfies[25]

$$\phi(\gamma) = (1-\gamma)\phi + \gamma\phi^{\text{comp}}.$$

In view of the observation that the complete network has the greatest bottleneck parameter across all financial networks, the above equality implies that $\phi(\gamma)$ is increasing in γ, establishing the following counterpart to Corollary 1:

Corollary 3. *Suppose that $\epsilon > \epsilon^*$, and consider a symmetric financial network for which $\phi > \bar{\phi}$. Then, the γ-convex combination of the network and the complete network is no more stable or resilient for all γ.*

This corollary implies that, in contrast to our results for the small shock regime, a more equal distribution of interbank liabilities cannot prevent the systemic collapse of the network in the presence of large shocks.

4.3 Multiple Shocks

The insights on the relationship between the extent of contagion and the structure of the financial network studied so far generalize to the case of multiple negative shocks.

Proposition 8. *Let p denote the number of negative shocks and let $\epsilon_p^* = n(a-v)/p$. There exist constants $y_p^* > \hat{y}_p > 0$, such that*

(a) If $\epsilon < \epsilon_p^*$ and $y > y_p^*$, then the complete network is the most stable and resilient financial network, whereas the ring network is the least resilient.

(b) If $\epsilon > \epsilon_p^*$ and $y > y_p^*$, then the complete and the ring financial networks are the least stable and resilient financial networks. Furthermore, if $p < n-1$, then there exists a δ-connected financial network that is strictly more stable than the complete and ring financial networks.

(c) If $\epsilon > \epsilon_p^*$ and $y \in (\hat{y}_p, y_p^*)$, then the complete network is the least stable and resilient financial network. Furthermore, the ring network is strictly more stable than the complete financial network.

[25] For any subset of banks S, we have $\sum_{i \in S, j \notin S} y_{ij}(\gamma) = (1-\gamma)\sum_{i \in S, j \notin S} y_{ij} + \gamma|S||S^c|/(n-1)$.

Parts (a) and (b) generalize the insights of Propositions 3 and 5 to the case of multiple shocks. The key new observation is that the critical threshold ϵ_p^* that defines the boundary of the small and large shock regimes is a decreasing function of p. Consequently, the number of negative shocks plays a role similar to that of the size of the shocks. More specifically, as long as the magnitude *and* the number of negative shocks affecting financial institutions are sufficiently small, more complete interbank claims enhance the stability of the financial system. The underlying intuition is identical to that behind Proposition 3: the more interconnected the financial network is, the better the excess liquidity of non-distressed banks are utilized in absorbing the shocks. On the other hand, if the magnitude *or* the number of shocks are large enough so that the excess liquidity in the financial system is not sufficient for absorbing the losses, financial interconnections serve as a propagation mechanism, creating a more fragile financial system. Furthermore, as in Proposition 5, weakly connected networks ensure that the losses are shared with the senior creditors of the distressed banks, protecting the rest of the system.

Part (c) of Proposition 8 contains a new result. It shows that in the presence of multiple shocks, the claims of the senior creditors in the ring financial network are used more effectively as a shock absorption mechanism than in the complete financial network. In particular, the closer the distressed banks in the ring financial network are to one another, the larger the loss their senior creditors are collectively forced to bear. This limits the extent of contagion in the network.[26]

As a final remark, we note that a multi-shock counterpart to Proposition 7 can also be established. In particular, if $m_{ij} < m^*$ for all i and j, then all banks in the financial network default at the face of p shocks of size $\epsilon > \epsilon_p^*$.

4.4 Non-Trivial Liquidation Proceeds

Our results thus far were restricted to the case in which the proceeds from liquidations are "trivial," i.e., $\zeta = 0$. We now show that our main results remain valid even if banks can recover a positive fraction $\zeta > 0$ of their projects' returns via liquidation.

Proposition 9. *Let $\epsilon_*(\zeta) = n(a-v) + \zeta A$ and $\epsilon^*(\zeta) = n(a-v) + n\zeta A$. Then, there exists $y^*(\zeta)$ such that for $y > y^*(\zeta)$,*

(a) If $\epsilon < \epsilon_*(\zeta)$, then the complete and the ring financial networks are, respectively, the most and the least stable and resilient financial networks.

(b) If $\epsilon > \epsilon^*(\zeta)$, then the complete and the ring networks are, respectively, the least stable and resilient networks, while any δ-connected network for small enough δ is strictly more stable and resilient.

(c) If $\epsilon_*(\zeta) < \epsilon < \epsilon^*(\zeta)$, then the complete network is strictly more stable and resilient than the ring network. Furthermore, if $\epsilon > \epsilon_*(\zeta) + \zeta A$, then there exists a δ-connected network which is strictly more stable and resilient than the complete network.

[26] In a related context, Alvarez and Barlevy (2013) show that the aggregate equity of the banking system with a ring network structure depends on the location of the shocks. Also see Barlevy and Nagaraja (2013) for an interesting connection between the problem of contagion in the ring financial network and the so-called circle-covering problem.

Part (a) corresponds to the small shock regime, in which the complete network outperforms all other regular financial networks and the ring network is the most fragile of all. Part (b), on the other hand, corresponds to our large shock regime results: for large enough shocks, the complete network becomes as fragile as the ring financial network,[27] whereas the presence of weakly connected components in the financial system guarantees that the losses are shared with the distressed bank's senior creditors, protecting the rest of the network.

Proposition 9 also establishes an intermediate regime, in which the complete network lies strictly between the ring and δ-connected financial networks in terms of stability and resilience. As part (c) shows, the threshold $\epsilon^*(\zeta)$ at which the complete network becomes the most fragile financial network no longer coincides with the threshold $\epsilon_*(\zeta) + \zeta A$ at which it starts underperforming δ-connected financial networks. In other words, even though both the small and large shock regimes exist regardless of the value of ζ, the phase transition between the two becomes smoother as ζ increases. Note that, as expected, the two thresholds coincide when $\zeta = 0$.

4.5 Size Heterogeneity

As mentioned earlier, our analysis has so far focused on regular financial networks in order to delineate the role of network structure on the financial system's fragility, while abstracting from the impact of asymmetries in size and leverage across banks. Focusing on the large shock regime, the next result illustrates that our key insights continue to hold even in the presence of such asymmetries.

Suppose that all assets and liabilities of bank i are scaled by a constant $\theta_i > 0$. In particular, its liabilities to its senior creditors and all other banks are equal to $\theta_i v$ and $y_i = \theta_i y$, respectively. Furthermore, i's short-term and long-terms returns are also scaled similarly as $\theta_i z_i$ and $\theta_i A$, respectively. As before, we assume that only one negative shock is realized and that $\zeta = 0$.

Proposition 10. *Suppose that bank j is hit with a negative shock $\epsilon > (a - v)\sum_{k=1}^{n}(\theta_k/\theta_j)$. Then,*

(a) Bank j defaults on its senior liabilities;

(b) All other banks also default if and only if $\widehat{m}_{ij} < \theta_i m^*$ for all banks i, where

$$\widehat{m}_{ij} = \theta_i + \sum_{k \neq j} \left(\frac{y_{ik}}{y_i}\right) \widehat{m}_{kj}. \tag{5}$$

In line with our earlier results, part (a) shows that a large enough shock guarantees that the senior creditors will bear some of the losses. Naturally, the corresponding threshold now depends on the relative size of the distressed bank: the greater the size of the bank, the smaller the threshold at which the senior creditors start to suffer.

[27] As in Proposition 8, the ring financial network may be more stable than the complete financial network in the presence of multiple shocks.

The second part of Proposition 10 is the counterpart to Proposition 7, establishing that the risk of systemic failures depends on the "size-adjusted" harmonic distances of other banks from the distressed bank j. Comparing (5) with (3) shows that the susceptibility of bank i to default not only depends on the intensity of the liabilities along the chains that connect j to i, but also on the size of all the intermediary banks that exist between the two. Indeed, if the financial network is connected, an increase in the relative size of any bank $k \neq i, j$ makes bank i more robust to a shock to bank j. This is a simple consequence of the fact that any such increase would raise i's distance from j. The effect of bank i's size on its own fragility is more subtle, as a greater θ_i increases both the distance $\widehat{m}_{ij}$ of bank i from the distressed bank j as well as the threshold $\theta_i m^*$. The two effects, however, are not proportional: whereas $\theta_i m^*$ increases linearly, $\widehat{m}_{ij}$ is sub-linear in θ_i. Consequently, increasing the relative size of a bank makes it more vulnerable to contagious defaults.

Finally, we remark that even though our results in this section were illustrated for an environment in which shocks can take only two values, similar results can be obtained for more general shock distributions.

5 Concluding Remarks

The recent financial crisis has rekindled interest in the relationship between the structure of the financial network and systemic risk. Two polar views on this relationship have been suggested in the academic literature and the policy world. The first maintains that the "incompleteness" of the financial network is a source of instability, as individual banks are overly exposed to the liabilities of a handful of financial institutions. According to this argument, a more complete financial network which limits the exposure of banks to any single counterparty would be less prone to systemic failures. The second view, in stark contrast, hypothesizes that it is the highly interconnected nature of the financial system that contributes to its fragility, as it facilitates the spread of financial distress and solvency problems from one institution to the rest in an epidemic-like fashion.

This paper provides a tractable theoretical framework for the study of the economic forces shaping the relationship between the structure of the financial network and systemic risk. We show that as long as the magnitude (or the number) of negative shocks is below a critical threshold, a more equal distribution of interbank liabilities leads to less fragility. In particular, all else equal, the sparsely connected ring financial network (corresponding to a credit chain) is the most fragile of all configurations, whereas the highly interconnected complete financial network is the configuration least prone to contagion. In line with the observations made by Allen and Gale (2000), our results establish that, in more complete networks, the losses of a distressed bank are passed to a larger number of counterparties, guaranteeing a more efficient use of the excess liquidity in the system in forestalling defaults.

We also show, however, that when negative shocks are larger than a certain threshold, the second view on the relationship between the structure of the financial network and the extent of contagion prevails. Now, completeness is no longer a guarantee for stability. Rather, in the presence of large shocks, financial networks in which banks are only weakly connected to one another are less prone to systemic failures. Such a "phase transition" is due to the fact that, the senior liabilities of banks, as well as the excess liquidity within the financial network, can act as shock absorbers. Weak interconnections guarantee that the more senior creditors of a distressed bank bear most of the losses and hence, protect the rest of the system against cascading defaults. Our model thus formalizes the robust-yet-fragile property of interconnected financial networks conjectured by Haldane (2009). More broadly, our results highlight the possibility that the same features that make a financial network structure more stable under certain conditions may function as significant sources of systemic risk and instability under another.

Our results indicate that the identification of systemically important financial institutions in the interbank network requires some care. In particular, some of the existing empirical analyses that rely on off-the-shelf and well-known measures of network centrality may be misleading, as such measures are relevant only if the interbank interactions are linear. In contrast, our analysis shows that, if the financial interactions exhibit non-linearities similar to those induced by unsecured debt contracts, the systemic importance of a financial institution is captured via its harmonic distance to other banks, suggesting that this new notion of network distance should feature in theoretically-motivated policy analyses.

Our model also highlights several possible avenues for policy interventions. From an *ex ante* perspective, a natural objective is to increase the stability and resilience of the financial system by regulating the extent and nature of interbank linkages. The important insight of our analysis is that such interventions have to be informed by the expected size of the negative shocks. For instance, imposing limits on the exposure between pairs of financial institutions motivated by the possibility of small shocks may be counterproductive if the shocks are in fact large. In addition, our characterization results in terms of the harmonic distance can serve as a guideline for *ex post* policy interventions once a shock is realized. In particular, our analysis suggests that injecting additional funds to or bailing out systemically important financial institutions that have a large impact on other entities within the network would contain the extent of contagion (even though such interventions may induce moral hazard-type concerns *ex ante*).

Another important dimension of policy analysis is the discussion of whether observed financial networks are likely to be (constrained) efficient. This is studied in some detail in the working paper version of our work (Acemoglu et al., 2013b). There we show that, in the presence of covenants that make the interbank interest rates contingent on the borrowers' lending decisions, the equilibrium pattern of interbank liabilities is efficient as long as the economy consists of two or three banks. However, in the presence of more than three banks, a specific (and new) type of externality arises: given that the contracts offered to a bank do not condition on the asset position of the banks to whom it lends, the borrowers do not internalize the effect of their decisions on their creditors' creditors. We show that this may lead to the formation of inefficient networks, embedding excessive counterparty risk from the viewpoint of social efficiency, both in the small and the large shock regimes. These results suggest that there may be room for welfare-improving government interventions at the network formation stage.

We view our paper as a first step in the direction of a systematic analysis of the broader implications of the financial network architecture. Several important issues remain open to future research.

First, this paper focused on the implications of a given network structure on the fragility of the financial system. A systematic analysis of the endogenous formation of financial networks and their efficiency and policy implications, along the lines of Acemoglu et al. (2013b) briefly discussed above, is an obvious area for future research.

Second, our focus was on a specific form of network interactions among financial institutions; namely, the spread of counterparty risk via unsecured debt contracts. In practice, however, there are other important types of financial interdependencies. In particular, (i) the fire sales of some assets by a bank may create distress on other institutions that hold similar assets; and (ii) withdrawal of liquidity by a bank (for example, by not rolling over a repo agreement or increasing the haircut on a collateral) may lead to a chain reaction playing out over the financial network. How the nature of these different types of financial interdependencies determine the relationship between underlying network structure and systemic risk remains an open question for future research.

Third, our model purposefully abstracted from several important institutional details of the banking system. For example, it eschewed other forms of interbank lending (such as repurchase agreements) as well as the differences between deposit-taking institutions, investment banks and other specialized financial institutions such as hedge funds. It also abstracted from the complex maturity structure of interbank liabilities which can create different types of contagion owing to a mismatch in the maturity of the assets and liabilities of a financial institution. Incorporating these important institutional realities is another obvious area of investigation.

Last but not least, a systematic empirical investigation of these and other types of network interactions in financial markets is an important area for research.

Appendix

A The Harmonic Distance

The harmonic distance defined in (3) provides a measure of proximity between a pair of banks in the financial network. According to this notion, two banks are closer to one another the more direct or indirect liability chains exist between them and the higher the face value of those liabilities are. This appendix studies some of the basic properties of the harmonic distance.

The first key observation is that the notion of harmonic distance is closely related to a discrete-state Markov chain defined over the network. In particular, let Q be a matrix whose (i, j) element is equal to the fraction of bank j's liabilities to i; that is, $q_{ij} = y_{ij}/y$. By construction, Q is a (row and column) stochastic matrix and hence, can be interpreted as the transition probability matrix of a Markov chain with the uniform stationary distribution. For this Markov chain, define the *mean hitting time from i to j* as the expected number of time steps it takes the chain to hit state j conditional on starting from state i. We have the following result:

Proposition A.1. *The harmonic distance from bank i to j is equal to the mean hitting time of the Markov chain from state i to state j.*

Thus, the harmonic distance provides an intuitive measure of proximity between any pair of banks: the longer it takes on average for the Markov chain to reach state j from state i, the larger the harmonic distance between the two corresponding banks in the financial network. The above observation enables us to identify the properties of the harmonic distance by relying on known results from the theory of Markov chains. For instance, given that expected hitting times in a Markov chain are non-symmetric, it is immediate that in general, $m_{ij} \neq m_{ji}$, even when the financial network is symmetric. This observation thus implies that the harmonic distance is not a notion of distance in its strictest sense. Nevertheless, it satisfies a weaker form of symmetry:

Proposition A.2. *Suppose that the financial network is symmetric. For any three banks i, j, and k,*

$$m_{ij} + m_{jk} + m_{ki} = m_{ik} + m_{kj} + m_{ji}. \tag{6}$$

An immediate implication of the above result is the following:

Corollary A.1. *If the financial network is symmetric, then there exists an ordering of banks such that if i is preceded by j, then $m_{ij} \leq m_{ji}$.*

The above ordering can be obtained by fixing an arbitrary reference bank k and ordering the rest of the banks according to the value of $m_{ik} - m_{ki}$. Such an ordering, however, is not necessarily unique due to the possibility of ties. Nevertheless, by Proposition A.2, the set of banks for which $m_{ij} = m_{ji}$ form an equivalence class, which implies that there exists a well-defined ordering of the equivalence classes, independent of the reference bank k (Lovász, 1993). More specifically, there exists a partition $(S_1, \dots, S_r)$ of the set of banks such that (i) $m_{ij} = m_{ji}$ for all $i, j \in S_t$; and (ii) if $t < t'$, then $m_{ij} < m_{ji}$

for all $i \in S_t$ and $j \in S_{t'}$. Therefore, in view of Proposition 7, banks in the higher equivalence classes are systemically more important as a large shock to them would lead to large chains of defaults. Yet, at the same time, such banks are not at much risk of contagious defaults if another bank is hit with a negative shock. In contrast, banks in the lower equivalence classes are prone to default due to contagion even though a negative shock to them would not lead to a large cascade of failures.

Proposition A.3 (triangle inequality)**.** *For any triple of banks i, j and k,*

$$m_{ij} + m_{jk} \geq m_{ik}.$$

Furthermore, the inequality is tight if and only if all liability chains from k to i pass through bank j.

The sum of harmonic distances of a given bank from all others is an invariant structural property of the financial network that does not depend on the identity of the bank:

Proposition A.4. *The quantity $\sum_{j\neq i} m_{ij}$ is a constant for all banks i in the financial network.*

We end this discussion by showing that the maximum pairwise harmonic distance between any two banks is minimized in the complete financial network. First, note that given the full symmetry in the complete financial network, m_{ij} is the same for all distinct pairs of banks i and j. Thus, equation (3) immediately implies that $m_{ij} = n-1$ for $i \neq j$. On the other hand, summing both sides of equation (3) over $j \neq i$ in an arbitrary financial network implies

$$\begin{aligned}\sum_{j\neq i} m_{ij} &= (n-1) + \sum_{j\neq i}\sum_{k\neq j} q_{ik} m_{kj} \\ &= (n-1) + \sum_{k\neq i}\sum_{j\neq k} q_{ik} m_{kj} - \sum_{k\neq i} q_{ik} m_{ki}.\end{aligned}$$

Given that the quantity $\sum_{j\neq k} m_{kj}$ does not depend on k, the above equality simplifies to $\sum_{k\neq i} q_{ik} m_{ki} = n-1$, implying that

$$\max_{k\neq i} m_{ki} \geq n-1.$$

Thus, the maximum pairwise harmonic distance is minimized in the complete financial network.

B Proofs

Notation

Let $Q \in \mathbb{R}^{n\times n}$ be the matrix whose (i,j) element is equal to the fraction of bank j's liabilities to i; that is, $q_{ij} = y_{ij}/y_j$. Whenever there is no risk of confusion, we use $y = [y_1, \dots, y_n]'$ to denote the vector of the banks' total liabilities to one another. Thus, rewriting equations (1) and (2) in matrix notation implies that a payment equilibrium is simply a pair of vectors (x, ℓ) that simultaneously solve

$$x = \big[\min\{Qx + e + \zeta\ell, y\}\big]^+ \tag{7}$$

$$\ell = \Big[\min\Big\{\frac{1}{\zeta}(y - Qx - e), A\Big\}\Big]^+, \tag{8}$$

where $e_j = c_j + z_j - v$ and $x_j = \sum_{i\neq j} x_{ij}$ is the total debt repayment of bank j to the rest of the banks.

Preliminary Lemmas

We start by stating two lemmas, which we will later utilize in the subsequent proofs. Given its simplicity, we have omitted the proof of Lemma B.1 below.

Lemma B.1. *Suppose that $\beta > 0$. Then,*

$$\Big|[\min\{\alpha, \beta\}]^+ - [\min\{\hat{\alpha}, \beta\}]^+\Big| \leq |\alpha - \hat{\alpha}|.$$

Furthermore, the inequality is tight only if either $\alpha = \hat{\alpha}$ or $\alpha, \hat{\alpha} \in [0, \beta]$.

Lemma B.2. *Suppose that (x, ℓ) is a payment equilibrium of the financial network. Then, x satisfies*

$$x = [\min\{Qx + e + \zeta A, y\}]^+. \tag{9}$$

Conversely, if $x \in \mathbb{R}^n$ satisfies (9), then there exists $\ell \in [0, A]^n$ such that (x, ℓ) is a payment equilibrium.

Proof. First, suppose that the pair (x, ℓ) is a payment equilibrium of the financial network. Thus, by definition, $\zeta\ell = [\min\{y - (Qx + e), \zeta A\}]^+$, which implies that

$$Qx + e + \zeta\ell = \max\left\{Qx + e, \min\{y, Qx + e + \zeta A\}\right\}.$$

Therefore,

$$\min\{Qx + e + \zeta\ell, y\} = \min\Big\{y, \max\left\{Qx + e, \min\{y, Qx + e + \zeta A\}\right\}\Big\}.$$

Simplifying the expression on the right-hand side above leads to

$$\min\{Qx + e + \zeta\ell, y\} = \min\{y, Qx + e + \zeta A\},$$

and hence, $x = [\min\{y, Qx + e + \zeta A\}]^+$, which is the same as (9).

To prove the converse, suppose that $x \in \mathbb{R}^n$ satisfies (9) and let $\ell = (1/\zeta)[\min\{y - (Qx + e), \zeta A\}]^+$. By construction, (8) is satisfied. Thus, to prove that (x, ℓ) is indeed a payment equilibrium, it is sufficient to show that the pair (x, ℓ) also satisfies (7). It is immedaite that

$$Qx + e + \zeta\ell = \max\left\{Qx + e, \min\{y, Qx + e + \zeta A\}\right\},$$

and therefore,

$$\begin{aligned}
[\min\{Qx + e + \zeta\ell, y\}]^+ &= [\min\left\{y, \max\left\{Qx + e, \min\{y, Qx + e + \zeta A\}\right\}\right\}]^+ \\
&= [\min\{y, Qx + e + \zeta A\}]^+ \\
&= x,
\end{aligned}$$

where the last equality is simply a consequence of the assumption that x satisfies (9). □

Proof of Proposition 1

Existence: In view of Lemma B.2, it is sufficient to show that there exists $x^* \in \mathbb{R}^n_+$ that satisfies $x^* = [\min\{Qx^* + e + \zeta A, y\}]^+$. Define the mapping $\Phi : H \to H$ as

$$\Phi(x) = \big[\min\{Qx + e + \zeta A, y\}\big]^+,$$

where $H = \prod_{i=0}^n [0, y_i]$. This mapping continuously maps a convex and compact subset of the Euclidean space to itself, and hence, by the Brouwer fixed point theorem, there exists $x^* \in H$ such that $\Phi(x^*) = x^*$. Lemma B.2 then implies that the pair (x^*, ℓ^*) is a payment equilibrium of the financial network, where $\ell^* = (1/\zeta)[\min\{y - (Qx^* + e), \zeta A\}]^+$. In particular, the collection of pairwise interbank payment $\{x^*_{ij}\}$ defined as $x^*_{ij} = q_{ij}x^*_j$ alongside liquidation decisions ℓ^*_i satisfy the collection of equations (1) and (2) for all i and j simultaneously. □

Generic Uniqueness: Without loss of generality, we restrict our attention to a connected financial network.[28] Suppose that the financial network has two distinct payment equilibria, denoted by (x, ℓ) and $(\hat{x}, \hat{\ell})$ such that $x \neq \hat{x}$.[29] By Lemma B.2, both x and $\hat{x}$ satisfy (9). Therefore, for any given bank i,

$$\begin{aligned} |x_i - \hat{x}_i| &= \big|[\min\{(Qx)_i + e_i + \zeta A, y_i\}]^+ - [\min\{(Q\hat{x})_i + e_i + \zeta A, y_i\}]^+\big| \\ &\leq |(Qx)_i - (Q\hat{x})_i| \,, \end{aligned} \tag{10}$$

where the inequality is a consequence of Lemma B.1. Summing both sides of the above inequality over all banks i leads to

$$\begin{aligned} \|x - \hat{x}\|_1 &\leq \|Q(x - \hat{x})\|_1 && (11) \\ &\leq \|Q\|_1 \cdot \|x - \hat{x}\|_1 && (12) \\ &= \|x - \hat{x}\|_1, \end{aligned}$$

where the last equality is due to the fact that the column sums of Q are equal to one. Consequently, inequalities (10)–(12) are all tight simultaneously. In particular, given that (10) is tight, and in view of Lemma B.1, for any given bank i either

$$(Qx)_i = (Q\hat{x})_i$$

or

$$0 \leq (Qx)_i + e_i + \zeta A \,,\ (Q\hat{x})_i + e_i + \zeta A \leq y_i. \tag{13}$$

Denoting the set of banks that satisfy (13) by B, it is immediate that, for all $i \in B$,

$$\begin{aligned} x_i &= e_i + \zeta A + (Qx)_i \\ \hat{x}_i &= e_i + \zeta A + (Q\hat{x})_i, \end{aligned}$$

[28] If the financial network is not connected, the proof can be applied to each connected component separately.

[29] The assumption that the two payment equilibria are distinct requires that $x \neq \hat{x}$. Note that if $x = \hat{x}$, then (8) immediately implies that $\ell = \hat{\ell}$.

and therefore,

$$(Qx)_i - (Q\hat{x})_i = x_i - \hat{x}_i \qquad \forall i \in B.$$

On the other hand, in view of the fact that $(Qx)_i = (Q\hat{x})_i$ for all $i \notin B$, we have

$$Q(x - \hat{x}) = \begin{bmatrix} x_B - \hat{x}_B \\ 0 \end{bmatrix},$$

and hence,

$$\|Q(x - \hat{x})\|_1 = \|x_B - \hat{x}_B\|_1. \tag{14}$$

Therefore, (12) holds as an equality only if $x_i = \hat{x}_i$ for all $i \notin B$, as it would otherwise violate (14). Hence,

$$Q_{BB}\big(x_B - \hat{x}_B\big) = x_B - \hat{x}_B, \tag{15}$$

where Q_{BB} is the submatrix of Q corresponding to the banks in B. On the other hand, the fact that the financial network is connected implies that Q, and hence, Q_{BB} are irreducible matrices.[30] Given that $x \neq \hat{x}$, equality (15) cannot hold unless B^c is empty. As a consequence,

$$\sum_{i=1}^{n} x_i = n\zeta A + \sum_{i=1}^{n} e_i + \sum_{i=1}^{n}\sum_{j=1}^{n} q_{ij}x_j$$

which implies $\sum_{i=1}^{n} e_i = -n\zeta A$, an equality that holds only for a non-generic set of parameters $z_1, \ldots, z_n$. Thus, the payment equilibria of the economy are generically unique. □

Proof of Proposition 2

Denote the set of banks that default on their senior debt by F, the set of banks that default but can pay their debts to the senior creditors by D, and the set banks that do not default by $S = (D \cup F)^c$. For any bank $i \in F$, we have

$$\pi_i + T_i = z_i + \zeta\ell_i + \sum_{j\neq i} x_{ij},$$

whereas for $i \in D$,

$$\pi_i + T_i = v.$$

On the other hand, for any bank $i \in S$ which does not default,

$$\pi_i + T_i = A - \ell_i + \zeta\ell_i + z_i - y + \sum_{j\neq i} x_{ij},$$

[30] An $n \times n$ matrix Q is said to be *reducible*, if for some permutation matrix P, the matrix $P'QP$ is block upper-triangular. If a square matrix is not reducible, it is said to be *irreducible*. If Q is a non-negative irreducible matrix with unit column sums, then all eigenvalues of any square submatrix of Q, say $\tilde{Q}$, lie within the unit circle, implying that equation $\tilde{Q}x = x$ has no non-trivial solutions. For more on this, see e.g., Berman and Plemmons (1979).

where $\zeta\ell_i$ is the proceeds that bank i obtains from liquidating its project (if any). Summing the above across all banks implies that the social surplus in the economy is equal to

$$\begin{aligned} u =&(A-y)|S| + v|D| + \sum_{i\notin D}(z_i + \zeta\ell_i) + \sum_{i\notin D}\sum_{j\neq i} x_{ij} - \sum_{i\in S}\ell_i \\ =&(A-y)|S| + \sum_{i=1}^{n}(z_i + \zeta\ell_i) + \sum_{i\in S}\sum_{j\neq i} y_{ij} - \sum_{i\in S}\ell_i, \end{aligned}$$

where the second equality is due to the fact that for $i \in D$ we have, $\sum_{j\neq i}(x_{ji} - x_{ij}) = z_i + \zeta\ell_i - v$. Further simplifying the above equality thus implies

$$\begin{aligned} u &= \sum_{i=1}^{n} z_i + |S|A + \zeta\sum_{i=1}^{n}\ell_i - \sum_{i\in S}\ell_i \\ &= n(a+A) - p\epsilon - (1-\zeta)\sum_{i=1}^{n}\ell_i, \end{aligned}$$

which completes the proof. □

Two Auxiliary Lemmas

Lemma B.3. *The number of bank defaults satisfies*

$$p \leq \sharp(\textit{defaults}) < \frac{p\epsilon}{a-v}$$

where p *is the number of realizations of negative shocks in the network.*

Proof. Given that the total interbank liabilities of each bank is equal to its total interbank claims and that $v > a - \epsilon$, any bank that is hit with a negative shock defaults. Hence, the lower bound is trivial. To obtain the upper bound, note that for any bank i that defaults but can meet its senior liabilities in full, we have

$$z_i + \sum_{j\neq i} x_{ij} = v + \sum_{j\neq i} x_{ji}.$$

Denoting the set of such banks by D and summing over all $i \in D$ imply

$$\sum_{i\in D} z_i + \sum_{i\in D}\sum_{j\neq i} x_{ij} = v|D| + \sum_{i\in D}\sum_{j\neq i} x_{ji}. \tag{16}$$

On the other hand, for any bank i that defaults on its senior liabilities (if such a bank exists), we have

$$\sum_{j\neq i} x_{ij} + z_i < v.$$

Summing over the set of all such banks, F, implies

$$\sum_{i\in F}\sum_{j\neq i} x_{ij} + \sum_{i\in F} z_i < v|F|. \tag{17}$$

Adding (16) and (17) leads to

$$p\epsilon - (a - v)\sharp(\text{defaults}) \geq \sum_{j \notin S} \sum_{i \in S} (y_{ij} - x_{ij}),$$

where S is the set of banks that do not default. By definition, the right-hand side of the above equality is strictly positive, proving that the number of defaults is strictly smaller than $p\epsilon/(a - v)$. □

Lemma B.4. *If $\epsilon < \epsilon_p^*$, then at least one bank does not default, where $\epsilon_p^* = n(a - v)/p$. On the other hand, if $\epsilon > \epsilon_p^*$, then at least one bank defaults to its senior creditors.*

Proof. Suppose that $\epsilon < \epsilon_p^*$ and that all banks default. Therefore,

$$z_i + \sum_{j \neq i} x_{ij} \leq v + \sum_{j \neq i} x_{ji},$$

for all banks i. Summing over i implies

$$na - p\epsilon \leq nv,$$

which contradicts the assumption that $\epsilon < \epsilon_p^*$.

To prove the second statement, suppose that $\epsilon > \epsilon_p^*$ and that no bank defaults on its senior liabilities. Thus,

$$z_i + \sum_{j \neq i} x_{ij} \geq v + \sum_{j \neq i} x_{ji},$$

for all banks i. Summing over i implies $na - p\epsilon \geq nv$, leading to a contradiction. □

Proof of Proposition 3

Proof of part (a). Without loss of generality assume that bank 1 is hit with the negative shock. By Lemma B.4, bank n does not default as it is the bank furthest away from the distressed bank. Moreover, as long as $y > y^* = (n - 1)(a - v)$, bank 1 — and consequently all banks — can meet their senior liabilities in full. This is due to the fact that $y + a - \epsilon > v$. Given that the set of banks that default form connected chain, say of length τ, the repayment of the last bank in default to its sole creditor satisfies

$$x_{\tau+1,\tau} = y + \tau(a - v) - \epsilon.$$

On the other hand, given that bank $\tau + 1$ does not default, we have

$$a + x_{\tau+1,\tau} > y + v.$$

As a result, $\tau \geq \epsilon/(a - v) - 1$, implying that the number of defaults reaches the upper bound established in Lemma B.3. Hence, the ring network is the least stable and least resilient financial network. □

Proof of part (b). By Lemma B.4, as long as $\epsilon < \epsilon^*$, there exists at least one bank that does not default. Given the full symmetry in the complete network, the $n-1$ banks that are not hit with the negative shock can thus meet their liabilities in full. Hence, the complete financial network is the most stable and most resilient regular financial network. □

Proof of part (c). Consider the financial network constructed as the γ-convex combination of the ring and the complete financial networks. Without loss of generality, we assume that bank 1 is hit with the negative shock. Define γ_d to be the value at which banks 1 through $d-1$ default while bank d is at the verge of default. At this value of γ, we have

$$x_1 = \left(\frac{\gamma_d}{n-1}\right)[(x_1 + \cdots + x_d) - x_1 + (n-d)y] + (1-\gamma_d)y + (a - v - \epsilon) \tag{18}$$

$$x_i = \left(\frac{\gamma_d}{n-1}\right)[(x_1 + \cdots + x_d) - x_i + (n-d)y] + (1-\gamma_d)x_{i-1} + (a - v) \tag{19}$$

for $2 \leq i \leq d$. Hence,

$$\begin{aligned}
\Delta_2 = x_2 - x_1 \;\; &= \;(1-\gamma_d)\Delta_1 - \left(\frac{\gamma_d}{n-1}\right)\Delta_2 + \epsilon \\
\Delta_i = x_i - x_{i-1} &= \;(1-\gamma_d)\Delta_{i-1} - \left(\frac{\gamma_d}{n-1}\right)\Delta_i,
\end{aligned}$$

where $\Delta_1 = x_1 - y$. Thus, for all $2 \leq i \leq d$,

$$\Delta_i = \beta^{i-1}(\Delta_1 + \epsilon/(1-\gamma_d)),$$

where for notational simplicity we have defined $\beta = (1 + \frac{\gamma_d}{n-1})^{-1}(1-\gamma_d)$. Given that $x_d = y$, the terms Δ_i must add up to zero, that is,

$$\sum_{i=1}^{d} \Delta_i = \left(\frac{1-\beta^d}{1-\beta}\right)\Delta_1 + \left(\frac{1-\beta^{d-1}}{1-\beta}\right)\left(\frac{\beta\epsilon}{1-\gamma_d}\right) = 0,$$

which immediately implies

$$\Delta_1 = -\left(\frac{1-\beta^{d-1}}{1-\beta^d}\right)\left(\frac{\beta\epsilon}{1-\gamma_d}\right).$$

Therefore,

$$x_i = y + \sum_{s=1}^{i}\Delta_s = y + \left(\frac{\beta^{d-1} - \beta^{i-1}}{1-\beta^d}\right)\left(\frac{\beta\epsilon}{1-\gamma_d}\right),$$

and as a result,

$$\sum_{i=1}^{d-1} x_i = (d-1)y + \left[\frac{d\beta^{d-1}(1-\beta) - (1-\beta^d)}{(1-\gamma_d)(1-\beta)(1-\beta^d)}\right]\beta\epsilon.$$

On the other hand, from (18) and (19), we have

$$\sum_{i=1}^{d-1} x_i = (d-1)y + \frac{d(a-v) - \epsilon}{\gamma_d(n-d)/(n-1)}.$$

Equating the above two equalities thus leads to

$$n(a-v)/\epsilon - 1 = \left[\frac{\beta^{d-1}(1-\beta)}{1-\beta^d}\right](n-d).$$

Therefore, the value γ_d at which $d-1$ banks default and bank d is at the verge of default must satisfy the above equality. For a fixed value of d, the right-hand side is increasing in β (and hence, decreasing in γ_d).[31] As a consequence, in order for the right-hand side to remain equal to the constant on the left-hand side, d has to decrease as γ increases. In other words, there will be weakly less defaults for higher values of γ. □

Proof of Proposition 4

We start with a lemma, which is a simple consequence of the definition of the payment equilibrium.

Lemma B.5. *Suppose that no bank defaults on its senior liabilities. Then, banks in sets s and d can meet their liabilities and default, respectively, if and only if*

$$(I - Q_{dd})^{-1}e_d < 0 \tag{20}$$

$$Q_{sd}(I - Q_{dd})^{-1}e_d + e_s \geq 0. \tag{21}$$

Proof of Proposition 4. First consider the original financial network $\{y_{ij}\}$, label the set of banks that default and can meet their liabilities in full by d and s, respectively. Lemma B.5 implies that inequalities (20) and (21) are satisfied. Next, consider the financial network $\{\tilde{y}_{ij}\}$ obtained by rewiring $\{y_{ij}\}$. By construction, $\tilde{Q}_{dd} = Q_{dd}$, which immediately implies that the rewired network also satisfies (20). Thus, the proof is complete once we show that it also satisfies the second inequality. We have

$$\begin{aligned}\tilde{Q}_{sd}(I - \tilde{Q}_{dd})^{-1}e_d + e_s &= \left((1-\gamma)Q_{sd} + \left(\frac{\gamma}{n-d}\right)\mathbf{1}_s\mathbf{1}_s'Q_{sd}\right)(I - Q_{dd})^{-1}e_d + e_s \\ &\geq \left(\frac{\gamma}{n-d}\right)\mathbf{1}_s\mathbf{1}_s'Q_{sd}(I - Q_{dd})^{-1}e_d + \gamma e_s,\end{aligned}$$

where the inequality is a consequence of the fact that $\{y_{ij}\}$ satisfies (21). Hence, using the identity $\mathbf{1}_d'Q_{dd} + \mathbf{1}_s'Q_{sd} = \mathbf{1}_d'$ implies,

$$\begin{aligned}\tilde{Q}_{sd}(I - \tilde{Q}_{dd})^{-1}e_d + e_s &\geq \left(\frac{\gamma}{n-d}\right)\mathbf{1}_s\mathbf{1}_d'e_d + \gamma e_s \\ &= \gamma\left(\frac{d(a-v)-\epsilon}{n-d} + a - v\right)\mathbf{1}_s,\end{aligned}$$

and hence,

$$\begin{aligned}\tilde{Q}_{sd}(I - \tilde{Q}_{dd})^{-1}e_d + e_s &\geq \left(\frac{\gamma}{n-d}\right)(n(a-v)-\epsilon)\mathbf{1}_s \\ &> 0,\end{aligned}$$

where the last inequality is a consequence of the fact that $\epsilon < \epsilon^*$. Thus, by Lemma B.5, none of the banks labeled s default. □

[31] This can be easily verified by noticing that $\beta^{d-1}(1-\beta) = (1-\beta^d)\big(1+\beta^{-1}+\cdots+\beta^{-(d-1)}\big)^{-1}$.

Proof of Corollary 1

Consider the regular financial network $\{y_{ij}\}$. By assumption, only the distressed bank, say bank j, defaults. This means that the bank is paid in full by its creditors; thus,

$$\begin{aligned} x_j &= [\min\{y + a - v - \epsilon, y\}]^+ \\ &> [\min\{y^* + a - v - \epsilon^*, y^*\}]^+ \\ &= 0, \end{aligned}$$

where the inequality is a consequence of the assumptions that $\epsilon < \epsilon^*$ and $y > y^*$. The fact that $x_j > 0$ implies that bank j can meet its liabilities to the senior creditors in full. Thus, by Proposition 4, any $(\{j\}, \{j\}, \gamma)$-rewiring of the financial network would lead to more weakly less defaults. Finally, noting that in the absence of financial contagion any γ-convex combination of a network with the complete network is essentially a $(\{j\}, \{j\}, \gamma)$-rewiring of that network completes the proof. □

Proof of Proposition 5

Proof of part (a). First consider the complete financial network. By Lemma B.4, the distressed bank defaults on its senior liabilities. Suppose that the remaining $n - 1$ banks do not default and that they can all meet their liabilities in full.[32] This would be the case only if

$$(n-2)\frac{y}{n-1} + (a - v) \geq y,$$

where we are using the fact that the distressed bank does not pay anything to its junior creditors. The above inequality, however, contradicts the assumption that $y > y^*$. Hence, all banks default, implying that the complete network is the least resilient and the least stable financial network.

Now consider the ring financial network and assume without loss of generality that bank 1 is hit with the negative shock. Once again by Lemma B.4, bank 1 defaults on its senior liabilities and hence, does not pay anything to its more junior creditor, bank 2. Thus, if some bank does not default, the length of the default cascade, τ, must satisfy

$$\tau(a - v) > y,$$

which in light of the assumption that $y > y^*$ guarantees that $\tau > n - 1$, leading to a contradiction. Hence, all banks default, implying that the ring financial network is the least stable and resilient financial network. □

Proof of part (b). Consider a δ-connected financial network where $\delta < \frac{a-v}{ny}$. By definition, there exists a partition (S, S^c) of the set of banks such that $y_{ij} \leq \delta y$ for all $i \in S$ and $j \in S^c$. Consequently,

[32] Note that given the symmetric structure of the complete network and the uniqueness of payment equilibrium, either all other banks default together or they all meet their liabilities fully at the same time.

$\sum_{j \notin S} y_{ij} \leq \delta y |S^c|$ for all banks $i \in S$, and therefore, the total liabilities of all banks in S to any bank $i \in S$ satisfies

$$\begin{aligned}\sum_{j \in S} y_{ij} &\geq y - \delta y |S^c| \\ &\geq y - (a - v).\end{aligned}$$

Therefore, if the negative shock hits a bank in S^c, all banks in S can still meet their liabilities in full. This is a simple consequence of the fact that $a - v + \sum_{j \in S} y_{ij} \geq y$, which guarantees that in the unique payment equilibrium of the financial network, no bank in S defaults. A similar argument shows that when the shock hits a bank in S, all banks in S^c can meet their liabilities in full. Thus, the financial network is strictly more stable and resilient than both the complete and the ring networks, in which all banks default. □

Proof of Proposition 6

Lemma B.6. *Suppose that bank j is hit with a shock $\epsilon > \epsilon^*$. If d denotes the set of all other banks that default, then*

$$(I - Q_{dd})^{-1}[(a - v)\mathbf{1}_d - yQ_{dj}] < 0. \tag{22}$$

Furthermore, if (22) is satisfied for a subset of banks d, then all banks in d default.

Proof. To prove the first statement, suppose that d and s denote the set of banks that default (excluding bank j) and can meet their liabilities in full, respectively. By definition,

$$x_d = Q_{dd} x_d + yQ_{ds}\mathbf{1}_s + (a - v)\mathbf{1}_d,$$

which implies

$$x_d = (I - Q_{dd})^{-1} \left[yQ_{ds}\mathbf{1}_s + (a - v)\mathbf{1}_d \right].$$

On the other hand, given that Q is a stochastic matrix, we have $Q_{ds}\mathbf{1}_s + Q_{dd}\mathbf{1}_d + Q_{dj} = \mathbf{1}_d$. Therefore, for all banks in d to default, it is necessary that

$$(I - Q_{dd})^{-1} \left[(a - v)\mathbf{1}_d - yQ_{dj} \right] < 0,$$

which proves the first statement.

To prove the second statement, suppose that for a given subset of banks d, inequality (22) is satisfied. Replacing for $Q_{dj} = \mathbf{1}_d - Q_{dd}\mathbf{1}_d - Q_{ds}\mathbf{1}_s$ immediately implies that

$$(I - Q_{dd})^{-1}[(a - v)\mathbf{1}_d + yQ_{ds}] < y\mathbf{1}_d.$$

In other words, even if no other bank outside of d (and the originally distressed bank j) defaults, still no bank in d can meet its liabilities in full. Therefore, regardless of the state of other banks, all banks in d default. □

Proof of Proposition 6 Consider the original financial network $\{y_{ij}\}$. By the first part of Lemma B.5, inequality (22) is satisfied. Next, consider the financial network $\{\tilde{y}_{ij}\}$ obtained by rewiring $\{y_{ij}\}$. By construction, $\tilde{Q}_{dd} = Q_{dd}$ and

$$\tilde{q}_{ij} = (1-\gamma)q_{ij} + \gamma/(n-1)$$

for any bank $i \neq j$. Consequently,

$$\begin{aligned}(I-\tilde{Q}_{dd})^{-1}[(a-v)\mathbf{1}_d - y\tilde{Q}_{dj}] =& (1-\gamma)(I-Q_{dd})^{-1}[(a-v)\mathbf{1}_d - yQ_{dj}] \\ &+ \gamma(a-v-y/(n-1))(I-Q_{dd})^{-1}\mathbf{1}_d.\end{aligned}$$

The first part of Lemma B.6 guarantees that the first term on the right-hand side above is negative. As for the second term, note that $(I-Q_{dd})^{-1}$ is an inverse M-matrix, and hence, is element-wise positive.[33] Furthermore, by assumption, $y > y^* = (n-1)(a-v)$, implying that the second term is also negative. Therefore, by the second part of Lemma B.6, all banks in d still default after the rewiring. □

Proof of Proposition 7

Proof of part (a). Let $m^* = y/(a-v)$. We prove the result by showing that if a bank i does not default, then $m_{ij} \geq m^*$. Given that the distressed bank j defaults on its senior liabilities, it is immediate that $x_j = 0$. Let $x_d \in \mathbb{R}^d$ denote the subvector of the equilibrium payment vector x corresponding to the banks in default, excluding the originally distressed bank j. Similarly, let $x_s \in \mathbb{R}^s$ denote the subvector corresponding to the banks that can meet their liabilities in full.[34] From the definition of a payment equilibrium, it is immediate that

$$x_d = Q_{dd}x_d + yQ_{ds}\mathbf{1}_s + (a-v)\mathbf{1}_d,$$

and as a result,

$$x_d = (I-Q_{dd})^{-1}\left(yQ_{ds}\mathbf{1}_s + (a-v)\mathbf{1}_d\right). \tag{23}$$

Furthermore, given that banks indexed s can meet their liabilities in full, it must be the case that

$$Q_{sd}x_d + yQ_{ss}\mathbf{1}_s + (a-v)\mathbf{1}_s \geq y\mathbf{1}_s.$$

Substituting for x_d from (23) implies

$$Q_{sd}(I-Q_{dd})^{-1}\mathbf{1}_d + \mathbf{1}_s \geq m^*\left[I - Q_{ss} - Q_{sd}(I-Q_{dd})^{-1}Q_{ds})\right]\mathbf{1}_s. \tag{24}$$

[33] A square matrix X is an *M-matrix* if there exist a non-negative square matrix B and a constant $r > \rho(B)$ such that $X = rI - B$, where $\rho(B)$ is the spectral radius of B. By Plemmons (1977, Theorem 2), the inverse of any non-singular M-matrix is element-wise non-negative. For more on M-matrices and their properties, see Berman and Plemmons (1979).

[34] With some abuse of notation, we use d and s to not only denote the set of defaulting and solvent banks, but also the size of the two sets, respectively. Hence, $s + d + 1 = n$.

On the other hand, by the definition of the harmonic distance (3), we have

$$m_{dj} = \mathbf{1}_d + Q_{dd} m_{dj} + Q_{ds} m_{sj} \tag{25}$$

$$m_{sj} = \mathbf{1}_s + Q_{sd} m_{dj} + Q_{ss} m_{sj}, \tag{26}$$

where $m_{dj} \in \mathbb{R}^d$ and $m_{sj} \in \mathbb{R}^s$ are vectors that capture the harmonic distances of the banks in default and the solvent banks to the distressed bank j, respectively. Solving for m_{dj} in (25) and replacing in (26) leads to

$$\begin{aligned} m_{sj} &= \mathbf{1}_s + Q_{sd}(I - Q_{dd})^{-1}\mathbf{1}_d + Q_{sd}(I - Q_{dd})^{-1} Q_{ds} m_{sj} + Q_{ss} m_{sj} \\ &\geq m^* \left[I - Q_{ss} - Q_{sd}(I - Q_{dd})^{-1} Q_{ds}\right] \mathbf{1}_s + Q_{sd}(I - Q_{dd})^{-1} Q_{ds} m_{sj} + Q_{ss} m_{sj}, \end{aligned}$$

where the inequality is a consequence of (24). Rearranging the terms thus implies

$$C m_{sj} \geq m^*(C\mathbf{1}_s),$$

where matrix $C \in \mathbb{R}^{s \times s}$ is given by

$$C = \left[I - Q_{ss} - Q_{sd}(I - Q_{dd})^{-1} Q_{ds}\right].$$

The proof is thus complete once we show that C^{-1} is an element-wise non-negative matrix, as it would immediately imply that $m_{sj} \geq m^* \mathbf{1}_s$. To establish this, first note that C is the Schur complement of the non-singular, M-matrix G defined as

$$G = \begin{bmatrix} I - Q_{ss} & -Q_{sd} \\ -Q_{ds} & I - Q_{dd} \end{bmatrix}.$$

By Berman and Plemmons (1979, exercise 5.8, p. 159), the Schur complement of any non-singular M-matrix is itself a non-singular M-matrix.[35] Thus, Theorem 2 of Plemmons (1977) guarantees that C^{-1} is element-wise non-negative, completing the proof. $\square$

Proof of part (b). Suppose that all banks default if bank j is hit with the negative shock, which means that $x_i < y$ for all banks $i \neq j$. On the other hand, from the definition of a payment equilibrium it is immediate that

$$x_i = a - v + \sum_{k \neq j} q_{ik} x_k,$$

for all $i \neq j$. Dividing both sides of the above equation by $a - v$ and comparing it with the definition of the harmonic distance over the financial network (3) implies that $x_i/(a - v) = m_{ij}$.[36] Thus, the fact that $x_i < y$ guarantees that $m_{ij} < m^*$, which completes the proof. $\square$

[35] For a proof, see Theorem 1 of Carlson and Markham (1979).

[36] Note that we are using the fact that for any j, equation (3) has a unique solution.

Proof of Lemma 1

Lemma B.7. *For any regular financial network,*

$$\frac{1}{8}n^2\phi^2 \leq 1-\lambda_2 \leq 2n\phi,$$

where λ_2 is the second largest eigenvalue of Q and ϕ is the bottleneck parameter.

Proof. Let Φ be the *conductance* of the graph, defined as

$$\Phi = \min_{S\subset N} \frac{Q(S,S^c)}{\min\{|S|,|S^c|\}},$$

where $Q(S,S^c) = \sum_{i\in S, j\notin S} q_{ij}$. It is easy to verify that

$$(n/2)\min\{|S|,|S^c|\} \leq |S||S^c| \leq n\min\{|S|,|S^c|\},$$

which implies that

$$\Phi/n \leq \phi \leq 2\Phi/n. \tag{27}$$

On the other hand, by Theorem 2 of Sinclair (1992),

$$\Phi^2/2 \leq 1-\lambda_2 \leq 2\Phi. \tag{28}$$

Combining (27) and (28) completes the proof. □

Proof of of Lemma 1. We first prove the lower bound. Let $M=[m_{ij}]$ denote the matrix of pairwise harmonic distances between the banks, and define $T = \mathbf{1}_n\mathbf{1}_n' + QM - M$. For any pair of banks $i\neq j$, we have

$$t_{ij} = 1 + \sum_{k=1}^{n} q_{ik}m_{kj} - m_{ij} = 0,$$

where the second equality is a consequence of the definition of the harmonic distance (3). This means that T is a diagonal matrix. Furthermore, $\mathbf{1}_n'T = n\mathbf{1}_n'$, which implies that all diagonal elements of T are equal to n, and as a consequence,

$$(I-Q)M = \mathbf{1}_n\mathbf{1}_n' - nI. \tag{29}$$

Solving for M, we have

$$M = n(\mathbf{1}_n\mathbf{1}_n'\mathrm{diag}(Z) - Z), \tag{30}$$

where $Z = \left(I - Q + \frac{1}{n}\mathbf{1}_n\mathbf{1}_n'\right)^{-1}$ and $\mathrm{diag}(X)$ is a diagonal matrix whose elements are the diagonal entries of X.[37] Therefore,

$$\begin{aligned} M\mathbf{1}_n &= n(\mathbf{1}_n\mathbf{1}_n'\mathrm{diag}(Z)\mathbf{1}_n - Z\mathbf{1}_n) \\ &= n(\mathrm{trace}(Z)-1)\mathbf{1}_n, \end{aligned}$$

[37] Since $I-Q$ is not invertible, equation (29) has infinitely many solutions. However, the restriction that $m_{ii}=0$ for all i uniquely determines the matrix M of pairwise harmonic distances.

where the second equality is a consequence of the fact that $Z\mathbf{1}_n = \mathbf{1}_n$. Consequently, for any bank i,

$$\frac{1}{n}\sum_{j\neq i} m_{ij} = \mathrm{trace}(Z) - 1 = \sum_{k=2}^{n} \frac{1}{1-\lambda_k}, \tag{31}$$

in which λ_k is the k-th largest eigenvalue of Q. The second equality above relies on the fact that the eigenvalues of Z are simply the reciprocal of the non-zero eigenvalues of $I - Q$. Thus, (31) implies

$$\max_{i\neq j} m_{ij} \geq \sum_{k=2}^{n} \frac{1}{1-\lambda_k} \geq \frac{1}{1-\lambda_2}.$$

Lemma B.7 now guarantees that $\max_{i\neq j} m_{ij} \geq 1/(2n\phi)$ establishing the lower bound in (4).

To establish the upper bound,[38] note that since Q is symmetric with 1 as its top eigenvalue, it can be written in spectral form as

$$Q = \frac{1}{n}\mathbf{1}_n\mathbf{1}_n' + \sum_{k=2}^{n} \lambda_k w_k w_k',$$

where $\{w_2, \dots, w_n\}$ are eigenvectors of Q corresponding to eigenvalues $\{\lambda_2, \dots, \lambda_n\}$ with lengths normalized to one. Consequently, $Z = \sum_{k=2}^{n} \frac{1}{1-\lambda_k} w_k w_k'$. Equation (30) then implies

$$m_{ij} = n\sum_{k=2}^{n} \frac{1}{1-\lambda_k}\left(w_{kj}^2 - w_{ki}w_{kj}\right),$$

for all pairs of banks i and j, and hence,

$$\begin{aligned} m_{ij} + m_{ji} &= n\sum_{k=2}^{n} \frac{1}{1-\lambda_k}(w_{ki} - w_{kj})^2 \\ &\leq \frac{n}{1-\lambda_2}\sum_{k=1}^{n}(w_{ki}^2 + w_{kj}^2 - 2w_{ki}w_{kj}) \\ &= \frac{2n}{1-\lambda_2}. \end{aligned}$$

The last equality above is a consequence of the observation that the collection of vectors $\{w_1, \dots, w_n\}$ form an orthonormal basis, and hence, it must be the case that $\sum_{k=1}^{n} w_{ki}^2 = 1$ for all banks i and that $\sum_{k=1}^{n} w_{ki}w_{kj} = 0$ for all $i \neq j$.[39] Given that the above inequality holds for any pairs of banks $i \neq j$, it is immediate that

$$\max_{i\neq j} m_{ij} \leq \frac{2n}{1-\lambda_2}. \tag{32}$$

Combining (32) with the lower bound in Lemma B.7 completes the proof. □

[38] The proof of the upper bound follows steps similar to those in, and generalizes, Lovász (1993, Theorem 3.1).

[39] By Horn and Johnson (1985, Theorem 2.1.4), if the rows of a square matrix X form an orthonormal basis, so do its columns.

References

Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2012), "The network origins of aggregate fluctuations." *Econometrica*, 80, 1977–2016.

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2013a), "The network origins of large economic downturns." NBER Working Paper No. w19230.

Acemoglu, Daron, Asuman Ozdaglar, and Alireza Tahbaz-Salehi (2013b), "Systemic risk and stability in financial networks." NBER Working Paper No. w18727.

Akram, Q. Farooq and Casper Christophersen (2010), "Interbank overnight interest rates-gains from systemic importance." Norges Bank Working Paper.

Allen, Franklin and Ana Babus (2009), "Networks in finance." In *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World* (Paul R. Kleindorfer and Yoram (Jerry) Wind, eds.), 367–382, Wharton School Publishing.

Allen, Franklin, Ana Babus, and Elena Carletti (2012), "Asset commonality, debt maturity and systemic risk." *Journal of Financial Economics*, 104, 519–534.

Allen, Franklin and Douglas Gale (2000), "Financial contagion." *Journal of Political Economy*, 108, 1–33.

Alvarez, Fernando and Gadi Barlevy (2013), "Mandatory disclosure and financial contagion." Working Paper.

Babus, Ana (2013), "The formation of financial networks." Discussion Paper 06-093, Tinbergen Institute.

Barlevy, Gadi and Haikady N. Nagaraja (2013), "Properties of the vacancy statistic in the discrete circle covering problem." FRB of Chicago Working Paper No. 2013-05.

Bech, Morten L. and Enghin Atalay (2010), "The topology of the federal funds market." *Physica A: Statistical Mechanics and its Applications*, 389, 5223–5246.

Bech, Morten L., James T. E. Chapman, and Rodney J. Garratt (2010), "Which bank is the "central" bank?" *Journal of Monetary Economics*, 57, 352–363.

Berman, Abraham and Robert J. Plemmons (1979), *Nonnegative Matrices in the Mathematical Sciences*. Academic Press, New York.

Bisias, Dimitrios, Mark Flood, Andrew Lo, and Stavros Valavanis (2012), "A survey of systemic risk analytics." *US Department of Treasury, Office of Financial Research*.

Blume, Larry, David Easley, Jon Kleinberg, Robert Kleinberg, and Éva Tardos (2011), "Network formation in the presence of contagious risk." *Proceedings of the 12th ACM Conference on Electronic Commerce*.

Brunnermeier, Markus K. and Martin Oehmke (2013), "Bubbles, financial crises, and systemic risk." volume 2, part B of *Handbook of the Economics of Finance*, 1221–1288, Elsevier.

Brunnermeier, Markus K. and Lasse Heje Pedersen (2005), "Predatory trading." *The Journal of Finance*, 60, 1825–1863.

Caballero, Ricardo J. and Alp Simsek (2013), "Fire sales in a model of complexity." *Journal of Finance*, 68, 2549–2587.

Cabrales, Antonio, Piero Gottardi, and Fernando Vega-Redondo (2013), "Risk-sharing and contagion in networks." Working Paper.

Carlson, David and Thomas L. Markham (1979), "Schur complements of diagonally dominant matrices." *Czechoslovak Mathematical Journal*, 29, 246–251.

Castiglionesi, Fabio, Fabio Feriozzi, and Guido Lorenzoni (2010), "Financial integration and liquidity crises." Working Paper.

Chung, Fan R. K. (1997), *Spectral Graph Theory*, volume 92. American Mathematical Society.

Cifuentes, Rodrigo, Gianluigi Ferrucci, and Hyun Song Shin (2005), "Liquidity risk and contagion." *Journal of the European Economic Association*, 3, 556–566.

Cohen-Cole, Ethan, Eleonora Patacchini, and Yves Zenou (2013), "Systemic risk and network formation in the interbank market." Working Paper.

Craig, Ben R., Falko Fecht, and Günseli Tümer-Alkan (2013), "The role of interbank relationships and liquidity needs." Working Paper.

Dasgupta, Amil (2004), "Financial contagion through capital connections: A model of the origin and spread of bank panics." *Journal of the European Economic Association*, 2, 1049–1084.

Denbee, Edward, Christian Julliard, Ye Li, and Kathy Yuan (2013), "Network risk and key players: a structural analysis of interbank liquidity." Working Paper.

Diamond, Douglas W. and Philip H. Dybvig (1983), "Bank runs, deposit insurance, and liquidity." *Journal of Political Economy*, 91, 401–419.

Eboli, Mario (2012), "A flow network analysis of direct balance-sheet contagion in financial networks." Working Paper.

Eisenberg, Larry and Thomas H. Noe (2001), "Systemic risk in financial systems." *Management Science*, 47, 236–249.

Elliott, Matthew, Benjamin Golub, and Matthew O. Jackson (2013), "Financial networks and contagion." Working Paper.

Farboodi, Maryam (2013), "Intermediation and voluntary exposure to counterparty risk." Working Paper.

Freixas, Xavier, Bruno M. Parigi, and Jean-Charles Rochet (2000), "Systemic risk, interbank relations, and liquidity provision by the central bank." *Journal of Money, Credit and Banking*, 32, 611–638.

Gai, Prasanna, Andrew Haldane, and Sujit Kapadia (2011), "Complexity, concentration and contagion." *Journal of Monetary Economics*, 58, 453–470.

Georg, Co-Pierre (2013), "The effect of the interbank network structure on contagion and common shocks." *Journal of Banking & Finance*, 37, 2216–2228.

Glasserman, Paul and H. Peyton Young (2013), "How likely is contagion in financial networks?" University of Oxford Department of Economics Discussion Paper 642.

Gofman, Michael (2011), "A network-based analysis of over-the-counter markets." Working Paper.

Gofman, Michael (2013), "Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions." Working Paper.

Haldane, Andrew G. (2009), "Rethinking the financial network." Speech delivered at the Financial Student Association in Amsterdam, The Netherlands, April 2009. www.bankofengland.co.uk/archive/Documents/historicpubs/speeches/2009/speech386.pdf.

Holmström, Bengt and Jean Tirole (1998), "Private and public supply of liquidity." *Journal of Political Economy*, 106, 1–40.

Horn, Roger A. and Charles R. Johnson (1985), *Matrix Analysis*. Cambridge University Press, New York, NY.

Kirkland, Stephen J. and Michael Neumann (2012), *Group Inverses of M-matrices and Their Applications*, volume 26. CRC Press.

Kiyotaki, Nobuhiro and John Moore (1997), "Credit chains." Working Paper.

Krishnamurthy, Arvind (2010), "Amplification mechanisms in liquidity crises." *American Economic Journal: Macroeconomics*, 2, 1–30.

Leitner, Yaron (2005), "Financial networks: Contagion, commitment, and private sector bailouts." *The Journal of Finance*, 60, 2925–2953.

Levin, David Asher, Yuval Peres, and Elizabeth Lee Wilmer (2009), *Markov Chains and Mixing Times*. American Mathematical Society, Providence, RI.

Lorenzoni, Guido (2008), "Inefficient credit booms." *The Review of Economic Studies*, 75, 809–833.

Lovász, László (1993), "Random walks on graphs: A survey." *Combinatorics, Paul Erdős is Eighty*, 2, 1–46.

Plemmons, Robert J. (1977), "M-matrix characterizations. I: Nonsingular M-matrices." *Linear Algebra and its Applications*, 18, 175–188.

Plosser, Charles I. (2009), "Redesigning financial system regulation." Speech at the New York University Conference "Restoring Financial Stability: How to Repair a Failed System". http://www.phil.frb.org/publications/speeches/plosser/2009/03-06-09_nyu-restoring-financial-stability.pdf.

Rochet, Jean-Charles and Jean Tirole (1996), "Interbank lending and systemic risk." *Journal of Money, Credit and Banking*, 28, 733–762.

Rotemberg, Julio J. (2011), "Minimal settlement assets in economies with interconnected financial obligations." *Journal of Money, Credit and Banking*, 43, 81–108.

Shin, Hyun Song (2008), "Risk and liquidity in a system context." *Journal of Financial Intermediation*, 17, 315 – 329.

Shin, Hyun Song (2009), "Securitisation and financial stability." *The Economic Journal*, 119, 309–332.

Sinclair, Alistair (1992), "Improved bounds for mixing rates of Markov chains and multicommodity flow." *Combinatorics, Probability and Computing*, 1, 351–370.

Soramäki, Kimmo and Samantha Cook (2013), "SinkRank: An algorithm for identifying systemically important banks in payment systems." *Economics: The Open-Access, Open-Assessment E-Journal*, 7.

Sorkin, Andrew Ross (2009), *Too Big to Fail: The Inside Story of How Wall Street and Washington Fought to Save the Financial System—and Themselves*. Viking, New York.

Vivier-Lirimont, Sébastian (2006), "Contagion in interbank debt networks." Working Paper.

Yellen, Janet L. (2013), "Interconnectedness and systemic risk: Lessons from the financial crisis and policy implications." Remarks at the American Economic Association/American Finance Association joint luncheon, San Diego, California.
http://www.federalreserve.gov/newsevents/speech/Yellen20130104a.pdf.

Zawadowski, Adam (2011), "Interwoven lending, uncertainty, and liquidity hoarding." Working Paper.

Zawadowski, Adam (2013), "Entangled financial systems." *The Review of Financial Studies*, 26, 1291–1323.

C Omitted Proofs (Not for Publication)

Proof of Proposition 8

Proof of part (a). First consider the complete financial network. If $\epsilon < \epsilon_p^*$, at least one bank does not default. Given the symmetry, all $n-p$ banks that are not hit with a negative shock do not default either, implying that the complete network is the most stable and resilient financial network in the face of small shocks.

Now consider the ring financial network and assume that p consecutive banks, labeled $i+1$ through $j = i+p$, are hit with negative shocks. An immediate observation is that all banks in default also form a connected chain, say of length $\tau \geq p$, the last of which is labeled $s = i + \tau$. In view of Lemma B.4, bank i does not default, as it is the bank furthest away from the realized shocks. As a result, as long as $y > y_p^* = (n-p)(a-v)$, in the unique payment equilibrium of the financial network, all banks can meet their senior liabilities v in full. This can be established by verifying that bank j — which is the bank facing the most amount of potential distress — can pay its senior debts. In particular,

$$x_{j,j-1} = y + (p-1)(a - \epsilon - v),$$

guaranteeing that $x_{j,j-1} + a - \epsilon > v$. Given that all banks can meet their senior liabilities, we have

$$x_{s+1,s} = y + \tau(a-v) - p\epsilon$$

where $s = i + \tau$ is the index of the last bank on the chain that defaults. On the other hand, given that $s+1$ does not default, we have $y \leq a - v + x_{s+1,s}$. As a result,

$$\tau = \left\lceil \frac{p\epsilon}{a-v} \right\rceil - 1 \geq \frac{p\epsilon}{a-v} - 1. \tag{33}$$

Hence, when shocks hit p consecutive banks on the credit chain, the number of bank failures reaches the upper bound established by Lemma B.3, implying that the ring network is the least resilient financial network. □

Proof of part (b). The proof follows a logic similar to that of Proposition 5. We first prove that if $\epsilon > \epsilon_p^*$, then the complete network is the least stable and resilient financial network. In particular, we show that all banks default. By Lemma B.4, the p distressed banks default on their senior liabilities. The remaining $n-p$ banks do not default only if

$$(n-p-1)\frac{y}{n-1} + (a-v) \geq y.$$

The above inequality, however, can hold only if $y < \hat{y}_p = (n-1)(a-v)/p \leq y_p^*$. Hence, the complete network is the least resilient and the least stable financial network as all n banks default.

We next show that if $\epsilon > \epsilon_p^*$, then all n banks in the ring network fail as well. Suppose not, and that there exists a bank that can pay all its creditors in full. On the other hand, by Lemma B.4, there is also a bank that defaults on its senior liabilities v. Consider the path on the ring network connecting bank j to bank l, such that (i) j defaults on its senior debt; (ii) l pays all its creditors in full; and (iii) all banks on the path default but can pay back their senior debt. Denote the length of the path connecting j to l by τ (see Figure 4), and suppose that there are h negative shocks realized on this path.

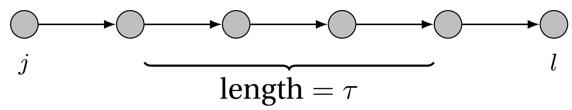

Figure 4. There are τ banks connecting j to l, all of which default, but can meet their senior liabilities.

Given that j does not pay anything to its junior creditor (which is the first bank on the path connecting it to l) and that l does not default, we have $(\tau+1)(a-v) - h\epsilon \geq y$, implying that

$$\begin{aligned}\tau &> \frac{y_p^* + h\epsilon_p^*}{a-v} - 1 \\ &= n - p - 1 + \frac{hn}{p}.\end{aligned}$$

On the other hand, the remaining $p-h$ shocks hit banks that are not on the path connecting j to l. Thus, the total number of defaults is at least $\tau + p - h$, implying

$$\sharp\text{defaults} > n - 1 + h\left(\frac{n}{p} - 1\right) \geq n - 1.$$

This, however, contradicts the assumption that at least one bank does not default.

Finally, consider a δ-connected financial network with the corresponding partition (S, S^c) such that $|S^c| = p$. Note that by definition, $\max\{y_{ij}, y_{ji}\} \leq \delta y$ for all $i \in S$ and $j \in S^c$. Therefore, for any bank $i \in S$, it must be the case that $\sum_{j \in S} y_{ij} \geq y - p\delta y$. On the other hand, in the case that all p negative shocks hit the banks in S^c, any bank i can meet its liabilities in full as long

$$a - v + \sum_{j \in S} y_{ij} \geq y.$$

Thus, as long as $\delta < (a-v)/(py)$, then no bank in S defaults, establishing that the given financial network is strictly more stable than the complete financial network. $\square$

Proof of part (c). An argument similar to the one invoked in the proof of part (b) shows that the when $\epsilon > \epsilon_p^*$ and $y > \hat{y}_p$, all banks in the complete network default. Therefore, the complete network is the least stable and resilient financial network.

To prove that the ring financial network is more stable than the complete network, we show that there exists a realization of the shocks for which at least one bank in the ring network does not default. In particular, consider the situation in which p consecutive banks, labeled 1 through p, are hit with negative shocks. By Lemma B.4, in the unique payment equilibrium, bank p defaults on its senior debt. Therefore, the length of the cascade of defaults following bank p, denoted by τ, satisfies

$$\tau(a-v) < y \leq (\tau+1)(a-v).$$

Thus, the number of defaults in the whole network is

$$\begin{aligned}\sharp\text{defaults} &= p+\tau \\ &< p + \frac{y_p^*}{a-v} \\ &= n,\end{aligned}$$

implying that at least one bank does not default. Hence, the ring network is strictly more stable than the complete network. □

Proof of Proposition 9

The proofs of parts (a) and (b) are similar to those of Propositions 3 and 5, respectively, and are thus omitted. To prove part (c), first consider the complete financial network and without loss of generality, assume that bank 1 is hit with the negative shock. It is easy to verify that, as long as $\epsilon_*(\zeta) < \epsilon < \epsilon^*(\zeta)$, the unique payment equilibrium is given by

$$\begin{aligned}(x_1,\ell_1) &= (a-v-\epsilon+\zeta A+y, A) \\ (x_i,\ell_i) &= \left(y, \frac{\epsilon-n(a-v)-\zeta A}{\zeta(n-1)}\right)\end{aligned}$$

Therefore, the total liquidation across the financial network satisfies

$$\sum_{i=1}^{n}\zeta\ell_i = \epsilon - n(a-v). \tag{34}$$

Next consider the ring financial network. Again, it is east to verify that if bank 1 is hit with the negative shock, then the total amount of liquidation across the financial network satisfies

$$\sum_{i=1}^{n}\zeta\ell_i = \tau\zeta A + [\epsilon-(\tau+1)(a-v)-\tau\zeta A]^+, \tag{35}$$

where $\tau = \lceil \epsilon/(a-v+\zeta A)\rceil - 1$ is the number of defaults. Comparing (34) with (35) then immediately implies that the extent of liquidation in the complete financial network is strictly smaller than the ring financial network. Hence, the former is strictly more stable and resilient than the latter.

Finally, consider the financial network depicted in Figure 2 with $q = 1$. Suppose that bank 1 is hit with a negative shock, which immediately implies that banks 1 and 2 default and liquidate their projects entirely, whereas all other banks can meet their liabilities in full. Consequently,

$$\sum_{i=1}^{n}\zeta\ell_i = 2\zeta A.$$

Comparing the above with (34) shows that as long as $\epsilon > \epsilon_*(\zeta) + \zeta A$, then the given 0-connected network is strictly more stable and resilient than the complete financial network. □

Proof of Proposition 10

Proof of part (a). The proof closely follows the proof of the second part of Lemma B.4. In particular, suppose that $\epsilon > (a-v)\sum_{k=1}^{n}\theta_k/\theta_j$, but all banks can meet their liabilities to the senior creditors in full. Thus, by the definition of the payment equilibrium,

$$z_i + \sum_{k\neq i} x_{ik} \geq \theta_i v + \sum_{k\neq i} x_{ki},$$

for all banks i. Summing over all i implies

$$a\sum_{i=1}^{n}\theta_i - \theta_j\epsilon \geq v\sum_{i=1}^{n}\theta_i,$$

which is a contradiction. Thus, the distressed bank j defaults on its liabilities to the senior creditors. □

Proof of part (b). In the presence of a large shock to bank j, all other banks default if and only if $x_i < \theta_i y$ for all i, where x_i's are the solutions to the following collection of equations:

$$x_i = \theta_i(a-v) + \sum_{k\neq j} q_{ik}x_k.$$

Comparing the above equation to (5), however, implies that $x_i = (a-v)\widehat{m}_{ij}$. Thus, all banks default if and only if $\widehat{m}_{ij} < \theta_i m^*$, completing the proof. □

Proof of Proposition A.1

Let $\{\chi_t\}_{t\geq 0}$ denote the discrete-time, discrete-space Markov chain with the transition probability matrix Q; that is, $\mathbb{P}(\chi_{t+1} = j|\chi_t = i) = q_{ij}$. Also, let τ_{ij} denote the number of time steps that it takes to visit state j for the first time; that is, $\tau_j = \min\{t \geq 0 : \chi_t = j\}$. Therefore, the mean hitting time of state j conditional on starting from state i is given by

$$\begin{aligned}
\mathbb{E}_i[\tau_j] &= \sum_{t=1}^{\infty} t\mathbb{P}(\tau_j = t|\chi_0 = i) \\
&= \sum_{t=1}^{\infty}\sum_{k=1}^{n} t\mathbb{P}(\tau_j = t, \chi_1 = k|\chi_0 = i) \\
&= \sum_{k=1}^{n} q_{ik}\sum_{t=1}^{\infty} t\mathbb{P}(\tau_j = t|\chi_1 = k),
\end{aligned}$$

implying that the mean hitting times satisfy the following fixed point equation:

$$\mathbb{E}_i[\tau_j] = 1 + \sum_{k=1}^{n} q_{ik}\mathbb{E}_k[\tau_j].$$

This equation, however, is identical to equation (3). Furthermore, given the argument in the proof of Lemma 1, the equation has a unique solution. Therefore, $\mathbb{E}_i[\tau_j] = m_{ij}$. □

Proof of Proposition A.2

Levin, Peres, and Wilmer (2009, Lemma 10.10) show that in any reversible Markov chain and for any three states i, j and k, we have $\mathbb{E}_i[\tau_j] + \mathbb{E}_j[\tau_k] + \mathbb{E}_k[\tau_i] = \mathbb{E}_j[\tau_i] + \mathbb{E}_k[\tau_j] + \mathbb{E}_i[\tau_k]$. On the other hand, by Proposition A.1, the harmonic distances in the financial network are equal to the mean hitting times in the corresponding Markov chain, establishing (6). $\square$

Proof of Corollary A.1

Pick an arbitrary bank k, and create an ordering of the rest of the banks according to the value of $m_{ik} - m_{ki}$. More specifically, let bank i appear before bank j if $m_{ik} - m_{ki} \geq m_{jk} - m_{kj}$. Proposition A.2, on the other hand, requires (6) to hold. Consequently, it must be the case that $m_{ij} \geq m_{ji}$. $\square$

Proof of Proposition A.3

Kirkland and Neumann (2012, Theorem 6.2.1) show that Markov chain mean hitting times satisfy the triangle inequality. Thus, by Proposition A.1, the harmonic distances satisfy the triangle inequality as well. $\square$

Proof of Proposition A.4

Equation (31) in the proof of Lemma 1 establishes that

$$\frac{1}{n}\sum_{j\neq i} m_{ij} = \sum_{k=2}^{n} \frac{1}{1-\lambda_k},$$

where $\{\lambda_2, \ldots, \lambda_n\}$ are the $n-1$ smallest eigenvalues of matrix Q (that is, excluding $\lambda_1 = 1$). Given that the right-hand side above is independent of i, it is immediate that the average harmonic distance from bank i to all other banks is an invariant property of the financial network. $\square$