



Systemic risk shifting in financial networks [☆]

Matthew Elliott ^{a,*}, Co-Pierre Georg ^{b,c}, Jonathon Hazell ^d

^a *Cambridge University, United Kingdom of Great Britain and Northern Ireland*

^b *University of Cape Town, South Africa*

^c *Bundesbank, Germany*

^d *Princeton, United States of America*

Received 26 June 2019; final version received 9 November 2020; accepted 11 November 2020

Available online 16 November 2020

Abstract

Banks face different but potentially correlated risks from outside the financial system. Financial connections can share these risks, but they also create the means by which shocks can be propagated. We examine this tradeoff in the context of a new stylized fact we present: German banks are more likely to have financial connections when they face more similar risks. We develop a model that can rationalize such behavior. We argue that such patterns are socially suboptimal and raise systemic risk, but can be explained by risk shifting. Risk shifting motivates banks to correlate their failures with their counterparties, even though it creates systemic risk.

© 2020 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

JEL classification: G21; G11; D85

Keywords: Financial networks; Asset correlation; Contagion

[☆] This paper subsumes “Endogenous Financial Networks: Efficient Modularity and Why Shareholders Prevent It.” We wish to thank our discussant Alireza Tahbaz-Salehi, as well as Daron Acemoglu, Agostino Capponi, Antonio Cabrales, Ben Golub, Piero Gottardi, Sanjeev Goyal, Matthew Jackson, Adam Szeidl, Alp Simsek, Xavier Vives and three anonymous referees, as well as numerous seminar and conference participants, for their helpful comments and discussions. Jörg Kalbfuss provided excellent research assistance. Elliott acknowledges funding from the Cambridge Endowment for Research in Finance and the European Research Council under the grant EMBED #757229. Hazell acknowledges funding from the Caltech Summer Research Fellowship programme. This paper is partly based on Hazell’s undergraduate dissertation. The views expressed in this paper are those of the authors and do not necessarily reflect the opinions of the Deutsche Bundesbank, the Eurosystem or its staff.

* Corresponding author.

E-mail address: mle30@cam.ac.uk (M. Elliott).

<https://doi.org/10.1016/j.jet.2020.105157>

0022-0531/© 2020 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

In the financial system, systemic risk comes from two sources. First, banks may be exposed to the same real investments. Second, banks have financial exposure to each other. In either case, shocks to the real economy can spread across the financial system. As emphasized by policymakers (e.g. Basel Committee on Banking Supervision, 2011), the interaction between real exposure and financial exposure is key for financial stability. Conventional wisdom argues banks should prefer counterparties with differing real exposure. Then, banks can share real risks across the financial system (Rochet and Tirole, 1996; Allen and Gale, 2000). But financial connections frequently play the opposite role. Banks often have similar real exposures to their financial counterparties. Financial connections then *concentrate* real risk. After a shock to the real economy, banks are doubly affected—first, through their direct exposure, and second through their indirect exposure via their financial counterparties. In the United States, various anecdotes from the 2008 financial crisis display this pattern. Many institutions had the same real exposure as their financial counterparties, which raised systemic risk:

- Among US commercial banks, over-the-counter (OTC) derivatives trading was concentrated at the five largest banks. These banks were financially exposed to each other, while at the same time facing the same real exposure to the housing market.¹
- Investment banks used monoline insurers to share their subprime mortgage risk, but these monoline insurers, through their asset management and collateralized debt obligation (CDO) arms, had the same real exposure to subprime mortgages.²

We first present a new stylized fact from the German banking system. Banks have similar real exposures to their financial counterparties. Therefore, if German commercial banks lend to one another in the interbank market, they tend to lend to similar non-financial firms. This fact may be surprising—financial connections do not diversify real exposure. Motivated by this stylized fact, we provide a model that can rationalize this behavior even though it generates systemic risk. We study a model with limited liability, real investments, and a financial network. We characterize socially efficient networks. In these networks, which minimize systemic risk, banks have different real exposures from their counterparties. Absent limited liability, banks have no incentive to deviate from the socially efficient network. But limited liability leads banks to deviate from social efficiency and engage in *systemic risk shifting*. Banks increase their equity values by having the same real exposures as their close financial counterparties. Then, banks fail at the same time as their counterparties, raising their values conditional on not failing, and, hence their equity values, but also increasing systemic risk.

The German commercial banking system is an ideal setting for studying the relationship between banks' financial and real exposures. Supervisory data from the German Credit Register records the near-universe of bank-bank and bank-firm lending. Interbank loans proxy for financial exposures, and commercial loans proxy for real exposures. Large interbank exposures, with long-term maturities, create substantial counterparty risk in the financial system.³ Bank-firm

¹ See, for example, the report of the 2011 Financial Crisis Inquiry Commission.

² SEC, "Risk Management Reviews of Consolidated Supervised Entities," internal memo to Erik Sirri and others, November 6, 2007, p. 3.

³ The share of interbank lending over equity is 17% for banks at the mean of our sample and 128% at the 99th percentile. Interbank loans have an average maturity of more than one year.

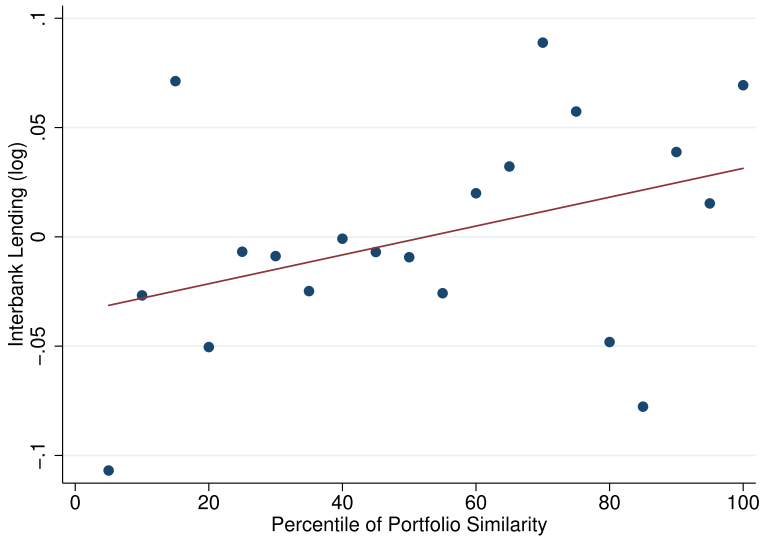


Fig. 1. Banks Have Similar Real Exposure to Their Financial Counterparties. *Notes:* the y-axis is log of net quarterly interbank lending between each pair of commercial banks in the German Credit Register for 2006-2014. The x-axis is the rank of how similar two banks’ commercial loan portfolios are in a quarter, based on the overlap of the firms to which each bank lends. We residualize both variables against time, lender and borrower fixed effects, and plot a binned scatter with 20 points.

lending is the largest class of real investments for most banks.⁴ Thus, Upper and Worms (2004) and others argue that shocks to firms or sectors can lead to systemic risk in the German financial system.

Fig. 1 illustrates our stylized fact. Banks with more similar real exposures tend to lend more to each other. The figure relates the similarity between two banks’ non-financial portfolios, and the size of their net interbank lending in a given quarter. We calculate a measure of each bank pair’s portfolio similarity, based on the overlap in the portfolio of firms to which they lend in a given quarter, and rank all bank pairs by this measure. We residualize by a set of fixed effects and then relate this measure to interbank lending between each bank pair. We present the result in a binned scatterplot. Portfolio similarity and interbank lending are significantly and positively related. The effect is large. When two banks move from the 25th to the 75th percentile of non-financial similarity, their net lending to one another increases by about 31% or EUR 13.8 million. The results are statistically significant, robust to restricting only to banks with national coverage, and hold after adding borrower-by-quarter, lender-by-quarter and borrower-lender fixed effects.⁵ Our main measure of real exposure comes from overlap in banks’ portfolios at the firm level, but our findings also hold with the analogous sectoral measure of similarity.

The general pattern that more similar nodes are more likely to link to each other in a network is known as *homophily*. Encapsulated in the proverb that “birds of a feather flock together,” this pattern has been noticed since at least the 16th century. Homophily is a robust feature in net-

⁴ The average bank in our sample has about ten times as much in loans to the real economy as in equity. Given the low rate of home ownership in Germany, residential lending is small for most banks.

⁵ Banks with national coverage exclude regional banks such as savings central banks (Landesbanken).

works.⁶ Our results document homophily in a new, and perhaps surprising, setting—the German banking system.

We develop a model in which banks choose their real exposures and financial connections that can rationalize this pattern, even though it is associated with systemic risk. The model has three key ingredients: real investments, a financial network, and limited liability. There is a set of financial institutions, termed banks. There is a set of real investments that generate random returns. Each bank holds a portfolio of the investments. Banks can hold any share in any investment. Banks also hold financial claims on one another. These claims form a financial network. Banks' market values are determined by the investment returns flowing to each bank. Banks have liabilities to external debt holders. The residual market value, after external debt, becomes equity value. In states of the world where market values are less than the value of external debt, banks fail. These failures trigger discontinuous falls in the value of financial claims, due to default costs. By limited liability, equity value is zero if banks fail.

Our model captures a tradeoff between two priorities: risk sharing, and minimizing systemic risk. By holding multiple real investments or financial claims on counterparties, banks share the risks associated with a given investment. But both investment portfolios and the financial network lead to systemic risk. If multiple banks hold the same investment, they may all simultaneously fail if the investment has a low return. If banks hold financial claims on one another, then a second bank may fail if a first bank's portfolio has a low return. Default costs amplify systemic risk by lowering the value of claims the second bank has with the first bank, when the first bank fails.

Our model makes two points. First, the pattern in the data—where banks have similar real exposures to their financial counterparties—is not socially efficient. We characterize the socially efficient networks and portfolios. Banks should hold different investments from their closest financial counterparties. This structure minimizes the likelihood that a bank's portfolio draws bad returns at the same time as its counterparties. Collectively, the banking system then absorbs fewer losses, which lowers the risk of bank failures. Our characterization reveals an additional subtlety. The social planner optimally partitions the financial network into groups of banks, with strong financial claims within groups, and weak claims between groups. This network prevents failures after relatively small falls in investment returns, and minimizes systemic risk after large falls in returns. The weak financial claims between groups prevent failures from spreading throughout the network.

Second, we show that the socially efficient outcome does not arise in equilibrium. Instead, consistent with the stylized fact, banks seek to have the same real exposure as their financial counterparties. The reason for this behavior is limited liability. The result of this behavior is greater systemic risk. Limited liability creates behavior that we and others term *systemic risk shifting*.⁷

We isolate the role of limited liability through two steps. In the first step, we show deviations from the socially efficient outcome cannot raise banks' expected market values—where market values are the sum of external debt and equity. In a deviation, banks can change either their financial claims or their investment portfolios. Deviating from the socially efficient outcome increases the expected number of failures. These failures create default costs, which lower the

⁶ See McPherson et al. (2001) for a survey.

⁷ See Acharya (2009) for an earlier use of the term.

market value of the deviating banks through the interconnectedness of the network. Thus, the socially efficient outcome can arise in equilibrium if banks maximize market value.

In the second step, we show banks can raise their *equity* values by deviating from the socially efficient outcome. Limited liability, which creates the difference between equity and market value, leads to socially inefficient behavior. With limited liability, banks increase their expected equity value by failing at the same time as their closest counterparties. Then, neither is affected by the default costs associated with the other's failure. Instead, external debtors lose value from counterparties' default costs. These deviations raise banks' market value when they do not fail, leading to higher equity value on average. Moreover these deviations raise systemic risk—the chance that a large share of the network fails. Consequently, there is systemic risk shifting. Changes in either portfolios or financial claims can lead to systemic risk shifting. Either change can raise equity value if banks become more likely to fail at the same time as their counterparties.

Related Literature

Our paper falls within the financial networks literature. The literature examining how the structure of exogenous financial networks affects systemic risk has expanded rapidly. Building on early important works, such as Allen and Gale (2000) and Freixas et al. (2000), this literature emphasizes that interconnections can spread contagion. Networks can generate systemic risk by facilitating the spread of relatively large shocks (Gai and Kapadia, 2010), or by interacting with various propagation mechanisms, such as bankruptcy costs (Elliott et al., 2014); uncertainty about banks' balance sheets (Caballero and Simsek, 2013; Alvarez and Barlevy, 2014); and fire sales (Cifuentes et al., 2005). Correlated exposures have also been shown to make a major contribution to systemic risk (Glasserman and Young, 2015) and will amplify the impact of the other mechanisms. Allen and Babus (2009) and Capponi (2016) survey this literature. Since these papers consider exogenous networks and fixed real exposures, they do not examine socially efficient network structures, nor whether individual banks might choose to deviate from the social optimum.

Most papers studying financial networks empirically focus on characterizing the network structure to identify the risk of contagion (see Upper and Worms (2004) and Craig and von Peter (2014) for an analysis of the German interbank network structure). Craig and Ma (2018) develop a structural model of the German interbank market to study the potential for systemic risk.⁸ A few papers empirically study contagion and systemic risk in financial networks (Furfine, 2003; Iyer and Peydro, 2011; Denbee et al., 2014; Anderson et al., 2019; Anderson et al., 2020). To the best of our knowledge, however, there is no prior empirical paper studying whether banks lend to counterparties with a similar loan portfolio.

A smaller literature focuses on socially efficient network structures and the endogenous formation of financial networks, in the presence of systemic risk (see, for example, Gofman, 2017). One major departure we make from this literature is by considering the joint choice of financial and real exposures. Our analysis is motivated by the stylized fact that banks are more likely to lend to other banks which have similar real exposures. Both financial and real exposure choices are required to explain this observation.⁹

⁸ Since we are interested in the creation of new links, we use the actual transfer of funds between two banks to denote interbank lending (i.e. a flow variable). In contrast, Craig and Ma (2018) are interested in credit risk and use the reported exposure from historical fund transfers (i.e. a stock variable) to denote interbank exposures.

⁹ Subsequent to us, Jackson and Pernoud (2019) also investigate this interaction. We compare their results to ours later in the paper.

Two papers close to ours, studying socially efficient network structures, are Cabrales et al. (2017) and Acemoglu et al. (2015a). Both examine social efficiency in the context of financial contagion. Acemoglu et al. (2015a) identify a key tradeoff facing the social planner. Denser connections prevent bankruptcies from small shocks, but facilitate the spread of contagion from large shocks. Therefore, highly connected networks are optimal if all shocks are relatively small. Conversely, if shocks are always large, then weak connections in the network can enhance stability by preventing the spread of contagion. Our paper builds on this analysis by considering socially efficient networks when shocks hitting the system can be *both* large or small with positive probability; and introduces an amplification mechanism via default costs.

Cabrales et al. (2017) study the same key tradeoff governing social efficiency as in our paper—how to limit contagion, while also allowing risk-sharing. They study considerably more general shock distributions than we do. Their first set of results can be viewed as a substantial and important generalization of the Acemoglu et al. (2015a) insight. They find that when shocks are sufficiently fat-tailed, the tradeoff is resolved by minimizing the chance of contagion and forming a maximally segmented network. When shocks are sufficiently thin-tailed, the tradeoff is resolved by maximizing risk-sharing in a fully integrated network structure. For intermediate shock distributions, their results are consistent with our findings. They find conditions under which the strength of connections in the network takes, at most, two values. We take this as suggestive that the forces we study apply more generally than our two point shock distributions. What we get from our more restrictive shock distribution, in combination with including costs of financial distress, is that the pattern in which these links are organized is crucial. We are able to characterize the socially optimal patterns. Links between different groups of highly connected banks must be weak to contain financial distress costs after one group fails.

A growing literature considers endogenous network formation.¹⁰ We discuss three papers closely related to ours. Each examines why privately and socially optimal behavior might differ. First, in Farboodi (2014), banks form links due to intermediation, since some banks have access to risky investment opportunities or funding opportunities, and others do not. Banks also intermediate to capture rents. Impressively, Farboodi (2014) succeeds in finding equilibrium networks, and the equilibrium structure she identifies matches empirically observed financial networks—there is a core of investment banks connected to each other, each with links to a set of commercial banks. In her model private behavior is socially inefficient because core banks behave in an excessively risky manner to capture intermediation rents. Second, Acemoglu et al. (2015b) also focuses on financial intermediation as the reason for network formation. Privately and socially efficient behavior may diverge because of a financial network externality—banks contract on a bilateral basis, and thus do not account for their role in creating a conduit that allows idiosyncratic shocks to develop into contagion. In our model, systemic risk also arises endogenously but from a different set of frictions, pertaining to risk shifting. The stylized fact we present, that banks correlate their real exposures with their largest counterparties is suggestive that this friction plays a role in practice, although other frictions may, of course, also matter. Third, Erol (2015) introduces government bailouts into a model of network formation among firms with bilateral exposures. Anticipating bailouts leads to “network hazard”—firms become less concerned about the choice of their counterparties’ own counterparties. In equilibrium, highly connected central

¹⁰ Papers considering related problems include Leitner (2005), Blume et al. (2011), Allen et al. (2012), Babus (2016), Zawadowski (2013), Eisert and Eufinger (2018), Erol and Vohra (2014), Di Maggio and Tahbaz-Salehi (2014), Wang (2014), Cohen-Cole et al. (2015), Galeotti et al. (2015), Cabrales et al. (2017), Erol and Ordoñez (2017), Galeotti et al. (2017), Kanik (2017), Bernard et al. (2017), Craig and Ma (2018), Chang and Zhang (2018) and Stanton et al. (2018).

firms emerge in the network. Unlike in our model, however, firms do not decide to correlate their real and financial exposures.

Outside the networks literature, the role of risk shifting in generating financial instability has been widely discussed in the context of a single-firm framework (e.g. Jensen and Meckling, 1976). Relatively little work considers risk shifting in a systemic setting. A first exception is Acharya (2009). We build on Acharya (2009) by combining portfolio choice of real investments with network choices of financial connections. This raises new questions and allows us to provide suggestive evidence consistent with risk shifting through portfolio choices in a network setting. A second exception is Farhi and Tirole (2012). In their model, banks face liquidity risk. Authorities can intervene to improve liquidity. They uncover a strategic complementarity. Each bank increases liquidity risk if other banks also have greater liquidity risk. Then, banks are likely to be illiquid at the same time, which prompts an intervention from the authorities. Our model also identifies a strategic complementarity, but the source is different: arising from limited liability instead of intervention by the authorities.

We structure the paper as follows. Section 2 presents our model. Section 3 examines socially efficient networks. Section 4 discusses equilibrium networks. Section 5 provides several robustness checks for the stylized fact that banks have similar real and financial exposure. Proofs are relegated to Appendix B.

2. The model

We introduce a model to rationalize the stylized fact we presented in the introduction. The model has three key ingredients. First, banks have financial claims on one another. Second, banks make real investments. Third, there is limited liability. With these ingredients, we ask whether it is socially efficient for banks to choose the same real exposures as their financial counterparties—as in the stylized fact—and whether socially efficient outcomes will arise in equilibrium.

2.1. Banks, investments, and financial claims

Banks. There is a set $N = \{1, \dots, n\}$ of financial institutions, which we refer to as banks. The economy lasts for three periods, $t = 0, 1, 2$. In period 0, each bank issues debt with face value \underline{v} to external debtors, that is, entities outside the financial system.

Investment Portfolio. In period 1, each bank chooses a portfolio of investments which yields a stochastic return. The return generated by bank i 's portfolio is p_i . There is one unit available for each of n types of investment. We let $\phi_i \in \Delta^n$ denote bank i 's portfolio, where Δ^n is the n -dimensional simplex, and use $\phi_{ik} \in [0, 1]$ to denote the share of i 's portfolio in investment k . Portfolios $\phi := \{\phi_i\}_i$ satisfies $\sum_i \phi_{ik} = 1$, so the return on all investments goes to some bank. We denote the set of possible portfolios by Φ .

Each investment type k generates a random return R_k . There is a portfolio maintenance cost $c > 0$ for each type of investment bank i makes. Thus, the return on i 's portfolio is

$$p_i = \sum_{k=1}^n (\phi_{ik} R_k - c I_{\phi_{ik} > 0}),$$

where $I_{\phi_{ik} > 0}$ is an indicator variable taking the value 1 if $\phi_{ik} > 0$ and 0 otherwise. The portfolio maintenance cost c captures monitoring or research costs related to each investment type, or the diluting of special expertise when banks hold more diversified portfolios.

Financial Network. Also in period 1, banks form financial claims on each others' investment portfolios. We denote these claims by a matrix \mathbf{A} where $A_{ij} \geq 0$ is the claim of bank i on the return generated by a counterparty bank j 's investments. Financial claims satisfy $\sum_i A_{ij} = 1$, so the returns generated by each portfolio go to some bank. Each bank must have weakly stronger claims on its own portfolio than any other banks, so $A_{ii} \geq A_{ji}$ for all i and all j . We can represent \mathbf{A} as a weighted, directed graph, where the banks are the nodes and the links are financial claims between banks. Thus, \mathbf{A} is a financial network.

In period 2, returns from the investments realize. The total investment returns flowing to bank i is its market value, v_i . Thus, the financial network \mathbf{A} and investment portfolios Φ chosen in period 1 jointly determine banks' realized market values in period 2.

Default. Banks can default. If $v_i < \underline{v}$ in period 2, then bank i is unable to repay its debt holders in full, which causes i to default. Then default costs lower the returns of i 's portfolio by an amount β . Default costs include costs associated with liquidating investments, possibly at a discount during financial turmoil; inefficient allocation of resources during default, and so forth.¹¹

Market Values. The market value of bank i in period 2 is

$$v_i = \sum_{j=1}^n A_{ij}(p_j - \beta I_{v_j < \underline{v}}). \tag{1}$$

Bank i 's market value depends on its financial claims, A_{ij} , and on counterparties' portfolio returns, p_j . If j defaults, then i suffers a fall in its market value proportionate to its claim on j 's portfolio. Letting emboldened variables represent vectors, we can rewrite equation (1) as

$$\mathbf{v} = \mathbf{A}(\mathbf{p} - \mathbf{b}(\mathbf{v})) \tag{2}$$

where $b(v_i) = \beta I_{v_j < \underline{v}}$. We show in the Online Appendix, Section OA1, that there exists a vector of market values \mathbf{v} which solves equation (2), and the set of solutions forms a complete lattice. The partial order is set inclusion across the set of banks that fail.¹² Thus, there is always a solution to equation (2) in which a minimal set of banks fail, such that in all other solutions a superset of these banks fail. Throughout, we focus on the solution in which this minimal set fails.

Risk Sharing. Banks benefit from risk sharing. Through risk-sharing they can avoid default, which would otherwise lower their market values. Banks share risks in two ways. First, banks hold financial claims on each others' portfolios, diversifying their real exposures. However, financial claims expose banks to counterparties' default costs. Second, banks can diversify their own portfolios directly. Again, this can help banks to avoid failure in some circumstances. However, more diversified portfolios incur larger portfolio maintenance costs. For a given bank, either diversified investments or financial claims may be more valuable. The size of default costs versus portfolio maintenance costs matters. The existing network of financial claims, and other banks' investment portfolios also matter.¹³

Systemic Risk. Financial claims and real investments also create systemic risk: several banks can fail at the same time. There are two mechanisms leading to systemic risk. First, banks might

¹¹ See Cifuentes et al. (2005), Gai and Kapadia (2010) or Caballero and Simsek (2013) for a more detailed treatment of fire sales in financial networks.

¹² Eisenberg and Noe (2001), Rogers and Veraart (2013), Elliott et al. (2014) and Acemoglu et al. (2015a) derive similar results.

¹³ Another important reason for financial connections, distinct from risk sharing, is intermediation (e.g. Farboodi, 2014; Acemoglu et al., 2015b).

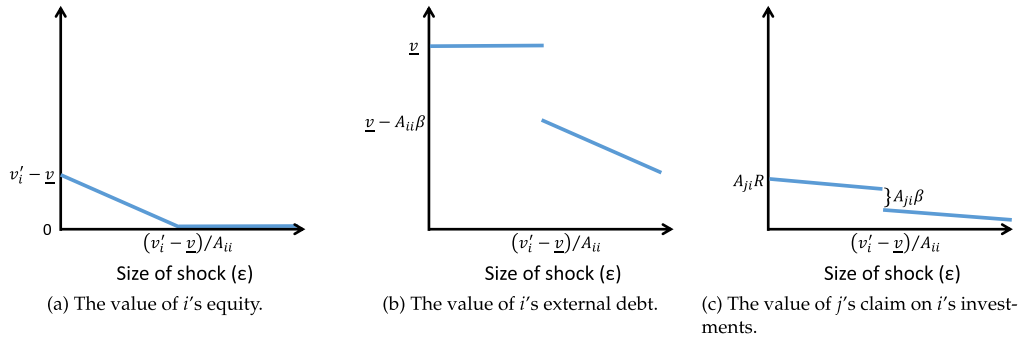


Fig. 2. Example—Equity and External Debt Value with Limited Liability. *Notes:* Let v_i' be the value of i prior to a shock ε to i 's investment. For shocks $\varepsilon < (v_i' - \underline{v})/A_{ii}$ bank i does not default and the losses due to the shock are spread among banks in proportion to their claims on the investment. The shock reduces the equity value of i (Panel (a)) and the value of a bank j with claims $A_{ji} < A_{ii}$ (Panel (c)). For larger shocks i fails, resulting in financial distress costs of β that further reduce the cashflows of i 's investment. These losses reduce both the value of external debt (Panel (b)) and the value of bank j 's claim on i 's investment (Panel (c)).

have the same real exposures. Several banks might hold the same investment. If that investment has a low return, several banks can fail. Second, banks might be exposed to one another through the financial system. Suppose a bank has financial claims on a counterparty. If the counterparty's portfolio has a low return, then its market value falls and it may default. If the counterparty defaults, then the bank's market value falls even further. Thus, default costs amplify real shocks across the financial network.

2.2. Equity value and limited liability

Let π_i be i 's equity value and δ_i be i 's value of external debt in period 2. There is limited liability, so

$$\pi_i = \max\{v_i - \underline{v}, 0\} \tag{3}$$

$$\delta_i = \min\{\underline{v}, v_i\}. \tag{4}$$

When $v_i \geq \underline{v}$, bank i does not default. External debtors receive the face value of their debt, \underline{v} . Bank i 's residual market value goes to equity holders. When $v_i < \underline{v}$, bank i defaults. Equity holders receive nothing and external debtors receive the market value v_i , which is less than the face value of their debt.

Example. A simple thought experiment shows how limited liability affects the value of contracts. Suppose that the return on i 's portfolio falls by some shock of size ε , while all other investment returns remain constant. After this shock no banks default. The market value of bank j , v_j , falls as j 's claims on i are now worth less—specifically, the value of j 's claims on i 's portfolio falls in value by $A_{ji}\varepsilon$. The equity value of bank i , π_i , also falls by $A_{ii}\varepsilon$. Since, by assumption i does not fail, the value of bank i 's external debt remains a constant \underline{v} .

Now suppose that ε is sufficiently large that after the shock, $v_i < \underline{v}$ and i defaults, but small enough that no other bank defaults. As bank i defaults, it incurs default costs that further decrease the return of its portfolio. Bank i 's equity holders receive nothing after i fails, due to limited liability. Bank j 's claim on i 's portfolio now falls by $A_{ji}(\varepsilon + \beta)$, due to the added impact of

counterparty default costs. The value of i 's external debt falls by $A_{ii}(\varepsilon + \beta)$, less i 's equity value absent the shock.

Finally, suppose that after the shock both i and j default, but no other banks default. In this case, j 's claim on i continues to fall by $A_{ji}(\varepsilon + \beta)$. However, i 's external debtors are exposed to j 's default costs. The debt value of i now falls by $A_{ii}(\varepsilon + \beta) + A_{ij}\beta$, less i 's equity value prior to the shock. The key features of these contracts are summarized in Fig. 2.¹⁴

2.3. Distribution of investment returns

We assume a simple process for the distribution of returns on investments. With probability r , a shock of value ε lowers the return on a single type of investment. The investment is selected uniformly at random, independent of banks' portfolios or financial claims. All investment types unaffected by the shock generate a return \bar{R} . If no shock arrives, then all investments return \bar{R} . To simplify algebra we also define $R = \bar{R} - c$.

Conditional on its occurrence, the shock is large with probability q and small with probability $(1 - q)$, that is:

$$\varepsilon = \begin{cases} \varepsilon_L > n(R - \underline{v}) & \text{with probability } q \\ \varepsilon_S \in (R - \underline{v}, n(R - \underline{v})] & \text{with probability } (1 - q). \end{cases} \tag{5}$$

We also assume that $\varepsilon_L \geq 2\varepsilon_S$. Further, small shocks are relatively common compared to large shocks. Specifically, we assume $0 < q < \frac{1}{n^2}$.

Discussion of shock process. A mix of small and large shocks creates a tradeoff between two priorities: first, risk sharing; and second, avoiding systemic risk. First, banks can share the risks associated with small shocks. These shocks are larger than the equity value of a single bank. Without financial claims or multiple investments, a bank will fail after a small shock hits. But these shocks are small enough that, under some networks and portfolios, banks can share risks and prevent any defaults. Second, large shocks create systemic risk. These shocks are large enough that at least one bank must default for any network and portfolio, and financial connections can serve to propagate defaults. Indeed, if large shocks are common, that is, q is large, then the benefits of risk sharing will be small and an unconnected financial system will be efficient. We assume large shocks are rare to create a tradeoff.¹⁵

2.4. Discussion of model assumptions

We close the section by discussing some features of our model in more detail.

Model of the Financial Network. We explore a tradeoff between risk sharing and minimizing systemic risk. Our model of financial claims is a simple way of capturing this tradeoff.¹⁶ In practice, many contracts feature this tradeoff. A real-world example, similar to our model, is an unsecured, bilaterally traded, over-the-counter swap. We briefly discuss two alternative models that do not permit such a tradeoff. First, a richer contracting space could eliminate the tradeoff.

¹⁴ Our model entertains two possibilities: (i) financial claims are written on the return generated by an investment after default costs have been subtracted, in which case financial claims are more senior than external debt; or (ii) financial claims are written on the return generated by investments, in which case financial claims are as senior as external debt.

¹⁵ In the Online Appendix, Section OA2 we generalize the shock distribution in several respects.

¹⁶ See Cabrales et al. (2017), Wang (2014) and Cabrales et al. (2015) for models of the financial network that also capture this tradeoff.

For example, suppose banks i and j write a contract that induces a cash flow from i to j if a small shock, but not a large shock, hits j 's portfolio. This contract can prevent j 's default after a small shock, but leaves i unaffected by the large shock. Thus, there is no tradeoff. Second, standard debt contracts, as modeled in the financial networks literature, do not allow risk sharing to prevent defaults. Their value does not fall until a bank defaults, and so no risks can be shared among banks prior to defaults (see, for example, Eisenberg and Noe, 2001). Of course, in practice, the value of debt contracts falls as banks' default risk rises. Therefore, in a distressed financial system, debt contracts may behave similarly to the contracts in our model.

External Debt is Fixed. For simplicity, we do not model banks' choice of external debt in period 0. Instead, the face value of external debt is fixed, and not contingent on the financial network or on portfolio choices. One could imagine, equivalently, that external debtors cannot observe the financial network when they write contracts.¹⁷ In equilibrium, external debtors should correctly anticipate the structure of the financial network. They should demand returns that compensate them for any risk associated with the equilibrium financial network and portfolios, and \underline{v} will be set accordingly. By taking the external debt of banks as given when banks form the financial network and choose portfolios, we do not model this step explicitly. However, our main results are all consistent with such behavior. We return to this issue at greater length in Section 4, when we discuss equilibrium behavior.

Relation of Model to Stylized Fact. Recall the stylized fact from the introduction: German banks with similar real investments have larger financial claims on each other. Our model lets banks choose both real investments and financial claims. Therefore, we can investigate what mechanisms might reproduce the stylized fact. We can ask whether the pattern is socially efficient, or whether the pattern generates systemic risk.

Our model of financial claims is consistent with German interbank lending, at least in a stylized sense. In the German data, we do not observe the exact contracts used by banks. We cannot see whether the contract is collateralized, a cashflow swap, a standard debt contract, or some more complicated instrument. Instead, we observe quarterly net transfers between banks. Given this ambiguity, we model a general contract, which summarizes banks' incentives to share risks.

3. Social efficiency

This section characterizes socially efficient networks and portfolios. Banks hold different portfolios of investments from their close financial counterparties. This finding contrasts with our stylized fact: in German data, banks hold similar investments to their close financial counterparties.

3.1. Social planner's problem

We start by setting up the social planner's problem.

Social planner. The social planner chooses the financial network and the set of investment portfolios at time $t = 1$, before investment returns are realized. The social planner maximizes the sum of expected equity and external debt value. The social planner therefore has a utilitarian objective function. However, as all agents in our model are risk-neutral and utility is transferable, an outcome maximizes this objective if and only if it is Pareto efficient.

¹⁷ Caballero and Simsek (2013) point out that inter-bank financial networks are highly opaque, especially to outsiders.

Participation constraint. We impose a participation constraint on the social planner. The planner must choose a network and a portfolio that does not decrease the expected market value of any bank below its autarky market value—that is, the expected market value if the bank has no financial claims and holds only a single asset. Since in autarky bank i fails following a large or small shock to its investment, the participation constraint implies

$$\mathbb{E}[v_i(\mathbf{A}, \phi)] \geq R - \frac{r}{n}(q\varepsilon_L + (1 - q)\varepsilon_S + \beta), \tag{6}$$

for all i . We denote the set of column stochastic and non-negative matrices satisfying these participation constraints by \mathcal{A} .

The participation constraint requires the social planner to give each bank approximately the same expected market value. We view the participation constraint as a minimal restriction on the planner’s choices that is also normatively appealing.¹⁸

We think it is unrealistic for the planner to allocate all the assets to a single bank, and without this restriction this provides a way in which the planner can avoid all but one of the banks from ever failing

Social Planner’s Problem. The social planner chooses a financial network $\mathbf{A} \in \mathcal{A}$ and set of portfolios $\phi \in \Phi$. Thus, the planner solves the following optimization problem:

$$\max_{\mathbf{A} \in \mathcal{A}, \phi \in \Phi} \mathbb{E} \left[\sum_{i \in N} \pi_i(\mathbf{A}, \phi) + \delta_i(\mathbf{A}, \phi) \right]. \tag{7}$$

We can simplify the social planner’s problem as follows:

Remark 1. The social planner’s problem is equivalent to minimizing the expected cost of defaults plus total portfolio maintenance costs, subject to the same constraints.

3.2. Social planner’s solution

To characterize the social planner’s solution, we will define a particular financial network and set of investment portfolios.

Clustered networks. First we define d^* to be the unique positive root of

$$d_i^2(R - \underline{v})\beta + d_i((R - \underline{v})(\varepsilon_L - \varepsilon_S) - \varepsilon_S\beta) + \varepsilon_S(n(R - \underline{v}) - \varepsilon_L) = 0. \tag{8}$$

For the rest of this section we abstract from the integer problems by assuming that

- (i) d^* is an integer; and
- (ii) n/d^* is an integer.

¹⁸ We intend the planner’s problem to provide a benchmark for what might be achievable via financial regulations. Without any participation constraint the planner would be incentivized to avoid default costs by allocating all claims on assets to a single bank. We view this as an unrealistic outcome for policy to obtain. While this motivates the inclusion of a participation constraint, an alternative would be to impose it on equity values instead of market values. That has some appeal given that we later consider banks maximizing their equity values as well as their market values. However, even when banks make individual choices to maximize shareholder value we expect debt holders to have a voice when financial regulations are set. By using market values for the participation constraint we capture this.




Fig. 3. A network in the class $\mathcal{A}^*(6)$, with 4 groups and 6 banks in each group.

Under the integer assumptions, we can partition the banks into n/d^* groups of d^* banks. We refer to each group as a *cluster*. Now we define the class of networks $\mathcal{A}^*(d^*)$. Letting G_i be the cluster that i belongs to

$$\mathcal{A}^*(d^*) := \left\{ \mathbf{A} \in \mathcal{A} : \begin{array}{ll} |G_i| = d^* & \text{for all } i \in N \\ A_{ji}^* = A_S := \frac{R-\underline{v}}{\varepsilon_S} & \text{for all } i, j : G_i = G_j \\ A_{ji}^* = A_W := \frac{R-\underline{v}}{\varepsilon_L + \beta d^*} & \text{for all } i, j : G_i \neq G_j \end{array} \right\}.$$

The class of networks $\mathcal{A}^*(d^*)$ is symmetric networks where organizations are partitioned into clusters of size d^* , with strong financial claims within clusters, and weaker financial claims between clusters. We say that a network $\mathbf{A} \in \mathcal{A}^*(d^*)$ is d^* -clustered. The network representation of claims for an $\mathbf{A} \in \mathcal{A}^*(6)$ is shown in Fig. 3.

Non-overlapping portfolios. We say that banks have *non-overlapping portfolios* if each bank’s portfolio consists of a single type of investment, and all banks make different types of investments.

Suppose that banks hold non-overlapping portfolios and the financial network is d^* clustered. Then, by construction, when a small shock hits any bank i ’s investment, $A_{ji}^* \varepsilon_S \leq \sum_{k \in N} A_{jk} R - \underline{v}$ so that no banks default. When a large shock hits a bank i ’s investment, banks in G_i default but banks outside G_i do not default. Thus, d^* banks fail following a large shock to i ’s investment and no banks fail following a small shock to any investment.

We can now state our first key result.

Proposition 1. *Under the maintained assumptions, namely integer conditions (i) and (ii) above and $q < \frac{1}{n^2}$, there exists an $\bar{r} > 0$ such that for all $r < \bar{r}$ a network \mathbf{A} and portfolios ϕ solve the social planner’s problem if $\mathbf{A} \in \mathcal{A}^*$ and portfolios are non-overlapping. Further, as $r \rightarrow 0$, a network \mathbf{A} and portfolios ϕ solve the social planner’s problem only if $\mathbf{A} \in \mathcal{A}^*$ and portfolios are non-overlapping.*

3.3. Intuition and proof outline

Intuitively, the networks and portfolios characterized in Proposition 1 are the optimal tradeoff between the two competing imperatives in our model. First, the social planner shares risks when

small shocks arrive, by preventing any defaults. Second, the social planner minimizes systemic risk when a large shock arrives.

There are several steps to proving the Proposition. First, we show that for r sufficiently small the planner will always choose to share risks through financial claims, instead of having banks hold portfolios of multiple investments. This simplifies the problem, allowing us to restrict attention to non-overlapping portfolios.

Second, we show that perfect risk sharing through financial claims, such that $A_{ij} = 1/n$ for all i, j , has fewer expected defaults than any network in which a small shock leads to at least one failure. With perfect risk sharing, banks only fail after large shocks. Large shocks are sufficiently rare that the social planner prefers to avoid a single failure when a small shock arrives than multiple failures when a large shock arrives.

Third, we simplify the problem using the participation constraints. Intuitively, as shocks become rare, the additional value of risk sharing is small. Consequently, the planner must give each bank i overall claims on projects that sum to approximately 1 (i.e., $\sum_j \sum_k A_{ij} \phi_{jk} \approx 1$) to satisfy the participation constraint. Thus, the set \mathcal{A} must be in the neighborhood of the set of doubly stochastic and non-negative matrices (i.e., networks A satisfying $\sum_i A_{ij} = \sum_i A_{ji} = 1$ and $A_{ij} \geq 0$).

The first three steps simplify the social planner’s problem. The problem reduces to choosing a non-negative and doubly stochastic network that minimizes the expected number of failures from large shocks, conditional on preventing failures from small shocks. While this problem is more tractable than our starting point, the space of possible networks remains large.

In the fourth step of the proof, we make progress by considering a simpler relaxed problem. The solution to this problem provides an upper bound on the planner’s objective, or, equivalently, a lower bound on the expected number of failures. We then show that (i) this bound is achieved by all networks in the class $\mathcal{A}^*(d^*)$ along with non-overlapping portfolios, and (ii) no other network portfolio pair achieves this bound.

In this simpler problem, we give up on simultaneously trying to minimize the number of failures following a large shock to any bank. Instead, we minimize the number of failures following a large shock to a given bank i , while making sure there is no failure after a small shock. We show that for all doubly stochastic network structures \mathbf{A} such that no small shock causes at least one bank to fail, at least d^* banks fail following a large shock to i .

We formalize the solution to the simpler problem in Lemma 2 in the Appendix. The proof of Lemma 2 works by combining inequalities. Let D_i be the set of banks that fail when a large shock hits i . Collectively, banks outside D_i cannot have overly large claims on banks in D_i , otherwise they would fail following a large shock. Banks within D_i can have stronger claims, but not so strong that they fail when a single small shock hits a bank in D_i . Minimizing the size of $|D_i|$ subject to these constraints implies minimizing d_i subject to the inequality

$$d_i^2(R - \underline{v})\beta + d_i((R - \underline{v})(\varepsilon_L - \varepsilon_S) - \varepsilon_S\beta) + \varepsilon_S(n(R - \underline{v}) - \varepsilon_L) \geq 0.$$

This implies that the number of failures following a large shock to i must be at least d^* , where d^* is the unique positive root of the left-hand side of the above inequality.

The more general problem of choosing a network structure that simultaneously minimizes the number of failures from a large shock to any investment is harder. One might suspect that the social planner faces a tradeoff, between minimizing the number of failures after a large shock to one investment, and minimizing the number of failures following a large shock to some other investment.

But in the fifth step, we show there is no tradeoff. We show any network $\mathbf{A} \in \mathcal{A}^*(d^*)$ with non-overlapping portfolios achieves d^* failures when a large shock hits *any* project, and no failures otherwise. As networks $\mathbf{A} \in \mathcal{A}^*(d^*)$ achieve the upper bound on the planner’s objective identified in Lemma 2, we have established the *if* part of Proposition 1.

In the sixth and final step, we establish the *only if* part of Proposition 1. We show that only networks $\mathbf{A} \in \mathcal{A}^*(d^*)$ are socially efficient. Obtaining the upper bound in Lemma 2 requires several inequalities to bind, which pins down the strength of each bank’s financial claims. Specifically, it can be shown that it is necessary but not sufficient for bank i to have d^* financial claims of size $(R - \underline{v})/\varepsilon_S$ and $n - d^*$ financial claims of size $(R - \underline{v})/(\varepsilon_L + d^*\beta)$. It turns out that the *only* way to wire the network so that each bank has financial claims with these values and the bound in Lemma 2 is achieved, is to arrange banks in a d^* -clustered network with non-overlapping portfolios.

3.4. Discussion

Proposition 1 is our first key finding. The social planner optimally chooses a network that is d^* -clustered with non-overlapping portfolios. We characterize socially efficient networks despite considering a large space of possible network structures and investment portfolios. We place minimal restrictions on the structure of the financial network. For example, some banks could have a few large financial counterparties, while other banks might have many small financial counterparties. Moreover, each bank can hold any type of investment.

Absence of homophily. The social planner chooses networks and portfolios without *homophily*. In a homophilous network, banks with similar real exposures would have similar financial exposures. By contrast, in the socially efficient network all banks are connected to one another, albeit with financial claims of varying strength. But there is no homophily: the social planner chooses non-overlapping portfolios. Further, while the choice of non-overlapping portfolios is intuitive, it is not obvious. For example, Cabrales et al. (2017) find that when banks face uncorrelated real exposures, but are heterogenous in the marginal shock distributions they face, homophily in the financial network is efficient—it is efficient for banks’ counterparties to face similar shock distributions.

Relation to stylized fact. In the introduction we present a new stylized fact from the German interbank system, alongside anecdotal evidence from the 2008 financial crisis. Real-world financial systems *do* exhibit homophily. Banks with similar real exposures have similar financial exposures, different from the socially efficient outcome in our model. We will study banks’ incentives to deviate from socially efficient networks, in Section 4.

Default costs and weak links between clusters. Proposition 1 also allows a better understanding of the role of default costs in amplifying systemic risk—and how the socially efficient network prevents such amplification from occurring. The relatively weak financial claims between clusters mitigate the impact of default costs. These claims are weak but positively valued, since for any $\mathbf{A}^* \in \mathcal{A}^*$ we have

$$0 < A_{ij}^* = A_W < A_S = A_{jk}^* \quad \text{for } G_i \neq G_j = G_k.$$

Weak links allow the social planner to prevent failures outside the cluster hit by the shock. Banks which do not initially fail after a large shock to another bank might subsequently fail—exposure to default costs could cause a second round of defaults. However, in the socially efficient network, the cumulative impact of the large shock and default costs does not cause failures beyond the cluster hit by the shock. Instead, the claims between clusters are too weak to transmit the

shock outside the cluster. Still, the weak claims enable risk-sharing between clusters after a small shock. Consequently fewer banks within a cluster are required in order to absorb the impact of a small shock without default. Therefore weak claims between clusters allow an optimal tradeoff between risk sharing and minimizing systemic risk.

It is suboptimal for clusters to be fully segmented, such that no bank holds claims on banks in another cluster. While full segmentation would prevent any part of the large shock from being transmitted outside the cluster, it would also prevent sharing of the small shock between clusters. To prevent failures following small shocks, clusters would then have to be larger, and more banks would fail following a large shock.¹⁹ Fully segmenting the network and allowing no claims between clusters, increases the expected number of failures.²⁰

Comparison with the literature. We now place our social planner's solution in the context of the literature. Our formal analysis assumes there can be a large or small shock with some probability. Then, neither the empty nor the complete network is optimal.

Suppose instead there is a single shock. If the shock is small, then a complete financial network in which $A_{ij} = 1/n$ is socially efficient. This network maximizes risk sharing. If the shock is large, then an empty network in which $A_{ii} = 1$ and $A_{ij} = 0$ for $i \neq j$. This network minimizes systemic risk. As the shock increases from being small to large in size according to our definition, the efficient network changes discontinuously. This discussion echoes results in Acemoglu et al. (2015a), and their argument that some financial networks may be 'robust yet fragile'.

Suppose instead there is a continuous distribution of shocks. Two important results from Cabrales et al. (2017) are that when the shock distribution is sufficiently fat-tailed, maximal segmentation of banks is optimal, while the maximal integration of banks is optimal when the distribution is sufficiently thin-tailed. Intuitively, in the fat-tailed case, large shocks are sufficiently common that the tradeoff between avoiding failures from small shocks and preventing systemic risk from large shocks is resolved by the corner solution that minimizes systemic risk. The opposite is true in the thin-tailed case. This provides an impressive generalization of the insight from Acemoglu et al. (2015a).

The distribution of shocks we study leads to an intermediate case—where it is not optimal to focus exclusively on either risk sharing or minimizing systemic risk. Cabrales et al. (2017) also consider this intermediate case. However, the absence of default costs from their model prevents the pattern of connections in socially optimal networks from being pinned down, although they are able to conclude that optimal networks feature links of only two different strengths. In this intermediate case, by modeling financial distress costs, we find that socially optimal networks are d^* -clustered. This yields the insight that weak, but non-zero, claims between clusters are optimal.

Extensions. Proposition 1 relies on some strong assumptions. In the Online Appendix, Section OA2, we show some ways in which the insight from Proposition 1 that weak links between clusters are optimal generalizes.

¹⁹ In different settings, Blume et al. (2011) and Erol and Vohra (2014) find that socially efficient networks are fully segmented.

²⁰ In the Online Appendix, Section OA3 we present some comparative static results on the size of clusters, number of clusters and strength of links between clusters.

4. Stability of the social planner’s solution

Let us take stock: in socially efficient financial networks and portfolios, banks hold different investments from their financial counterparties. Further, we will see in this section that the social planner’s solution is a stable outcome in a decentralized financial system where each bank maximizes its expected market value. But in German data, banks tend to lend to the same firms and industries as their largest financial counterparties.

The analysis in the later half of this section achieves a reconciliation. In the presence of limited liability, banks seek to hold investments similar to those of their close financial counterparties, though this pattern is not socially efficient and generates systemic risk. This behavior is *systemic risk shifting*.²¹

4.1. Bank choices

We allow each bank to choose its portfolio $\phi_i \in \Delta^n$. We also suppose banks can engage in bilateral trades. In these trades we permit the banks to exchange their financial claims on each others’ portfolios.

Definition 1. Claims $A' \in \mathcal{A}$ can be reached from A through a *feasible bilateral trade* between i and j if for all banks $k \in N$:

- (i) $A'_{ik} + A'_{jk} = A_{ik} + A_{jk}$ for all k
- (ii) $A_{kl} = A'_{kl}$ for all $k \neq i, j$ and all l .

Bilateral trading is intended to capture decentralized trade in inter-bank markets.^{22,23} We assume that each bank’s level of external debt, \underline{v} , holds constant when banks change either portfolios or financial claims.²⁴

4.2. If banks maximize market value, the social planner’s solution is stable

We begin by assuming that banks seek to maximize their expected market value, summing both debt and equity value. In this case, we will see that the social planner’s solution is stable. Thus, the difference between market and equity value, due to limited liability, is crucial.

A network portfolio pair is stable if two conditions are met. First, no bank can unilaterally change its portfolio to strictly increase its expected market value. Second, no bank can change its portfolio, and simultaneously agree a bilateral trade with another bank to change its financial claims, so that these changes strictly increase its expected market value and do not lower the other bank’s expected market value.

²¹ We also show in the Online Appendix, Section OA4 that there is a formal sense in which the planner’s solution favors debt holders over equity holders. This creates a basic tension between the interests of the equity holders and efficient outcomes. Intuitively, debt holders’ interests are well aligned with the planner’s interests because debt holders like to avoid failures.

²² For related studies of network formation financial markets, see Farboodi (2014); Erol and Vohra (2014); Erol (2015); or Acemoglu et al. (2015b).

²³ A large share of inter-bank trading is bilateral or over-the-counter, as opposed to centralized or exchange-traded (BIS, 2015).

²⁴ Subsection 4.3 discusses the robustness of our conclusions to this assumption.

Definition 2. A network portfolio pair (\mathbf{A}, ϕ) is *stable* if and only if

- (i) for all $\phi' = (\phi'_i, \phi_{-i})$ such that $\phi'_i \in \Delta^n$, $\mathbb{E}[v_i(\mathbf{A}, \phi')] \leq \mathbb{E}[v_i(\mathbf{A}, \phi)]$; and
- (ii) for all $\phi' = (\phi'_i, \phi_{-i})$ such that $\phi'_i \in \Delta^n$ and for every pair of banks $i, j \in N$, if for all feasible bilateral trades between i and j yielding claims \mathbf{A}' either

$$\mathbb{E}[v_i(\mathbf{A}', \phi')] \leq \mathbb{E}[v_i(\mathbf{A}, \phi)] \quad \text{or} \quad \mathbb{E}[v_j(\mathbf{A}', \phi')] < \mathbb{E}[v_j(\mathbf{A}, \phi)].$$

In our stability definition, banks unilaterally change portfolios but bilaterally adjust their financial claims. Thus, a network portfolio pair is unstable if a single bank wants to change its portfolio holding, keeping fixed the portfolios of the other banks. We interpret this as an implicit market clearing condition—if there is positive excess demand for an investment, then the network portfolio pair is unstable. While we think it is natural to model portfolio choices as being unilateral, our financial claims are inherently bilateral and require a contract to be written between two banks. Hence, we require a pair of banks to find a mutually valuable change in financial claims for the network portfolio pair to be unstable. In effect, the market for investments is “competitive”, whereas the market for financial claims is “game theoretic”.

Under our definition of stability, banks’ incentives are not generally aligned with the social planner. Bank i wants to implement a change in ϕ_i that increases its market value $\mathbb{E}[v_i(\mathbf{A}, \phi)]$, even if other banks’ expected market values decline by more than i ’s market value increases. Similarly, banks i and j will engage in trades that raise their joint market value, even if these trades cause other banks’ expected market values to decline by more than i and j gain. Nevertheless, the social planner’s solution is stable.

Proposition 2. *There exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ all socially efficient network portfolio pairs (A^*, ϕ^*) are stable.*

The potential for systemic risk creates strong externalities that are not being internalized. Still, there are no deviations from the planner’s solution that raise market value. The reason why the social planner’s solution remains stable in the presence of these externalities is that the connectedness of the socially efficient network helps align incentives. Changes that impose losses on third parties cause new failures. Through the connectedness of the socially optimal network, these losses affect the banks which deviate from the socially efficient network portfolio pair. Thus, despite the rich set of available deviations in either portfolios and/or financial claims, there is stability.

Example. Consider a socially efficient network portfolio pair (A^*, ϕ^*) , and label investments so that $\phi_{ii}^* = 1$ for all i . Suppose now that bank i decides to change its portfolio. Suppose it still makes investments only of one type, but now instead of choosing investments of a different type from all other banks, i chooses the same investment type as bank j . Suppose further that j is in the same cluster as i ($j \in D_i(A^*)$). Such a change lowers the market value of other banks. Now when the small shock hits the investment type j both i and j suffer direct losses, causing all the banks in their cluster to fail. The expected losses of a bank k from this event is

$$\underbrace{\left(\frac{r(1-q)}{n} \right)}_{\text{Prob. small shock hits investment } j} \left(\underbrace{A_{ki}^* \varepsilon_S}_{\text{Small shock to investment } i} + \underbrace{A_{kj}^* \varepsilon_S}_{\text{Small shock to investment } j} \right)$$

$$+ \underbrace{\left(\sum_{l \in D_i} A_{kl}^* \beta \right)}_{\text{Additional defaults}} \Big). \tag{9}$$

Before i 's deviation, the expected loss to k from the small shock hitting bank i 's investment or bank j 's investment was a smaller amount:

$$\begin{aligned} & \underbrace{\left(\frac{r(1-q)}{n} \right)}_{\text{Prob. small shock hits investment } i} \underbrace{\left(A_{ki}^* \varepsilon_S \right)}_{\text{Small shock to investment } i} \\ + & \underbrace{\left(\frac{r(1-q)}{n} \right)}_{\text{Prob. small shock hits investment } j} \underbrace{\left(A_{kj}^* \varepsilon_S \right)}_{\text{Small shock to investment } j} . \end{aligned} \tag{10}$$

The change in i 's portfolio choice causes failures that would otherwise not have occurred, and these losses are passed around the banks in the system. Comparing equations (10) and (9) shows how these additional losses are absorbed by the different banks. Crucially though, bank i accrues additional losses of

$$\left(\frac{r(1-q)}{n} \right) \left(\sum_{l \in D_i} A_{il}^* \beta \right) > 0.$$

As such, even though the change in portfolio choice by i imposes large negative externalities on other banks that i does not internalize, the deviation also reduces i 's expected value making the deviation unprofitable.

4.3. If banks maximize equity value, the social planner's solution is not stable

By contrast, when banks maximize expected equity value, the social planner's solution is no longer stable. Limited liability creates a wedge between equity and market value. In a precise sense, limited liability leads equilibrium behavior to differ from the socially efficient outcome.

Suppose that banks maximize expected equity value, that is

$$\mathbb{E}[\pi_i(\mathbf{A}, \phi)] = \mathbb{E} \left[\max \left\{ v_i(\mathbf{A}, \phi) - \underline{v}, 0 \right\} \right]. \tag{11}$$

We adjust the definition of stability accordingly.

Definition 3. A network portfolio pair (\mathbf{A}, ϕ) is *stable under limited liability* if and only if

- (i) for all $\phi' = (\phi'_i, \phi_{-i})$ such that $\phi'_i \in \Delta^n$, $\mathbb{E}[\pi_i(\mathbf{A}, \phi')] \leq \mathbb{E}[\pi_i(\mathbf{A}, \phi)]$; and
- (ii) for all $\phi' = (\phi'_i, \phi_{-i})$ such that $\phi'_i \in \Delta^n$, and for every pair of banks $i, j \in N$, if for all feasible bilateral trades between i and j yielding claims \mathbf{A}' either

$$\mathbb{E}[\pi_i(\mathbf{A}', \phi')] \leq \mathbb{E}[\pi_i(\mathbf{A}, \phi)] \quad \text{or} \quad \mathbb{E}[\pi_j(\mathbf{A}', \phi')] < \mathbb{E}[\pi_j(\mathbf{A}, \phi)].$$

Proposition 3. For any socially efficient network portfolio pair (A^*, ϕ^*) with $d^* \geq 2$, (A^*, ϕ^*) is not stable under limited liability. Further,

- (i) each bank i has a unilateral deviation from (A^*, ϕ^*) to (A^*, ϕ') for $\phi' = (\phi_{-i}^*, \phi'_i)$ that raises its expected equity value.
- (ii) for each bank i there exists a bank j such that i and j have a feasible bilateral trade from (A^*, ϕ^*) to (A', ϕ^*) that strictly increases the expected equity value of both i and j .

Proposition 3 shows that the social planner's solution is not stable under limited liability. Banks have valuable deviations that raise the probability of a bank's failure, but also raise the bank's expected equity value.

Examples. To gain intuition, consider an example of a portfolio deviation that raises expected equity value. Suppose that bank i changes its portfolio. Instead of holding a portfolio consisting only of investments of type i , bank i holds a portfolio consisting only of investments of type j , the same investment as another bank in i 's cluster. This change causes i to fail when a small shock hits investment j . But the failure does not lower i 's equity value. By limited liability, i 's equity value after this deviation is zero. Prior to the change, in the social planner's solution, i 's equity value is also zero. Further, when a shock now hits investment type i , no banks are affected and i has positive equity value, whereas before it had zero equity value. Overall, i 's expected equity value increases. Importantly, i is more likely to fail after the deviation. But the default costs affect i 's external debt holders and not i 's equity holders, due to limited liability.

As a second example, banks can deviate with bilateral trades. Consider banks i and j in different clusters. Consider a feasible bilateral trade that increases i 's claims on banks within i 's cluster, while eliminating i 's claims on j 's cluster. This trade increases the equity value of bank i when the small shock hits j 's cluster— i is no longer exposed to the small shock. Further, bank i 's equity value does not change when a small shock hits its own cluster. Though i now fails, its equity value is zero due to limited liability. Prior to the deviation, i 's equity value is also zero when a small shock hits its cluster, though i does not fail. So after the deviation, i raises both its probability of failure and its expected equity value—by taking advantage of limited liability.

Our stability concept entertains general deviations, in which banks can adjust both their portfolios and their financial claims. However, the proposition shows that restricted deviations still render the social planner's solution unstable. There are valuable deviations in which banks hold fixed their portfolios and only change their financial claims. There are also valuable deviations in which banks hold fixed their financial claims and only change their portfolios. Therefore, the instability result does not rely on the interaction between banks' real and financial exposure. Banks can adjust *either* their financial *or* their real claims, to make a profitable deviation.

Role of assumption: face value of external debt is fixed. We assume that banks take the face value of external debt as fixed when contemplating deviations. However, for our results, this is consistent with external debt being endogenously determined and reflecting the risks debt-holders expect to face before financial claims and portfolios are chosen.²⁵ We expect external debtors to anticipate and adjust the cost of their debt to account for risks associated with the financial network formed and portfolios chosen—while leaving this step unmodeled. Our results on the social planner's solution, the stability of these efficient network and portfolio pairs when banks maximize their expected market values, and the instability of these efficient network and portfolio pairs when banks maximize their expected equity values, all require us to consider only the level of external debt that would obtain when external debt holders anticipate the social planner's solution will be played. As banks are symmetric in the social planner's solution, our

²⁵ Implicitly, this requires that external debt holders cannot issue debt conditional on networks and portfolios.

assumption that external debt is \underline{v} for all banks is consistent with this. Of course, the return required by debt holders will be different under other stable network and portfolio pairs. In combination with the huge space of possible networks portfolio pairs, this presents a further difficulty characterizing the set of stable networks. We do not attempt such a characterization, but this is an important consideration for future work.

4.4. Systemic risk shifting

Limited liability encourages equity holders to raise the probability of default and deviate from social efficiency. This behavior, often termed risk shifting, is well known (Jensen and Meckling, 1976). However more novel to our setting, risk shifting has consequences for systemic risk—this behavior is *systemic risk shifting*. Banks deviate from the social planner’s solution in order to fail at the same time as their counterparties. Therefore, these deviations raise the probability that a large part of the system will fail at the same time, after either small or large shocks. Systemic risk increases.

Remark 2 clarifies the nature of systemic risk shifting.

Remark 2. Letting $\Pr(x)$ denote the probability of the event x , bank i ’s equity value can be rewritten as

$$\begin{aligned} \mathbb{E}[\pi_i] = & \Pr(v_i \geq \underline{v}) \left(\sum_j \sum_k A_{ij} \mathbb{E}[p_j | v_i \geq \underline{v}] - \underline{v} \right) \\ & - \beta \Pr(v_i \geq \underline{v}) \left(\sum_{j \in N} A_{ij} \Pr(v_j < \underline{v}) \right) \\ & + \beta \sum_{j \in N} A_{ij} \text{Cov}[I_{v_j < \underline{v}}, I_{v_i < \underline{v}}]. \end{aligned} \tag{12}$$

The first term on the right-hand side of equation (12) is simply the market value of bank i ’s when i does not default. If the default costs were zero, this would be the only term in equation (12). This term captures well-understood risk shifting incentives of bank i stemming from limited liability. Bank i would like to increase its market value when it does not fail, but does not care about its market value conditional on failing. The second term shows that the equity value of i is reduced by the failure of its counterparties in the states in which i does not fail. Also, note that portfolio shares and maintenance costs enter portfolio returns p_j , so do not directly appear in equation (12).

The third term in equation (12) captures systemic risk shifting. i ’s equity value is greater if it fails at the same time as its counterparties. Intuitively, suppose i fails in different states from its counterparties. Then, counterparties’ default costs subtract from i ’s equity value in these states, lowering expected equity value. By contrast, suppose i fails at the same time as its counterparties. Then i ’s equity value is zero. Due to limited liability, counterparties’ default costs do not subtract from i ’s equity value. Instead, counterparty default costs lower the value of i ’s external debt. This force enables banks’ deviations from the social planner’s solution to raise equity value. Moreover, if banks fail at the same time as their counterparties, systemic risk is higher.

Our model lets banks adjust either their investment portfolios or their financial claims. Proposition 3 shows that banks can use either portfolios or financial claims to make deviations that

raise equity value. Systemic risk shifting is the reason. By changing either portfolios or financial claims, banks increase the likelihood of failing at the same time as their counterparties.

Relation to stylized fact. Recall the stylized fact from the introduction. German banks have investments similar to those of their close financial counterparties. Recall, also, the social planner’s solution. In contrast to the stylized fact, banks hold investments different from those of their financial counterparties, in the socially efficient outcome. Systemic risk shifting is a reconciliation. Systemic risk shifting means banks find it privately beneficial to have the same real exposure as their financial counterparties. This pattern is not socially efficient, and generates systemic risk.

4.5. Systemic risk shifting in other networks and portfolios

We now show that the valuable deviations available to banks, starting from the social planner’s solution, are also available in many other networks and portfolios. Systemic risk shifting matters more generally.

Lemma 1. *For any network portfolio pair (A, ϕ) , bank i can increase its expected equity value by deviating to (A, ϕ') if*

- (i) *there is an investment type k such that $A_{ii}\phi_{ik} > 0$ and $v_i(A, \phi) \geq \underline{v}$ conditional on a small shock hitting investments of type k ;*
- (ii) *there is an investment type l such that $v_i(A, \phi) < \underline{v} + (1 - q)A_{ii}\phi_{ik}\varepsilon_S$ conditional on a small shock hitting investments of type l .*

We prove Lemma 1 constructively. The deviation that raises equity value is to a portfolio ϕ'_i such that $\phi'_{ij} = \phi_{ij}$ for all $j \neq k, l$, $\phi'_{ik} = 0$ and $\phi'_{il} = \phi_{il} + \phi_{ik}$. Under this deviation, bank i ’s value is unaffected when a shock hits a bank $j \neq k, l$. When a small shock hits assets of type l , bank i ’s equity value falls by at most $(1 - q)A_{ii}\phi_{ik}\varepsilon_S$. Bank i ’s profits can fall no further due to limited liability. On the other hand, when a shock hits assets of type k , i ’s equity value now increases by at least $A_{ii}\phi_{ik}\varepsilon_S$. When a large shock hits investments of type l and k the analysis is similar. Overall, expected equity value increases.

Lemma 1 underscores the importance of systemic risk shifting. Equity holders frequently face incentives to hold the same investments as their largest counterparties, while increasing systemic risk. Suppose, for example, that bank j is a large counterparty of bank i and bank j ’s portfolio consists of investments of type l . Then, conditional on a small shock hitting investment type l , bank i will suffer direct losses before any bankruptcy costs of $A_{ij}\phi_{jl}\varepsilon_S$. A sufficient condition for bank i ’s value to then be less than $\underline{v} + (1 - q)A_{ii}\phi_{ik}\varepsilon_S$ when a small shock hits assets of type l is then that $(A_{ij}\phi_{jl} + A_{ii}\phi_{il})\varepsilon_S > R - \underline{v}$. This condition is easier to satisfy when bank j is a large counterparty of bank i , bank j ’s portfolio consists of investments of type l , and bank i also has large holdings in investments of type l .

Systemic risk shifting in other papers. In a non-network, two-bank setting, Acharya (2009) identifies the same systemic risk shifting mechanism, and coins the term. Subsequent to us, Jackson and Pernoud (2019) build on our foundations to investigate further the interaction between real exposures and financial connections in a more general network setting, with more general contractual forms determining the financial interdependencies between banks. The phenomenon we study persists in this setting absent limited liability because there is still a complementarity between the returns of banks and their counterparties—banks want to avoid failure when their

counterparties earn high returns, and fail when counterparties make large losses. This further emphasizes the importance of considering banks' joint problem of choosing financial connections and exposures to the real economy when designing regulations.

5. Stylized fact

In this section, we flesh out the stylized fact we presented in the introduction—that German commercial banks lend more to each other when they have more similar real exposures. The relationship is statistically significant, quantitatively large, and robust across numerous specifications.

To test our hypothesis that banks lend more to banks with similar real exposures, we use quarterly data on all of the roughly 170 commercial banks in Germany from 2006/Q1 to 2014/Q4.²⁶ As discussed at some length in the introduction, among German commercial banks interbank loans represent a substantial and systemically important financial connection, while loans to firms are an important and substantial asset class generating meaningful real exposures.²⁷ More generally, lending to the corporate sector is a growing threat to financial stability worldwide, due to rapid growth in leveraged lending.²⁸ Our focus on corporate lending reflects these concerns.

We use quarterly data from the German large credit registry. This data contains detailed information about bank-firm lending collected under the German Banking Act, which requires all financial institutions located in Germany to report the end-of-quarter value of every loan to a firm based in Germany, provided the loan exceeded EUR1.5 million during the preceding quarter. The credit registry is extensive and covers well over 90% of all interbank loans and more than 70% of all bank-firm loans. The interbank data contain net liquidity flows between banks in each quarter, but does not report the exact financial instrument used (i.e. loan or bond).

Using these data, we first define our dependent variable $\log(\text{Amount}_{ij,t})$ as the log of the amount of interbank lending from bank i to bank j at date t , normalized by the total interbank lending of bank i at date t . Our main independent variable is a measure of portfolio similarity between lender bank and borrower bank.²⁹ This can be measured as follows: consider every bank-firm loan as a vector along a basis in a vector space where each direction corresponds to a specific asset class. An investment of EUR10 million in asset A is then defined as a vector of length 10 along direction A . The portfolio of each bank can then be represented by a vector in this K dimensional vector space (where K is the number of firms in the dataset, in our case around $K = 230,000$). The similarity of two portfolios can then easily be expressed through

²⁶ The German banking system has about 1,900 banks at present. Crucially, the German banking system has a three-pillar structure with, in addition to the commercial banks, about 1,000 cooperative banks and about 450 savings banks (as well as a number of subsidiaries of foreign banks and banks with a special purpose). Cooperative and savings banks do not usually access the entire interbank market, but rather transact through their respective head institutions. By focusing on commercial banks only, we avoid these institutional differences.

²⁷ While we cannot rule out the possibility that some of these loans are syndicated, syndicated loans make up a small fraction of all loans to non-financial firms in Germany and even fewer of these are in syndicates with other German banks. More details on the syndicated loan market in Europe can be found [here](#).

²⁸ See, for example, FOMC (2018), BoE (2018), IMF (2018) or BIS (2018).

²⁹ Only a fraction of all commercial banks in Germany are publicly listed. Consequently, market-based measures of correlation like CoVaR (Adrian and Brunnermeier, 2016) can be computed only for a subset of banks in our sample. However, as Abbassi et al. (2017) show, market-based measures of interdependence are a good proxy for correlations obtained from bank-level credit register data.

the Euclidean distance (i.e. the length of the distance vector between the two banks' portfolio vectors):

$$Distance_{ij,t} = \sqrt{\sum_{k=1}^K (x_{i,t}^k - x_{j,t}^k)^2} \tag{13}$$

where $x_{i,t}^k$ is bank i 's holding of asset k in quarter t . Two banks have a high portfolio similarity if they have a low Euclidean distance. In some specifications we use the normalized Euclidean distance, which we compute using portfolio weights instead of the exposures themselves: $w_{i,t}^k = x_{i,t}^k / X_{i,t}$, where $X_{i,t}$ is the total bank-firm lending at date t . For robustness, we consider a second measure of the similarity between banks' real exposures. We calculate the cosine similarity between two portfolios as

$$CosineSimilarity_{ij,t} = \frac{\sum_{k=1}^K x_{i,t}^k x_{j,t}^k}{\sqrt{\sum_{k=1}^K (x_{i,t}^k)^2} \sqrt{\sum_{k=1}^K (x_{j,t}^k)^2}} \tag{14}$$

as a robustness check for our main specifications.

For our main specification, we focus on the intensive margin, using an unbalanced panel with the amount of additional interbank lending from bank i to bank j in period t as the panel variable. Our main specification is

$$\log(Amount_{ij,t}) = \beta_{it} + \beta_{jt} + \beta_{ij} + \beta_t + \gamma \log(Distance_{ij,t}) + \varepsilon_{ij,t}.$$

$\log(Amount_{ij,t})$ is the log of the amount of interbank lending from bank i to bank j normalized by the total amount of interbank lending by bank i at date t . $\log(Distance_{ij,t})$ is the Euclidean distance between two banks' portfolios. We also include a quarterly time-fixed effect, annually varying borrower and lender fixed effects, as well as a borrower-lender pair fixed effect to account for unobserved heterogeneity. This is an extremely strong set of fixed effects, meaning that time-varying bank controls add very little explanatory value and do not change our results qualitatively.³⁰ Since the large credit registry only reports *exposures*, but not flows, we compute the amount of new credit provided from i to j using the difference of the exposure from i to j in t , minus the exposure from i to j in $t - 1$ (i.e. in the previous quarter). We have used the convention that a negative flow from i to j is a positive flow from j to i so that we can take the log of the amount as the main dependent variable. Using this convention, banks lend more to similar banks if the sign of γ is negative.

Results are shown in Table 2 in Appendix A for our main result and in Table 4 for a robustness check using cosine similarity as our measure of the similarity between banks' real exposures. In all specifications, banks lend more to partners with similar exposures to the real economy. The effects are statistically significant and also quantitatively substantial. In all specifications, the relationship remains significant and large. We saturate our estimation with a large set of

³⁰ We also added controls from detailed monthly balance sheet information, using the Bundesbank's balance sheet statistics (BISTA). These controls did not change our results qualitatively. Results for this specification are available upon request.

fixed effects to control for unobserved heterogeneity.³¹ We always include a time-fixed effect to control for trends. In model (1) we also include lender and borrower fixed effects. In model (2) we use yearly varying lender and borrower fixed effects, which controls for a substantial amount of unobserved variation and, additionally, a borrower-lender pair fixed effect. In model (3) we furthermore include a borrower-lender fixed effect, i.e. one fixed effect per borrower-lender pair.³²

In Panel A we compute the normalized Euclidean distance based on banks' co-investment in firms. Most banks only have few firms in common, however, since the total number of firms is relatively large compared to the number of firms commercial banks invest in. We find strong evidence for our hypothesis in all three specifications. Next, in Panel B, we use the lagged non-normalized Euclidean distance as the explanatory variable. This allows us to investigate whether our results are being driven only by banks' different portfolio sizes. As results remain statistically significant, homophily is present in banks' portfolio compositions, rather than just their portfolio sizes.

As previously discussed, our theory predicts a positive correlation between firms' real and financial exposure, without asserting causality. Thus, we can formulate our empirical question via the co-movement of variables, without having to establish claims about causality. Irrespectively, it is interesting to investigate whether lagged Euclidean distance explains banks' propensity to link to each other. Panel C reports on this exercise. Finally, in Panel D, as an alternative measure of how dissimilar banks' portfolios are, we classify loans into one of 74 sectors and apply the Euclidean distance measure in this space.

In all four estimations we find strong empirical evidence for our hypothesis that banks tend to lend to banks that are similar to themselves. Our results are qualitatively robust to further robustness checks, including: (i) using an alternative measure of similarity based on the cosine similarity of both banks' portfolios as shown in Panel A of Table 4; and (ii) adding a set of bank-specific control variables to control for observed heterogeneity. Our results are economically significant. As we report in the introduction, when the Euclidean distance between two banks decreases from the 75th to the 25th percentile, i.e. when their non-financial similarity increases, their net lending grows by roughly 26%.³³

In Table 3 we consider the extensive margin of loan creation. We create a binary variable that takes a value of one if banks i and j have an interbank loan with one another at date t , and zero otherwise. The results from this balanced panel are reported in Panel A. In Panel B we explicitly study the creation of new loans. The dependent variable is $ENTRY_{ij,t}$ which equals one if there exists an interbank loan from bank i to bank j at date t that did not exist on date $t - 1$. We find overwhelming support for our hypothesis.

³¹ In particular, this accounts for the possibility that a lender is a core bank, while the borrower is a periphery bank. Since we even include time-varying lender- and borrower fixed effects, we even allow for the possibility that this feature changes over time.

³² Even though we have winsorized the bank-firm and bank-bank exposures at the 5% and 95% levels, our results could still be driven by a few outliers. Fig. 1 from the introduction shows that this is not the case. In this figure we plot the residuals from regressing the dependent and independent variable on time-varying borrower and lender fixed-effects, as well as borrower-lender fixed-effects.

³³ This is based on model (1), Panel A from Table 2 and the percentile variables presented in the summary statistics (Table 1).

6. Concluding remarks

This paper makes two contributions. First, we presented a new stylized fact from the German banking system. Banks have similar real exposures to their financial counterparties. If German commercial banks lend to one another in the interbank market, they tend to lend to similar non-financial firms. Second, we rationalized this behavior and showed it generates systemic risk. We presented a model with limited liability, real investments, and a financial network. We characterized socially efficient networks. In these networks, which minimize systemic risk, banks have different real exposures from their counterparties. Absent limited liability, banks have no incentive to deviate from the socially efficient network. But limited liability leads banks to deviate from social efficiency and engage in *systemic risk shifting*. Banks increase their equity values by having the same real exposures as their close financial counterparties. Then, banks fail at the same time as their counterparties, raising their expected equity value but also increasing systemic risk.

Policymakers emphasize jointly considering financial and real exposures when minimizing systemic risk (Basel Committee on Banking Supervision, 2011). Much anecdotal evidence points in this direction. We provide theory and evidence showing that the joint behavior of real and financial exposures matters for systemic risk.

We close with a caveat. Our model predicts only a correlation between banks' real and financial exposures and is silent on causality—it does not matter whether banks adjust their real exposures in light of their financial exposure, or adjust their financial exposure in response to real exposures. We have, in effect, considered both alternatives. If we changed the model so that banks could only engage in unilateral deviations to change the correlation of their real exposures with their counterparties, the socially efficient outcome would be stable when banks maximize their values, but not when they maximize their equity values under limited liability. If we changed the model so that banks could only engage in bilateral trades to change their financial exposures while holding their real exposures fixed, the socially efficient outcome would be stable when banks maximize their values, but not when they maximize their equity values under limited liability as long as failure costs are not too high. Regardless of these details, the stylized fact presented in Section 5 is consistent with our model.

Appendix A. Tables

Table 1
Summary statistics for dependent and independent variables.

	N	mean	median	sd	p25	p75
$\log(\text{Amount}_{ij,t})$	33,891	-6.247527	-5.973928	3.310431	-8.446989	-3.763935
$\log(\text{Distance}_{ij,t})$	33,891	14.4612	14.7249	1.352575	13.8095	15.3281
$\log(\text{NormDistance}_{ij,t})$	33,891	0.2540417	0.186755	0.1677311	0.1317385	0.3378451
$\log(\text{SectorDistance}_{ij,t})$	33,891	15.61501	16.13923	1.373794	15.23235	16.4708

Note: $\log(\text{Amount})_{ij,t}$ is the log of the amount of interbank lending from bank i to bank j normalized by the total amount of interbank lending by bank i at date t . $\log(\text{Distance}_{ij,t})$ is the log of the Euclidean distance of portfolio choice of banks i and j . $\log(\text{Distance}_{ij,t-1})$ is the lagged log of the Euclidean distance of portfolio choice of banks i and j . $\log(\text{NormDistance}_{ij,t-1})$ is the lagged log of the Euclidean distance of the relative portfolio weight (the amount lent from bank i to firm k normalized by the total portfolio size of bank i) of banks i and j . And $\log(\text{SectorDistance}_{ij,t-1})$ is the lag of the log of the Euclidean distance of portfolio choices of banks i and j based on their investment not in individual firms, but in sectors of the economy.

Table 2

Interbank lending and portfolio distance – intensive margin.

PANEL A – OVERLAP IN FIRM EXPOSURE – NORMALIZED DISTANCE

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{NormDistance}_{ij,t})$	-1.509*** (-4.53)	-2.965*** (-4.58)	-0.696* (-1.70)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	33,885	33,745	33,048
R^2 (adjusted)	0.247	0.276	0.491

PANEL B – OVERLAP IN FIRM EXPOSURE – DISTANCE

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{Distance}_{ij,t})$	-0.608*** (-9.14)	-0.892*** (-10.97)	-0.121** (-2.40)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	33,885	33,745	33,048
R^2 (adjusted)	0.417	0.448	0.592

PANEL C – OVERLAP IN FIRM EXPOSURE – LAGGED DISTANCE

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{Distance}_{ij,t-1})$	-0.533*** (-7.59)	-0.759*** (-10.02)	-0.0559 (-0.75)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	17,854	17,595	17,034
R^2 (adjusted)	0.428	0.472	0.596

PANEL D – OVERLAP IN LAGGED SECTORAL EXPOSURE

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{SectorDistance}_{ij,t})$	-0.672*** (-10.07)	-0.716*** (-10.56)	-0.186 (-1.24)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	17,854	17,595	17,034
R^2 (adjusted)	0.438	0.477	0.596

Note: The dependent variable is always $\log(\text{Amount})_{ij,t}$, the log of the amount of interbank lending from bank i to bank j normalized by the total amount of interbank lending by bank i at date t . The independent variables are: (i) $\log(\text{NormDistance}_{ij,t})$, the log of the Euclidean distance of the relative portfolio weight (the amount lent from bank i to firm k normalized by the total portfolio size of bank i) of banks i and j ; (ii) $\log(\text{Distance}_{ij,t})$, the log of the Euclidean distance of portfolio choice of banks i and j ; and (iii) $\log(\text{Distance}_{ij,t-1})$, the lagged log of the Euclidean distance of portfolio choice of banks i and j ; (iv) $\log(\text{SectorDistance}_{ij,t-1})$, the lag of the log of the Euclidean distance of portfolio choices of banks i and j based on their investment not into individual firms, but into sectors of the economy. Standard errors are always clustered at the borrower and lender levels.

Table 3
Interbank lending and portfolio distance – extensive margin.

PANEL A – EXTENSIVE MARGIN

$1_{ij,t}$	(1)	(2)	(3)
$\log(NormDistance_{ij,t-1})$	-0.307*** (-8.12)	-0.385*** (-6.31)	-0.109*** (-3.13)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	134,400	134,400	134,400
R^2 (adjusted)	0.154	0.193	0.335

PANEL B – NEW CREATION OF LINKS

$ENTRY_{ij,t}$	(1)	(2)	(3)
$\log(NormDistance_{ij,t-1})$	-0.102*** (-5.98)	-0.0683*** (-3.45)	-0.0156 (-0.71)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	134,400	134,400	134,400
R^2 (adjusted)	0.025	0.038	0.050

Note: In Panel A, the dependent variable is $1_{ij,t}$, a dummy variable that equals one if there is a loan from i to j at date t and zero otherwise. In Panel B the dependent variable is $ENTRY_{ij,t}$, a dummy variable that equals one if a loan is created between banks i and j at date t that did not exist at date $t - 1$, and zero otherwise. The independent variable is in both cases $\log(NormDistance_{ij,t})$, the log of the Euclidean distance of the relative portfolio weight (the amount lent from bank i to firm k normalized by the total portfolio size of bank i) of banks i and j . Standard errors are always clustered at the borrower and lender levels.

Table 4
Interbank lending and portfolio distance – robustness checks.

PANEL A – INTENSIVE MARGIN

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{CosineSimilarity}_{ij,t-1})$	0.106*** (6.16)	0.0995*** (4.82)	-0.0297 (2.31)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	26,174	25,983	25,488
R^2 (adjusted)	0.28	0.309	0.500

PANEL B – EXTENSIVE MARGIN

$1_{ij,t}$	(1)	(2)	(3)
$\log(\text{CosineSimilarity}_{ij,t-1})$	0.891* (1.95)	1.398*** (2.68)	-0.0470 (-0.17)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	79,873	79,806	79,691
R^2 (adjusted)	0.182	0.219	0.350

PANEL C – FREQUENTLY INTERACTING PAIRS ONLY

$\log(\text{Amount}_{ij,t})$	(1)	(2)	(3)
$\log(\text{NormDistance}_{ij,t})$	-1.427*** (-4.16)	-3.081*** (-4.39)	-0.589 (-1.56)
Time FE	Yes	Yes	Yes
Lender + Borrower FE	Yes	-	-
Time-varying Lender + Borrower FEs	No	Yes	Yes
Borrower-Lender FE	No	No	Yes
N	31,303	31,163	30,466
R^2 (adjusted)	0.239	0.269	0.486

Note: In Panel A, the dependent variable is $\log(\text{Amount})_{ij,t}$, the log of the amount of interbank lending from bank i to bank j normalized by the total amount of interbank lending by bank i at date t . In Panel B, the dependent variable is $1_{ij,t}$, a dummy variable that equals one if there is a loan from i to j at date t and zero otherwise. In Panel B the dependent variable is $\text{ENTRY}_{ij,t}$, a dummy variable that equals one if a loan is created between banks i and j at date t that did not exist at date $t-1$, and zero otherwise. The independent variable is in both cases $\log(\text{CosineSimilarity}_{ij,t})$, the log of the cosine similarity of the relative portfolio weight (the amount lent from bank i to firm k normalized by the total portfolio size of bank i) of banks i and j . In Panel C we have restricted the sample on *infrequent linkages only*, i.e. on borrower-lender pairs that have a connection with one another at most 20% of the time. The independent variable is $\log(\text{NormDistance}_{ij,t-1})$ the lagged log of the Euclidean distance of the relative portfolio weight (the amount lent from bank i to firm k normalized by the total portfolio size of bank i) of banks i and j . The computation is conditional on two banks having at least one exposure in common. Standard errors are always clustered at the borrower and lender levels.

Appendix B. Omitted proofs

B.1. Proof of Remark 1

Using equations (3) and (4) we have that

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in N} \pi_i + \delta_i \right] &= \mathbb{E} \left[\sum_{i \in N} \max\{v_i - \underline{v}, 0\} + \min\{\underline{v}, v_i\} \right] \\ &= \sum_{i \in N} \sum_{j \in N} A_{ij} \mathbb{E}[p_j] - \beta \sum_{i \in N} \sum_{j \in N} A_{ij} \mathbb{E}[I_{v_j < \underline{v}}] \end{aligned}$$

Noting that $\mathbb{E}[p_j] = \bar{R} - r\mathbb{E}[\varepsilon]/n - \sum_k cI_{\phi_{ik} > 0}$; and that $\sum_{i \in N} A_{ij} = 1$ because the dependency matrix is column stochastic, it follows that:

$$\begin{aligned} \mathbb{E} \left[\sum_{i \in N} \pi_i + \delta_i \right] &= n\bar{R} - r\mathbb{E}[\varepsilon] - \sum_j \sum_k cI_{\phi_{jk} > 0} - \beta \mathbb{E} \left[\sum_{j \in N} I_{v_j < \underline{v}} \right] \\ &= nR - r\mathbb{E}[\varepsilon] - \mathbb{E}[\text{cost. of defaults}] - \underbrace{\sum_j \sum_k cI_{\phi_{jk} > 0}}_{\text{portfolio costs}} \end{aligned}$$

Since all but the final two term of the above equation are exogenously given, the social planner maximizes $\mathbb{E} \left[\sum_{i \in N} \pi_i + \delta_i \right]$ by minimizing the expected cost of defaults less portfolio maintenance costs. \square

B.2. Proof of Proposition 1: preliminaries

We begin by stating and then proving a Lemma. We define a network to be η -doubly stochastic if it is column stochastic and $\sum_j A_{ij} \in [1 - \eta, 1 + \eta]$ for all i .

Lemma 2. *If banks have non-overlapping portfolios, there exists a $\bar{\eta} > 0$ such that for all $\eta < \bar{\eta}$, in all η -doubly stochastic network structures \mathbf{A} for which no small shock always causes at least one bank to fail, at least d^* banks fail following a large shock to i .*

Let $\mathcal{A}(\eta)$ be the set of non-negative, η -doubly stochastic, n -by- n matrices. Let $D_i(\mathbf{A})$ be the set of organizations that fail following a large shock to i , and set $d_i(\mathbf{A}) = |D_i(\mathbf{A})|$. Now consider the following problem,

$$\text{P1: } \min_{\mathbf{A} \in \mathcal{A}(\eta)} d_i(\mathbf{A}) \text{ subject to } A_{jk} \varepsilon_S \leq (1 + \eta)R - \underline{v} \text{ for all } j, k \in N$$

As \mathbf{A} must be η -doubly stochastic, the equity value of each bank absent a shock is at least $(1 - \eta)R - \underline{v}$ and at most $(1 + \eta)R - \underline{v}$. As, by definition, each organization $j \notin D_i$ does not fail following a large shock to i , an upper bound on the losses absorbed by banks not in D_i after a large shock to i is then $(n - d_i)((1 + \eta)R - \underline{v})$. The remaining losses must be absorbed by the remaining banks. Thus, collectively organizations in D_i incur losses of at least $\varepsilon_L + d_i \beta - [n((1 + \eta)R - \underline{v}) - (1 + \eta)d_i(R - \underline{v})]$, and so

$$\sum_{j \in D_i} \left[\sum_{k \in D_i} A_{jk} \beta + A_{ji} \varepsilon_L \right] \geq \varepsilon_L + d_i \beta - [n((1 + \eta)R - \underline{v}) - d_i((1 + \eta)R - \underline{v})]. \quad (15)$$

It follows immediately from the constraints in P1 that

$$\sum_{j \in D_i} A_{ji} \leq \frac{d_i(R(1 + \eta) - \underline{v})}{\varepsilon_S}, \quad (16)$$

and

$$\sum_{j \in D_i} \sum_{k \in D_i} A_{jk} \leq \frac{d_i^2(R(1 + \eta) - \underline{v})}{\varepsilon_S}. \quad (17)$$

Combining inequalities (15), (16) and (17)

$$f(d_i) := d_i^2 \left[(1 + \eta)R\beta - \underline{v}\beta \right] + d_i \left[((1 + \eta)R - \underline{v})(\varepsilon_L - \varepsilon_S) - \varepsilon_S\beta \right] + \varepsilon_S[n((1 + \eta)R - \underline{v}) - \varepsilon_L] \geq 0.$$

For all η sufficiently small, as $\varepsilon_L > n(R - \underline{v})$, the constant term of this quadratic equation is always negative; and the quadratic coefficient is always positive. It follows from the quadratic formula that $f(d_i)$ always has exactly one positive real root, which we denote by $d_i^*(\eta)$. Moreover, as $f(d_i)$ increases in d_i at $d_i^*(\eta)$, for values of $d_i < d_i^*(\eta)$ the constraint is violated ($f(d_i) < 0$). Thus, $\lceil d_i^*(\eta) \rceil$ is the unique solution to P1.

Now consider the value of $d_i^*(\eta)$ as $\eta \rightarrow 0$. By the quadratic formula the positive root of the inequality, $d_i^*(\eta)$, is continuous in η . Moreover, by Remark 2 in Section OA3 of the Online Appendix, $d_i^*(\eta)$ decreases in R and so also decreases in η . Thus, for all $\eta > 0$, $d_i^*(\eta) < d_i^*$. Finally, as by assumption d_i^* is an integer, for all η sufficiently small, $\lceil d_i^*(\eta) \rceil = d_i^*$. Hence, for all η sufficiently small, the unique solution to P1 is d_i^* .

We now argue that the solution to P1 provides a lower bound on the number of failures that will hit following a large shock to i , subject to there being no small shock that always causes at least one bank to fail. Bank j 's equity value in any state of the world in which a small shock hits k 's project is at most

$$\sum_l A_{jl}R - A_{jk}\varepsilon_S - \underline{v} = ((1 + \eta)R - \underline{v}) - A_{jk}\varepsilon_S.$$

The constraint that no small shock always causes at least one bank to fail therefore implies that $A_{jk} \leq ((1 + \eta)R - \underline{v})/\varepsilon_S$ for all j , and all k . This is the constraint imposed in P1, and so the solution to P1 provides a lower bound on the number of failures that must be incurred following a large shock to i , if a network is chosen that avoids there being a failure for sure after a small shock hits any bank. \square

B.3. Proof of Proposition 1

There are five steps to the proof: First, we show that the planner will always choose non-overlapping portfolios. Second, we show that there can be no small shock that results in one or more failures in a socially optimal network. Third, we show that the participation constraints mean that for all $\eta > 0$, there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ the social planner must choose

a η -doubly stochastic network. Fourth, we apply Lemma 2 to show that all d^* -clustered networks with separate shocks are socially optimal for all $r > 0$ sufficiently small. Finally, we show that all socially optimal networks are d^* -clustered with non-overlapping portfolios for $r > 0$ sufficiently small.

Step 1: Non-overlapping portfolios.

Consider any feasible portfolio ϕ and network A , where ϕ does *not* satisfy the non-overlapping portfolios condition. This generates claims for each bank i on assets of type k equal to $\sum_j A_{ij}\phi_{jk}$. Consider now the alternative network portfolio pair ϕ' and network A' . We let ϕ' be the non-overlapping portfolios in which each bank i only makes investments of type i . Thus, $\phi'_{ii} = 1$ and $\phi'_{ij} = 0$ for all i and all $j \neq i$. For all bank pairs i, k , we set $A'_{ik} = \sum_j A_{ij}\phi_{jk}$. Note that

$$\sum_i A'_{ik} = \sum_i \sum_j A_{ij}\phi_{jk} = \sum_j \sum_i A_{ij}\phi_{jk} = \sum_j \phi_{jk} = 1.$$

Thus the network A' is non-negative and column stochastic, and thus feasible.

Suppose then the planner chooses the network portfolio pair (A', ϕ') instead of (A, ϕ) . Suppose first that $r = c = 0$. Then (A', ϕ') and (A, ϕ) , by construction, generate exactly the same market values for all banks in all states of the world. Thus, the planner would be indifferent between them. Suppose this leads to a value of V for the planner's objective. Now consider the case in which $r = 0$, but $c > 0$. In this case, the planner's objective decreases by a strictly positive amount $\zeta > 0$.

Now we consider $r > 0$ and $c > 0$. By Remark 1, the social planner seeks to minimize $\beta \mathbb{E} \left[\sum_{j \in N} I_{v_j < \underline{v}_j} \right] + \sum_j \sum_k c I_{\phi_{ik} > 0}$. The first term in the objective can differ under the network portfolio pairs (A', ϕ') and (A, ϕ) as in the two cases different numbers of banks may fail in different states of the world. However, the value of this difference in the planner's objective is bounded from above by $rn\beta$ (supposing there are never any failures in one case and that all banks fail following any shock in the other case). However, for all $c > 0$ there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ we have $rn\beta < \zeta$. Thus, for all r sufficiently small, the planner always chooses non-overlapping portfolios. For the remainder of the proof we use this result and restrict attention to non-overlapping portfolios.

Step 2: No small shocks cause any failures. We show that the planner always chooses a network satisfying this by demonstrating that any network in which a small shock causes at least one failure is dominated by the complete network in which $A_{ij} = 1/n$ for all i and all j .

Under the complete network, the average value of a bank (assuming non-overlapping portfolios and recalling that $R = \bar{R} - c$) is

$$R - (1 - q)r\varepsilon_S/n - qr(\varepsilon_L + n\beta)/n.$$

The first term is the average project return absent any shocks, the second term represents losses from small shocks, which never result in a failure, and the third term represents losses from large shocks which always result in n failures.

An upper bound (that is not achievable) on the average value of a bank when there is a one project that when hit by a small shock always results in at least one failure, is given by the average value that would result assuming there are never any failures from any shocks to other projects and there is exactly one failure following a small or large shock to this project. This upper bound is

$$R - (1 - q)r\varepsilon_S/n - qr\varepsilon_L/n - r\beta/n^2.$$

The first term is the average returns absent any shocks, the second term represents direct losses from small shocks, the third term represents direct losses from large shocks, and the fourth term represents losses from a project that results in one failure when it is hit by a small (or large) shock. This fourth term is given by the product of the probability with which a shock hits the project in question (r/n) and the average loss per bank from a failure (β/n). Thus, a sufficient condition for the complete network to be preferred is that $q < 1/n^2$, which holds by assumption.

Step 3: Row stochasticity. We now show that for r sufficiently small the planner will choose an approximately doubly stochastic network. We do so in two steps. First, we show that there exists a row stochastic network satisfying the participation constraints for all $r \in [0, 1]$. We then show that for all $\eta > 0$ there exists a $\bar{r} > 0$ such that for $r < \bar{r}$, any network that is not η -doubly stochastic will violate a participation constraint.

It is straightforward to verify that any network $\mathbf{A} \in \mathcal{A}^*$ satisfies the participation constraints for all $r \in [0, 1]$. The participation constraints require that

$$v_i(\mathbf{A}) \geq R - \frac{r}{n} (q\varepsilon_L + (1 - q)\varepsilon_S + \beta). \tag{18}$$

This is satisfied for any network $\mathbf{A} \in \mathcal{A}^*$ because $q < 1/d^*$. For any $\mathbf{A} \in \mathcal{A}^*$,

$$v_i(\mathbf{A}) = R - \frac{r}{n} (q\varepsilon_L + (1 - q)\varepsilon_S + qd^*\beta).$$

By inspection, the participation constraint (equation (18)) is continuous in r and $\lim_{r \rightarrow 0} v_i(\mathbf{A}) = \sum_j A_{ij}R \geq R$ for all i . This is equivalent to $\sum_j A_{ij} \geq 1$ for all i . As, by column stochasticity, $\sum_i \sum_j A_{ij} = n$, if there is any i such that $\sum_j A_{ij} > 1$ there must exist a k such that $\sum_j A_{kj} < 1$ implying that at least one participation constraint will be violated in the limit. Thus, in the limit, satisfying all the participation constraints requires the network to be doubly stochastic. Further, by the continuity of the participation constraint in r , for all $\eta > 0$ there exists a $\bar{r} > 0$ such that for $r < \bar{r}$, any network that is not η -doubly stochastic will violate a participation constraint.

Step 4: d^* -clustered networks obtain the bound. By Step 2, for all socially optimal networks there are never any failures following a small shock. By step 3, for all $\eta > 0$, there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ the planner must choose a network that is η -doubly stochastic. Thus, there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ the lower bound on the expected number of failures that must occur when a large shock hits a given organization i that we found in Lemma 2 holds in all solutions to the social planner’s problem. Specifically, Lemma 2 shows that in a socially optimal network there must be at least d^* failures following a large shock to any project for all $r > 0$ sufficiently small.³⁴ We show now that all networks $\mathbf{A} \in \mathcal{A}^*$ achieve this bound simultaneously for all i when r is sufficiently small.

The class of networks \mathcal{A}^* is defined by partitioning the banks into groups such that

- (i) $A_{ji} = \frac{R-v}{\varepsilon_S}$ if $G_i = G_j$.
- (ii) $A_{ji} = \frac{R-v}{\varepsilon_L + \beta d^*}$ if $G_i \neq G_j$.
- (iii) $|G_i| = d^*$ for all $i \in N$.

³⁴ The lower bound d^* , given by the positive root of $f(d_i)$ as defined in equation (18), was obtained by minimizing the number of failures when a large shock hits bank i without regard to how many banks fail following a large shock to other banks. Thus, any network can at best achieve the lower bound of d^* failures following a large shock to all banks i while having no failures otherwise, and any such network will be socially optimal.

If a small shock hits a bank i (and no other bank is hit by a small shock at the same time) then j 's equity value will be $R - \underline{v} - A_{ji}\varepsilon_S$. Substituting in the possible values of A_{ji} above, if $j \in G_i$ then this equity value is weakly positive and j does not fail while, if $j \notin G_i$ then this equity value is strictly positive and j does not fail. Thus, there are no failures following a small shock if there are separate shocks. If a large shock hits any bank i then j fails if $j \in G_i$, but if $j \notin G_i$ then j 's equity value is $R - \underline{v} - A_{ji}\varepsilon_L - A_{ji}d^*\beta$, which is weakly positive and so j does not fail. Thus, the lower bound characterized in Lemma 2, is obtained and there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ networks $\mathbf{A} \in \mathcal{A}^*$ are socially optimal.

Step 5: Only d^* -clustered networks obtain the bound. We now show that only networks $\mathbf{A} \in \mathcal{A}^*$ can be socially optimal for r sufficiently small.

By step 2 all socially optimal networks must have no failures following any small shock, and by Lemma 2, for r sufficiently small all socially optimal networks must have exactly d^* failures following every large shock. The proof of Lemma 2 first placed an lower bound on the number of failures for η -doubly stochastic networks of $d^*(\eta)$, and then showed that $\lceil d^*(\eta) \rceil = d^*$ for all η sufficiently small. Obtaining the lower bound $d^*(\eta)$ required inequalities (15), (16) and (17) to all bind. By inequality (16) we must have $\sum_{j \in D_i} A_{ji} = d_i(R(1 + \eta) - \underline{v})/\varepsilon_S$. As there are no failures following a small shock, we also know that $A_{jk} \leq (R(1 + \eta) - \underline{v})/\varepsilon_S$ for all j, k , and so we must have $A_{jk} = (R(1 + \eta) - \underline{v})/\varepsilon_S$, for all $j, k \in D_i$ and for all $i \in N$.

There then exists a $\bar{\eta} > 0$ such that for all $\eta < \bar{\eta}$ the following argument holds: If $j, k \in D_i$ then $A_{jk}\varepsilon_L > R(1 + \eta) - \underline{v}$. This implies that j fails following a large shock to k , and so $j \in D_k$. Hence, $D_k \supseteq D_i$. As $i \in D_i$, we conclude that $i \in D_k$. But then, if $j, i \in D_k$ then $A_{ji}\varepsilon_L > R - \underline{v}$ and j also fails following a large shock to i . Thus, $D_i \subseteq D_k$. Combining set inclusions we conclude that $D_i = D_k$ for all $k \in D_i$. Hence, to achieve the lower bound for all $i \in N$, the set of banks N must be partitioned into disjoint subsets such that when a large shock hits the investment of any bank in the set, all banks in the set default.

As inequality (15) in the proof of Lemma 2 must bind to achieve the bound, the losses collectively absorbed by banks $j \notin D_i$ after a large shock to investment i are $(n - d^*(\eta))(R(1 + \eta) - \underline{v})$. As each of these banks is not in D_i , the most losses any one of them can absorb is $R(1 + \eta) - \underline{v}$ (otherwise they would fail following a large shock to i and would be in D_i). Thus, for all $j \notin D_i$,

$$\left(A_{ji}\varepsilon_L + \sum_{k \in D_i} A_{jk}\beta \right) = R(1 + \eta) - \underline{v}. \tag{19}$$

Similarly, to prevent more than d^* failures when a large shock hits bank $h \in D_h$,

$$\left(A_{jh}\varepsilon_L + \sum_{k \in D_h} A_{jk}\beta \right) = R(1 + \eta) - \underline{v}. \tag{20}$$

Rearranging equation (20), and as $D_i = D_h$ for all $h \in D_i$,

$$A_{jh} = \left(\frac{R(1 + \eta) - \sum_{k \in D_i} A_{jk}\beta - \underline{v}}{\varepsilon_L} \right),$$

for all $j \notin D_i$ and for all $h \in D_i$. Thus, $A_{jk} = A_{jh}$ for all $j \notin D_i$ and for all $k, h \in D_i$, and so

$$A_{jh} = \frac{R(1 + \eta) - \underline{v}}{\varepsilon_L + \beta d^*}, \tag{21}$$

for all $j \notin D_i$ and for all $h \in D_i$.

Combining the conditions we have found, for all η sufficiently small, any network \mathbf{A} that achieves the lower bound of $d^*(\eta)$ failures must satisfy:

- (i) $A_{ij} = \frac{R(1+\eta)-v}{\varepsilon_S}$ if $D_i = D_j$.
- (ii) $A_{ij} = \frac{R(1+\eta)-v}{\varepsilon_L + \beta d^*}$ if $D_i \neq D_j$.
- (iii) $|D_i| = d^*(\eta)$ for all $i \in N$.

As there are, by assumption, no integer problems, $d^* = \lim_{\eta \rightarrow 0} d^*(\eta)$ and n/d^* are integers. Moreover, as argued in the proof of Lemma 2, $d^*(\eta)$ is continuous in η and converges to d^* from below as $\eta \rightarrow 0$. Thus, for all η sufficiently small $\lceil d^*(\eta) \rceil = d^*$.

B.4. Proof of Proposition 2

Let ω_{Lk} and ω_{Sk} denote the states of the world in which a large shock and small shock hit investments of type k respectively. Consider a socially optimal network portfolio pair (A^*, ϕ^*) , and without loss of generality let $\phi_{ii}^* = 1$ for all i . Define the set of network portfolio pairs \mathcal{D} by the pairs consisting of networks A' that can be reached via a feasible trade between i and j from A^* , and portfolios $\phi' = (\phi'_{-i}, \phi'_i)$ such that $\phi'_i \in \Delta^n$. The following four lemmas are helpful.

Lemma 3. For all $(A', \phi') \in \mathcal{D}$ if $\phi'_{ik} > 0$ for $k \neq i$ then

- (i) If $k \in D_i$ and $j \notin D_i$ then all banks $D_i \setminus \{i\}$ fail in state of the world ω_{Sk} , while if $k \in D_i$ and $j \in D_i$ then all banks $D_i \setminus \{i, j\}$ fail and at least one of i and j fail in state of the world ω_{Sk} .
- (ii) If $k \notin D_i$ then all banks $N \setminus \{i, j\}$ fail and i and/or j fail in state of the world ω_{Lk} .
- (iii) If $k \in D_i$ and $A' = A^*$ then all banks $D_i \setminus \{i\}$ fail in state of the world ω_{Sk} .
- (iv) If $k \notin D_i$ and $A' = A^*$ then all banks N fail following a large shock to assets of type k .

Lemma 4. For all $(A', \phi') \in \mathcal{D}$, in state of the world ω_{Lk}

- (i) If $k \notin D_i$ and $k \notin D_j$, then banks D_k fail. If $k \notin D_i$ but $k \in D_j$, then banks $D_j \setminus \{j\}$ fail and j and/or i fail.
- (ii) If $k \in D_i \setminus \{i\}$ and $j \notin D_i$, then banks $D_i \setminus \{i\}$ fail and i and/or j fail.
- (iii) If $k \in D_i \setminus \{i\}$ and $j \in D_i$, then banks $D_i \setminus \{i, j\}$ fail and i and/or j fail.
- (iv) If $k \in D_i \setminus \{i\}$ and $A' = A^*$ banks D_i fail.

Lemma 5. For all $\eta > 0$ and all $(A', \phi') \in \mathcal{D}$ there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ if $v_i((A', \phi')) \geq v_i((A^*, \phi^*))$ and $v_j((A', \phi')) \geq v_j((A^*, \phi^*))$, then $1 - \eta < \sum_m A'_{jm} < 1 + \eta$ and $1 - \eta < \sum_m A'_{im} < 1 + \eta$.

We refer to new counterparty losses as losses incurred in a state of world ω that would not have been incurred in the same state of the world ω for (A^*, ϕ^*) .

Lemma 6. If $\phi'_{ik} > 0$, $k \notin D_i$ and $v_i((A', \phi')) + v_j((A', \phi')) \geq v_i((A^*, \phi^*)) + v_j((A^*, \phi^*))$, then the new expected joint counterparty losses of i and j under (A', ϕ') are at least

$$\frac{2(d^*A_S + (n - 2d^*)A_W)\beta r q}{n}$$

Unilateral Deviations We begin by considering unilateral deviations so that $(A', \phi') = (A^*, \phi')$. Note that

$$\begin{aligned} \mathbb{E}[v_i(A^*, \phi')] &= \bar{R} - r \left(\frac{d^* A_S + (n - d^*) A_W}{n} \right) (q \varepsilon_L + (1 - q) \varepsilon_S) \\ &\quad - \beta \sum_l A_{il}^* \Pr[v_l < \underline{v} | A^*, \phi'] - \sum_k c I_{\phi_{ik} > 0}. \end{aligned}$$

We therefore have that

$$\begin{aligned} \mathbb{E}[v_i(A^*, \phi')] - \mathbb{E}[v_i(A^*, \phi^*)] &\leq \beta \sum_l A_{il}^* \Pr[v_l < \underline{v} | A^*, \phi^*] \\ &\quad - \beta \sum_l A_{il}^* \Pr[v_l < \underline{v} | A^*, \phi'] \end{aligned} \tag{22}$$

The inequality will hold with equality when i 's portfolio ϕ'_i contains only one asset type, and otherwise with strict inequality. Equation (22) show that the profitability of the deviation can be evaluated by considering the expected counterparty losses that i incurs. Following a deviation there must exist an asset type $k \neq i$ such that $\phi_{ik} > 0$. We break this down into two cases, $k \in D_i$ and $k \notin D_i$.

If $k \in D_i$ then by part (iii) of Lemma 3, all banks $D_i \setminus \{i\}$ fail following a small shock to assets of type k . Thus, the deviation generates new expected counterparty losses for i of $\frac{A_S(d^* - 1)\beta r(1 - q)}{n}$, where $A_S(d^* - 1)\beta$ are the losses incurred conditional on a small shock hitting assets of type k , and $r(1 - q)/n$ is the probability of this happening. By parts (i) and (iv) of Lemma 4, as for socially optimal network portfolio pairs, banks D_l fail when a large shock hits a bank $l \neq i$ for any $\phi'_i \in \Delta^n$. Thus, the most that i can save in expected counterparty losses are the losses associated with the failures that occur when a large shock hits assets of type i under (A^*, ϕ^*) . The maximum possible value of these savings to i is $\frac{A_S d^* \beta r q}{n}$. Thus, the deviation is unprofitable because

$$\frac{A_S \beta (d^* - 1) r (1 - q)}{n} > \frac{A_S \beta r q d^*}{n},$$

where the inequality follows from $q < 1/n^2$.

If $k \notin D_i$ then by part (iv) of Lemma 3 all banks N will fail following a large shock to assets of type k . This generates new counterparty losses for i of $\frac{(A_S d^* + A_W(n - 2d^*))\beta r q}{n}$. Again, an upper bound on what can be saved from a deviation is $\frac{A_S \beta r q d^*}{n}$ and so, as

$$\frac{(A_S d^* + A_W(n - 2d^*))\beta r q}{n} > \frac{A_S d^* \beta r q}{n},$$

the deviation is unprofitable.

Bilateral Deviations

We will use v'_i in place of $v_l(A', \phi')$ and v_i^* in place of $v_l(A^*, \phi^*)$ for all l to save on notation. We will show that the overall change in expected values for banks i and j when moving from the network portfolio pair (A^*, ϕ^*) to (A', ϕ') decreases, so that at least one bank must be made worse off by any possible deviation.

$$\mathbb{E}[v'_i + v'_j] - \mathbb{E}[v_i^* + v_j^*]$$

$$\begin{aligned}
 &\leq \mathbb{E} \left[\sum_l \sum_k (A'_{il} + A'_{jl}) \phi'_{lk} p_k \right] - \mathbb{E} \left[\sum_l \sum_k (A^*_{il} + A^*_{jl}) \phi^*_{lk} p_k \right] \\
 &\quad - \sum_l (A'_{il} + A'_{jl}) \beta \Pr[v_l < \underline{v} | A', \phi'] + \sum_l (A^*_{il} + A^*_{jl}) \beta \Pr[v_l < \underline{v} | A^*, \phi^*] \\
 &= \sum_l (A^*_{il} + A^*_{jl}) \beta (\Pr[v_l < \underline{v} | A^*, \phi^*] - \Pr[v_l < \underline{v} | A', \phi']) \tag{23}
 \end{aligned}$$

This first inequality holds with equality when i 's portfolio ϕ'_i contains only one asset type and otherwise with strict inequality. The next equality holds because in all feasible bilateral trade $A'_{il} + A'_{jl} = A^*_{il} + A^*_{jl}$. Inequality (23) shows that for there to be a jointly profitable deviation for banks i and j , they must reduce their joint expected counterparty losses.

CASE 1: $D_i = D_j$

By Lemma 4 the only possible failures that might be prevented when a large shock hits assets of type $k \neq i$, is a single failure of either i or j , and that is only possible for $k \in D_i$. The joint expected value of preventing all such failures to i and j is $2A_S \beta r q (d^* - 1)/n$ —banks i and j jointly save costs of $2A_S \beta$ when a large shock hits a bank $k \in D_i \setminus \{i\}$, and the probability of such a shock hitting is $r q (d^* - 1)/n$. In addition, banks i and j can jointly save at most $2A_S d^* \beta r q/n$ when a large shock hits bank i . This places an upper bound on the possible value of the deviation of

$$\frac{2A_S \beta r q (2d^* - 1)}{n}.$$

We show now that, in fact, in any profitable deviation banks D_i will fail following a large shock to assets of type $k \in D_i \setminus \{i\}$. For either i or j to survive in state of the world ω_{Lk} , its value in this state of the world has to increase compared to its value under (A^*, ϕ^*) . Without loss of generality, suppose that bank i survives in this state of the world (as $j \in D_i$ the argument for j is identical). Then, as $k \in D_i \setminus \{i\}$

$$v_i(A', \phi' | \omega_{Lk}) \geq \underline{v} = R - A^*_{ik} \varepsilon_S.$$

We also know, under the assumption that i avoids failure, that

$$v_i(A', \phi' | \omega_{Lk}) \leq R \sum_m A'_{im} - A'_{ik} \varepsilon_L - \beta \sum_{l \in D_i \setminus \{i\}} A'_{il},$$

where the inequality comes from the lower bound on portfolio maintenance costs of c .

Thus, we must have

$$R \sum_m A'_{im} - A'_{ik} \varepsilon_L - \beta \sum_{l \in D_i \setminus \{i\}} A_{il} \geq R - A^*_{ik} \varepsilon_S.$$

By Lemma 5, for all $\eta > 0$ there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ we have $\sum_m A'_{im} \leq (1 + \eta)$. Hence, for $r < \bar{r}$, we need

$$A_S - A'_{ik} \geq \frac{A_S (\varepsilon_L - \varepsilon_S) - \eta R}{\varepsilon_L}.$$

As $A_S (\varepsilon_L - \varepsilon_S) > 0$, there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ in all profitable bilateral trades, $A_S (\varepsilon_L - \varepsilon_S) > \eta R$. Thus, for all $r < \bar{r}$ we have that $A'_{ik} < A_S$. As the trade must be feasible, this implies that $A'_{jk} > A_S$, and, hence j now fails following a small shock to k . This has a joint expected cost to i and j of at least

$$\frac{2A_S\beta r(1-q)}{n} > \frac{2A_S\beta r q(2d^* - 1)}{n}.$$

Thus there is no profitable deviation in which banks D_i do not all fail when a large shock hits a bank $k \in D_i \setminus \{i\}$.

The only possible remaining gains from a deviation require reducing i and j 's joint counterparty failure costs when a large shock hits i . Suppose then that there are less than d^* failures when a large shock hits assets of type i . We have just shown that if the deviation is profitable and $j \in D_i$, all banks D_i fail when a large shock hits assets of type j . But if $\phi'_{ii} = \phi^*_{ii} = 1$, then a large shock hitting j is equivalent to a large shock hitting i —in this case banks i and j are symmetric. Thus, for all profitable deviations in which $\phi'_{ii} = \phi^*_{ii} = 1$, banks D_i must all fail when a large shock hits bank i . Hence, if $\phi'_{ii} = \phi^*_{ii} = 1$ there is no scope for i and j avoiding any of the joint counterparty failure costs they incur under (A^*, ϕ^*) in a profitable deviation by deviating to a network portfolio pair (A', ϕ^*) , and so all such deviations are unprofitable.

We can therefore restrict attention to $\phi'_{ii} < 1$, which implies that there exists an asset type $k \neq i$ such that $\phi'_{ik} > 0$. Then, by Lemma 3(i), if $k \in D_i$ then, following a small shock to k , banks $D_i \setminus \{i, j\}$ as well as i and/or j fail. This generates new expected joint losses for i and j of at least

$$\frac{2A_S\beta(d^* - 1)r(1-q)}{n} > \frac{2A_Sd^*\beta r q}{n}. \tag{24}$$

The inequality just requires that $q < (d^* - 1)/(2d^* - 1)$ which holds as $q < 1/n^2 < 1/3 \leq (d^* - 1)/(2d^* - 1)$. Thus, the deviation is jointly unprofitable.

Suppose now $\phi'_{ik} > 0$ and $k \notin D_i$. Then, by Lemma 6 there will be expected losses of at least

$$\frac{(2(d^*A_S + (n - 2d^*)A_W)\beta r q)}{n} \geq \frac{2A_S\beta r q d^*}{n}.$$

Thus, the deviation is again jointly unprofitable.

CASE 2: $D_i \neq D_j$

By Lemma 4 parts (i) and (ii), banks i and j jointly incur at least the same losses in state of the world ω_{Lk} for $k \neq i$ following a deviation as they do under (A^*, ϕ^*) . Thus, an upper bound on the expected joint losses that can be avoided by a deviation is $(A_S + A_W)d^*\beta r q/n$ (the expected losses associated with a large shock hitting assets of type i). Moreover, by Lemma 4 part (i) if a large shock hits bank j , then banks $D_i \setminus \{i\}$ will fail and i and/or j will fail. Thus, if $\phi'_{ii} = 1$, then again by symmetry of i and j , banks $D_i \setminus \{i\}$ will fail and i and/or j will fail following a large shock to assets of type i . This generates expected joint losses of $(A_S + A_W)d^*\beta r q/n$ and removes any scope for the deviation being profitable.

We can therefore restrict attention to $\phi'_{ii} \neq 1$. Hence, there exists a $k \neq i$ such that $\phi'_{ik} > 0$. By Lemma 3(i) if $k \in D_i$ then following a small shock to assets of type k at least banks $D_i \setminus \{i\}$ fail. This generates expected losses for i and j of at least

$$\frac{(A_S + A_W)(d^* - 1)\beta r(1-q)}{n} > \frac{(A_S + A_W)d^*\beta r q}{n},$$

and so the deviation is again unprofitable.

If $\phi'_{ik} > 0$ for $k \notin D_i$ then by Lemma 6 there will be expected losses of at least

$$\frac{(2(d^*A_S + (n - 2d^*)A_W)\beta r q)}{n} \geq \frac{2A_S\beta r q d^*}{n}.$$

Thus, the deviation is again jointly unprofitable.

B.5. Proof of Lemma 3

Proof. Part (i) In state of the world ω_{Sk} , an upper bound on the values of banks can then be found by assuming no banks fail (or equivalently, setting $\beta = 0$). Thus, an upper bound on the value of a bank l such that $l \in D_k$ and $l \neq i, j$ in state ω_{Sk} is

$$\begin{aligned} v_l(A', \phi' | \omega = \omega_{Sk}) &\leq \sum_m \sum_h A'_{lm} \phi_{mh} p_h - c \\ &= R - A_S \varepsilon_S - A_S \phi'_{ik} \\ &= \underline{v} - A_S \phi'_{ik} < \underline{v}. \end{aligned} \tag{25}$$

The first inequality holds because we are assuming no banks fail. The first equality is because $R = \bar{R} - c$, and under any feasible trade between i and j , $A'_{lm} = A^*_{lm}$ for all m . The second equality holds by the definition of A_S . The final inequality follows from $\phi'_{ik} > 0$. Thus, all banks l such that $l \in D_k$ and $l \neq i, j$ fail.

There are then two cases to consider. If $j \notin D_i$ then we have already shown that all banks $D_i \setminus \{i\}$ fail. If $j \in D_i$ then an upper bound on the average value of banks i and j is

$$\begin{aligned} \frac{v_i(A', \phi' | \omega = \omega_{Sk}) + v_j(A', \phi' | \omega = \omega_{Sk})}{2} &\leq \sum_m \sum_h \frac{(A'_{im} + A'_{jm}) \phi_{mh} p_h}{2} - c \\ &= R - A_S \varepsilon_S - A_S \phi'_{ik} \varepsilon_S \\ &= \underline{v} - A_S \phi'_{ik} \varepsilon_S < \underline{v} \end{aligned}$$

The first inequality holds because the minimum possible average portfolio maintenance costs are c and because we are assuming no banks fail. The first equality is because $R = \bar{R} - c$, and under any feasible trade between i and j , $A'_{il} + A'_{jl} = A^*_{il} + A^*_{jl}$ for all l . Thus, either bank i fails, or bank j fails or both fail.

Part (ii)

In state of the world ω_{Lk} , an upper bound on the values of banks can then be found by assuming no banks fail (or equivalently, setting $\beta = 0$). Thus, for a bank $l \in D_k$ and $l \neq j$

$$\begin{aligned} v_l(A', \phi' | \omega = \omega_{Lk}) &\leq \sum_m \sum_h A'_{lm} \phi_{mh} p_h - c \\ &= R - A_S \varepsilon_L - A_W \phi'_{ik} \\ &= \underline{v} - A_S (\varepsilon_L - \varepsilon_S) - A_W \phi'_{ik} < \underline{v} \end{aligned} \tag{26}$$

The first equality is because $R = \bar{R} - c$ and under any feasible trade between i and j , $A'_{lm} = A^*_{lm}$ for all m . The second equality holds by the definition of A_S , and the final inequality holds because $\varepsilon_L > \varepsilon_S$ and $\phi'_{ik} > 0$. Thus, all banks in D_k except possibly j fail in state ω_{Lk} after a deviation to the network portfolio pair (A', ϕ') .

There are now two cases to consider. First, suppose that $j \notin D_k$. Then, by the above argument all d^* banks in D_k fail in state of the world ω_{Lk} . As in a feasible bilateral trade $A'_{il} + A'_{jl} = A^*_{il} + A^*_{jl}$ for all l , an upper bound on the average value of banks i and j in state ω_{Lk} is then

$$\begin{aligned} &\frac{v_i(A', \phi' | \omega = \omega_{Lk}) + v_j(A', \phi' | \omega = \omega_{Lk})}{2} \\ &\leq \sum_m \sum_h \frac{(A'_{im} + A'_{jm}) \phi'_{mh} p_h}{2} - \sum_{m \in D_k} \frac{(A'_{im} + A'_{jm}) \beta}{2} - c \end{aligned}$$

$$\begin{aligned}
 &= R - A_W \varepsilon_L - d^* A_W \beta - \left(\frac{A_W + A_S}{2} \right) \phi'_{ik} \varepsilon_L \\
 &= \underline{v} - \left(\frac{A_W + A_S}{2} \right) \phi'_{ik} \varepsilon_L < \underline{v}.
 \end{aligned}$$

Thus, i and/or j must fail.

Suppose now instead that $j \in D_k$. We then know for sure only that banks $D_k \setminus \{j\}$ fail. Nevertheless, bank i or j still fails in state ω_{Lk} . Now we have

$$\begin{aligned}
 &\frac{v_i(A', \phi' | \omega = \omega_{Lk}) + v_j(A', \phi' | \omega = \omega_{Lk})}{2} \\
 &\leq \sum_m \sum_h \frac{(A'_{im} + A'_{jm}) \phi'_{mh} p_h}{2} - \sum_{m \in D_k \setminus \{j\}} \frac{(A'_{im} + A'_{jm}) \beta}{2} - c \\
 &= R - \left(\frac{A_W + A_S}{2} \right) \varepsilon_L - \left(\frac{A_W + A_S}{2} \right) (d^* - 1) \beta \\
 &\quad - \left(\frac{A_W + A_S}{2} \right) \phi'_{ik} \varepsilon_L \\
 &= \underline{v} - \frac{A_S(\varepsilon_L - \varepsilon_S + (d^* - 2)\beta)}{2} - \left(\frac{A_W + A_S}{2} \right) (\beta + \phi'_{ik} \varepsilon_L) \\
 &< \underline{v}
 \end{aligned}$$

The second equality uses the definitions of both A_S and A_W . Thus, again, i and/or j fail.

An upper bound on the value of a bank $l \notin D_k \cup \{i, j\}$ in state ω_{Lk} is then

$$\begin{aligned}
 v_l(A', \phi' | \omega_{Lk}) &\leq R - A_W \varepsilon_L - A_W d^* \beta - A_W \phi'_{ik} \varepsilon_L \\
 &= \underline{v} - A_W \beta - A_W \phi'_{ik} \varepsilon_L < \underline{v}.
 \end{aligned} \tag{27}$$

Thus, all banks $l \neq i, j$ fail, and at least one of i and j fails.

Part (iii)

From part (i) we have already seen that when a large shock hits asset of type k (for which $\phi'_{ik} > 0$), and $k \in D_i \setminus \{i\}$, all banks $D \setminus \{i\}$ fail if $j \notin D_i$. As $A' = A^*$ there is no trade, and so it is without loss of generality to let i 's trade partner be $j \notin D_i$, and we can apply the result from part (i).

Part (iv)

From part (ii) when a large shock hits assets of type k (for which $\phi'_{ik} > 0$), and $k \notin D_i \setminus \{i\}$, all banks except possibly one of i and j fail. However, if $A' = A^*$, then an upper bound on j 's value, which also applies to i 's value, is

$$\begin{aligned}
 v_i(A^*, \phi' | \omega_{Lk}) &\leq R - A_W \varepsilon_L - (A_W(n - d^*) + A_S(d^* - 1))\beta - A_W \phi'_{ik} \varepsilon_L \\
 &= \underline{v} - (A_W(n - 2d^*) + A_S(d^* - 1))\beta - A_W \phi'_{ik} \varepsilon_L < \underline{v}.
 \end{aligned}$$

Thus, if $A' = A^*$, then following the deviation to (A^*, ϕ') with $\phi'_{ik} > 0$, all banks will fail in state of the world ω_{Lk} . \square

B.6. Proof of Lemma 4

Proof. Consider a large shock to bank $k \neq i$. In this case, an upper bound on the value of a bank l such that $l \in D_k$ and $l \neq i, j$ in state ω_{Lk} is

$$\begin{aligned}
 v_l(A', \phi'|\omega_{Lk}) &\leq \sum_m \sum_h A'_{lm} \phi_{mh} p_h - c \\
 &= R - A_S \varepsilon_L \\
 &= \underline{v} - A_S(\varepsilon_L - \varepsilon_S) < \underline{v}.
 \end{aligned}
 \tag{28}$$

The first inequality is obtained by setting $\beta = 0$. The first equality holds because for any feasibly bilateral trade between i and j $A'_{lm} = A^*_{lm}$ for all m . The final equality is from the definition of A_S . Thus, banks other than i and j always fail when a large shock hits a bank in their cluster. Thus, for $k \notin \{D_i \cup D_j\}$ all banks D_k fail in state of the world ω_{Lk} . This proves part (i) for all cases except $D_i \neq D_j$ and $k \in D_j$.

If $j \in D_i$, an upper bound on the average value of banks i and j in state ω_{Lk} for $k \in D_i \setminus \{i\}$ is

$$\begin{aligned}
 \frac{v_i(A', \phi'|\omega_{Lk}) + v_j(A', \phi'|\omega_{Lk})}{2} &\leq \sum_m \sum_h \frac{(A'_{im} + A'_{jm}) \phi_{mh} p_h}{2} - c \\
 &= R - A_S \varepsilon_L \\
 &= \underline{v} - A_S(\varepsilon_L - \varepsilon_S) < \underline{v}.
 \end{aligned}
 \tag{29}$$

The reasoning underlying this series of equations is very similar to above. The only notable difference is that we no longer have $A'_{im} = A^*_{im}$ for all m , but for all feasible bilateral trades between i and j we do have $A'_{im} + A'_{jm} = A^*_{im} + A^*_{jm}$ for all m . Thus, if $j \in D_i$ and $k \in D_i \setminus \{i\}$ at least one of i and j must fail in state of the world ω_{Lk} . Thus, there are at least $d^* - 1$ failures, and these failures include i and/or j . This proves part (iii).

Now suppose that $j \notin D_i$. Then, an upper bound on the average value of banks i and j if large shock hits a bank $k \in D_i \cup D_j$ for $k \neq i$ is

$$\begin{aligned}
 &\frac{v_i(A', \phi'|\omega_{Lk}) + v_j(A', \phi'|\omega_{Lk})}{2} \\
 &\leq \sum_m \sum_h \frac{(A'_{im} + A'_{jm}) \phi_{mh} p_h}{2} - \sum_{l \in D_k \setminus \{i, j\}} \sum_m \frac{(A'_{im} + A'_{jm}) \beta}{2} - c \\
 &= \frac{R - A_S \varepsilon_L - A_S(d^* - 1)\beta}{2} + \frac{R - A_W \varepsilon_L - A_W(d^* - 1)\beta}{2} \\
 &= \underline{v} - \frac{A_W(d^* - 2)\beta - A_S(\varepsilon_L - \varepsilon_S) - (A_S - A_W)(d^* - 1)\beta}{2} < \underline{v}
 \end{aligned}$$

This series of equations follows a similar logic to that of the equations above, but uses the definitions of both A_S and A_W when moving to the last equality. As the average value of i and j is lower than \underline{v} , at least one of i and j must again fail. Thus, there are at least d^* failures, and these failures include i and/or j . This completes the proof of part (i) and also proves part (ii).

Finally, suppose that $A' = A^*$ and consider a large shock to a bank $k \in D_i \setminus \{i\}$. For all banks $l \in D_i$

$$\begin{aligned}
 v_l(A^*, \phi'|\omega_{Lk}) &\leq R - A_S \varepsilon_L \\
 &= \underline{v} - A_S(\varepsilon_L - \varepsilon_S) < \underline{v}.
 \end{aligned}$$

Thus, all banks in D_i fail. \square

B.7. Proof of Lemma 5

Proof. In the state of the world in which no shock hits any asset types, under the network portfolio pair (A^*, ϕ^*) bank i 's value is equal to R . A lower bound on bank i 's expected value under the network portfolio pair (A^*, ϕ^*) is therefore $R - r\ell$, where ℓ is the (finite) maximum losses that bank i incurs in any state of the world.

For a network portfolio pair $(A', \phi') \in \mathcal{D}$ an upper bound on the expected value of bank i is $\sum_m A'_{im} R$. Thus, a necessary condition for the deviation to the network portfolio pair $(A', \phi') \in \mathcal{D}$ to increase i 's expected value (i.e., for $v_i((A', \phi')) > v_i((A^*, \phi^*))$) is that

$$\sum_m A'_{im} R \geq R - r\ell.$$

Fix any $\eta > 0$, then $\sum_m A'_{im} \geq 1 - r\ell/R$. Set $\bar{r} = \eta R/\ell$. Note that $\bar{r} > 0$ because ℓ is finite and $R > 0$. Thus, for $r < \bar{r}$ we have $\sum_m A'_{im} \geq 1 - \eta$.

Repeating the exercise for j , we conclude that $\sum_m A'_{jm} \geq 1 - \eta$. Moreover, as the bilateral trade that results in A' is feasible $\sum_m (A'_{jm} + A'_{im}) = \sum_m (A^*_{jm} + A^*_{im}) = 2$. Hence, $\sum_m A'_{jm} = 2 - \sum_m A'_{im}$ and we have $2 - \sum_m A'_{im} \geq 1 - \eta$, which can be rearranged to give $\sum_m A'_{im} \leq 1 + \eta$. By a symmetric argument we also have $\sum_m A'_{jm} \leq 1 + \eta$. \square

B.8. Proof of Lemma 6

Proof. As $\phi'_{ik} > 0$ and $k \notin D_i$, by Lemma 3(ii) there will be at least $n - 1$ failures following a large shock to bank k , and only bank i or j might survive. If, instead, all n banks failed, there would be new expected joint counterparty losses of

$$\frac{2(d^* A_S + (n - 2d^*) A_W) \beta r q}{n} \tag{30}$$

It only remains to show that for any profitable deviation with $\phi'_{ik} > 0$ and $k \notin D_i$, then either (i) all n banks fail in state of the world ω_{Lk} or (ii) if only $n - 1$ banks fail in state of the world ω_{Lk} , then there is at least one new failure after a large shock to some asset $l \neq k$.

We will first show that there is no jointly profitable deviation in which bank j survives in state of the world ω_{Lk} , without incurring sufficiently many new expected joint counterparty losses. For j to survive, j must have a value of $v_j(A', \phi' | \omega_{Lk}) \geq \underline{v}$. However, by Lemma 5 and Lemma 3(i)

$$v_j(A', \phi' | \omega_{Lk}) \leq R(1 + \eta) - A'_{ji} \phi'_{ik} \varepsilon_L - A'_{jk} \varepsilon_L - (1 + \eta - A'_{jj}) \beta.$$

To show that this implies $v_j(A', \phi' | \omega_{Lk}) < \underline{v}$ we will put a lower bound on j 's losses in this state of the world by showing that, for the deviation to be jointly profitable, we must have $A'_{jl} \leq A^*_{jl}$ for all $l \in D_j$ and $\sum_{m \in D_l} A'_{jm}/d^* \leq A_W$ for all $l \notin D_j \cup \{i\}$ and either $\sum_{m \in D_i} A'_{jm}/d^* \leq A_W$ or else $A'_{jm} \leq A_W$ for all $m \in D_i \setminus \{i\}$.

First, towards a contradiction, suppose that in a jointly profitable deviation $A'_{jl} > A^*_{jl}$, for $l \in D_j \setminus \{i\}$. Then, when a small shock hits assets of type l ,

$$\begin{aligned} v_j(A', \phi' | \omega_{Sl}) &\leq (1 + \eta)R - A^*_{jl} \varepsilon_S - (A'_{jl} - A^*_{jl}) \varepsilon_S \\ &= \underline{v} + \eta R - (A'_{jl} - A^*_{jl}) \varepsilon_S \end{aligned}$$

By Lemma 5, for r sufficiently small $R\eta < (A'_{jl} - A^*_{jl}) \varepsilon_S$ and, hence, $v_j(A', \phi' | \omega_{Sl}) < \underline{v}$. This means that j fails following a small shock to assets of type l . This generates additional expected new joint counterparty losses for i and j of at least $(A_S + A_W) \beta r (1 - q)/n$. By equation

(22) the deviation is profitable only if it reduces expected joint counterparty losses. The maximum possible reduction in expected joint counterparty losses from the trade is $2A_S\beta r q/n < 2(A_S + A_W)\beta r q/n < (A_S + A_W)\beta r(1 - q)/n$. Thus, the deviation is unprofitable, which is a contradiction. We must therefore have $A'_{jl} \leq A^*_{jl}$ for all $l \in D_j \setminus \{i\}$.

Now consider a large shock to assets of type l for $l \notin \{D_i \cup D_j\}$. By Lemma 4 part (i) banks D_l fail. Thus,

$$\begin{aligned} v_j(A', \phi' | \omega_{Ll}) &\leq (1 + \eta)R - A'_{jl}\varepsilon_L - \sum_{m \in D_l} A'_{jm}\beta \\ &= \underline{v} + \eta R - (A'_{jl} - A^*_{jl})\varepsilon_L - \sum_{m \in D_l} (A'_{jm} - A^*_{jm})\beta \end{aligned}$$

Let $\ell \in \operatorname{argmax}_{l \notin \{D_j \cup D_i\}} \sum_{m \in D_l} A'_{jm}$ and let $h \in \operatorname{argmax}_{l \in D_\ell} A'_{jl}$. If $\sum_{m \in D_\ell} A'_{jm}/d^* > A_W$ or $\sum_{m \in D_\ell} A'_{jm}/d^* = A_W$ and $A'_{jh} > A_W$, then $v_j(A', \phi' | \omega_{Lh}) < \underline{v}$. If $h = k$, this implies that j fails in state ω_{Lk} , which is a contradiction. We can therefore restrict attention to $h \neq k$. We have already seen that banks D_h and bank j fail in state ω_{Lh} . Further, all banks $m \notin D_h \cup \{i, j\}$ then have values of at most

$$v_m(A', \phi' | \omega_{Lh}) \leq R - A_W\varepsilon_L - A_W(d^* + 1)\beta < \underline{v},$$

and so also fail. Thus, there are at least $n - 1$ failures in state of the world ω_{Lh} , and this generates additional new counterparty losses for i and j of

$$\frac{2((d^* - 1)A_S + (n - d^*)A_W)\beta r q}{n} > \frac{2(A_S\beta r q)}{n}.$$

Thus, the deviation is unprofitable. Hence, we must have that for all $l \notin \{D_i \cup D_j\}$, $A'_{jl} \leq A_W$.

Now consider a large shock to assets of type $h \in \operatorname{argmax}_{l \in D_i \setminus \{i\}} A'_{jl}$. By Lemma 4 part (ii), all banks $D_i \setminus \{i\}$ fail, as do at least one of i and j . If bank i fails, then for j not to also fail we need that either $\sum_{l \in D_i} A'_{jl}/d^* \leq A_W$, or if $\sum_{l \in D_i} A'_{jl}/d^* > A_W$ then $A'_{jh} \leq A_W$. If j does fail, then by the same logic as above, there will be a cascade of failures in which all banks fail, exceeding the bound on expected new counterparty losses (30). Suppose then that j fails but not i in state ω_{Lh} . As j fails, along with banks $D_i \setminus \{i\}$, all banks $m \in D_j \setminus \{j\}$ have value

$$v_m(A', \phi' | \omega_{Ll}) \leq R - A_W\varepsilon_L - (A_W(d^* - 1) + A_S)\beta < \underline{v},$$

and so also fails. But then for any bank $h \notin \{D_i \cup D_j\}$

$$v_h(A', \phi' | \omega_{Ll}) \leq R - A_W\varepsilon_L - A_W(2d^* - 1)\beta < \underline{v},$$

and so all these banks also fail.

Combining the above possibilities, we must have that bank i fails in state ω_{Lk} and either $\sum_{l \in D_i} A'_{jl}/d^* \leq A_W$, or if $\sum_{l \in D_i} A'_{jl}/d^* > A_W$ then $A'_{jh} \leq A_W$.

Recall that for j not to fail when a large shock hits k we need

$$R(1 + \eta) - A'_{ji}\phi'_{ik}\varepsilon_L - A'_{jk}\varepsilon_L - (1 + \eta - A'_{jj})\beta \geq R - A_W\varepsilon_L - d^*A_W\beta. \tag{31}$$

By Lemma 5, η can be made arbitrarily small by requiring r to be sufficiently small. Thus, the only possibility for satisfying this inequality is if $A'_{jk} < A_W$, while $\sum_l A'_{jl} = 1 + \eta$. However, as shown above, we must have $A'_{jl} \leq A^*_{jl}$ for any $l \in D_j$, or any $l \notin \{D_i \cup D_j\}$. Moreover, if

$\sum_{l \in D_i} A'_{jl}/d^* \leq A^*_{jl}$ then there is no remaining scope for reducing A'_{jk} . Thus, we must have $\sum_{l \in D_i} A'_{jl}/d^* > A^*_{jl}$, which implies that $A'_{jl} \leq A^*_{jl}$ for all $l \in D_i \setminus \{i\}$. Thus, the only possibility for reducing A'_{jk} by an amount x , while $\sum_l A'_{jl} = 1 + \eta$, is to set $A'_{ji} = A^*_{ji} + \eta + x$. The inequality becomes increasingly slack as x increases, so the best chance of satisfying it is to set $x = A^*_{jk}$. Even in this case, we will see that inequality (31) is violated. We then have

$$v_j(A', \phi' | \omega_{Lk}) \leq R(1 + \eta) - (A^*_{ji} + A^*_{jk})\phi'_{ik}\varepsilon_L - (1 + \eta - A_S)\beta \tag{32}$$

$$\leq R(1 + \eta) - 2A_W\phi'_{ik}\varepsilon_L - (1 + \eta - A_S)\beta \tag{33}$$

Thus, as $d^*A_S + (n - d^*)A_W = 1$, for $v_j(A', \phi' | \omega_{Lk}) \geq \underline{v}$ we need

$$\phi'_{ik} \leq \frac{1}{2} - \frac{A_S(d^* - 1)\beta + (n - 2d^*)A_W\beta - (R - \beta)\eta}{2A_W\varepsilon_L}.$$

Thus, by Lemma 5, there exists a $\bar{r} > 0$ such that for all $r < \bar{r}$ we require $\phi'_{ik} < 1/2$.

To realize the potential gains from the deviation that can make it profitable, we need to avoid at least one bank in the set $D_i \setminus \{i\}$ failing when a large shock hits bank i . This requires that $A_S\phi'_{ii}\varepsilon_L \leq A_S\varepsilon_S$, which requires $\phi'_{ii} \leq \frac{\varepsilon_S}{\varepsilon_L} \leq 1/2$.

If ϕ'_{ii} and ϕ'_{ik} are both less than half, then bank i must hold a third asset q . But by Lemma 4 part (ii), there are $n - 1$ failures when a large shock hits q , so we have exceeded the bound on new expected joint counterparty losses (30). Thus, there is no profitable deviation without either (i) n banks failing in state of the world ω_{Lk} or (ii) $n - 1$ banks failing in state of the world ω_{Lk} and an extra failure after a large shock hits some other asset $l \neq k$.

The argument showing there is no profitable deviation in which i does not fail when a large shock hits k follows the same steps, and the full details are omitted for brevity. First, to prevent a large cascade of at least $n - 1$ additional failures that would make the deviation unprofitable, banks i 's average dependence of a set of banks $D_l \neq D_i$ cannot increase. Second, to prevent additional failures when a small shock hits a bank in D_i that would make the deviation unprofitable, either i 's average dependency on banks D_i must not increase, or else $A'_{il} \leq A_S$ for all $l \in D_i \setminus \{i\}$. If i 's average dependency on banks D_i does not increase, then there is no scope for having $A'_{ik} < A_W$, and thus i fails in state ω_{Lk} as claimed. If, instead, $A'_{il} \leq A_S$ for all $l \in D_i \setminus \{i\}$, then the only way to have $A'_{ik} < A_W$ is to increase A'_{ii} . However, if this is done by enough to prevent i from failing in state ω_{Lk} , then at least banks $D_i \setminus \{i\}$ fail in state of the world ω_{Li} , and the deviation is again jointly unprofitable. \square

B.9. Proof of Proposition 3

Proof. Part i. One can easily verify that the network portfolio pair (A^*, ϕ^*) satisfies the conditions of Lemma 1. Investment k in Lemma 1 corresponds to investment i in the this proof. Investment l in Lemma 1 is the investment corresponding to any bank in i 's cluster. Thus, by that Lemma, there exists a change in bank i 's portfolio that is profitable.

Part ii. Consider now a deviation from (A^*, ϕ^*) to (A', ϕ^*) , reached via a bilateral trade between i and $j \notin D_i$. Specifically, let $A'_{ik} = A_S + A_W$ for all $k \in D_i$, $A'_{ik} = 0$ for all $k \in D_j$, $A'_{jk} = A_S + A_W$ for all $k \in D_j$, $A'_{jk} = 0$ for all $k \in D_i$, and otherwise set $A'_{ik} = A^*_{ik}$ and $A'_{jk} = A^*_{jk}$. It is easy to verify that this trade is feasible.

We show now that this trade is strictly profitable for i . Strict profitability for j then follows by symmetry. First, observe that the value of all banks remains the same in states of the world ω_{Sl} and ω_{Li} for $l \notin \{D_i \cup D_j\}$.

Next, observe that in states of the world ω_{Sk} and ω_{Lk} for $k \in D_i$ we have under the socially optimal network portfolio pair that $\pi_i(A^*, \phi^*|\omega_{Lk}) = \pi_i(A^*, \phi^*|\omega_{Sk}) = 0$. Thus, by limited liability, $\pi_i(A', \phi^*|\omega_{Lk}) \geq \pi_i(A^*, \phi^*|\omega_{Lk})$ and $\pi_i(A', \phi^*|\omega_{Sk}) \geq \pi_i(A^*, \phi^*|\omega_{Sk}) = 0$.

Moreover, under the socially optimal network portfolio pair, in state of the world ω_{Lk} for $k \in D_j$, $\pi_i(A^*, \phi^*|\omega_{Lk}) = 0$. This follows from the definition of A_W . Thus, by limited liability, we must have $\pi_i(A', \phi^*|\omega_{Lk}) \geq 0$.

The only remaining states of the world to consider are ω_{Sk} for $k \in D_j$. Following the trade,

$$v_j(A', \phi^*|\omega_{Sk}) \leq v_j(A^*, \phi^*|\omega_{Sk}) - A_W \varepsilon_S = \underline{v} - A_W \varepsilon_S < \underline{v}$$

so bank j will now fail. This implies that for banks $l \in D_j \setminus \{j\}$

$$v_l(A', \phi^*|\omega_{Sk}) \leq v_l(A^*, \phi^*|\omega_{Sk}) - A_S \beta = \underline{v} - A_S \beta < \underline{v},$$

and so the failure of j precipitates a cascade of failures in which all banks $D_k = D_j$ fail. However, banks $l \notin D_j$ do not fail. To see that these banks survive, note that if the set of banks D_j , and only the banks D_j fail, we have

$$v_l(A', \phi^*|\omega_{Sk}) \geq R - A_W \varepsilon_S - d^* A_W \beta = \underline{v} + A_W (\varepsilon_L - \varepsilon_S) > \underline{v},$$

where the first inequality is tight for all banks $l \notin D_j \setminus \{i\}$ and the second equality follows from the definition of A_W .

Consider then the value of bank i in this state of the world. We have seen that in states of the world ω_{Sk} for $k \in D_j$, banks D_j fail given (A', ϕ^*) , while no banks failed for (A^*, ϕ^*) . However, because of the trade, none of the losses from these failures are incident on i —that is, $A'_{ik} = 0$ for all $k \in D_j$. Thus, we have

$$v_i(A', \phi^*|\omega_{Sk}) = R > R - A_W \varepsilon_S = v_i(A^*, \phi^*|\omega_{Sk}).$$

Thus, the deviation from (A^*, ϕ^*) to (A', ϕ^*) is strictly profitable for i , and by symmetry, strictly profitable for j . \square

B.10. Proof of Remark 2

We have

$$\begin{aligned} & \mathbb{E}[\pi_i] \\ &= \sum_{\omega \in \Omega} \psi(\omega) \left(\sum_{j \in N} A_{ij} (p_j(\omega) - \beta I_{v_j(\omega) < \underline{v}}) - \underline{v} \right) (1 - I_{v_i(\omega) < \underline{v}}) \\ &= \sum_{j \in N} A_{ij} [\mathbb{E}[p_j | v_i \geq \underline{v}] \Pr(v_i \geq \underline{v}) - \beta (\Pr(v_j < \underline{v}) - \mathbb{E}[I_{v_j < \underline{v}} I_{v_i < \underline{v}}])] - \underline{v} \Pr(v_i \geq \underline{v}) \\ &= \sum_{j \in N} A_{ij} [\mathbb{E}[p_j | v_i \geq \underline{v}] \Pr(v_i \geq \underline{v}) \\ &\quad - \beta (\Pr(v_j < \underline{v})(1 - \Pr(v_i < \underline{v})) - \text{Cov}[I_{v_j < \underline{v}}, I_{v_i < \underline{v}}])] \\ &\quad - \underline{v} \Pr(v_i \geq \underline{v}) \\ &= \Pr(v_i \geq \underline{v}) \left(\sum_{j \in N} A_{ij} [\mathbb{E}[p_j | v_i \geq \underline{v}] - \beta \Pr(v_j < \underline{v})] - \underline{v} \right) \end{aligned}$$

$$+ \beta \sum_{j \in N} A_{ij} \text{Cov} [I_{v_j < \underline{v}}, I_{v_i < \underline{v}}]. \quad \square$$

B.11. Proof of Lemma 1

We prove the result constructively by finding a unilateral portfolio deviation that is always profitable under the maintained assumptions of Lemma 1. Consider the deviation by bank i from portfolio ϕ_i to ϕ'_i in which i sets $\phi'_{ij} = \phi_{ij}$ for all $j \neq k, l$, $\phi'_{ik} = 0$ and $\phi'_{il} = \phi_{il} + \phi_{ik}$. Since the deviation is unilateral, we have $\phi'_j = \phi_j$ for all $j \neq i$. We let ω_{Lk} and ω_{Sk} denote the states of the world in which a large and small shock, respectively, hits assets of type k .

The equity value of all banks is weakly greater in states of the world ω_{Ln} and ω_{Sn} for $n \neq l, m$, and the state of the world in which no shock hits. Firm i pays the same or fewer portfolio maintenance costs, since i is no longer paying the maintenance cost on asset k . Thus, in these states, p_j is weakly higher for all banks j , hence all banks' equity values weakly increase.

Therefore, a sufficient condition for i 's equity value to increase is

$$\begin{aligned} E [\pi_i (A, \phi')] &> E [\pi_i (A, \phi)] \\ \implies (1 - q) &([\pi (A, \phi'|\omega_{Sk}) - \pi (A, \phi|\omega_{Sk})] + [\pi (A, \phi'|\omega_{Sl}) - \pi (A, \phi|\omega_{Sl})]) \\ + q &([\pi (A, \phi'|\omega_{Lk}) - \pi (A, \phi|\omega_{Lk})] + [\pi (A, \phi'|\omega_{Ll}) - \pi (A, \phi|\omega_{Ll})]) > 0 \end{aligned} \quad (34)$$

that is, expected profits, conditional on large or small shocks hitting k or l , must increase after the deviation. We will calculate the value of each square bracketed term to show that inequality (34) does indeed have to be satisfied in our constructed trade.

We have

$$\pi (A, \phi'|\omega_{Sk}) - \pi (A, \phi|\omega_{Sk}) \geq A_{ii} \phi_{ik} \varepsilon_S. \quad (35)$$

After the deviation to ϕ' , bank i weakly lowers its portfolio maintenance costs, since it no longer pays a portfolio maintenance cost for holding asset k . Bank i 's market value increases directly by $A_{ii} \phi_{ik} \varepsilon_S$, since bank i no longer has a portfolio share in asset k . Bank i 's market value in this state may increase further, since after the deviation, no banks that do not fail under ϕ fail under ϕ' when a small shock hits k ; but some banks that fail under ϕ may no longer fail under ϕ' . Thus, fewer failure costs are deducted from i 's market value in this state.

We have

$$\pi (A, \phi'|\omega_{Sl}) - \pi (A, \phi|\omega_{Sl}) > -(1 - q) A_{ii} \phi_{ik} \varepsilon_S. \quad (36)$$

We have $\pi (A, \phi'|\omega_{Sl}) \geq 0$ by limited liability. We have $\pi (A, \phi|\omega_{Sl}) < (1 - q) A_{ii} \phi_{ik} \varepsilon_S$ according to assumption (ii) of the Lemma.

We have

$$\pi (A, \phi'|\omega_{Lk}) - \pi (A, \phi|\omega_{Lk}) \geq 0. \quad (37)$$

We know that $\pi (A, \phi'|\omega_{Lk})$ cannot be strictly lower than $\pi (A, \phi|\omega_{Lk})$, since i is less exposed to the large shock to investment type k after the deviation. But i could fail even after the deviation, in which case equation (37) equals zero.

We have

$$\pi (A, \phi'|\omega_{Ll}) - \pi (A, \phi|\omega_{Ll}) > -(1 - q) A_{ii} \phi_{ik} \varepsilon_S. \quad (38)$$

We have $\pi(A, \phi'|\omega_{LI}) \geq 0$ by limited liability. We have $\pi(A, \phi|\omega_{LI}) < (1 - q) A_{ii} \phi_{ik} \varepsilon_S$ according to assumption (ii) of the Lemma.

Summing and rearranging inequalities (35)-(38) yields

$$\begin{aligned} & (1 - q) \left([\pi(A, \phi'|\omega_{Sk}) - \pi(A, \phi|\omega_{Sk})] + [\pi(A, \phi'|\omega_{Sl}) - \pi(A, \phi|\omega_{Sl})] \right) \\ & + q \left([\pi(A, \phi'|\omega_{Lk}) - \pi(A, \phi|\omega_{Lk})] + [\pi(A, \phi'|\omega_{LI}) - \pi(A, \phi|\omega_{LI})] \right) \\ & > (1 - q) (A_{ii} \phi_{ik} \varepsilon_S - (1 - q) A_{ii} \phi_{ik} \varepsilon_S) - q (1 - q) A_{ii} \phi_{ik} \varepsilon_S \\ & = (1 - q) (A_{ii} \phi_{ik} \varepsilon_S - A_{ii} \phi_{ik} \varepsilon_S) = 0 \end{aligned}$$

so inequality (34) is satisfied and the constructed deviation is profitable.

Appendix C. Supplementary material

Supplementary material related to this article can be found online at <https://doi.org/10.1016/j.jet.2020.105157>.

References

- Abbassi, P., Brownlees, C., Hans, C., Podlich, N., 2017. Credit risk interconnectedness: what does the market really know? *J. Financ. Stab.* 29, 1–12.
- Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2015a. Systemic risk and stability in financial networks. *Am. Econ. Rev.* 105, 564–608.
- Acemoglu, D., Ozdaglar, A.E., Tahbaz-Salehi, A., 2015b. Systemic risk in endogenous financial networks, available at SSRN 2553900.
- Acharya, V.V., 2009. A theory of systemic risk and design of prudential bank regulation. *J. Financ. Stab.* 5, 224–255.
- Adrian, T., Brunnermeier, M., 2016. CoVaR. *Am. Econ. Rev.* 106, 1705–1741.
- Allen, F., Babus, A., 2009. Networks in finance. In: *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World*, p. 367.
- Allen, F., Babus, A., Carletti, E., 2012. Asset commonality, debt maturity and systemic risk. *J. Financ. Econ.* 104, 519–534.
- Allen, F., Gale, D., 2000. Financial contagion. *J. Polit. Econ.* 108, 1–33.
- Alvarez, F., Barlevy, G., 2014. Mandatory disclosure and financial contagion. Tech. rep., Working Paper. Federal Reserve Bank of Chicago.
- Anderson, H., Erol, S., Ordoñez, G., 2020. Interbank Networks in the Shadows of the Federal Reserve Act, available at SSRN 3367203.
- Anderson, H., Paddrik, M., Wang, J.J., 2019. Bank networks and systemic risk: evidence from the national banking acts. *Am. Econ. Rev.* 109, 3125–3161.
- Babus, A., 2016. The formation of financial networks. *Rand J. Econ.* 47, 239–272.
- Basel Committee on Banking Supervision, 2011. Basel III: A global regulatory framework for more resilient banks and banking systems. Technical Note. Bank for International Settlements.
- Bernard, B., Capponi, A., Stiglitz, J.E., 2017. Bail-ins and bail-outs: Incentives, connectivity, and systemic stability. Tech. rep. National Bureau of Economic Research.
- BIS, 2015. Semiannual OTC derivatives statistics. Tech. rep. Bank for International Settlements.
- BIS, 2018. BIS Quarterly Review, September 2018. Tech. rep. Bank for International Settlements.
- Blume, L., Easley, D., Kleinberg, J., Kleinberg, R., Tardos, É., 2011. Network formation in the presence of contagious risk. In: *Proceedings of the 12th ACM Conference on Electronic Commerce*. ACM, pp. 1–10.
- BoE, 2018. Financial Policy Committee statement from its meeting - 3 October 2018. Tech. rep. Bank of England.
- Caballero, R.J., Simsek, A., 2013. Fire sales in a model of complexity. *J. Finance* 68, 2549–2587.
- Cabrales, A., Gale, D., Gottardi, P., 2015. Financial Contagion in Networks. *Oxford Handbook on the Economics of Networks*. Forthcoming.
- Cabrales, A., Gottardi, P., Vega-Redondo, F., 2017. Risk sharing and contagion in networks. *Rev. Financ. Stud.* 30, 3086–3127.

- Capponi, A., 2016. Systemic risk, policies, and data needs. In: *Optimization Challenges in Complex, Networked and Risky Systems*, INFORMS, pp. 185–206.
- Chang, B., Zhang, S., 2018. Endogenous market making and network formation, available at SSRN 2600242.
- Cifuentes, R., Ferrucci, G., Shin, H.S., 2005. Liquidity risk and contagion. *J. Eur. Econ. Assoc.* 3, 556–566.
- Cohen-Cole, E., Patacchini, E., Zenou, Y., 2015. Static and dynamic networks in interbank markets. *Netw. Sci.*, 98–123.
- Craig, B., Ma, Y., 2018. Intermediation in the Interbank Lending Market. Columbia University. mimeo.
- Craig, B., von Peter, G., 2014. *J. Financ. Intermed.* 23, 322–347.
- Denbee, E., Julliard, C., Li, Y., Yuan, K., 2014. Network Risk and Key Players: A Structural Analysis of Interbank Liquidity.
- Di Maggio, M., Tahbaz-Salehi, A., 2014. Financial intermediation networks. In: *Columbia Business School Research Paper*.
- Eisenberg, L., Noe, T.H., 2001. Systemic risk in financial systems. *Manag. Sci.* 47, 236–249.
- Eisert, T., Eufinger, C., 2018. Interbank networks and backdoor bailouts: benefiting from other Banks' Government guarantees. *Manag. Sci.* In press.
- Elliott, M., Golub, B., Jackson, M.O., 2014. Financial networks and contagion. *Am. Econ. Rev.* 104, 3115–3153.
- Erol, S., 2015. Network hazard and bailouts. Tech. rep., working paper.
- Erol, S., Ordoñez, G., 2017. Network reactions to banking regulations. *J. Monet. Econ.* 89, 51–67.
- Erol, S., Vohra, R., 2014. Network Formation and Systemic Risk, Second Version. PIER Working Paper.
- Farboodi, M., 2014. Intermediation and voluntary exposure to counterparty risk. Tech. rep., Working Paper. University of Chicago.
- Farhi, E., Tirole, J., 2012. Collective moral hazard, maturity mismatch, and systemic bailouts. *Am. Econ. Rev.* 102, 60–93.
- FOMC, 2018. Minutes of the Federal Open Market Committee. Tech. rep. Federal Reserve Bank.
- Freixas, X., Parigi, B.M., Rochet, J.-C., 2000. Systemic risk, interbank relations, and liquidity provision by the central bank. *J. Money Credit Bank.*, 611–638.
- Furfine, C., 2003. *J. Money Credit Bank.* 35, 111–128.
- Gai, P., Kapadia, S., 2010. Contagion in financial networks. *Proc. R. Soc. A, Math. Phys. Eng. Sci.*, rspa20090410.
- Galeotti, A., Ghiglino, C., Goyal, S., 2015. Financial Linkages, Portfolio Choice and Systemic Risk.
- Galeotti, A., Golub, B., Goyal, S., 2017. Targeting interventions in networks. preprint. arXiv:1710.06026.
- Glasserman, P., Young, H.P., 2015. How likely is contagion in financial networks? *J. Bank. Finance* 50, 383–399.
- Gofman, M., 2017. Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions. *J. Financ. Econ.* 124, 113–146.
- IMF, 2018. Global Financial Stability Report, October 2018. Tech. rep. International Monetary Fund.
- Iyer, R., Peydro, J.-L., 2011. Interbank contagion at work: evidence from a natural experiment. *Rev. Financ. Stud.* 24, 1337–1377.
- Jackson, M.O., Pernoud, A., 2019. What makes financial networks special? Distorted investment incentives, regulation, and systemic risk measurement, distorted investment incentives. In: *Regulation, and Systemic Risk Measurement*.
- Jensen, M.C., Meckling, W.H., 1976. Theory of the firm: managerial behavior, agency costs and ownership structure. *J. Financ. Econ.* 3, 305–360.
- Kanik, Z., 2017. Can Rescues by Banks Replace Costly Bail-Outs in Financial Networks? Available at SSRN 3057615.
- Leitner, Y., 2005. Financial networks: contagion, commitment, and private sector bailouts. *J. Finance* 60, 2925–2953.
- McPherson, M., Smith-Lovin, L., Cook, J.M., 2001. Birds of a feather: homophily in social networks. *Annu. Rev. Sociol.* 27, 415–444.
- Rochet, J.-C., Tirole, J., 1996. Interbank lending and systemic risk. *J. Money Credit Bank.* 28, 733–762.
- Rogers, L.C., Veraart, L.A., 2013. Failure and rescue in an interbank network. *Manag. Sci.* 59, 882–898.
- Stanton, R., Walden, J., Wallace, N., 2018. Mortgage loan flow networks and financial norms. *Rev. Financ. Stud.* 31, 3595–3642.
- Upper, C., Worms, A., 2004. Estimating bilateral exposures in the German interbank market: is there a danger of contagion? *Eur. Econ. Rev.* 48, 827–849.
- Wang, J.J., 2014. Distress Dispersion and Systemic Risk in Networks, available at SSRN.
- Zawadowski, A., 2013. Entangled financial systems. *Rev. Financ. Stud.*