

Illiquidity Spirals in Coupled Over-the-Counter Markets *

Christoph Aymanns[†]
University of St. Gallen

Co-Pierre Georg[‡]
University of Cape Town
Deutsche Bundesbank

Benjamin Golub[§]
Harvard

September 6, 2018

Abstract

We model intermediaries trading economically coupled assets, each asset in its own over-the-counter market—e.g., secured debt and the underlying collateral. Incentives to provide liquidity in one market are increasing in counterparties' activity in both markets. The intermediaries' activity is thus the outcome of a game of strategic complements on two coupled trading networks. We model a crisis as an exogenous change to network structure, as well as the exogenous exit of some intermediaries. This causes an *illiquidity spiral* across the two networks. We find that in coupled networks, in contrast to uncoupled ones, illiquidity spirals can be so severe that liquidity vanishes *discontinuously* as we vary the shock. Liquidity can be improved if one of the two OTC markets is replaced by an exchange, or if the two OTC markets have more links in common.

Keywords: market liquidity, funding liquidity, over-the-counter markets

JEL Classification: G21, G23, D85

*For very helpful comments, we thank our discussants Gilles Chemla, Adam Copeland, Maryam Farboodi, Bernard Herskovic, and John Nash; also (in random order) Semyon Malamud, Andrei Shleifer, Yu An, Rama Cont, Itay Goldstein, Darrell Duffie, Elizabeth Leicht, Ben Hambly, Robin Greenwood, Chaojun Wang, Martin Oehmke, Max Bruche, David Scharfstein, Selman Erol, Stephen Morris, Ana Babus, and John Campbell. We are also grateful to seminar participants (Copenhagen Business School, Imperial College Business School, HEC Paris, Columbia University, the University of Toronto Fields Institute, the Bank of Canada, Deutsche Bundesbank, UPE, Oxford), and conference participants (FIRS 2017, 2017 Warwick Business School Frontiers of Finance Conference, 2017 Cowles Conference on General Equilibrium and its Applications, Conference on Network Science and Economics at the Washington University in St. Louis, 2016 ECB-IMF Workshop on Money Markets, Monetary Policy Implementation and Market Infrastructures, UECE Lisbon Meetings in Game Theory and Applications 2015). The views expressed in this paper do not necessarily reflect the views of Deutsche Bundesbank. Christoph Aymanns would like to acknowledge generous support from the LSE Systemic Risk Centre.

[†]E-Mail: christoph.aymanns@gmail.com

[‡]E-Mail: co-pierre.georg@bundesbank.de

[§]E-Mail: bgolub@fas.harvard.edu

1 Introduction

Many important financial markets are *over-the-counter* (OTC): market participants trade with one another in decentralized transactions, rather than through an exchange. Due to search frictions, traders cannot reoptimize their counterparties instantly, and so trading relationships are persistent.¹ As a result, the performance of the market—most importantly, how successfully it facilitates access to liquidity for its users—depends on the network of these trading relationships.² In particular, an intermediary’s provision of liquidity depends on the willingness of its own counterparties to keep trading, and access to liquidity is not uniform within realistic financial networks.³

An important subtlety is that one OTC market may be interdependent, or *coupled*, with another one due to the nature of the financial instruments traded in the two markets. The essence of the coupling is that an intermediary’s circumstances in one market influence its behavior in another—for example, because access to liquidity in one market is complementary to trading in the other. A natural example of coupling comes from the market for *repo* lending: short-term lending of cash secured by collateral such as bonds or other assets.⁴ Brunnermeier and Pedersen (2009) argue, in the context of a centralized market, that such a coupling is an important aspect in the fragility of liquidity provision. The complementarity is as follows: when the bond market is less liquid, traders are less willing to take a bond as collateral, so fewer loans are extended at a given price. Conversely, repo funding is used to finance collateral (e.g., bond) purchases, so illiquidity in the repo market reduces liquidity in the collateral market.

This paper studies coupled OTC markets and asks, in this context, the basic questions about liquidity and its fragility. In particular: How severe is the potential evaporation of liquidity during crisis times, and how does this depend on the network structures of the markets? Modeling strategic liquidity provision as a game in two coupled trading networks, our main

¹Examples of OTC markets include the market for repurchase agreements (repo), the interbank market, and markets for credit default swaps and other derivatives. On the search frictions in OTC markets, see, for example, Duffie, Gârleanu, and Pedersen (2007); Di Maggio, Kermani, and Song (2017); Aquilina and Suntheim (2017).

²How the network among intermediaries affects various financial outcomes is the focus of an active literature on financial networks: see, e.g., Gofman (2014); Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015); Cabrales, Gottardi, and Vega-Redondo (2017); Di Maggio, Kermani, and Song (2017); Elliott and Hazell (2017); Elliott, Golub, and Jackson (2014); Erol and Vohra (2017); Farboodi (2017); Malamud and Rostek (2017); Wang (2016).

³Acharya, Afonso, and Kovner (2017), for example, show that during the 2008 crisis, domestic and foreign banks had differential access to the market for asset-backed commercial paper. Copeland, Martin, and Walker (2014), document substantial heterogeneity in access even to tri-party repo funding in late 2008.

⁴Repo has become a major source of funding liquidity for banks, money market funds, security lenders, and others. Copeland, Martin, and Walker (2014) estimate the sum of all repo outstanding on a typical day in July and August 2008 to be \$6.1 trillion. Less liquid, non-government bond repo outstanding was approximately 500 billion Euros in the EU ($\approx 10\%$ of the entire market) and ≈ 500 bn USD (mostly agency MBS collateral) in the US; see Baklanova, Copeland, and McCaughrin (2015) and ICMA (2016) for details.

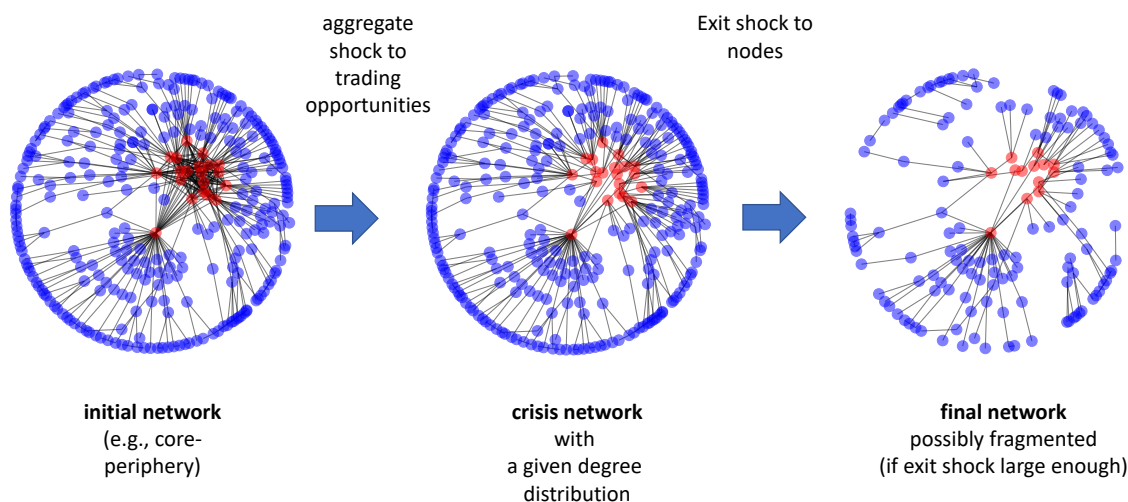


Figure 1: Illustrative networks of trading opportunities (for repo or collateral market) in three different states. *Left*: pre-crisis network with a core-periphery structure. *Middle*: crisis network, with some links disrupted. This is modeled as a random network with a given degree distribution. (The network shown has a power law degree distribution.) *Right*: Network after the exit shock, which causes some intermediaries to withdraw at random.

contribution is to analyze circumstances when these markets are particularly fragile, and identify phenomena occurring in such games that have no analogue when there is only one trading network.⁵ We focus on how coupling can exacerbate illiquidity spirals; which networks are particularly susceptible to this; and which interventions by policymakers can best improve the stability and resilience of markets.⁶

Our model takes the perspective of a macroprudential stress test: an analyst anticipates a crisis scenario, and asks what will happen to liquidity.⁷ We now describe it, using the repo/collateral case for concreteness. There are two OTC markets: one for repo and one for collateral. The market participants are called *intermediaries*. The timing of the model is illustrated in Figure 1. In each market, a *crisis network* is realized, whose links specify the potential trading partners of each intermediary in each network. This realization of the two networks is random from the perspective of the analyst, which captures the fact that some trading opportunities may become inactive due to an aggregate shock occurring at a time of crisis and will not be the same as the networks of exposures measured during normal times.⁸ (Appendix A presents evidence that during the 2008 financial crisis, links present in normal times became inoperative.) Each intermediary makes a binary decision in each market: whether to be *active* and provide liquidity to

⁵We borrow the framing from Malamud and Rostek (2017) who ask analogous questions in a different context.

⁶An extensive networks literature has focused on a planner intervening in a network; for recent examples, see Gofman (2017), Fainmesser and Galeotti (2017), and Galeotti et al. (2017).

⁷The importance of macroprudential stress tests for policymakers has been highlighted for example by Constandino (2017), Anderson et al. (2018) and Aymanns et al. (2018).

⁸For instance, certain counterparty relationships may be effectively shut down due to bilateral exposure limits or other events. See, e.g., Di Maggio et al. (2017); Perignon et al. (2017).

its counterparties. We assume payoffs such that, outside of a special circumstance that we will discuss in a moment, an intermediary is willing to be active in each market if and only if it has access to liquidity in both markets. (Section 2.2 discusses the reasons underlying the coupling or complementarity we have sketched.) Note that liquidity is a local attribute: whether an intermediary has access to each market depends on whether the *particular* intermediaries it trades with choose to be active in the market.⁹ As for the “special circumstance,” some intermediaries, chosen uniformly at random, are hit with an idiosyncratic *exit shock*, and if a node is hit in this way, it is inactive regardless of others’ behavior. The fraction of idiosyncratically shocked intermediaries is called *size* of the exit shock, and it is the key parameter in the macroprudential stress test: varying it is a way of probing how fragile the system is to random failures.

Formally, we study a strategic game of liquidity provision among the intermediaries in the crisis network, conditional on the exit shock realization, and analyze its Nash equilibria. Our liquidity measure is simply the equilibrium number of intermediaries willing to provide repo and buy collateral.¹⁰ The policymaker conducting the macroprudential stress test wishes to assess the expected equilibrium liquidity provision from an *ex ante* perspective. Thus, our calculations focus on the expected liquidity measure, with expectations taken over the crisis network realization (drawn according to certain flexible families of distributions) and the identities of the intermediaries who experience the exit shock. We plot this quantity as a function of the size of the exit shock in order to probe how sensitive liquidity outcomes are to idiosyncratic exit realizations. A key force in the model is a sort of network-driven feedback loop. If some intermediaries are forced to exit the market due to the realization of an exogenous shock, those who are dependent on them for market access—in either market—exit as well, and this contagion propagates through the system. This is an illiquidity spiral in the coupled networks.

Our analysis of the model yields four main results. First, coupling between OTC markets exacerbates these illiquidity spirals: they are more severe, compared to a hypothetical case in which the networks are not coupled. Second, they are more severe in a qualitatively stark way: coupling creates the potential for sudden market freezes, where liquidity vanishes *discontinuously* as we increase the size of the exit shock only slightly—a phenomenon that does not occur without coupling. Third, the fragility of coupled OTC markets is reduced when the two networks have more links in common—that is, when an intermediary’s counterparties in one network are

⁹The investment bank BNP Paribas explained its decision to close down two of its funds by saying “The complete evaporation of liquidity in certain market segments of the US securitization market has made it impossible to value certain assets fairly regardless of their quality or credit rating. [...] Traders are reluctant to bid on securities backed by risky mortgages because they are difficult to sell on.” (“BNP Paribas Freezes Funds as Loan Losses Roil Markets,” Bloomberg.com, August 9, 2007, as cited in Acharya, Gale, and Yorulmazer (2011)).

¹⁰In general, there may be multiple equilibria, which, as we will show, can be ranked by their liquidity. Throughout this paper, we focus on the equilibrium that permits the *maximum* possible liquidity provision. Our results are unchanged if we were to measure liquidity by the number of intermediaries active in *either* market.

more likely to be counterparties in the other. Fourth, a financial system in which repo and collateral are both traded OTC is significantly less resilient to an exogenous shock than a system in which collateral (or repo) is traded on an exchange, providing everyone with access to trading opportunities with everyone else. Thus, regulators can greatly increase stability by introducing an exchange, even if they can do so in only one market.¹¹

To establish these results, we begin by extending the standard supermodular models used to study networked economies (Elliott et al., 2014; Acemoglu et al., 2015) to the coupled-network case. In Section 3, we reduce the characterization of the maximal liquidity measure to the study of a certain graph-theoretic notion—*mutually stable sets* in the final (post-shocks) network, and give a simple algorithm for tracking the illiquidity spirals and computing the liquidity measure. Armed with this algorithm, we can study how the expected liquidity measure depends on the network structure. To build intuition, Section 4 considers stylized star and core-periphery networks.¹² These cases offer the simplest introduction to a key force in the model, the role of *fragile connectors*. A fragile connector is an intermediary whose connectedness to others is fragile in one market, but which provides liquidity to many intermediaries in the other. In a core-periphery network, a fragile connector would be an intermediary that is in the periphery of, say, the collateral market and the core of the repo network: if it loses its access to the core in the repo network, it is not able to function, and all its peripheral neighbors are disabled in the collateral network. Increasing overlap between the two hubs makes this configuration less likely. Note that fragile connectors are distinctive to the multilayer setting: in a single-layer network, a node on the brink of being disconnected is unlikely to be pivotal to many others' connectedness, because (by hypothesis) it has few active counterparties. But in a multilayer network, a node can indeed be nearly disconnected in one network and crucial to the connectivity of the other.

For our main results, we model the crisis network as a random network with a given degree distribution, which can be specified quite flexibly; this permits idiosyncratic variation in links while allowing the policymaker doing the macroprudential stress test to capture data or beliefs about the network structure.¹³ Our approach is motivated by data from the recent financial crisis, which shows that the distribution of crisis exposures is not well-described by very simple network shapes such as core-periphery networks, yet is also asymmetric and has non-

¹¹This conclusion holds even in a model which abstracts away from important benefits of an exchange, such as centralizing clearing.

¹²Core-periphery networks consist of some core intermediaries (all interconnected among themselves), and peripheral intermediaries, each connected only to a few core intermediaries. See, for example, Craig and Von Peter (2014) and the overview of the literature given in Glassermann and Young (2016a).

¹³An emerging literature studies contagion in financial networks using supervisory data, see for example Greenwood, Landier, and Thesmar (2015) and Duarte and Eisenbach (2015). These papers, however, do not consider the case of coupled financial networks.

trivial structure in its degree distribution (Appendix A discusses this in more detail). Concretely, we can fix a full joint distribution of degrees (i.e. specify what fraction of intermediaries have each possible number of counterparties in each network), and then study networks that are random conditional on that data.

In such networks, we investigate *market freezes* in the repo and collateral markets as an extreme outcome of an illiquidity spiral. In particular, we study the equilibrium liquidity measure as we vary the size of the exit shock, increasing it from a low to a high level. As this variation occurs, the subset of intermediaries withdrawing due to the exit shock increases in size. In a market freeze, liquidity in both markets evaporates entirely *and abruptly* as we increase the shock severity. That is, starting from a healthy amount of market and funding liquidity, the exit of a very small set of additional intermediaries may result in the total freeze of the repo and collateral markets. This stark amplification of an exit shock is caused by the coupling of the repo and collateral markets, and is a robust outcome in “generic-looking” random networks, as opposed to very particular ones. For the case of large random networks we can obtain analytical results identifying where this abrupt transition occurs. We compare this case with the alternative environment in which only one of the networks is OTC. When the other market is replaced by a centralized exchange, markets may still freeze but liquidity evaporates continuously. That is, starting from any initial condition, the withdrawal of an additional intermediary may only result in a small decrease in equilibrium liquidity. Furthermore, the number of intermediaries that have to withdraw for equilibrium liquidity to vanish is always larger when at least one market is centralized; that is, the market is more robust in that case. These results indicate that the network structure of the repo and collateral markets has an important impact on the resilience of funding and market liquidity. In this sense, our results provide a novel perspective on the illiquidity spirals in the markets for short-term secured debt that were at the heart of the global financial crisis.

While we illustrate our results in the repo market, the applicability of our model is not restricted to this example. As [Acharya, Schnabl, and Suarez \(2013\)](#) point out, banks have been moving an increasing share of their assets off their balance sheets using special purpose vehicles (SPVs), peaking at \$1.3 trillion in 2007. These SPVs issued commercial paper that was backed, for example, by mortgages and bought by a variety of financial institutions. Crucially, the SPVs were endowed with explicit and implicit liquidity guarantees, forcing banks to take them back onto their balance sheets in times of crisis. These liquidity guarantees create a coupling between the different ABCP markets since banks use the same \$1 guarantee to underwrite several guarantees simultaneously. Once one of the guarantees is activated, the bank cannot use the same liquidity for another guarantee, changing its ability to trade in other networks.

We discuss the related literature, and our contributions relative to it, extensively in Section 7. To be brief, our main claim is that we offer a methodological contribution in showing how the theory of games in coupled networks can be applied to analyze important financial markets. Though network models have been used extensively to study over-the-counter markets and other trading relationships, to the best of our knowledge, ours is the first paper to focus on the financial implications of *coupled* networks.¹⁴ The conclusions emphasized above—e.g., on market freezes, the benefits of centralizing at least one market, and the benefits of greater overlap between two OTC networks—show that the theory has implications for policy-relevant questions. At a technical level, we develop ideas from literatures on coupled networks, which were often analyzed using heuristics, to rigorously study the conditions under which market freezes occur; here, both the conditions we give and the proofs are new.

2 Model

2.1 Game of liquidity provision

There is a set $N = \{1, \dots, n\}$ of intermediaries. Intermediaries may trade bilaterally with other intermediaries in a *repo* and a *collateral* market, $\mu \in \{R, C\}$. In the repo market, a repo seller provides financing to a repo buyer against a security as collateral. In the collateral market, intermediaries trade the security.

Each intermediary can trade in a given market only with a subset of other intermediaries, which can depend on the market. The set of trading relationships in market μ during crisis times is taken as exogenous and described by a directed network \mathcal{G}_μ . A directed network \mathcal{G} is a set of nodes $V(\mathcal{G})$ together with a set $E(\mathcal{G})$ of directed links, i.e., ordered pairs (i, j) with $i, j \in N$, which we often write as $i \rightarrow j$. In the repo market, a link $i \rightarrow j$ in \mathcal{G}_R means that i provides repo financing (a secured loan) to j , providing cash in exchange for a security and a promise to repurchase it at a later date. We describe such a link by saying that i provides *funding liquidity* to j . In the collateral market, the link $i \rightarrow j$ in \mathcal{G}_C means that i purchases collateral from j . Therefore, we say that i provides *market liquidity* to j . We assume there are no self-links $i \rightarrow i$. Both markets share the same node set: $V(\mathcal{G}_R) = V(\mathcal{G}_C) = N$. We will call such a pair $(\mathcal{G}_C, \mathcal{G}_R)$ a *multilayer network*.

We focus on repo as the prime example of secured lending, but our model can be translated to any secured lending market. More generally, though we use the repo market terminol-

¹⁴Trading in single-layer networks is studied in various setups. In addition to references above, see also, for example, Gai, Haldane, and Kapadia (2011), Duffie, Malamud, and Manso (2014), Condorelli, Galeotti, and Renou (2016).

ogy throughout, the same primitives can be used to study any two trading networks that are coupled through the behavior of their nodes. See Section 2.2 for more on this.

It will be useful to define the set of an intermediaries trading partners as its neighborhoods.

Definition 2.1 (Neighborhoods). The *in-neighborhood* of intermediary i in market μ , i.e. in network \mathcal{G}_μ , is the set of intermediaries whose directed links point to i ,

$$K_{i,\mu}^- = \{j \mid j \rightarrow i \in E(\mathcal{G}_\mu)\},$$

and can be interpreted as the set of intermediaries providing liquidity to i . The in-degree of intermediary i in market μ is the size of the in-neighborhood $d_{i,\mu}^- = |K_{i,\mu}^-|$.

Analogously, we define the set of intermediaries that obtain liquidity from i as the out-neighborhood $K_{i,\mu}^+$ by replacing $j \rightarrow i$ with $i \rightarrow j$ in the above definition. The out-degree $d_{i,\mu}^+$ is defined as the number of these intermediaries.

Intermediaries play a strategic game of complete information.¹⁵ Let $a^R \in \{0, 1\}$ and $a^C \in \{0, 1\}$ correspond to intermediaries' decisions of whether to be *active* in each market, with the pair (a_i^R, a_i^C) being intermediary i 's *action*.¹⁶ An intermediary's payoff will depend on the actions of its counterparties and the realization of an exogenous, exit shock w_i to the intermediary. The outcome $w_i = 1$ is a good, or business-as-usual, shock and $w_i = 0$ is a bad shock.¹⁷ We often refer the fraction of intermediaries hit by a bad shock $1 - \sum_i w_i / n$ as the size or the severity of the shock.

To describe the payoffs, it will be useful to define an auxiliary variable $S_i^\mu = \sum_{j \in K_{i,\mu}^-} a_j^\mu$, the number of active counterparties in i 's in-neighborhood in market μ , which depends on \mathbf{a}_{-i} (because i is not an element of its own neighborhood). Then an intermediary's payoff is

$$u_i(\mathbf{a}) = \begin{cases} \pi(S_i^R, S_i^C) - c(w_i) & \text{if } a_i^R = a_i^C = 1 \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

Here $\pi(\cdot)$ is a function describing the payoffs of operating with the given levels of neighborhood activity (these depend on the network and neighbors' decisions) and $c(w_i)$ describes the costs

¹⁵This complete-information assumption amounts to assuming that the intermediaries know who their counterparties are, and whether they are active. The assumption is not essential, and can be relaxed using recent results in the theory of network games. But it considerably streamlines the analysis and puts the focus on the planner's uncertainty over how a crisis will affect the network.

¹⁶As usual, \mathbf{a} refers to the profile of all actions, i.e. $\mathbf{a} = (a_i^R, a_i^C)_{i \in N}$ and \mathbf{a}_{-i} refers to the profile of all actions other than i 's own.

¹⁷Examples of bad shocks include news about markets that trigger intermediaries' internal risk limits, or the bankruptcy of an intermediary (e.g., as a consequence of the discovery of fraud).

of operating, which depends on one's own shock. We take a reduced-form approach to the actual trade: payoffs represent the consequences of intermediaries' being active; these payoffs can be derived from a more detailed description of intermediaries' activities.¹⁸

The critical assumptions imposed on the ingredients of the payoff are as follows.

1. The function π is increasing in each argument, capturing that the returns to operating are increasing in the levels of neighbors' liquidity provision in both networks.
2. An intermediary facing at least one counterparty in each market is willing to operate given a good shock realization: $\pi(1, 1) > c(1)$.
3. An intermediary having no counterparty in either market is unwilling to operate in either market, even given the good shock realization: $\pi(0, 1) < c(1)$ and $\pi(1, 0) < c(1)$.
4. An intermediary having a bad shock realization is unwilling to operate: $\pi(S_i^R, S_i^C) < c(0)$ for all values of (S_i^R, S_i^C) .

Under these assumptions, an intermediary's best response to its in-neighbors' actions can be summarized as follows:

$$\mathcal{R}_i(\mathbf{a}_{-i}, w_i) = \begin{cases} (1, 1) & \text{if } S_i^\mu \geq 1 \text{ for } \mu = R, C \text{ and } w_i = 1 \\ (0, 0) & \text{otherwise.} \end{cases} \quad (2)$$

Note that the intermediary takes the same action in both markets in any best response, and so as a shorthand in discussing equilibria and deviations, it will be without loss to let the variable y_i denote the best-response action of i in both markets, and we will often call it simply i 's best response. It can be written as

$$y_i = w_i \cdot B(S_i^R(\mathbf{a}_{-i}) \cdot S_i^C(\mathbf{a}_{-i})), \quad (3)$$

where we define the operator $B(x) = 1$ if $x > 0$ and $B(x) = 0$ otherwise.

The fact that the threshold level of activity that is required for intermediary i to operate is equal to 1 is not essential. What is essential to the analysis we will use, which is based on the theory of supermodular games, is that intermediary i 's level activity (which needn't be binary) is weakly increasing in others' levels of activity, and at some threshold,¹⁹ which may depend on a combination of others' activities, the intermediary shuts down. Analogues of our results can

¹⁸Section 2.2 below discusses some financial foundations for the strategic structure we assume.

¹⁹In particular, replacing 1 by a generalized threshold in the best-response function leads to a very similar analysis, and detailed calculations are available upon request.

be derived in that richer environment, but the insights are seen most sharply in the special case we study here.

2.2 Interpretation of intermediary behavior

In the previous section we outlined a simple model for intermediaries in the repo and collateral markets. In this section we explain the key economic forces behind the intermediaries' best responses. Intermediaries operate in the market as middlemen: their reason for participation is to collect spreads from intermediating repo and collateral trades. Intermediaries have an incentive to keep their inventory of cash and collateral small and their repo exposure flat. This allows the intermediary to maximize the return it can earn on its capital. Given that intermediaries do not hold large inventories of cash or collateral, *the repo and collateral markets are effectively complements*, as liquidity in one market facilitates intermediation in the other: access to the repo market allows intermediaries to fund collateral purchases, while access to the collateral market ensures that an intermediary can liquidate collateral in case one of its repo counterparties defaults. It is this complementarity between repo and collateral markets that is captured, in a very simple way, in our best response. We discuss it in more detail below.

Note that the game we have defined focuses on the intermediaries, and leaves out the sources and destinations of funding liquidity and collateral that are external to the intermediaries. These fundamental traders provide the basic demand and supply that drives the chains of lending and re-lending we study. In the context of the game, they can be modeled as non-strategic nodes who always set $a_i^C = a_i^R = 1$. Because they do not change the fundamental connectedness properties of our networks, we omit them here. They are, however, discussed more fully in Appendix I.

We now discuss the financial considerations behind intermediaries' constraints and decisions in more detail. Intermediaries face a regulatory capital constraint: they must hold an amount of capital no less than some percentage of total risky assets. This creates inventory costs, and an intermediary thus has incentives to hold a small inventory of the collateral asset relative to the gross volume of collateral and repo that it intermediates.²⁰

We assume that intermediaries have incentives to operate at or near the regulatory capital constraint (DeAngelo and Stulz, 2015). This, along with the fact that they hold small inventories (and therefore do not have much collateral to liquidate at will) means that they seek financing in order to provide repo to another intermediary or to purchase collateral, rather than using their own cash. We also assume that the intermediary will lend in the repo market only if it has

²⁰Empirical evidence for dealer inventory management in equity and derivatives markets can be found for example in Hendershott and Menkveld (2014) or Abad et al. (2016).

access to the collateral market for immediate liquidation of counterparties' collateral in case they default.²¹

We close by illustrating an intermediary's activity in the repo and collateral market. Suppose an intermediary extends a repo loan of size ΔA_R to one of its counterparties. Increasing repo lending leaves the capital ratio unchanged²² but requires an equal increase in reverse repo borrowing $\Delta L_R = \Delta A_R$.²³ Note that the intermediary requires access to the collateral market where it would liquidate the collateral in case of counterparty default on this transaction. In other words, extending a repo loan requires access to both funding liquidity and market liquidity. If an intermediary lacks access to one of the two, it is either simply unable to perform the transaction, or faces prohibitive costs.²⁴

Now suppose, instead, that an intermediary purchases an amount ΔA_C of the collateral asset from one of its counterparties. The intermediary obtains reverse repo to fund the initial purchase: $\Delta L_R = \Delta A_C$. This transaction leads to a temporary expansion of the intermediary's balance sheet. The intermediary reverses this expansion by selling an amount ΔA_C of the collateral asset to one of its other counterparties and by repaying the reverse repo loan. In summary, purchasing the collateral asset requires access to both funding liquidity and market liquidity.

The exit shock vector w may be interpreted as the realization of uncertainty in the value of an illiquid asset on the intermediaries' balance sheet as the crisis unfolds. A shock realization $w_i = 0$ corresponds to an adverse shock that forces intermediary i to cease its intermediation activity and become inactive. Without the adverse shock, if an intermediary can intermediate transactions without violating any of the above assumptions, its intermediation activity is profitable. This motivates the best response functions of the liquidity provision game: As long as intermediary i has enough trading partners active in the repo market (to provide funding liquidity), and enough, possibly different, trading partners active in the market for collateral (to provide market liquidity), and as long as it has not experienced an exit shock, the intermediary's

²¹This is motivated by risk-management concerns: if intermediaries held on to collateral in such cases, they would be exposed to fundamental risk in the value of the collateral (see, for example [Oehmke \(2014\)](#)).

²²Assuming that repo lending implies a zero risk weight for the intermediary's capital requirement.

²³Realistically, the repo buyer requires a *haircut* on the collateral, i.e. will provide a fraction $\alpha \in [0, 1]$ of the collateral's face value as funding. The haircut insures the repo buyer against the risk of a devaluation of the collateral. The repo buyer will also require interest ($r \geq 0$) per unit of repo financing. For simplicity, we assume that the collateral is safe, i.e. that $\alpha = 1$ (though the collateral may still be illiquid), and that there is no discounting, ($r = 0$).

²⁴Note the similarity of our model to [Acharya et al. \(2011\)](#). They consider a model in which debt has a short maturity relative to the collateral and buyers of the collateral require collateralized funding to purchase the asset if it needs to be liquidated. [Acharya et al. \(2011\)](#) show that, in the presence of liquidation costs, the debt-bearing capacity of a collateral asset is determined by the future liquidation value of the collateral, which in turn is determined by its future debt-bearing capacity. Here, we consider a simplified version of this model with a binary debt-bearing capacity depending on the availability of funding and market liquidity in the neighborhood of a given intermediary.

best response is to be active in both markets.²⁵

Our model of intermediaries' decisions to be active can be embedded in a more detailed model, which relates them to payoffs for the intermediary. Appendix I sketches such a model, deriving payoffs richer than those in (1) above, but giving rise to the same best responses. An intermediary is active in a market if it makes profits from intermediating (repo or collateral). These profits are equal to a spread charged from trading the collateral, which depends on its access to trading partners, less a cost of doing business.

3 Equilibrium for general networks

In the following we will consider arbitrary directed networks of trading relationships, \mathcal{G}_R and \mathcal{G}_C , for the repo and collateral markets.

3.1 Definition and existence of equilibrium

Intermediaries play a game of complete information given a realization of the exit shock. An *equilibrium* in this game is a fixed point of the best-response correspondence \mathcal{R} , whose components are the functions \mathcal{R}_i defined in (3).

To study these, we first introduce some notation and terminology. First, for any two vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$ define the pointwise ordering

$$\mathbf{x} \leq \mathbf{y} \iff x_i \leq y_i \text{ for all } i \in N.$$

With this ordering, we have a game of strategic complements: the best response of each agent is monotone in its argument (holding the exit shock w fixed). Thus, a standard application of Tarski's fixed-point theorem shows that equilibria exist and are nicely ordered. Indeed, they form a complete lattice, containing a greatest element (a maximal equilibrium that pointwise dominates every other) and a least element. This lattice structure permits a rich theory of comparative statics: see [Milgrom and Roberts \(1990\)](#) and [Zhou \(1994\)](#).²⁶

Following the rationale in [Elliott et al. \(2014\)](#), we make a best-case assumption, focusing on the maximum equilibrium, to derive a lower bound on the cascade size. We call the maxi-

²⁵Note that choosing to be active in a market does not imply that actual transactions occur. Instead the outcome of the liquidity provision game constrains potential transactions in these markets.

²⁶For similar arguments in other financial network applications, see [Eisenberg and Noe \(2001\)](#) or [Elliott et al. \(2014\)](#).

imum equilibrium \mathbf{y}^* . A natural equilibrium liquidity measure is then

$$\mathcal{L}(\mathbf{y}^*) = \frac{1}{n} \sum_i \mathcal{Y}_i^*, \quad (4)$$

which is simply the fraction of intermediaries that are active (i.e., willing to provide liquidity) in equilibrium.

3.2 Equilibrium, network structure and shocks

3.2.1 Simple examples

Let us briefly consider conditions under which the maximum equilibrium will look very simple, i.e. will be $\mathbf{y}^* = \mathbf{0}$ or $\mathbf{y}^* = \mathbf{1}$. If the exit shock vector takes the form $\mathbf{w} = \mathbf{0}$, the equilibrium will be $\mathbf{y}^* = \mathbf{0}$. Equally, for an arbitrary exit shock vector but in the absence of links between the intermediaries, the equilibrium will be $\mathbf{y}^* = \mathbf{0}$. If the exit shock vector takes the form $\mathbf{w} = \mathbf{1}$ and all intermediaries have at least one incoming link in both \mathcal{G}_C and \mathcal{G}_R , in the maximum equilibrium, all intermediaries will choose to be active and $\mathbf{y}^* = \mathbf{1}$.

3.2.2 Characterization and algorithm

In general, the equilibrium of the liquidity provision game depends on two factors: the network structure of the over-the-counter markets during the crisis and the realization of the exit shock. We will examine this dependence.

Suppose that, after the crisis networks have been realized, their structure satisfies the following assumption.

Assumption 3.1. All intermediaries have at least one incoming edge in \mathcal{G}_R and \mathcal{G}_C .

In this case, for an exit shock realization $\mathbf{w} = \mathbf{1}$, the maximal equilibrium is $\mathbf{y}^* = \mathbf{1}$. We can think of this as a *pre-(idiosyncratic)-shock* situation, in which no assets have lost any value; our assumption guarantees that no intermediaries withdraw from the crisis networks.²⁷ Starting from such a baseline, we can consider equilibrium outcomes for more interesting realizations of the exit shock, i.e. for shock vectors with some $w_i = 0$, to probe the fragility of liquidity in the crisis networks. We refer to this regime as the *post-(idiosyncratic)-shock* regime. In the following we will characterize how the post-shock equilibrium depends on the structure of the crisis networks \mathcal{G}_C and \mathcal{G}_R .

²⁷This result is stated and proved as Lemma 3 in Appendix C.1, using some terminology we introduce later in this section. The assumption is without loss of generality in that if it is not satisfied we can simply remove the intermediaries not satisfying it from network.

Let W denote the set of intermediaries for which $w_i = 0$. Let $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$ denote the networks after all the edges corresponding to the intermediaries in W (i.e., those that have received a bad shock, $w_i = 0$) have been removed. To link the equilibrium outcome to the structure of $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$, first define a stable subset of nodes in a network \mathcal{G} .

Definition 3.1 (Stable subset). In a network $\mathcal{G} = (V, E)$, a subset $V' \subset V$ is *stable* if, for each $i \in V'$, there is a $j \in V'$ such that $(j, i) \in E$. That is, every node in V' has an incoming edge from V' .

We now make an analogous definition for coupled networks.

Definition 3.2 (Mutually stable subset). Let \mathcal{G}_R and \mathcal{G}_C be directed graphs. A mutually stable subset of the coupled network $(\mathcal{G}_R, \mathcal{G}_C)$ is a set V' of nodes that is stable in \mathcal{G}_μ for $\mu \in \{R, C\}$.

The existence and size of mutually stable subsets is closely related to the existence of a maximal equilibrium in our game.

Proposition 1 (Maximal equilibrium and mutually stable subsets). *In the maximal equilibrium \mathbf{y}^* , the set of active intermediaries ($y_i^* = 1$) equals the maximal mutually stable subset of $(\mathcal{G}_R(W), \mathcal{G}_C(W))$.*

In other words, if we take a stable subset of $(\mathcal{G}_R(W), \mathcal{G}_C(W))$ that is maximal under set inclusion and set their activity level to 1, we get the maximal equilibrium. That this is an equilibrium follows from the form of \mathcal{R} : First, none of these intermediaries have been shocked, as the shocked ones were removed from $(\mathcal{G}_R(W), \mathcal{G}_C(W))$. Second, all intermediaries in such a set have, by definition of a mutually stable subset, at least one incoming link in both networks from other intermediaries in the set, so, by definition of \mathcal{R} , is an equilibrium for all of them to be active. To complete the proof of Proposition 1, we must show that no equilibrium has a set of active intermediaries that is larger than the set of those active in \mathbf{y}^* . To this end, take \mathbf{y} satisfying $\mathcal{R}(\mathbf{y}) = \mathbf{y}$ and observe that this set of intermediaries (by definition of \mathcal{R}) is mutually stable. Thus it must be contained in the maximal mutually stable set.

The above gives a graph-theoretic description of the maximal equilibrium. We can also give a description in terms of fixed-point theory. By the supermodular structure of the game, the maximum lattice point can be found by starting from the maximum feasible actions $\mathbf{y} = (1, \dots, 1)$, and repeatedly applying the best response function, \mathcal{R} (Milgrom and Roberts, 1990). Algorithm 1 below makes this precise.

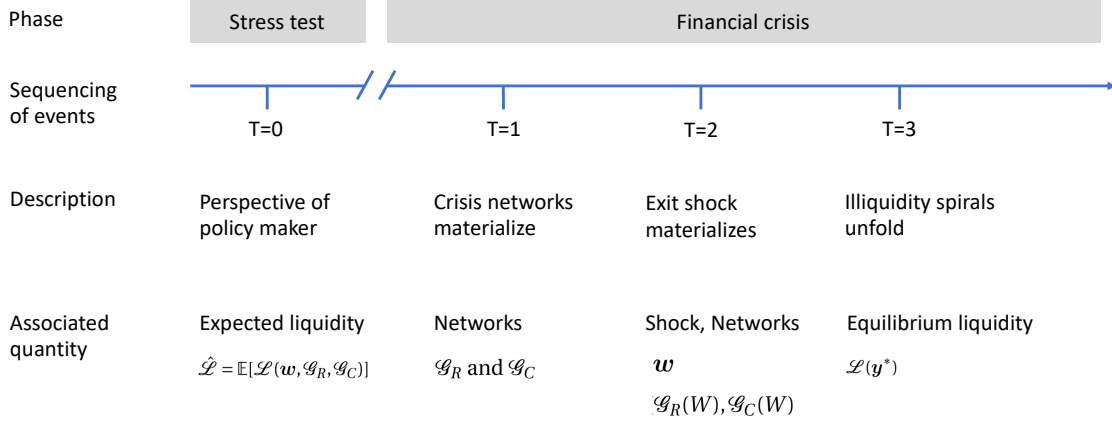


Figure 2: Macroprudential stress testing and crisis timeline.

Algorithm 1 Algorithm to compute the equilibrium (greatest fixed point of \mathcal{R}).

```

y ← 1
while y ≠  $\mathcal{R}(\mathbf{y}, w)$  do
    y ←  $\mathcal{R}(\mathbf{y}, w)$ 
end while
return y

```

3.3 Macroprudential stress testing and ex-ante expected liquidity

As mentioned in the introduction, our analysis of liquidity takes the perspective of a macroprudential stress test (e.g. [Constâncio \(2017\)](#), [Anderson et al. \(2018\)](#)) in which a policy analyst anticipates a crisis scenario and asks what will happen to liquidity. Figure 2 clarifies this sequencing of events and places the quantities we have developed so far into the context of the macroprudential stress test.

Conceptually our analysis can be divided into two phases. A stress testing phase ($T = 0$) and a crisis phase ($T = 1, \dots, 3$). In the stress testing phase, using only information available at time $T = 0$, a policy analyst anticipates what occurs in the three events that constitute a crisis. First, he takes a view on the structure of the networks \mathcal{G}_C and \mathcal{G}_R that emerge at $T = 1$ due to the aggregate shock that initiates the crisis. Second, he considers the idiosyncratic exit shocks that materialize at $T = 2$ and lead to the withdrawal of intermediaries and the post-shock networks $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$. Finally, given $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$ illiquidity spirals unfold and the equilibrium liquidity \mathcal{L} is realized.

From the perspective of the policy analyst, both the shock and network realizations are unknown. Thus, for the policy analyst, the natural measure of liquidity is not \mathcal{L} but its ex-




Figure 3: Example networks

ante expectation given distributions over shocks and crisis networks $w \sim F_w$ and $\mathcal{G}_\mu \sim H_\mu$, i.e.

$$\hat{\mathcal{L}} = \mathbb{E}[\mathcal{L}(w, \mathcal{G}_R, \mathcal{G}_C)]. \quad (5)$$

In what follows, our object of interest will be the expected equilibrium liquidity measure $\hat{\mathcal{L}}$ and its dependency on the distributions $w \sim F_w$ and $\mathcal{G}_\mu \sim H_\mu$, reflecting different beliefs of the policy analyst over crises scenarios.

4 Two examples: Star and core-periphery networks

To build intuition for the mechanics of illiquidity spirals and to preview our main results, we will illustrate our outcomes when the crisis networks—both for repo and collateral—come from canonical classes of simple networks: star and core-periphery networks. Some of our central findings have manifestations in some form even in these simple examples. First, we will see that networks in which the two layers (repo and collateral) have a more overlapping link structure—with the same intermediaries more likely to be at the center of both—are *more* resilient. Second, making one market centralized improves the robustness of the system.

These results apply to star networks and networks with a *core-periphery* structure, which is a stylized representation of patterns often seen in empirical studies of the exposures generated in over-the-counter markets (see [Abad et al. \(2016\)](#) and [Craig and Von Peter \(2014\)](#)).

In the following we will discuss the example of the star network in detail and, to remain concise, will comment only briefly on our results for core-periphery networks. Our full treatment of core-periphery networks can be found in [Appendix B](#).

4.1 Star network

Suppose both the repo network \mathcal{G}_R and the collateral network \mathcal{G}_C are star networks. We will study the behavior of illiquidity spirals, in which those who exit following a shock cause other intermediaries to withdraw from the market, and characterize the resulting equilibrium liquidity measure. We will then explain how the calculations anticipate our main findings, which hold for much richer network structures.

Figure 3 illustrates a star network. The network consists of two types of nodes: a single hub node (blue rectangle) and a set of peripheral nodes (yellow circles). The hub node has bi-directional links to all peripheral nodes but peripheral nodes are not connected to each other. Thus, the OTC markets $\mu \in \{R, C\}$ are characterized by a partition of the set of intermediaries N into a single hub intermediary $B_{H,\mu}$ and set of peripheral intermediaries $B_{P,\mu} = N \setminus B_{H,\mu}$.

Let us now consider an exit shock profile $w^j \in \mathbb{R}^N$ in which intermediary j , chosen uniformly at random, receives an adverse shock, i.e.

$$w_i^j = \begin{cases} 0 & \text{if } i = j, \\ 1 & \text{otherwise.} \end{cases}$$

The probability that intermediary j receives an adverse shock is $P(w^j) = 1/n$. The following results for the post-shock liquidity measure will be computed by averaging over all shock profiles of this form.

We also consider the realization of the network labels—that is, the identity of the hub in each network—to be random and equally likely. Thus, we have a (simple) random multilayer network, and we will condition on realizations of the identities of the hub and the periphery as random variables.

To compute the post-shock liquidity measure it is sufficient to consider two cases. In the first case the hub intermediaries in the repo and collateral markets are different, i.e. $B_{H,C} \neq B_{H,R}$. In the second case the hub intermediaries are the same, i.e. $B_{H,C} = B_{H,R}$.

Proposition 2 (Coupled star networks – post-shock liquidity measure). *Consider two star networks \mathcal{G}_C and \mathcal{G}_R . Let $E[\cdot | \cdot]$ denote the conditional expectation operator. The conditional expected post-shock liquidity measures are*

$$E[\mathcal{L} | B_{H,C} \neq B_{H,R}] = \frac{(n-2)(n-1)}{n^2},$$

$$E[\mathcal{L} | B_{H,C} = B_{H,R}] = \frac{(n-1)^2}{n^2}.$$

Given a probability $q = P(B_{H,C} = B_{H,R})$ the expected post-shock liquidity measure for the coupled star networks is

$$\begin{aligned}\hat{\mathcal{L}}^{s-s} &:= E[\mathcal{L}] = qE[\mathcal{L} \mid B_{H,C} = B_{H,R}] + (1 - q)E[\mathcal{L} \mid B_{H,C} \neq B_{H,R}] \\ &= q \frac{(n-2)(n-1)}{n^2} + (1 - q) \frac{(n-1)^2}{n^2}.\end{aligned}$$

For the proof see Appendix C.2.

First, note that the conditional post-shock equilibrium liquidity is always less than $n - 1$, the liquidity in case of no additional withdrawals after the adverse shock. We call the phenomenon that additional intermediaries withdraw following an exit shock an *illiquidity spiral*. Second, note that the equilibrium liquidity is smaller when the repo and collateral markets do not share the same hub node. This suggests that OTC markets that are less similar are less resilient to exit shocks. For the star network the intuition for this result is very simple. In coupled star networks only the failure of a hub can trigger the withdrawal of additional intermediaries. If the exit shock hits each intermediary with equal probability, then it is more likely that a hub will be hit if the star networks do not share the same hub.

In later sections we will extend this observation to a richer class of networks, by showing that stability is increasing in the structural overlap of the coupled networks: the likelihood that a given link in one network also exists in the other. In fact, even for the star network, one could think of the probability q as a measure of this overlap: as q increases, it becomes more likely that links overlap.

We next turn to the question of how the stability of the system changes when one of the OTC markets is replaced by a centralized exchange. In a centralized exchange all intermediaries can trade with all other intermediaries. Therefore, we model a centralized exchange as a fully connected (complete) network. Suppose that the collateral market is replaced by a complete network. There can never be contagion “through” the complete network, because no node is critical to connectivity within it. In other words, the mutually stable subsets of the pair of networks $(\mathcal{G}_R, \mathcal{G}_C)$ are simply the stable subsets of \mathcal{G}_R . Therefore, the expected post-shock liquidity $\hat{\mathcal{L}}^{s-c}$ in the star-complete configuration is the same as in the star-star configuration conditional on the two networks sharing the same hub intermediary, i.e. $B_{H,C} = B_{H,R}$.

Proposition 3 (Coupled star and complete networks – post-shock liquidity measure). *Consider a star network \mathcal{G}_C and a complete network \mathcal{G}_R . The post-shock liquidity measure is*

$$\hat{\mathcal{L}}^{s-c} = E[\mathcal{L}] = \frac{(n-1)^2}{n^2}.$$

Furthermore, the post-shock liquidity in the star-complete configuration always exceeds the post-shock liquidity in the star-star configuration if $q > 0$:

$$\hat{\mathcal{L}}^{s-c} > \hat{\mathcal{L}}^{s-s}.$$

The latter assertion follows immediately from the fact that $E[\mathcal{L} \mid B_{H,C} \neq B_{H,R}] < E[\mathcal{L} \mid B_{H,C} = B_{H,R}] = \hat{\mathcal{L}}^{s-c}$ and the fact that $\hat{\mathcal{L}}^{s-s}$ is a convex combination of these two quantities. Hence, the coupling of OTC repo and collateral markets leads to a reduction of the expected equilibrium liquidity relative to the benchmark case of a centralized collateral market.

4.2 Core-periphery network

Core-periphery networks (see Figure 3 for a stylized example), are often used in models of over-the-counter markets since they depict a segmented dealer-client structure in a very simple way.²⁸ Nodes in the stylized core-periphery network we study are partitioned into two sets: a set of *core* nodes and a set of *peripheral* nodes. A core node is connected to all other core nodes and a subset of peripheral nodes via bi-directional links. A peripheral node is connected only to a single core node via a bi-directional link.

We will only outline our main results here, deferring a full treatment of illiquidity spirals in core-periphery networks to Appendix B. As mentioned above, our results for the coupled star network carry over to coupled core-periphery networks. First, as the two coupled networks become similar (i.e. overlapping), the extent of illiquidity spirals is reduced. In the context of core-periphery networks this means that networks in which nodes are in the core or periphery of both networks are more stable than networks in which nodes are in the core of one but in the periphery of the other network. Second, if one of the core periphery networks is replaced by a centralized exchange, equilibrium liquidity is always improved.

Our analysis of core-periphery networks goes beyond our discussion of star networks in that we also consider exit shocks in which more than just one node receives an adverse shock. In particular, anticipating our analysis in Section 5, we study expected equilibrium liquidity as a function of the fraction of nodes that received an adverse exit shock (i.e. the size of the exit shock). We find that equilibrium liquidity first decreases quickly for small shock sizes but then decreases more slowly as shock sizes become larger (see Figure 4). Importantly, in a core periphery network, equilibrium liquidity never vanishes entirely if there exists a subset of nodes that are in the core of both networks.

²⁸See Wang (2016) for a recent model of trading in core periphery networks.

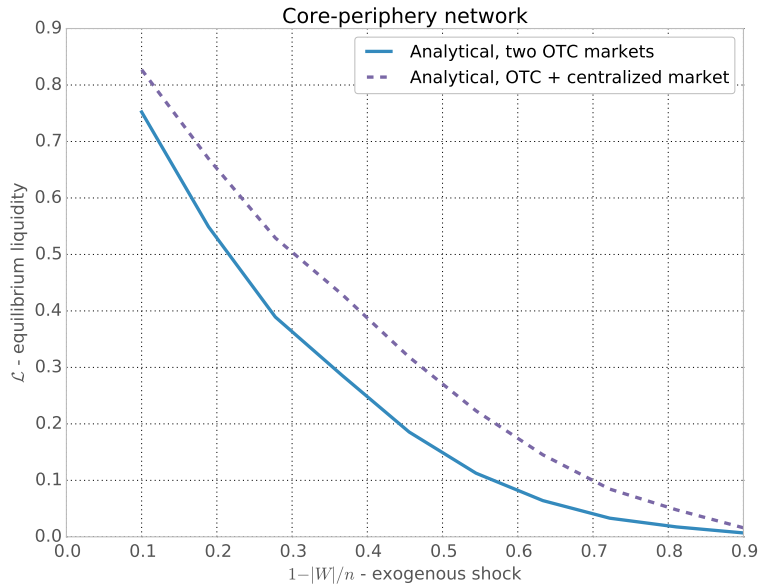


Figure 4: Let W denote the set of all intermediaries hit with an exit shock. Equilibrium liquidity as a function of the fraction of intermediaries $1 - |W|/n$ that withdraw from the repo and collateral markets following an exit shock in a core periphery network. The core periphery network is determined by the fractions ($n_{cc} = 0, n_{cp} = 2, n_{pc} = 2, n_{pp} = 50$). Continuous line: analytical approximation for the case of repo and collateral OTC networks. Dashed line: analytical approximation for the case of OTC repo market and centralized collateral market.

5 Random networks with given degree distributions

In our discussion of star and core-periphery networks above, we assumed that the structure of the crisis networks is not only quite simple, and the shape of the network is known in advance (up to relabelings). As we have argued in the introduction, one of the essential aspects of a crisis is that the network of trading opportunities may be considerably changed relative to its normal configuration by an aggregate shock, and the analyst conducting the macroprudential stress test faces uncertainty over this change. In particular, some trading relationships may become inactive. Therefore, in this section, we consider the (random) trading network in effect after such a change. (Section 6 contains some more discussion of modeling considerations.)

Our model of the crisis network allows us to parameterize key features of the resulting linkages while at the same time modeling the uncertainty in potential trading links from the perspective of the policymaker. In particular, we consider in- and out-degrees of the nodes (recall the definitions from Section 2.1) that follow an arbitrary joint distribution; links are random conditional on nodes' degrees. This sort of model makes the study of contagions especially tractable, allowing us to examine fragility without parametric assumptions on the shape of the

degree distribution. After carrying out this exercise, we further build on it by using it to “add noise” to the structured core-periphery networks studied in the previous section (see Appendix G).

The examples studied in the prior section illustrate that the exit shock degrades liquidity provision. A key question we will focus on is whether (from the ex ante perspective of the analyst) the random crisis network is robust or fragile: whether liquidity degrades gradually as the a crisis gets worse, or whether it exhibits extreme sensitivity. The device we use to probe this is to study how sensitive the network is to the size of an exit shock w that removes some nodes uniformly at random. In particular, a configuration is said to be fragile if the total liquidity is very sensitive to the size of the exit shock, and thus can collapse following a small number of nodes being removed or disabled.

As we will see, under certain conditions, there is, indeed, extreme sensitivity. In other words, the randomness in the crisis network we study here makes the fragility of coupled networks more acute. This demonstrates most clearly the financial contagion phenomena that occur in multilayer networks, but which have no single-layer counterpart. It also highlights to policymakers the sorts of crisis network configurations that raise the greatest concern. In addition, the results in this section reinforce two insights first seen in the pure core-periphery model. First, our results here imply that making one network centralized improves stability, uniformly across crisis states. Second, if the two crisis networks are more similar, in the sense of having more overlapping links, then stability is improved.

5.1 Random network models

For the rest of this section, the networks of interest are the ones induced by the aggregate shock, i.e. the crisis networks. Let $\mathbf{d}_\mu^+ = (d_{i,\mu}^+)_{i=1}^n$ and $\mathbf{d}_\mu^- = (d_{i,\mu}^-)_{i=1}^n$ be sequences of non-negative integers representing the out-degree and in-degree, respectively, of an intermediary $i \in N$ in market $\mu \in \{R, C\}$, where as before $n = |N|$. In Appendix C.3, we impose some technical conditions on these degree sequences that make random graphs generated from them well-behaved.²⁹ Let $G_\mu(n, \mathbf{d}_\mu^+, \mathbf{d}_\mu^-)$ be the set of graphs on n nodes with degree sequences \mathbf{d}_μ^+ and \mathbf{d}_μ^- . A random network \mathcal{G}_μ (for an n which is left implicit in the notation) is then a draw from $G_\mu(n, \mathbf{d}_\mu^+, \mathbf{d}_\mu^-)$ uniformly at random. Our simulations deal with these graphs directly. In our analytical results, we study the limit of large networks, $n \rightarrow \infty$. In this limit, we assume our technical assumptions impose that these degree sequences are consistent with joint distributions $(p_{jk,\mu})_{j,k}$ for $\mu \in \{R, C\}$, where $p_{jk,\mu}$ is the fraction of intermediaries with in-degree j and out-degree k in

²⁹These assumptions ensure, for example, that the first and second moment of the degree distribution remain bounded in the limit $n \rightarrow \infty$ and that the sum of all out-degrees matches the sum of all in-degrees.

network μ . We fix these degree distributions throughout the section.

We now make the following assumption:

Assumption 5.1. For $\mu \in \{R, C\}$ and all $k \geq 0$, we have $p_{0k, \mu} = 0$.

In other words, we assume that each node has at least one incoming link in both networks with probability 1. This is the analog of Assumption 3.1 in the random networks model, and thus all nodes are in a maximal stable set, and indeed in a mutually stable set (recall Lemma 3 in Appendix C.1). We also impose:

Assumption 5.2. The random networks \mathcal{G}_R and \mathcal{G}_C are *independent* realizations of G_R and G_C , respectively.

This implies that, for example, intermediary i 's out-degree in \mathcal{G}_R is independent of its out-degree in \mathcal{G}_C . In Section 5.5, we relax this assumption.

5.2 Liquidity when one network is complete: the case of a centralized collateral market

It is convenient to start with a case where one market, say the collateral market, is centralized. This reverses the order of presentation relative to the previous section, but the one-network case is necessary here to introduce key ideas for the coupled-network case. Formally, let $\bar{\mathcal{G}}_C$ denote the complete network, and assume for this subsection that $\mathcal{G}_C = \bar{\mathcal{G}}_C$. This is equivalent to the study of liquidity in one network, \mathcal{G}_R ; the completeness of the other network means that illiquidity spirals occur only through \mathcal{G}_R .

Let the repo market correspond to a random network \mathcal{G}_R drawn from G_R . Now we discuss the exit shocks that we use to probe the fragility of the network.³⁰ There is a pre- and a post-shock state. The pre-shock state corresponds to $w = 1$, in which nobody receives bad shocks. In the post-shock state, a fraction $1 - x$ of intermediaries chosen uniformly at random receive an adverse shock $w_i = 0$. These intermediaries withdraw from both markets. We call $1 - x$ the *size* of the exit shock or the severity of the shock. Recall that $\hat{\mathcal{L}}^*(x)$ is the expected liquidity of the maximal equilibrium from the perspective of the policy analyst conducting the macroprudential stress test as defined in Section 3.3. How does $\hat{\mathcal{L}}^*(x)$ vary as function of the shock size, and in particular is it very sensitive to the exact shock size at some values of x ?

³⁰Recall that \mathcal{G}_R represents the configuration after an aggregate shock.

5.2.1 The giant component

In the present case of a complete collateral network, equilibrium liquidity in any graph \mathcal{G}_R corresponds to the maximal stable set in \mathcal{G}_R . This follows from Proposition 1 along with the fact that $\mathcal{G}_C = \overline{\mathcal{G}_C}$, so that mutually stable sets are simply those sets in \mathcal{G}_R . The characterization of the maximal stable set in \mathcal{G}_R , in turn, is reducible to the study of a certain kind of *giant component* in the network \mathcal{G}_R . We now build up the definition of this object.

It is important to note that $\mathcal{G}_R(W)$, the network with the shocked intermediaries stripped of their edges, is equal in distribution as a draw from G_R with a different (suitably thinned) degree distribution. Thus, all arguments here apply equally to the shocked and unshocked cases and we drop the argument W for readability.

Definition 5.1 (Strongly and weakly connected subsets). In a network $\mathcal{G} = (V, E)$, a subset $V' \subset V$ is:

- (a) *strongly connected* if, for any nonempty, proper subset $V'' \subseteq V'$, there is an edge from V'' to $V' \setminus V''$. (Note this implies an edge exists in the other direction as well.)
- (b) *weakly connected* if for any nonempty, proper subset $V'' \subseteq V'$, there is an edge *between* V'' and $V' \setminus V''$ in one direction or the other.

A *strong* (resp., *weak*) connected component is defined to be a maximal strongly (resp., weakly) connected subset with more than one node.

Now, fix a single network μ and its associated degree distribution $(p_{jk,\mu})_{j,k}$. Degree sequences $\mathbf{d}_\mu^+ = (d_{i,\mu}^+)_{i=1}^n$ and $\mathbf{d}_\mu^- = (d_{i,\mu}^-)_{i=1}^n$ are drawn from that distribution, satisfying the technical assumptions of Appendix C.3. Let γ_n (resp. $\tilde{\gamma}_n$) be the fraction of nodes in a maximum-cardinality strongly (resp., weakly) connected component. Let ρ_n be the fraction of nodes in nonsingleton weakly connected components other than the maximum-cardinality one.³¹ Standard results about random graphs under our assumptions are summarized as follows (see Cooper and Frieze (2004) for details):

Lemma 1. Under our maintained technical assumptions,

- (a) $\rho_n \rightarrow 0$ always, so that there is at most one strongly or weakly component of nonnegligible size;
- (b) γ_n tends, with probability one, to a constant $c \geq 0$ that depends only on $(p_{jk,\mu})_{j,k}$;
- (c) $\tilde{\gamma}_n$ tends, with probability one, to the same constant c .

³¹If there are several components of maximum cardinality, select an arbitrary one.

If $c > 0$, the largest strong (weak) component is said to be the *giant strong (weak) component* (of size c) in the random graph. Otherwise we say there is no giant component.

5.2.2 Characterization of liquidity

We have seen above that equilibrium liquidity corresponds to the maximal stable set in \mathcal{G}_R . It can now be deduced easily from the fact above that, asymptotically, c is the fraction of nodes in the maximal stable set in the random graph with n nodes. Thus c is the equilibrium liquidity. In the shocked regime, the degree distribution governing \mathcal{G}_R is different, and thus corresponds to a different c . In other words, in our setup of randomly disabling a fraction $1 - x$ of intermediaries, the giant component size c depends on the shock size $1 - x$. We can now state the main result for the case where one market is centralized:

Proposition 4. *Let \mathcal{G}_R and \mathcal{G}_C be drawn as described at the start of this section. In particular let $\mathcal{G}_C = \bar{\mathcal{G}}_C$ be a complete network. There exists a value r_c such that the expected liquidity of the maximal equilibrium vanishes smoothly at shock size $1 - r_c$.*

- (1) for all $x \in [0, r_c]$, there is no liquidity: $\hat{\mathcal{L}}^*(x) = 0$;
- (2) for all $x \in (r_c, 1]$, liquidity is positive: $\hat{\mathcal{L}}^*(x) > 0$;
- (3) The transition between the regimes is smooth: $\hat{\mathcal{L}}^*(r_c^-) = \hat{\mathcal{L}}^*(r_c)$.

After the reductions we have gone through above, the proof of Proposition 4 is a straightforward application of known results on the giant component in directed networks (see Newman (2002) and Cooper and Frieze (2004)). These immediately yield the smooth transition of liquidity regimes in Proposition 4. Because an illiquidity spiral can never propagate through the complete network, all other choices of \mathcal{G}_C must lead to a smaller critical shock size.³²

The technique behind the analysis of c as a function of x is expressing the probability that a random node is in the giant component via a fixed-point equation. For instance, it can be seen that a node is in the giant weak component if and only if at least one neighbor is. The resulting fixed-point equation characterizes the size of the giant component and its behavior as we vary x . For the full proof of Proposition 4, see Appendix C.4.

5.3 Stress in coupled over-the-counter markets

We now turn to the case where both networks are over-the-counter, with nontrivial network structures in each. Let the repo (resp., collateral) market correspond to a random network \mathcal{G}_R

³²An exception to this is the special case when \mathcal{G}_C is a copy of \mathcal{G}_R , as we will discuss in Section 5.5. However, given the independence of \mathcal{G}_R and \mathcal{G}_C this event has vanishing probability as $n \rightarrow \infty$.

(resp., \mathcal{G}_C) drawn from G_R (resp., G_C). As before, these are the networks after the aggregate shock, and we then consider the effect of an exit shock by comparing the pre- and post-exit shock states. The pre-shock state corresponds to $w = 1$, in which nobody receives bad shocks. In the post-shock state, a fraction $1 - x$ of intermediaries chosen uniformly at random receive an adverse shock $w_i = 0$. These intermediaries withdraw from both markets. As before, we call $1 - x$ the *size* of the exit shock. Let $\hat{\mathcal{L}}^*(x)$ be the expected liquidity of the maximal equilibrium. How does $\hat{\mathcal{L}}^*(x)$ vary as function of the shock size?

We will use the notions introduced above to state a key assumption:

Assumption 5.3. Fix $(p_{jk,\mu})_{j,k}$; select, uniformly at random, of a fraction $1 - x$ of nodes and remove all their edges.³³ Let $c(x)$ be the size of the giant strong component in the resulting graph. *Giant-component concavity* holds if c is concave in x over the range where $c(x) > 0$.

We are now in a position to state our main result on liquidity in coupled random networks.

Proposition 5. *Consider random networks \mathcal{G}_R and \mathcal{G}_C drawn as described above, with a degree distribution satisfying giant-component concavity. Also, let $\bar{\mathcal{G}}_C$ denote the complete network.*

- (A) *There is a value $x_c \in [0, 1]$ such that the expected liquidity of the maximal equilibrium has a discontinuity at shock size $1 - x_c$. That is:*
- (1) *for all $x < x_c$, there is no liquidity: $\hat{\mathcal{L}}^*(x) = 0$ for large enough n ;*
 - (2) *there is a constant $L > 0$, such that for all $x \geq x_c$, the liquidity is bounded below by L : i.e., $\hat{\mathcal{L}}^*(x) \geq L$ for large enough n .*
- (B) *The critical shock size for the case of a centralized collateral market is always greater than the critical shock size for the case of an OTC collateral market: $1 - x_c < 1 - r_c(\mathcal{G}_R, \bar{\mathcal{G}}_C)$.*

Proposition 5 states that there are two liquidity regimes, and the one that obtains depends on the size of the exit shock. If the shock is sufficiently small, i.e. less than $1 - x_c$ in size, both repo and collateral markets are liquid. However, if the shock size increases beyond its critical value by an arbitrarily small amount, liquidity in both markets vanishes. This is the *market freeze* regime. The transition between the liquid and the frozen market regime is *discontinuous* at the critical shock size: starting from a liquid market, the withdrawal of a very small measure of additional intermediaries is amplified through the coupled structure of the repo and collateral markets to the extent that all liquidity disappears entirely. Formally, this is reflected in the fact that at x_c , liquidity goes from a strictly positive value to 0.

³³If technical assumptions in the appendix that apply to the original graph, they also apply to the graph obtained from this process, called the “percolated” graph, so the lemma about giant components still holds.

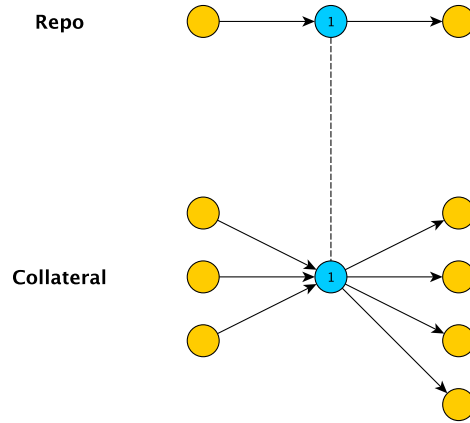


Figure 5: Illustrative example of a fragile connector: Intermediary 1 is a stylized example of a fragile connector. It receives liquidity in the repo market from only a single intermediary and is therefore susceptible to the withdrawal of this critical intermediary. At the same time intermediary 1 is the sole provider of liquidity to a number of intermediaries in the collateral market – it acts as a connector in the market for collateral. Thus if intermediary 1 were to withdraw it would lead to a loss of access to liquidity for a large number of intermediaries in the collateral market. The existence of such fragile connectors is crucial in the mechanism underlying Proposition 5.

If shock size is increased, more intermediaries withdraw exogenously, and some intermediaries lose all their counterparties, being forced to withdraw as well. The question is when and why this should result in a discontinuous loss of liquidity. Intuitively, the transition from the liquid to the frozen regime is discontinuous because the complementary nature of the repo and collateral markets produces “fragile connectors” (see Fig. 5 for an illustrative example). An intermediary with few counterparties in the repo market is fragile since it can easily lose access to liquidity. If the same intermediary is an important intermediary of liquidity in the collateral market, it will be a fragile connector. The withdrawal of such an intermediary becomes more likely as the shock size is increased and once it occurs, can have devastating consequences for liquidity in both markets. Note that this result holds for any random network that satisfies the assumptions made in Appendix C.3. It does not hold for arbitrary networks, however. One counterexample comes from stylized core-periphery networks discussed in Section 4. For these networks, a shock can never spread through the core and hence fragile connectors do not exist.

While we have shown here that discontinuities can occur in certain networks as they grow large, our result also has important implications for networks of smaller size. In particular, if these networks display a discontinuity in the limit, there exists a regime in which they are susceptible to full collapse after the removal of a single node. For details and a numerical analysis of this point, see the Online Appendix.

We now discuss the proof of Proposition 5. First, observe liquidity is positive in the $n \rightarrow \infty$

limit if and only if there is a mutually stable set in the two networks that comprises a positive fraction of nodes. In our proof of this result in the appendix, we relate this to a notion from random graph theory called the *mutual giant component* (Buldyrev et al., 2010). Just like the giant component in one graph, this turns out to be unique if it exists. As a result of this key simplification, the fraction of intermediaries in a mutual giant component again satisfies a fixed-point condition: a node’s probability of being in it depends on its neighbors’ probability of being in it. This is analogous to, but more complicated than, the fixed-point equation we discussed above in the one-network case. The degree sequences \mathbf{d}_μ^+ and \mathbf{d}_μ^- (for both values of μ) and the size of the exit shock $1 - x$ enter this equation, in determining how many neighbors can link one to it, and the probability that one is not in it due to an exit shock.³⁴ Thus, the first step is to reduce the study of liquidity to the study of this fixed-point equation. Once it is written down, it becomes possible to study how the solutions depend on the shock size, $1 - x$, and to understand what features of the degree distribution lead to a discontinuous changes. Indeed, our main technical contribution relative to the prior literature is to state general conditions on the degree distribution such that the discontinuity occurs: the main ingredient is giant-component concavity.

It is useful to compare the result to Proposition 4. This result shows that, when one of the two markets is complete—e.g., if it corresponds to a centralized exchange—the transition from the liquid to the frozen market regime is no longer abrupt but *smooth*. In addition, as noted in Proposition 5(B), the transition always occurs at a larger shock size in the presence of a centralized exchange. Here, liquidity is less sensitive to the withdrawal of a single intermediary and can only vary smoothly with the size of the exit shock: a sudden market freeze is not possible. Comparing the two propositions emphasizes the stabilizing effect of a centralized exchange on liquidity in the presence of exit shocks.

The main difference between the case when both networks are genuinely over-the-counter case and the centralized collateral case is the absence of fragile connectors. Since in the complete collateral network all intermediaries receive and provide liquidity to each other, there can be no contagion through the complete network. While an intermediary may be fragile in the repo market, its withdrawal cannot lead to a large number of further withdrawals in the collateral market. The absence of fragile connectors therefore removes the amplification effect that

³⁴Closely related to this fixed-point equation is the fact that we can define two interacting branching processes occurring in the two networks, which in distribution, reflect the extended network neighborhood around a typical node. The degree distributions \mathbf{d}_μ^+ and \mathbf{d}_μ^- and the size of the exit shock $1 - x$ determine the branching probabilities. If, starting from a randomly chosen intermediary, the branching process goes extinct, a node cannot form part of the giant component. This reasoning allows us to specify a system of coupled equations whose greatest solution yields $\mathcal{L}^*(x)$. Appendix C.4 develops this further. Amini et al. (2013) and Elliott et al. (2014), and Buldyrev et al. (2010) use such techniques to characterize giant components, though none of these papers characterize outcomes for arbitrary degree distributions in coupled networks.

results from the complementarity of the repo and collateral markets. This leads to a smooth transition and an increased critical shock size.

5.4 Liquidity for binomial and power law degree distributions

At a qualitative level, the results so far have shown that fragility is a robust property of the systems in question for a large set of degree distributions. For concreteness, we illustrate the results in Propositions 5 and 4 by considering particular degree distributions: binomial (Erdős-Rényi) and scale-free (power law) degree distribution.




Figure 6: Equilibrium liquidity s^* as a function of the fraction of intermediaries $1 - x$ that withdraw from the repo and collateral markets following an exit shock in an Erdős-Rényi network with average degree $\lambda = 5$.

Erdős-Rényi: This is the simplest type of random network, corresponding to random matching of counterparties. This type of graph is obtained by letting each directed link exist with a given probability q . We hold the average in- and out-degree $\lambda = nq$ fixed as n varies. Here, due to the independence of in- and out-degrees the joint degree distribution factorizes into $p_{jk} = p_j p_k$ with $p_j = p_k$ and we have

$$p_k = \binom{n-1}{k} q^k (1-q)^{n-k-1}.$$




Figure 7: Equilibrium liquidity s^* as a function of the fraction of intermediaries $1 - x$ that withdraw from the repo and collateral markets following an exit shock in a scale free network.

Scale free: A more realistic random graph structure models the wide heterogeneity in degrees. As for the Erdős-Rényi networks we assume that the in- and out-degrees are independent, such that the joint degree distribution factorizes into $p_{jk} = p_j p_k$. We take $p_j = p_k$ and $p_k = C_\mu k^{-\alpha}$ for $\alpha \in (2, 3]$ and $k > 1$. The constant that normalizes the degree distribution is $C = 1/(\zeta(\alpha) - 1)$, where $\zeta(\cdot)$ is the Riemann zeta function. The exponent α determines how dispersed the degree distribution is; for $\alpha < 2$, the variance of the degree distribution diverges.

Illustrating the results: We solve for the liquidity measure of the maximal equilibrium $\hat{\mathcal{L}}^*(x)$ numerically; the detailed calculations can be found in Appendix E. For each degree distribution, we also compute $\hat{\mathcal{L}}^*(x)$ when the collateral market is replaced by a complete network. In Fig. 6 and 7 we present the results for the binomial and scale free degree distributions, respectively. The findings of Propositions 5 and 4 are apparent. First, in the case of two coupled OTC markets $(\mathcal{G}_R, \mathcal{G}_C)$ there is a discontinuous transition from the liquid to the frozen market regime. Second, when the collateral market is replaced by a complete network, yielding the pair $(\mathcal{G}_R, \bar{\mathcal{G}}_C)$, the transition is smooth and occurs at a greater shock size. Our results are robust to the choice of parameters as long as the degree distributions satisfy the requirements laid out at the beginning of this section.

It is clear that the discontinuities are stark, and under our parameters, the transition for the power-law case is steeper and happens for a smaller shock size, even though average degree

is actually lower.

5.5 Correlations between repo and collateral networks

Assumption 5.2 imposed that \mathcal{G}_R and \mathcal{G}_C are independent draws from their respective distributions. Under this assumption, an intermediary's counterparties in the repo market are uncorrelated with its counterparties in the collateral market. In many cases however, the presence of a trading relationship between two intermediaries in a given market is correlated with their being linked in another market. Here we discuss how this affects our results.

Let \mathcal{G}_R and \mathcal{G}_C be random networks with the same degree distribution. Then, as $n \rightarrow \infty$ $|E(\mathcal{G}_C)| = |E(\mathcal{G}_R)| = M$. Define the *overlap* measure of network similarity as:

$$\omega = \frac{\#\{i \rightarrow j \mid i \rightarrow j \in E(\mathcal{G}_C) \wedge i \rightarrow j \in E(\mathcal{G}_R)\}}{M}.$$

If \mathcal{G}_R and \mathcal{G}_C are independent, as $n \rightarrow \infty$ the fraction of overlapping edges vanishes and $\omega = 0$. If \mathcal{G}_R is a copy of \mathcal{G}_C , all edges overlap and $\omega = 1$. Let \mathcal{G}_C and \mathcal{G}_R be two Erdős-Rényi random networks with overlap ω . How are the results in Proposition 5 affected by different levels of overlap?

Let us discuss some extreme cases to build intuition. The independent-networks case that has been a focus of our results is essentially the $\omega = 0$ case, because with $n \rightarrow \infty$ and finite degrees, the probability of two independently drawn neighborhoods overlapping tends to 0. Now, for the other extreme, consider $\omega = 1$, so that $\mathcal{G}_R = \mathcal{G}_C$. The maximal equilibrium when $\omega = 1$ is the same as the maximal equilibrium when the collateral network is a complete network. This is because, as in the case where \mathcal{G}_C is complete, the mutually stable subsets are exactly the stable subsets of \mathcal{G}_R . Thus for $\omega = 1$, the results in Proposition 4 apply while for $\omega = 0$, the results in Proposition 5 apply.

Varying the overlap parameter ω then interpolates between these two extremes. For values of overlap, ω , which are not too high, the sensitivity to shocks is still quite extreme, as in the independent-networks case covered by Proposition 5. For higher values, the dependence of liquidity on x is much smoother. Numerical calculations and a heuristic calculation, both detailed in Appendix F, show that the transition occurs at an overlap of approximately $\omega_c = 2/3$ if \mathcal{G}_C and \mathcal{G}_R are Erdős-Rényi random networks.

Thus, as the two over-the-counter markets become more similar, i.e. as intermediaries share more counterparties across markets, liquidity becomes more resilient to exit shocks. This may appear counterintuitive at first, since it seemingly contradicts notion that a diversified set of counterparties protects against random shocks to one's counterparties. However, the com-

plementary nature of repo and collateral markets that makes diversification here harmful rather than helpful and leads instead to an amplification of exit shocks through the over-the-counter markets.

We also present $\hat{\mathcal{L}}^*(x)$ explicitly for two levels of overlap $\omega = \{0.2, 0.8\}$; see Fig. 13 in Appendix F. As expected for $\omega = 0.2$, we observe a discontinuous transition from the liquid to the frozen regime, while for $\omega = 0.8$, we observe a continuous transition.³⁵

6 Discussion

In this section, we discuss a number of our modeling assumptions, examining their robustness as well as natural extensions and questions suggested by the results.

Binary liquidity-provision decisions

We interpret the model laid out in Section 2.1 as a liquidity-provision game between capital- and cash-constrained intermediaries. Importantly, rather than modeling trades in the repo and collateral markets in detail, we model in reduced-form a single decision about whether to provide liquidity in each market. The graph formed by active intermediaries in the maximal equilibrium can be interpreted as the set of *potential* trades. We abstract away from volume and frequency of trade along a given link. Moreover, an intermediary either decides to provide liquidity to all of its counterparties or none. In particular, as long as there is at least one trading partner active in the repo market, and one, possibly different, trading partner active in the market for collateral, intermediary i is willing to be active in both markets.

In a more detailed model that endogenized decisions over which links to keep “open” or “active,” there would be two new forces relative to our model. On the one hand, an intermediary may sometimes choose to stay partly operational in one market if it has that option, while it would have shut down if given a binary option between staying active and shutting down. In this sense our best-response function is a conservative approximation of the best-response function in a richer model: it resolves ambiguity in favor of shutting down. The second force goes in the opposite direction: other intermediaries’ possibility of remaining partly operational may make a given intermediary more likely to remain (partly) operational after a given shock. Thus, there is no easy comparison between our model and one with richer, link-by-link decisions about the willingness to trade. Because the strategic dynamics in this model are already intricate, we view

³⁵We compute $\hat{\mathcal{L}}^*(x)$ both via our approximate method and numerically via algorithm 1. Note that the heuristic solution approximates the numerical solution quite well, though there are clear finite size effects for the numerical solution (we used $n = 2000$).

the current model as a useful simplification.

However, the basic structure of a supermodular game would carry over to a suitable version of the more general link-by-link game. Let $a_{ij}^\mu \in \{0, 1\}$ denote the action of intermediary i vis-a-vis intermediary $j \in K_{i,\mu}^+$. Then one could impose the following constraint: $\sum_{j \in K_{i,\mu}^+} a_{ij}^\mu \leq c_{i,\mu} + \sum_{j \in K_{i,\mu}^-} a_{ji}^\mu$. That is, the liquidity that intermediary i provides to its counterparties is constrained, up to a level $c_{i,\mu}$, by the liquidity it receives from its counterparties. Now the best-responses—how an intermediary withdraws liquidity—would be more complicated, reflecting, for example, intermediary-level heterogeneity such as counterparty risk. The game can still be defined to be supermodular, and so one can again define a maximal equilibrium. The propagation of shocks, however, can be quite different due to the two forces we have sketched. On the one hand, cascades can start from smaller shocks, because intermediaries can react in a less extreme way. On the other hand, sensitivity to parameters will also be less extreme, because liquidity will be withdrawn gradually and a shock that might have triggered full shutdown of a set of intermediaries before will only trigger a milder contagion in a subset of those intermediaries. Such extensions would be interesting to explore in future work.

Intermediaries facing tight constraints

We motivated the intermediaries' best responses by tight capital and cash-in-advance constraints. Intermediaries earn more from their intermediation activity by increasing leverage and reducing cash reserves to boost returns until regulatory constraints become tight. This will be optimal if the intermediary does not anticipate the adverse shock. One way to relax the assumption of tight constraints is to change the best response in Eq. (2) for some intermediaries, making it optimal for them to operate even without active counterparties. These intermediaries can be thought of having slack constraints and would provide liquidity in both markets irrespective of their neighbors' actions, increasing equilibrium liquidity provision. However, this slackness parameter will be endogenous, and is likely to tighten in a crisis scenario.

Network structure

We studied the game of liquidity provision in three types of networks: stylized star and core-periphery networks as well as random networks. Star and core-periphery networks allowed us to develop an intuition for illiquidity spirals and exhibit simple versions of our main results. We then moved on to crisis networks viewed as random networks with an arbitrary degree distribution.

While some results are common across the two settings, other predictions differ markedly.

In view of this, it is worth commenting further on the crisis network model, what it represents, and how it relates to observables. Recall that a directed link represents an unobservable opportunity to trade—in contrast to a measured exposure network that reflects an aggregation of trades across time.³⁶ One way of interpreting the random realization of links in our model is as uncertainty over which trading opportunities are realized, even in normal times—e.g., which counterparties are available to trade when they are needed (Duffie et al. (2007) emphasize the “by appointment” nature of OTC trade in many markets, suggesting that frictions in finding trading opportunities are considerable). Empirically, measured exposure networks can have rich degree distributions, as we discuss in Appendix A; our model is designed to capture these flexibly, along with the uncertainty in trading opportunities just discussed.

Our most important reason for taking the random network approach, however, is that we are concerned with behavior during crises. In such times, the network of trading opportunities changes considerably from its normal configuration and is therefore inherently uncertain from the perspective of a policy maker conducting a macroprudential stress test prior to the crisis. This is the case even if trading opportunities in normal times are *not* uncertain from the analyst’s perspective.

As a robustness check, we use numerical calculations to study a class of network structures that interpolate between core-periphery and random networks with a given degree distribution (see the Online Appendix). These can be viewed as core-periphery networks with noise—i.e., ones whose structure is perturbed by a random shock. Our results show that even a slight perturbation of this form (i.e., slight planner uncertainty) gives rise to the systemic fragility we have identified. Indeed, the transition to the fragile regime (with market freezes) is caused by relatively small reductions in the links among the most connected nodes. This further underlines the stark impact of aggregate shocks even when they are small. It also motivates empirical research into the structure of trading networks both in normal and crisis times.

Also in the Online Appendix we show that the abrupt evaporation of liquidity observed for random networks is not limited to large networks. Even in small networks, liquidity may evaporate abruptly, if particularly critical intermediaries fail. Such critical intermediaries tend to be intermediaries that are sole liquidity providers to intermediaries that themselves provide liquidity to many other intermediaries.

³⁶For some challenges involved in inferring the network structure from observational data, see Golub and Jackson (2010) and Chierichetti, Liben-Nowell, and Kleinberg (2011).

7 Relation to the Literature

Market freezes in theory. In the liquidity-provision game between intermediaries outlined in the previous section, an intermediary's choice to be active in either market determines whether it provides market and funding liquidity to its trading partners. While we make simple reduced-form assumptions on liquidity provision in order to focus attention on network considerations, our model relates to an active literature on liquidity hoarding as a source for financial market freezes. This literature considers the precise mechanisms of liquidity provision in more detail. Banks in [Gale and Yorulmazer \(2013\)](#) choose to hoard liquidity, even if there is a willing borrower in the market, because of a precautionary or a speculative motive. [Heider, Hoerova, and Holthausen \(2015\)](#) show that interbank markets can break down and banks start hoarding liquidity if banks have private information about their assets and adverse selection is prevalent. [Bond and Leitner \(2015\)](#) show that a freeze in the market for an asset can arise when traders hold inventories of similar assets and their leverage constraints are tight. Empirical evidence corroborates these theoretical models. [Ashcraft, McAndrews, and Skeie \(2011\)](#) and [Acharya and Merrouche \(2013\)](#), for example, show that banks in the US and UK indeed were hoarding liquidity during the global financial crisis. Liquidity in an OTC market is, unlike in centralized markets, *local*: due to frictions, a situation can exist in such markets where some market participants have liquidity supply while others have liquidity demand, but no trade ensues because they don't have a link.³⁷ The ensuing over-the-counter network structure is then the result of the individual decisions by market participants whether or not to be active in the market.

Empirical evidence of financial market freezes. A number of authors empirically study the fragility of repo markets during the 2007/2008 financial crisis. [Gorton and Metrick \(2012\)](#) argue that a central aspect of the crisis was a system-wide run on short-term collateralized debt, and in particular on certain non-government bond repo markets. [Krishnamurthy, Nagel, and Orlov \(2014\)](#) show that, while the tri-party repo market has been more stable during the crisis than bilateral repo, markets for asset-backed commercial paper (ABCP) experienced a significant contraction.³⁸ This finding is mirrored by [Copeland, Martin, and Walker \(2014\)](#), who document substantial heterogeneity in access even to tri-party repo funding in late 2008.

Illiquidity spirals. In the literature on secured lending, the theoretical papers most closely related to ours are [Brunnermeier and Pedersen \(2009\)](#) and [Acharya, Gale, and Yorulmazer \(2011\)](#).

³⁷One reason why market participants with liquidity supply do not provide liquidity to those with liquidity demand is search frictions in over-the-counter markets.

³⁸[Covitz, Liang, and Suarez \(2012\)](#) study the fragility of asset-backed commercial paper markets. Our model naturally extends to ABCP markets.

Our model of illiquidity spirals in repo and collateral markets explicitly takes into account the OTC network structure of these markets. This sets us apart from [Brunnermeier and Pedersen \(2009\)](#) who study the feedback between market and funding liquidity in centralized markets. Our model shows that the network structure alone, abstracting from haircut and pricing feedback, can lead to a significant amplification of exogenous shocks in collateral and repo markets. [Acharya, Gale, and Yorulmazer \(2011\)](#) show how a bank's ability to obtain secured funding depends on the risk and liquidation value of the collateral and how this dependency leads to a feedback between collateral and debt markets mediated by the debt capacity (essentially, quantity) offered. In contrast to these works, we show that, given the complementarity of collateral and secured debt markets the over-the-counter nature of these markets is sufficient to generate a feedback between market and funding liquidity that amplifies exogenous shocks.

Our paper is also related to [Martin, Skeie, and von Thadden \(2014\)](#) who model repo runs arising from pure coordination failure in a dynamic model. They show that repo markets in which haircuts cannot be adjusted, such as a centralized or tri-party repo market, can be more fragile than bilateral repo markets in which haircuts can be adjusted. We contribute to this debate by showing that if at least one OTC market is replaced by a centralized exchange, the resilience of liquidity improves. This suggests at least two opposing effects that need to be taken into account when judging the merits of centralized exchanges: the flexibility of haircuts and adverse network effects.

Financial over-the-counter networks. Most empirical studies of OTC markets focus on financial exposure networks. Approximate core-periphery structures are often observed (see [Di Maggio, Kermani, and Song \(2017\)](#) on the inter-dealer corporate bond market and [Li and Schürhoff \(2018\)](#) on the market for municipal bonds). Much contemporary interbank lending is secured, making the interbank market another relevant reference case. [Craig and Von Peter \(2014\)](#) show, for example, that the German interbank market generally has a core-periphery structure with noise. In the international context, [Gabrieli and Georg \(2014\)](#) show that the Euroarea interbank market follows a core-periphery structure less closely, with large international banks connecting the different national core-periphery networks. Because of the sensitivity of trading outcomes even to small aggregate shocks, our result suggest the examination of networks of trading opportunities (as distinct from networks of exposures), especially during crises or times of volatility.

Contagion in financial networks. A large literature studies contagion in financial networks that ensues when the default of one financial institution causes the subsequent default of other financial institutions (see, for example, [Acemoglu, Ozdaglar, and Tahbaz-Salehi \(2015\)](#), [Elliott,](#)

Golub, and Jackson (2014), Zawadowski (2013), and Farboodi (2017), as well as Glassermann and Young (2016b) for an extensive overview). Burkholz, Leduc, Garas, and Schweitzer (2016) study cascading failures in a multiplex network. Firms in their model have a core and a subsidiary business unit who are each exposed to possible contagion within their respective network of business relationships. A default of the core business unit will lead to a default of the subsidiary, but not necessarily vice versa. Burkholz, Leduc, Garas, and Schweitzer (2016) observe a similar amplification mechanism between network layers to ours and find that the extent of the amplification is sensitive to the strength of the coupling between two two layers. While our model could be interpreted as a contagion model, contagion occurs via the intermediaries' decision variables to withdraw from markets rather than via actual defaults. This is a key difference to the existing literature on contagion in financial networks and in line with the empirical evidence from the 2008 financial crisis: even the default of a large intermediary, such as Lehman Brothers, did not trigger many subsequent defaults,³⁹ while it is likely to have led to a freeze in markets for short-term collateralized debt.

Multilayer network theory. There is a growing literature in applied mathematics and physics on coupled, or *multilayer*, networks. A seminal paper is Buldyrev, Parshani, Pau, Stanley, and Havlin (2010), and since then there have been a variety of applications, including, e.g., to the question of whether firms should spin off subsidiary units—see, e.g., Burkholz, Leduc, Garas, and Schweitzer (2016). In terms of the theory of coupled networks, our contribution is to formulate general conditions on the degree distributions of each of the trading networks that yield the stark conclusions (e.g., about discontinuous collapses in coupled markets). These conditions are satisfied by standard network structures, but the general conditions were not previously known. The relation to of coupled networks to certain games of strategic complements that we identify may also be of independent interest.

Games on networks. Our paper is also related to the networks literature in economic theory, especially contagion and games on networks. Papers such as Blume (1993) and Ellison (1993) first suggested that local interaction, modeled via a network structure, can be used to study the likelihood that various equilibria would be played and how an economy may reach an equilibrium. Whereas these early papers focused on noisy heuristic adjustment procedures, Morris (1997, 2000) studied games with standard (no-noise) solution concepts and related networks to games of incomplete information. The latter paper's results, applied to a network game, show when a network can support heterogeneous actions, and what conditions result in equilibria

³⁹An exception is the Reserve Primary Fund, who, due to a large exposure to Lehman Brothers, filed for bankruptcy the day after Lehman Brothers filed for bankruptcy.

such as (in our context) “everyone withdraws.” Jackson and Yariv (2007) and Galeotti, Goyal, Jackson, Vega-redondo, and Yariv (2009) developed this sort of model to accommodate random networks described by a degree distribution. Our approach has much in common with this theoretical literature on games in networks. We also use the structure of supermodular games (as in Milgrom and Roberts (1990)) to identify benchmark equilibria, and look at their structure for large random networks. The main innovation relative to these papers on games in networks is that we study multilayer networks, and analyze how the multilayer aspect of their structure affects the best-response structure of the game, especially when the underlying networks are random. Equilibria depend more sharply on the parameters on the network than has been reported previously, due to the discontinuities discussed above. Thus, our paper relates to the network games literature broadly, and offers new game-theoretic implications arising from multilayer network structures.

8 Conclusion

We develop a model of intermediary liquidity provision in over-the-counter repo and collateral markets and study the maximal equilibria of the resulting complete-information game in two coupled networks from the perspective of a policy analyst conducting a macroprudential stress test. The coupling occurs through the complementary nature of repo and collateral trading and is reflected in the best responses. In particular, as long as there is at least one trading partner active in the repo market, and one, possibly different, trading partner active in the market for collateral, an intermediary is willing to be active in both markets.

The presence of *fragile connectors*—intermediaries that are on the brink of isolation in one market and critical intermediaries in the other—makes equilibrium liquidity fragile. In particular, the withdrawal of such a fragile connector can lead to a sudden market freeze. Even in the absence of fragile connectors, the complementary nature of repo and collateral markets can amplify exit shocks and lead to illiquidity spirals. Replacing at least one OTC market by a centralized exchange reduces the extent of illiquidity spirals and improves the resilience of liquidity. Increasing overlap—making the two networks more similar—also increases resilience.

Moving forward, a natural next step would be to extend the model to account for incomplete information. In this setting, the realization of the shock profile or the parameters entering intermediaries’ best response functions may only be partially known to the agents. This would, for example, allow us to study how liquidity is affected by changes in intermediaries’ beliefs about the distribution of the exogenous shock. In turn, this would permit a study of the realistic phenomenon that liquidity may evaporate when bad news is published, if that news

coordinates beliefs in a suitable way (Angeletos and Werning, 2006; Golub and Morris, 2017a,b; Morris and Yildiz, 2017).

Another natural set of theoretical questions concerns the resilience of the network and how the potential abruptness of liquidity evaporation changes individual nodes' incentives to invest in protective measures—e.g., due diligence. Public goods problems in networks have received a great deal of study (Bramoullé and Kranton, 2007; Allouch, 2015; Elliott and Golub, 2018), but the distinctive externalities associated with multilayer phenomena have yet to be explored.⁴⁰

Our model raises a number of issues for policymakers. First, we illustrate the potential fragility of liquidity in over-the-counter markets and show how it may be reduced by moving towards centralized exchanges. Second, our results highlight the importance of better measurement and empirical and study of the structure of these markets, in particular with respect to fragile connectors.

References

- ABAD, J., I. ALDASORO, C. AYMANNIS, M. D'ERRICO, L. F. ROUSOVA, P. HOFFMANN, S. LANGFIELD, AND T. ROUKNY (2016): "Shedding light on dark markets: first insights from the new EU-wide OTC derivatives dataset," *European Systemic Risk Board Occasional Paper Series*.
- ACEMOGLU, D., A. OZDAGLAR, AND A. TAHBAZ-SALEHI (2015): "Systemic Risk and Stability in Financial Networks," *American Economic Review*, 105, 564–608.
- ACHARYA, V., G. AFONSO, AND A. KOVNER (2017): "How do global banks scramble for liquidity? Evidence from the asset-backed commercial paper freeze of 2007," *Journal of Financial Intermediation*, 30, 1–34.
- ACHARYA, V. AND O. MERROUCHE (2013): "Precautionary hoarding of liquidity and interbank markets: Evidence from the subprime crisis," *Review of Finance*, 17, 107–160.
- ACHARYA, V., P. SCHNABL, AND G. SUAREZ (2013): "Securitization Without Risk Transfer," *Journal of Financial Economics*, 107, 515–536.
- ACHARYA, V. V., D. GALE, AND T. YORULMAZER (2011): "Rollover risk and market freezes," *The Journal of Finance*, 66, 1177–1209.
- ALLOUCH, N. (2015): "On the private provision of public goods on networks," *Journal of Economic Theory*, 157, 527–552.
- AMINI, H., R. CONT, AND A. MINCA (2013): "Resilience to contagion in financial networks," *Mathematical finance*.
- ANDERSON, R., C. BABA, J. DANIELSSON, U. S. DAS, H. KANG, AND M. SEGOVIANO (2018): "Macroprudential stress tests and policies: Searching for Robust and Implementable Frameworks," Report, International Monetary Fund.

⁴⁰There is also the additional subtlety of how the discontinuities we identify interact with discontinuities in network density that come from strategic effects—see, e.g., Golub and Livne (2010).

- ANGELETOS, G.-M. AND I. WERNING (2006): “Crises and Prices: Information Aggregation, Multiplicity, and Volatility,” *American Economic Review*, 96, 1720–1736.
- AQUILINA, M. AND F. SUNTHEIM (2017): “Liquidity in the UK corporate bond market: evidence from trade data,” *The Journal of Trading*, 12, 67–80.
- ASHCRAFT, A., J. MCANDREWS, AND D. SKEIE (2011): “Precautionary Reserves in the Interbank Market,” *Journal of Money, Credit, and Banking*, 43, 311–348.
- ATHREYA, K. B. AND P. JAGERS (2012): *Classical and modern branching processes*, vol. 84, Springer Science & Business Media.
- AYMANN, C., J. D. FARMER, A. M. KLEINNIJENHUIS, AND T. WETZER (2018): “Models of financial stability and their application in stress tests,” *Handbook of Computational Economics*, 4, 329–391.
- BAKLANOVA, V., A. M. COPELAND, AND R. MCCAUGHRIN (2015): “Reference guide to US repo and securities lending markets,” *FRB of New York Working Paper No. FEDNSR740*.
- BARABÁSI, A.-L. AND R. ALBERT (1999): “Emergence of scaling in random networks,” *science*, 286, 509–512.
- BLUME, L. E. (1993): “The Statistical Mechanics of Strategic Interaction,” *Games and Economic Behavior*, 5, 387–424.
- BOND, P. AND Y. LEITNER (2015): “Market run-ups, market freezes, inventories, and leverage,” *Journal of Financial Economics*, 115, 155–167.
- BRAMOULLÉ, Y. AND R. KRANTON (2007): “Public goods in networks,” *Journal of Economic Theory*, 135, 478–494.
- BRITTON, T., S. JANSON, A. MARTIN-LÖF, ET AL. (2007): “Graphs with specified degree distributions, simple epidemics, and local vaccination strategies,” *Advances in Applied Probability*, 39, 922–948.
- BRUNNERMEIER, M. K. AND L. H. PEDERSEN (2009): “Market liquidity and funding liquidity,” *Review of Financial Studies*, 22, 2201–2238.
- BULDYREV, S. V., R. PARSHANI, G. PAU, H. E. STANLEY, AND S. HAVLIN (2010): “Catastrophic cascade of failures in interdependent networks,” *Nature*, 464, 1025–1028.
- BURKHOLZ, R., M. V. LEDUC, A. GARAS, AND F. SCHWEITZER (2016): “Systemic Risk in multiplex networks with asymmetric coupling and threshold feedback,” *Physica D: Nonlinear Phenomena*, 323–324, 64–72.
- CABRALES, A., P. GOTTARDI, AND F. VEGA-REDONDO (2017): “Risk sharing and contagion in networks,” *The Review of Financial Studies*, 30, 3086–3127.
- CHIERICHETTI, F., D. LIBEN-NOWELL, AND J. M. KLEINBERG (2011): “Reconstructing patterns of information diffusion from incomplete observations,” in *Advances in Neural Information Processing Systems*, 792–800.
- CHOI, S., A. GALEOTTI, AND S. GOYAL (2017): “Trading in networks: theory and experiments,” *Journal of the European Economic Association*, jvw016.
- CONDORELLI, D., A. GALEOTTI, AND L. RENOU (2016): “Bilateral Trading in Networks,” *The Review of Economic Studies*, Forthcoming.
- CONSTÂNCIO, V. (2017): “Macroprudential stress-tests and tools for the non-bank sector,” Prepared remarks at the esrb annual conference, frankfurt am main, 22 september 2017, European Central Bank.
- COOPER, C. AND A. FRIEZE (2004): “The size of the largest strongly connected component of a random digraph with a given degree sequence,” *Combinatorics, Probability and Computing*, 13, 319–337.

- COPELAND, A., A. MARTIN, AND M. WALKER (2014): "Repo runs: Evidence from the tri-party repo market," *The Journal of Finance*, 69, 2343–2380.
- COVITZ, D. M., N. LIANG, AND G. SUAREZ (2012): "The evolution of a financial crisis: Collapse of the asset-backed commercial paper market," *The Journal of Finance*, 68, 815–848.
- CRAIG, B. AND G. VON PETER (2014): "Interbank tiering and money center banks," *Journal of Financial Intermediation*, 23, 322–347.
- DEANGELO, H. AND R. STULZ (2015): "Liquid-claim production, risk management, and bank capital structure: Why high leverage is optimal for banks," *Journal of Financial Economics*, 116, 219–236.
- DI MAGGIO, M., A. KERMANI, AND Z. SONG (2017): "The Value of Trading Relationships in Turbulent Times," *Journal of Financial Economics*, forthcoming.
- DUARTE, F. AND T. EISENBACH (2015): "Fire-Sale Spillovers and Systemic Risk," Staff report no. 645, Federal Reserve Bank of New York.
- DUFFIE, D., N. GÁRLEANU, AND L. H. PEDERSEN (2007): "Valuation in over-the-counter markets," *The Review of Financial Studies*, 20, 1865–1900.
- DUFFIE, D., S. MALAMUD, AND G. MANSO (2014): "Information Percolation in Segmented Markets," *Journal of Economic Theory*, 153, 1–32.
- EISENBERG, L. AND T. H. NOE (2001): "Systemic risk in financial systems," *Management Science*, 47, 236–249.
- ELLIOTT, M. AND B. GOLUB (2018): "A network approach to public goods," *Journal of Political Economy*, 126, forthcoming.
- ELLIOTT, M. AND J. HAZELL (2017): "Financial Networks: Efficient Modularity and Why Shareholders Prevent It," Working paper, University of Cambridge.
- ELLIOTT, M. L., B. GOLUB, AND M. O. JACKSON (2014): "Financial networks and contagion," *American Economic Review*, 104, 3115–3153.
- ELLISON, G. (1993): "Learning, Local Interaction, and Coordination," *Econometrica*, 61.
- EROL, S. AND R. VOHRA (2017): "Network Formation and Systemic Risk," Working paper.
- FAINMESSER, I. AND A. GALEOTTI (2017): "Pricing network effects," *Review of Economic Studies*, 83, 165–198.
- FARBOODI, M. (2017): "Intermediation and Voluntary Exposure to Counterparty Risk," Princeton University Working Paper, available at SSRN: ssrn.com/abstract=2535900.
- GABRIELI, S. AND C.-P. GEORG (2014): "A Network view on Interbank Liquidity," Discussion Paper 44/2014, Deutsche Bundesbank.
- GAI, P., A. G. HALDANE, AND S. KAPADIA (2011): "Complexity, Concentration, and Contagion," *Journal of Monetary Economics*, 58, 453–470.
- GALE, D. AND T. YORULMAZER (2013): "Liquidity Hoarding," *Theoretical Economics*, 8, 291–324.
- GALEOTTI, A., B. GOLUB, AND S. GOYAL (2017): "Targeting Interventions in Networks," Available at ssrn.com/abstract=3054353.
- GALEOTTI, A., S. GOYAL, M. O. JACKSON, F. VEGA-REDONDO, AND L. YARIV (2009): "Network Games," *Review of Economic Studies*, 77, 218–244.

- GLASSERMANN, P. AND H. P. YOUNG (2016a): "Contagion in Financial Networks," *Journal of Economic Literature*, 54, 779–831.
- (2016b): "Contagion in Financial Networks," *Journal of Economic Literature*, 54, 779–831.
- GOFMAN, M. (2014): "A Network-Based Analysis of Over-the-Counter Markets," Working paper, Wisconsin School of Business.
- (2017): "Efficiency and stability of a financial architecture with too-interconnected-to-fail institutions," *Journal of Financial Economics*, 124, 113–146.
- GOLUB, B. AND M. O. JACKSON (2010): "Using selection bias to explain the observed structure of Internet diffusions," *Proceedings of the National Academy of Sciences*, 107, 10833–10836.
- GOLUB, B. AND Y. LIVNE (2010): "Strategic Random Networks: Why Social Networking Technology Matters," [Ssrn.com/abstract=1577209](http://ssrn.com/abstract=1577209).
- GOLUB, B. AND S. MORRIS (2017a): "Expectations, Networks and Conventions," Available at SSRN: <https://ssrn.com/abstract=2979086>.
- (2017b): "Higher-Order Expectations," Available at SSRN: <http://ssrn.com/abstract=2979089>.
- GORTON, G. AND A. METRICK (2012): "Securitized banking and the run on repo," *Journal of Financial Economics*, 104, 425–451.
- GOYAL, S. AND F. VEGA-REDONDO (2007): "Structural holes in social networks," *Journal of Economic Theory*, 137, 460–492.
- GREENWOOD, R., A. LANDIER, AND D. THESMAR (2015): "Vulnerable Banks," *Journal of Financial Economics*, 115, 471–485.
- HEIDER, F., M. HOEROVA, AND C. HOLTHAUSEN (2015): "Liquidity hoarding and interbank market rates: The role of counterparty risk," *Journal of Financial Economics*, 118, 336–354.
- HENDERSHOTT, T. AND A. J. MENKVELD (2014): "Price pressures," *Journal of Financial Economics*, 114, 405–423.
- ICMA (2016): "European Repo Market Survey," Tech. rep., ICMA.
- JACKSON, M. O. AND L. YARIV (2007): "Diffusion of Behavior and Equilibrium Properties in Network Games," *American Economic Review*, 97, 92–98.
- KRISHNAMURTHY, A., S. NAGEL, AND D. ORLOV (2014): "Sizing Up Repo," *The Journal of Finance*, 69, 2381–2417.
- LI, D. AND N. SCHÜRHOFF (2018): "Dealer Networks," *The Journal of Finance*, forthcoming.
- MALAMUD, S. AND M. ROSTEK (2017): "Decentralized exchange," *American Economic Review*, 107, 3320–3362.
- MANEA, M. (2015): "Intermediation and resale in networks," Tech. rep., Working paper.
- MARTIN, A., D. SKEIE, AND E. VON THADDEN (2014): "Repo Runs," *The Review of Financial Studies*, 27, 957–989.
- MILGROM, P. AND J. ROBERTS (1990): "The economics of modern manufacturing: Technology, strategy, and organization," *The American Economic Review*, 511–528.
- MORRIS, S. (1997): "Interaction Games: A Unified Analysis of Incomplete Information, Local Interaction and Random Matching," Santa Fe Institute Working Paper No. 97-08-072E.
- (2000): "Contagion," *Review of Economic Studies*, 67, 57–78.

- MORRIS, S. AND M. YILDIZ (2017): “Crises: Equilibrium Shifts and Large Shocks,” Mimeo., MIT.
- NEWMAN, M. (2010): *Networks: an introduction*, Oxford University Press.
- NEWMAN, M. E. (2002): “Spread of epidemic disease on networks,” *Physical review E*, 66, 016128.
- OEHMKE, M. (2014): “Liquidating illiquid collateral,” *Journal of Economic Theory*, 149, 183–210.
- PERIGNON, C., D. THESMAR, AND G. VUILLEMEY (2017): “Wholesale Funding Dry-Ups,” *The Journal of Finance*, forthcoming.
- SON, S.-W., G. BIZHANI, C. CHRISTENSEN, P. GRASSBERGER, AND M. PACZUSKI (2012): “Percolation theory on interdependent networks based on epidemic spreading,” *EPL (Europhysics Letters)*, 97, 16006.
- WANG, C. (2016): “Core-periphery trading networks,” .
- ZAWADOWSKI, A. (2013): “Entangled Financial Systems,” *Review of Financial Studies*, 26, 1291–1323.
- ZHOU, L. (1994): “The set of Nash equilibria of a supermodular game is a complete lattice,” *Games and Economic Behavior*, 7, 295–300.

A The structure of financial networks during crises

For a policy analyst conducting a macroprudential stress test of the type outlined in Section 3.3, what is relevant is the structure of the networks of trading opportunities during crisis times. Our model of this network is a random network with a given degree distribution that can be flexibly specified by the modeler. This section documents some empirical observations that explain why this model was chosen.⁴¹

The network of trading opportunities is always latent and therefore difficult to measure empirically. However, a natural starting point is the network of observed exposures during the normal times. In Figure 8, we plot the cumulative degree distribution of overnight interbank exposures aggregated over a reserve maintenance period (RMP)⁴² among euro area banks. The first observation is that this network has a different distribution than one would see in a core-periphery network of the same size.⁴³ Rather than having a sharp distinction in degree between

⁴¹The final subsection of Section 6 discusses some more general considerations on the identification of this network.

⁴²To produce this plot, we used the data from Gabrieli and Georg (2014), who study the network of euro area overnight interbank loans obtained from the Target2 payment system. A reserve maintenance period corresponds roughly to a month and is the natural unit of aggregation as banks are required to hold minimum average reserves over this period commensurate with their deposits. The overnight interbank market is the main mechanism how these reserves are redistributed among banks.

⁴³Consider a core-periphery network with n_C core nodes and n_P peripheral nodes, each linked to one core node; we would expect a bi-modal degree distribution with one peak around degree one (corresponding to the peripheral nodes) and another around degree $n_C - 1 + n_P/n_C$ (corresponding to the core nodes). With peripheral nodes connecting to more core nodes, the distribution would have larger support, but there would still be a break in the degree distribution corresponding to the structural difference between the core and periphery.




Figure 8: Cumulative degree distribution of overnight interbank loans in Europe from the Target2 payment system in the reserve maintenance period starting in September 2008 compared with a stylized core-periphery network of the same size.

core and periphery, the distribution in Figure 8 appears heavy-tailed.⁴⁴ Figure 8 suggests that stylized core-periphery networks cannot capture the rich structure of observed exposures: a model with a more flexible degree distribution is needed, even for network structure during normal times, which is consistent with our modeling approach.

Our second observation, illustrated by Figure 9, is that this degree distribution thinned out substantially during the height of the global financial crisis in September 2008. We plot the distribution of changes in the number of neighbors a bank has from a given RMP relative to the RMP immediately prior. We label RMPs by the month they start in (throughout our sample, this label identifies an RMP uniquely). A negative mean indicates that connections are cut and the network is thinned out, while a positive mean implies more connections are being created. The mean of the change of a bank's number of neighbors in July 2008 is 1.96, while in September 2008 it is -1.43 and in July it is -0.25 .⁴⁵ That the network continued thinning out is in line with the continuing crisis in the euro area that turned into the euro area sovereign debt crisis in early

⁴⁴Note that, for confidentiality reasons, the CDF depicted must be censored at degree 50, but the figure is qualitatively unchanged without the censoring.

⁴⁵Since the Target2 payment system became operational in the last country of the euro area only in May 2008, we do unfortunately not have data for any date before then.

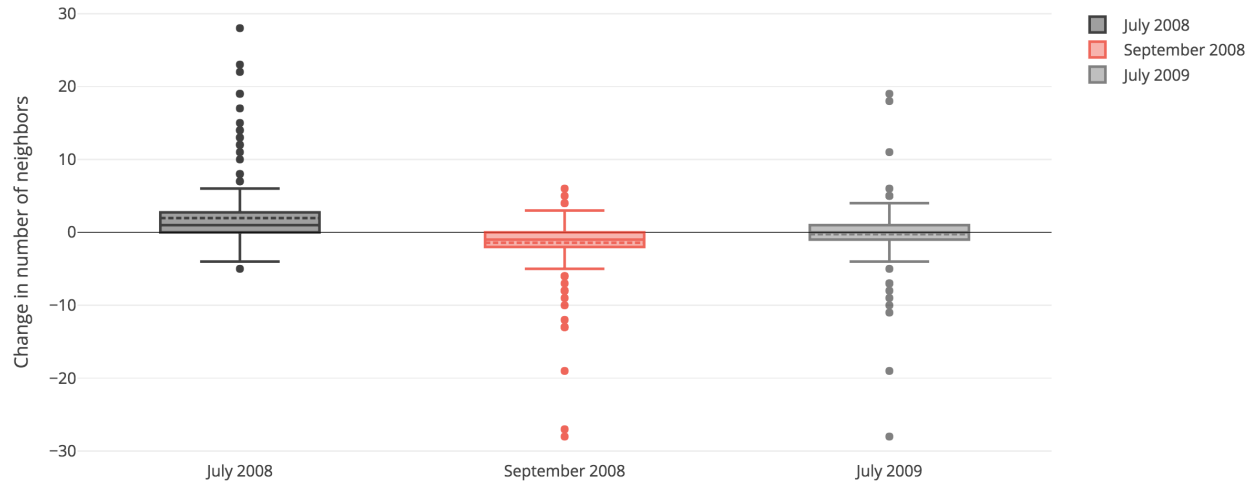


Figure 9: Box plot of the change of the number of a bank’s neighbors in the euro area overnight interbank network aggregated on reserve maintenance period basis. Changes are computed from one reserve maintenance period to the next. A negative change implies the a given bank has lost neighbors.

2010.

Our final observation is that even insofar as exposure networks pre-crisis are well-approximated by simple benchmarks such as core-periphery structures (as discussed, e.g., in (Craig and Von Peter, 2014; Li and Schürhoff, 2018)), the shocks to network structure during a crisis can change the situation considerably. Around the insolvency of the US investment bank Lehman Brothers in September 2008, Gabrieli and Georg (2014) show that about 52% of the links in the euro area overnight interbank market violated the stylized core-periphery structure and a large fraction of links changed. The random network models we use in Section 5 are designed to be flexible enough to capture important features of network structure that are consequences of this “thinning out.”

B Detailed treatment of core-periphery networks

In this section we provide a full treatment of our core-periphery results outlined in Section 4.2. We will show that, as for star networks, an adverse shock leads to the withdrawal of additional intermediaries and that the extent of this amplification depends on how many intermediaries are peripheral in one network and central in the other.

Nodes in the stylized core-periphery network we will study are partitioned into two sets: a set of *core* nodes and a set of *peripheral* nodes. A core node is connected to all other core nodes and a subset of peripheral nodes via bi-directional links. A peripheral node is connected only

to a single core node via a bi-directional link (see Fig. 3). Let $\Omega : N \rightarrow \{cc, cp, pc, pp\}$ denote the map that assigns each intermediary to either the core or the periphery of networks \mathcal{G}_C and \mathcal{G}_R . For example, an intermediary that is in the core of \mathcal{G}_C and \mathcal{G}_R would be labelled cc while an intermediary that is in the core of \mathcal{G}_R and the periphery of \mathcal{G}_C would be labelled cp .

We study a random way of generating core-periphery networks: we fix the total number of intermediaries with each label $n_{tp} = \#\{i \mid \Omega(i) = tp\}$, and generate core-periphery networks uniformly at random consistent with this. Thus, the parameter vector $\mathbf{n} = (n_{cc}, n_{cp}, n_{pc}, n_{pp})$ fully determines the distribution of coupled core-periphery networks we will study in this section.

We parameterize a shock profile w by the vector $\mathbf{m} = (m_{cc}, m_{cp}, m_{pc}, m_{pp})$ whose elements are the number of intermediaries of a particular type with $w_i = 1$. In other words, the shock profile gives the number of shocked intermediaries of each type. The shocks are drawn uniformly at random subject to these shock sizes. We denote the complement vector, i.e. the number of intermediaries with $w_i = 0$, by $\overline{\mathbf{m}} = (\overline{m}_{cc}, \overline{m}_{cp}, \overline{m}_{pc}, \overline{m}_{pp})$. The shock size is given by $\overline{m}_W = \sum_{tp} \overline{m}_{tp}$.

For an intermediary i in the core of both networks, let \mathcal{P}_i^R be the set of unshocked peripheral neighbors of i in \mathcal{G}_R , i.e.

$$\mathcal{P}_i^R = \{j \in \mathcal{G}_R \mid i \rightarrow j \wedge w_j = 1\}.$$

Let the random variable k_R (respectively, k_C) denote the size of the set \mathcal{P}_i^R (respectively, \mathcal{P}_i^C). Given the distribution of network and shocks, the expected values of k_R and k_C will be,

$$\begin{aligned} \overline{k}_R(\mathbf{m}) &= E[k_R \mid m_{pc}, m_{pp}] = \frac{m_{pc} + m_{pp}}{n_{cc} + n_{cp}}, \\ \overline{k}_C(\mathbf{m}) &= E[k_C \mid m_{cp}, m_{pp}] = \frac{m_{cp} + m_{pp}}{n_{cc} + n_{pc}}. \end{aligned}$$

We can derive an expression for the expected post-shock liquidity measure by averaging over all parameterizations of the shock profile. For this we first compute an expression for the expected number of additional intermediaries that will withdraw conditional on the withdrawal of an intermediary of a particular type, e.g. cc .

Lemma 2 (Type conditional spill-over). Let $z(\mathbf{m}) \in [0, 1]$ denote the expected fraction of peripheral intermediaries in \mathcal{G}_C that are not connected to a shocked intermediary (i such that $w_i = 0$) in the core of \mathcal{G}_R . Then the expected number of additional intermediaries that will withdraw conditional on the withdrawal of an intermediary of a particular type is given by

- cc (core-core): $a_{cc} = \overline{k}_R + \overline{k}_C z$,




Figure 10: Shared peripheral intermediaries in a core periphery network

- cp (core-periphery): $a_{cp} = \bar{k}_R$,
- pc (periphery-core): $a_{pc} = \bar{k}_C z$,
- pp (periphery-periphery): $a_{pp} = 0$.

To develop intuition for this result, let us consider the effects of the withdrawal of an intermediary given its type, i.e. cc , cp , pc or pp . If a core intermediary withdraws in a given network, only its peripheral neighbors will withdraw. If a peripheral intermediary withdraws, no additional intermediary will withdraw. If a cc intermediary withdraws, its peripheral neighbors in both networks will withdraw. If a cp (pc) intermediary withdraws, only peripheral neighbors in \mathcal{G}_R (\mathcal{G}_C) will withdraw. Finally, if a pp intermediary withdraws, no additional intermediary withdraws.

The number of additional intermediaries that withdraw following the withdrawal of a particular intermediary would be easy to compute if core intermediaries in \mathcal{G}_C and \mathcal{G}_R did not share peripheral intermediaries: for a core intermediary, this would simply be the number of peripheral intermediaries that rely on it in either network. However, if $\sum_i w_i$ becomes large relative to the network size n , some withdrawing core intermediaries are quite likely to share periphery intermediaries. Fig. 10 illustrates this situation. Here the withdrawing core intermediaries labeled 1 and 2 share a peripheral neighbor labeled 5. Treating the amplification effect of 1 and 2 as independent would result in the double-counting of intermediary 5. We can correct for this double-counting by appropriately scaling the expected number of peripheral neighbors of a core node in one of the two networks. In Lemma 2, the correction consists of scaling \bar{k}_C by the

fraction (z) of peripheral intermediaries in \mathcal{G}_C that are not connected to a core intermediary in \mathcal{G}_R with $w_i = 0$. The following result gives the consequence for aggregate activity of the effect quantified in Lemma 2:

Proposition 6 (Core-periphery – post shock liquidity measure). *The expected number of withdrawing intermediaries given $\bar{\mathbf{m}}$ and \mathbf{n} is given by*

$$A(\bar{\mathbf{m}}) = (1 + a_{cc})\bar{m}_{cc} + (1 + a_{cp})\bar{m}_{cp} + (1 + a_{pc})\bar{m}_{pc} + (1 + a_{pp})\bar{m}_{pp}.$$

The expected number of intermediaries providing liquidity conditional on shock size \bar{m}_W is given by

$$\hat{\mathcal{L}}^{cp-cp}(\bar{m}_W) = E[\mathcal{L} \mid \bar{m}_W] = 1 - \frac{1}{n} \sum_{\bar{\mathbf{m}}} A(\bar{\mathbf{m}}) P(\bar{\mathbf{m}}, \mathbf{n})$$

where $P(\cdot, \cdot)$ is the multivariate hypergeometric distribution with parameters \mathbf{n} and $\bar{\mathbf{m}}$.

It is possible to derive an exact expression for the quantity $z(\mathbf{m})$ defined in Lemma 2, which figures in a_{cc}, a_{cp}, a_{pc} and a_{pp} . Cleaner expressions are obtained using the approximation $z(\mathbf{m}) \approx 1 - (m_{cc} + m_{cp}) / (n_{cc} + n_{cp})$ for simplicity.⁴⁶ This approximation performs well in numerical experiments and is used in our numerical results, discussed below for core-periphery networks.

The quantity z can be thought of as a measure of overlap between the two core periphery networks. As z goes to zero, the core nodes in the repo and collateral networks share an increasing fraction of peripheral neighbors and thus become increasingly overlapping. Lemma 2 and Proposition 6 show that as z decreases, i.e. overlap increases, the expected post shock liquidity measure also increases. The intuition for this result follows directly from Lemma 2: only intermediaries that are in the core of at least one network can cause the withdrawal of additional intermediaries. Furthermore, the extent to which additional intermediaries withdraw depends on how many peripheral neighbors failing intermediaries in the cores of the networks share. If they share many peripheral neighbors, the illiquidity spiral due to the network coupling is dampened.

Finally, suppose that in a coupled core-periphery network parameterized by $\mathbf{n} = (n_{cc}, n_{cp}, n_{pc}, n_{pp})$, the collateral market is replaced by a complete network, i.e. a centralized exchange. This special case is nested in our parameterization of coupled core-periphery networks—it corresponds to

⁴⁶Note that $m_{cc} + m_{cp}$ is simply the number of core intermediaries in \mathcal{G}_R that survive and $n_{cc} + n_{cp}$ is the total number of core intermediaries in \mathcal{G}_R . Hence, we approximate z by the fraction of core intermediaries in \mathcal{G}_R that fail. Computing the exact expression involves keeping track of all network configurations and their probabilities and is complicated by the dependencies introduced by sampling without replacement. However, computing this exact expression for z does not add any substantial insights.

the vector $\mathbf{n}' = (n'_{cc}, 0, n'_{pc}, 0)$. Since the number of core and periphery nodes in the repo market should not be changed in going from \mathbf{n} to \mathbf{n}' , we require that $n'_{cc} = n_{cc} + n_{cp}$ and $n'_{pc} = n_{pc} + n_{pp}$.

Proposition 7 (Core-periphery and centralized market – post shock liquidity measure). *Let \mathcal{G}_R be a core periphery network and \mathcal{G}_C be a complete network. Given \mathbf{n}' and $\bar{\mathbf{m}}$, the expected number of withdrawing intermediaries is*

$$A(\bar{\mathbf{m}}) = (1 + \bar{k}_R)\bar{m}_{cc} + \bar{m}_{pc}.$$

The expected number of intermediaries providing liquidity conditional on shock size \bar{m}_W is

$$\hat{\mathcal{L}}^{cp-c}(\bar{m}_W) = E[\mathcal{L} | \bar{m}_W] = 1 - \frac{1}{n} \sum_{\bar{\mathbf{m}}} A(\bar{\mathbf{m}}) P(\bar{\mathbf{m}}, \mathbf{n}')$$

where P is defined as above. For a fixed shock size \bar{m}_W and $n_{cp} + n_{pp} > 0$, the post-shock liquidity in case of a centralized collateral market is always greater than post-shock liquidity in the pure core-periphery case

$$\hat{\mathcal{L}}^{cp-c}(\bar{m}_W) > \hat{\mathcal{L}}^{cp-cp}(\bar{m}_W).$$

In other words, when a collateral market with some peripheral nodes is replaced by a centralized exchange, post shock liquidity is always improved. The intuition for this result is very similar to the intuition for Proposition 3 where we show a similar result for star networks. If one of the two networks is replaced by a complete network (i.e. a centralized exchange), then no contagion can pass through this network. This is equivalent to setting $z(\mathbf{m}) = 0$ in Lemma 2. Thus the above result will hold, irrespective of the exact functional form of $z(\mathbf{m})$, provided that $z(\mathbf{m}) > 0$ for some network configurations when both markets are core-periphery. This is the case for the scenarios we are considering here.⁴⁷

To illustrate the size of this effect and the comparative statics of the liquidity measure for different shock sizes, we numerically evaluate the post shock liquidity measure in Fig. 4 for an example with $(n_{cc} = 0, n_{cp} = 2, n_{pc} = 2, n_{pp} = 50)$ using the approximation $z(\mathbf{m}) \approx 1 - (m_{cc} + m_{cp}) / (n_{cc} + n_{cp})$.⁴⁸

⁴⁷The case when there are just two nodes is an exception. However, in this case the notion of a core and a periphery are not well defined. Therefore, we exclude this corner case from our analysis.

⁴⁸To speed up the calculation, rather than summing over the entire probability space, $E[\mathcal{L} | \bar{m}_W]$ is approximated by its Monte Carlo average.

C Proofs

C.1 General facts

Lemma 3. Suppose Assumption 3.1 is satisfied. Then in a maximal equilibrium, $y_i^* = 1$ for all i ; equivalently, the maximal stable set consists of all nodes.

Proof. Fix a \mathcal{G} and define $\widehat{\mathcal{G}}$ to be the same graph with all edges reversed. Note in this graph each node has at least one outgoing edge. If we follow an arbitrary path, continuing by some out-edge at each step, then we will eventually reach a cycle of more than one node (since there are no self-edges by assumption). From this it follows that any directed path ends in a strong component. Thus (remembering the fact that $\widehat{\mathcal{G}}$ is \mathcal{G} with edges reversed) it follows that *any node in \mathcal{G} has a directed path leading to it from a strong component in \mathcal{G}* . Now, note that both this set and the path together constitute a stable set. There is such a set for each node, and they are all stable; thus their unions stable. This proves the result. \square

C.2 Star network

Proof. Proposition 2. First consider the case where $B_{H,C} \neq B_{H,R}$. Suppose that the hub intermediary in either market is hit by the adverse shock. The withdrawal of the hub intermediary forces all peripheral intermediaries to stop providing liquidity as they are fully dependent on the hub intermediary. Thus, in equilibrium $y_i^* = 0$ for all i and $\mathcal{L} = 0$. Now suppose a peripheral intermediary is hit by the adverse shock. Its withdrawal will not affect any other intermediaries since their provider of liquidity (the hub intermediary) has not been affected. Thus, in equilibrium $y_i^* = 1$ for all intermediaries that did not receive the adverse shock and $\mathcal{L} = n - 1$.

The probability that the adverse shock hits the hub intermediary in either market is $P(w^j \wedge j = B_{H,\mu}) = 2/n$. Conversely, the probability that the adverse shock does not hit the hub intermediary is $P(w^j \wedge j \neq B_{H,\mu}) = (n-2)/n$. Combing this with the above, the expected equilibrium liquidity measure is

$$\begin{aligned} E[\mathcal{L} \mid B_{H,C} \neq B_{H,R}] &= P(w = B_{H,\mu})\mathcal{L}(w = B_{H,\mu}) + P(w \neq B_{H,\mu})\mathcal{L}(w \neq B_{H,\mu}) \\ &= \frac{1}{n} \left[\frac{2}{n} \cdot 0 + \frac{n-2}{n} (n-1) \right] = \frac{(n-2)(n-1)}{n^2}. \end{aligned}$$

Now consider the case where $B_{H,C} = B_{H,R}$. Clearly, if the hub intermediary is hit by the adverse shock, in equilibrium $\mathcal{L} = 0$. This occurs with probability $1/n$. Conversely, if a peripheral intermediary is hit by the adverse shock, in equilibrium $\mathcal{L} = n - 1$. The expected equilibrium

measure is then given by

$$E[\mathcal{L} \mid B_{H,C} = B_{H,R}] = \frac{1}{n} \left[\frac{1}{n} 0 + \frac{n-1}{n} (n-1) \right] = \frac{(n-1)^2}{n^2}.$$

□

C.3 Preliminaries: theory of random networks

C.3.1 The configuration model

In the following we introduce random networks which are drawn uniformly at random conditional on a degree distribution. A standard device for generating and analyzing these graphs is the *configuration model*. All of the concepts introduced below apply equally to both markets $\mu \in \{R, C\}$. To avoid notional clutter, we drop the subscript μ for now.

For each n , let $\mathbf{d}_n^+ = (d_{i,n}^+)_{i=1}^n$ and $\mathbf{d}_n^- = (d_{i,n}^-)_{i=1}^n$ be sequences of non-negative integers representing the out-degrees and in-degrees, respectively, of intermediaries $i \in N$, where as before $n = |N|$ is the cardinality of the set of intermediaries. Note that all out-edges must have a corresponding in-edge, therefore $\sum_i^n d_{i,n}^+ = \sum_i^n d_{i,n}^-$. For a given n , denote the empirical distribution of degrees by

$$p_{jk,n} := \frac{1}{n} \#\{i \in N \mid d_{i,n}^+ = j, d_{i,n}^- = k\}.$$

The n in the subscript distinguishes this *empirical* distribution, associated with a given n , from an asymptotic distribution that is independent of the particular population size n , which we will introduce below.

Given \mathbf{d}_n^+ and \mathbf{d}_n^- satisfying the consistency condition noted above between total in- and out-degrees, let $G(n, \mathbf{d}_n^+, \mathbf{d}_n^-)$ be the set of graphs with degree sequences $\mathbf{d}_n^+, \mathbf{d}_n^-$. We define a standard device, the configuration model, for drawing graphs uniformly from this set:

Definition C.1 (Bollobás configuration model—see, e.g., [Amini et al. \(2013\)](#)). Consider a set of nodes $N = \{1, \dots, n\}$ and degree sequences $\mathbf{d}_n^+ = (d_{i,n}^+)_{i=1}^n$ and $\mathbf{d}_n^- = (d_{i,n}^-)_{i=1}^n$. Define for each node i a set of incoming and outgoing *half-edges*, H_i^- and H_i^+ , respectively. The set of all incoming and outgoing half edges is denoted by H^- and H^+ , respectively. A random directed multigraph $\tilde{\mathcal{G}}(n, \mathbf{d}_n^+, \mathbf{d}_n^-)$ drawn from the configuration model is then induced in the obvious way from a matching of all incoming half-edges H^- to outgoing half-edges H^+ drawn uniformly at random from the set of all such matchings. This multigraph may contain self-edges or multiple edges between two nodes. A graph without self-edges or multiple edges is a simple graph, and we condition on realizations that yield simple graphs. It is a standard fact that the resulting random variable is a draw uniformly at random from $G(n, \mathbf{d}_n^+, \mathbf{d}_n^-)$.

We are only interested in simple graphs and therefore need to impose conditions on the degree sequences to ensure that the probability of self-edges and multiple edges vanishes. We follow [Amini et al. \(2013\)](#) and [Britton et al. \(2007\)](#) and impose some standard conditions to ensure this, which turn out to yield this and other useful technical properties. The conditions are easiest to impose on an *infinite tuple* $((\mathbf{d}_n^+, \mathbf{d}_n^-))_{n=1}^\infty$ of pairs of degree sequences. The first set of conditions on an infinite tuple is:

Assumption C.1. For each n , \mathbf{d}_n^+ and \mathbf{d}_n^- are sequences of non-negative integers such that $\sum_i^n d_{i,n}^+ = \sum_i^n d_{i,n}^-$ and, for some joint distribution $(p_{jk})_{j,k \geq 0}$, over in- and out-degrees

1. $p_{jk,n} \rightarrow p_{jk}$ for every $j, k \geq 0$ as $n \rightarrow \infty$,
2. $\lambda := \sum_{j,k} p_{jk} j = \sum_{j,k} p_{jk} k \in (0, \infty)$,
3. $\sum_{i=1}^n (d_{i,n}^+)^2 + (d_{i,n}^-)^2 = O(n)$.

Note that conditions (2) and (3) imply that the average degree and the second moment degrees cannot diverge as the network becomes large. When this condition holds, we say it holds and the infinite tuple of degree sequences is *consistent with the joint distribution* $(p_{jk})_{j,k \geq 0}$.

We follow [Cooper and Frieze \(2004\)](#) and further require that the infinite tuple is *proper*. The technical assumptions comprising this definition require that a quantity akin to the degree sequence's second moment must grow much slower with the network size than the maximum degree of the sequence. This ensures that, while the maximum degree may go to infinity, the degree sequence does not become too dispersed:

Assumption C.2 (Proper degree sequences, [Cooper and Frieze \(2004\)](#)). Let Δ_n denote the maximum degree. Then

1. Let $\rho_n = \max\left(\sum_{j,k} \frac{j^2 k p_{jk,n}}{\lambda_n}, \sum_{j,k} \frac{k^2 j p_{jk,n}}{\lambda_n}\right)$. If $\Delta_n \rightarrow \infty$ with n then $\rho_n = o(\Delta_n)$.
2. $\Delta_n \leq \frac{n^{1/12}}{\log n}$.

We call an infinite tuple that satisfies Assumptions [C.1](#) and [C.2](#) *well-behaved*.

So far we have dealt with a single network. Now we extend our formalism to deal with two networks. To do this, we consider *two* infinite tuples $((\mathbf{d}_{C,n}^+, \mathbf{d}_{C,n}^-))_{n=1}^\infty$ and $((\mathbf{d}_{R,n}^+, \mathbf{d}_{R,n}^-))_{n=1}^\infty$, which are always well-behaved and are viewed as random variables.

Now note that our assumption that the two networks \mathcal{G}_R and \mathcal{G}_C are drawn independently ([Assumption 5.2](#)) implies $d_{iC,n}^+, d_{iR,n}^+$ and $d_{iC,n}^-, d_{iR,n}^-$ and $d_{iC,n}^+, d_{iR,n}^+$ and $d_{iC,n}^-, d_{iR,n}^-$ for all $i \in N$. In other words, the in- (out-) degree of an intermediary in the repo market gives no information about its in- (out-) degree in the collateral market. The networks $\mathcal{G}_{R,n}$ and $\mathcal{G}_{C,n}$ are independent, uniform draws from the sets $G(n, \mathbf{d}_{n,R}^+, \mathbf{d}_{n,R}^-)$ and $C(n, \mathbf{d}_{n,C}^+, \mathbf{d}_{n,C}^-)$, respectively.

C.3.2 Equilibrium and the mutual giant component

We can now establish a connection between the equilibrium liquidity defined in Section 3 and certain asymptotic properties of the random graphs defined above.

For each μ , we fix a joint distribution $(p_{jk,\mu})_{j,k \geq 0}$ over in- and out-degrees and well-behaved infinite tuples of degree sequences consistent with this distribution.

Definition C.2 (Giant out-component). Define $S_{\mu,n}$ to be any largest-cardinality strong component $S'_{\mu,n}$ and all nodes reachable by following a directed path out from $S'_{\mu,n}$.⁴⁹ Then the sequence of graphs is said to have a giant out-component if $S_{\mu,n}$ is well-defined for all large enough n and

$$\lim_{n \rightarrow \infty} \frac{1}{n} |S_{\mu,n}| \rightarrow s_\mu > 0.$$

It can be shown that, under the technical assumptions made above, if the sequence has a giant component, then asymptotically $S_{\mu,n}$ is unique—see [Cooper and Frieze \(2004\)](#). Suppressing the n index, we denote the subgraph of \mathcal{G}_μ associated with it by $GC_o(\mathcal{G}_R)$.

Definition C.3 (Mutual giant out-component). For large enough n , the mutual giant out-component $M = MGC_o(\mathcal{G}_R, \mathcal{G}_C)$ is defined to be a maximum-cardinality mutually stable subset of both $GC_o(\mathcal{G}_R)$ and $GC_o(\mathcal{G}_C)$.

The size of the mutual giant out-component and the equilibrium liquidity measure are then related as follows.

Lemma 4. Let \mathbf{y}^* be an equilibrium for $\mathcal{G}_R, \mathcal{G}_C$ and a shock profile \mathbf{w} as given in Section 3, with W being the set of shocked ($w_i = 0$) nodes. Then

$$\mathcal{L}(\mathbf{y}^*) = \frac{1}{n} \sum_i y_i^* \geq \frac{1}{n} |MGC_o(\mathcal{G}_R(W), \mathcal{G}_C(W))|.$$

In the limit of large networks we obtain

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_i y_i^* \rightarrow \frac{1}{n} |MGC_o(\mathcal{G}_R(W), \mathcal{G}_C(W))|,$$

The size of the mutual giant out-component is a lower bound on the size of the maximal mutually stable set, and thus the number of active intermediaries in equilibrium. It is only a lower bound since there may exist small, mutually stable components outside the mutual giant out-component. However, as the network becomes large, results from [Cooper and Frieze \(2004\)](#) imply that the relative size of these small mutually stable components vanishes; the reason is

⁴⁹Every nonempty proper subset has an edge to its complement or from it.

that even the weak components in either of the two markets which are not part of the giant out-component have a negligible size as $n \rightarrow \infty$ (recall Section 5.2.1). Therefore, in the limit of large networks the size of the mutual giant out-component is sufficient to compute the equilibrium liquidity. In the following we will discuss how the mutual giant out-component can be found.

C.3.3 A branching process approximation of equilibrium liquidity

In this section we will invoke results from the theory of branching processes and probability generating functions to compute the size of the giant mutual out-component (see Cooper and Frieze (2004) and Buldyrev et al. (2010)). We will first characterize the giant out-component in a single market (i.e., one graph without any coupling) and will then proceed to derive the size of the mutual giant out-component.

Computing the giant out-component The distribution of the out-degree of the terminal node of a randomly chosen link in a large graph is, in the $n \rightarrow \infty$ limit, given by

$$p_k^+ := \sum_j \frac{j}{\lambda} p_{jk},$$

where p_{jk} is the joint distribution for in- and out-degrees (see for example Cooper and Frieze (2004) and Newman (2010)). Note that the out-link is j times more likely to end up at a node with in-degree j . The average degree λ enters as a normalizing constant.⁵⁰

Suppose one starts to explore the network from a randomly chosen link via a breadth-first search algorithm. How many intermediaries can one reach by following only out-going links? In a random network model, this exploration process can be approximated by a standard branching process where the number of offspring (i.e. outgoing links) of any node is distributed according to $(p_k^+)_{k=0}^\infty$.⁵¹ Let $H(z) := \sum_k p_k^+ z^k$ denote the corresponding probability generating function. Recall the following useful result on the extinction probability of a branching process (see Athreya and Jagers (2012)):

Lemma 5. The probability f that the branching process defined by $(p_k^+)_{k=0}^\infty$ goes extinct is the smallest solution to $f = H(f)$.

⁵⁰The distribution for the in-degree of the terminal node of a randomly chosen link can be defined similarly but is not of interest for us.

⁵¹Of course this only corresponds to the number of intermediaries explored if the breadth-first search does not turn back on itself and does not re-explore parts it has already seen. The assumption that this does not occur is usually referred to as the requirement that the network is “locally tree-like”, i.e. that there are no short cycles. Hence the application of the branching process is indeed an approximation. However, under our maintained technical assumptions, Cooper and Frieze (2004) show rigorously that this approximation is indeed valid.

Then, the size of the giant out-component is given by a simple corollary of Lemma 5 (see for example Newman (2010)) that we summarize in the following lemma.

Lemma 6. Given the probability f that the branching process defined by $(p_k^+)_{k=0}^\infty$ goes extinct, the fraction of nodes in the giant out-component is

$$g(f) := 1 - \sum_{jk} p_{jk} f^k.$$

This follows from the fact that the probability that a random node with k outgoing links is not in the giant out-component is simply f^k .

Equilibrium liquidity and the mutual giant out-component Now that we have established how to compute the size of the giant out-component for a single network we can proceed to derive the size of the mutual giant component. The following derivation builds on results of Buldyrev et al. (2010). From now on we will associate each of the quantities introduced in Section C.3.3 with the collateral or repo markets via the subscripts C or R respectively. For example, $H_C(z)$ will be the probability generating function of the out-degree process for the network corresponding to the collateral market.

Consider the following coupled branching process—first, with $x = 1$, i.e., no shock. Choose a link at random in \mathcal{G}_R and follow it the node it goes into. Since \mathcal{G}_C and \mathcal{G}_R are independent by assumption, the intermediary we reach will be in the giant out-component of \mathcal{G}_C with probability s_C (the fraction of nodes in the giant out-component in the collateral network).

If the node is not in the giant out-component of \mathcal{G}_C the branching process will not continue further. We may equivalently assume that in that case the node reached has no out-neighbors. As discussed by Newman (2010) the branching process is identical in distribution to one in which we “thin” the degree distribution of the collateral network as follows:

$$\hat{p}_{jk,R}(s_C) := \underbrace{\sum_{l=j}^{\infty} \sum_{m=k}^{\infty} p_{lm,R}}_A \underbrace{\binom{l}{j} (1-s_C)^{l-j}}_B \underbrace{s_C^j \binom{m}{k} (1-s_C)^{m-k} s_C^k}_C. \quad (6)$$

The distribution $\hat{p}_{jk,R}(s_C)$ is the joint distribution of in- and out-degrees of the repo network after a fraction $1 - s_C$ of nodes has been removed uniformly at random. The transformed degree distribution consists of three terms: A , B and C . A corresponds to the initial probability that a random node has in-degree l and out-degree m . B is the probability that j out of initially l in-links are present after thinning. Finally, C is the probability that k out of initially m out-links are present after thinning. An equivalent argument can be made for the repo market. Again

we obtain a transformed degree distribution $\hat{p}_{jk,C}(s_R)$. Similarly, we define the transformed distributions for the out-degree process $\hat{p}_{k,R}^+(s_C)$ and $\hat{p}_{k,C}^+(s_R)$.

Now, suppose a fraction $1 - x$ of nodes, selected uniformly at random, withdraws from the markets due to the exit shock. We call $1 - x$ the size of the exit shock. The final size of the mutual giant out-component (and thus liquidity in the maximal equilibrium) is then determined by the branching process on the residual networks $\mathcal{G}_R(W)$ and $\mathcal{G}_C(W)$ after the withdrawal of the shocked intermediaries. Let $\hat{\mathcal{L}}^*(x)$ be the expected liquidity of the maximal equilibrium.

Lemma 7. Given the degree-distributions $p_{jk,\mu}$ for $\mu \in \{R, C\}$ and a shock of size $1 - x$, the size of the giant out-component in the repo (collateral) network s_R^* (s_C^*) is the greatest solution to

$$\begin{aligned} s_R &= xg_R(f_R, s_C) = x \left(1 - \sum_{jk} \hat{p}_{jk,R}(s_C) f_R^k \right), \\ f_R &= H_R(f_R, s_C) = \sum_k \hat{p}_{k,R}^+(s_C) f_R^k, \\ s_C &= xg_C(f_C, s_R) = x \left(1 - \sum_{jk} \hat{p}_{jk,C}(s_R) f_C^k \right), \\ f_C &= H_C(f_C, s_R) = \sum_k \hat{p}_{k,C}^+(s_R) f_C^k. \end{aligned} \tag{7}$$

Liquidity is then

$$\hat{\mathcal{L}}^*(x) := s^* = xg_R(xg_C(s^*)).$$

To see this, note that the expression $xg_R(xg_C(s^*))$ is monotonically increasing in s (see Lemma 11) in $[0, 1]$. Then by Tarski's fixed point theorem a maximum fixed point s^* exists. Note that the realization of the shock bounds the size of the giant out-component, and thereby equilibrium liquidity, from above by x . This is simply because a fraction of $1 - x$ of intermediaries withdraw from the markets due to the exit shock realization.

Suppose now that one of the markets is replace by a centralized exchange so that we can replace the corresponding network by a complete network. What is the size of the mutual giant out-component?

Lemma 8. Let \mathcal{G}_R be a random network. Let $\bar{\mathcal{G}}_C$ be a complete network. Given a shock of size $1 - x$, the size of the giant out-component in the repo network s_R^* is the greatest solution to

$$\begin{aligned} s_R &= g_R(f_R, x) = x \left(1 - \sum_{jk} \hat{p}_{jk,R}(x) f_R^k \right), \\ f_R &= H_R(f_R, x) = \sum_k \hat{p}_{k,R}^+(x) f_R^k. \end{aligned}$$

Liquidity is then

$$\hat{\mathcal{L}}^*(x) := s_R^* = g_R(f_R, x).$$

Thus, if the collateral network is replaced by a complete network, the size of the mutual giant out-component is simply the size of the giant out-component of the repo network taken on its own (we have discussed the study of its size above). To see this, first note that if $\bar{\mathcal{G}}_C$ is complete there will be no contagion through $\bar{\mathcal{G}}_C$. All intermediaries in $\bar{\mathcal{G}}_C$ are active except those that are not in the giant out-component of the repo market. Therefore it is not necessary to compute the size of the giant out-component in the collateral network via a branching process as in Lemma 7. The greatest fixed point exists by the same argument as in the proof of Lemma 7.

C.4 Proofs of random network results

In the following we will prove Propositions 4 and 5. Our main contribution is to provide conditions on the degree distributions of the random networks for which Propositions 5 holds. We will first provide a sketch of the proof of Proposition 4 since this is a standard result from the literature (see Cooper and Frieze (2004)) and is useful for the subsequent proofs of Proposition 5.

For the proof of Proposition 4 we will use standard properties of a generic probability generating function (pgf) that we summarize in the following remark.

Remark 1. A generic pgf $f(s) = \sum_i p_i s^i$ has the following properties:

- (i) $f(0) = p_0$,
- (ii) $f(1) = 1$,
- (iii) $f'(1) = df/ds(1) > 0$ (increasing),
- (iv) $d^2f/ds^2 > 0$ (convex) for $s > 0$.

Therefore $s^* = f(s^*)$ has a solution $s^* < 1$ if $f'(1) > 1$. Otherwise only the trivial solution $s^* = 1$ exists. $s^* = 0$ is not a solution if $p_0 > 0$. Note that the solution s^* is continuous in the slope $f'(1)$, i.e. as $f'(1) \rightarrow 1$ we have that $s^* \rightarrow 1$.

We illustrate some graphical intuition for this proof in Fig 11 in Appendix D.

Proof. Proposition 4. Recall that for $x = 1$ we have $H(z) = \sum_k p_k^+ z^k$ with $p_k^+ = \sum_j j p_{jk} / \lambda$. It can be shown, see for example Newman (2010) or Cooper and Frieze (2004), that after a fraction

$1 - x$ of intermediaries are removed uniformly at random from the network, the pgf of the out-degree distribution becomes $\hat{H}(z, x) = H(1 - x + xz)$. From remark 1 we know that $f = H(f) = 1$ if $dH/dz(1) = H'(1) \leq 1$ and $f < 1$ if $H'(1) > 1$. When $f = 1$ the size of the giant out-component vanishes, i.e. $g(1) = 0$. If $f < 1$ the size of the giant out-component is $g(f) > 0$, i.e. the giant out-component exists. Thus we need to ask at which x_c the derivative of the pgf becomes $\hat{H}'(1) = 1$. Note that $\hat{H}'(1) = xH'(1)$. Thus

$$x_c = \frac{1}{H'(1)} = \frac{\lambda}{\sum_{jk} p_{jk} jk}. \quad (8)$$

Since f is continuous as the derivative $\hat{H}'(1)$ changes, it is also continuous in x which determines $\hat{H}'(1)$. Note that Assumption 5.1 ensures that in the absence of an exit shock there exists a giant out-component of positive size. This concludes the proof. \square

Lemma 9. Let $f(x)$ denote the smallest solution $f = H(f, x)$. Then $f(x)$ is continuous, monotonically decreasing in x for $x \in [0, 1]$.

Proof. Lemma 9. $f(x) = H(f(x), x)$ is continuous follows from the proof of Proposition 4. To show that $f(x)$ is monotonically decreasing we use the result from the proof of Proposition 4 that $x \in [0, x_c]$: $f = 1 \implies df/dx = 0$. Now consider what happens when $x \in (x_c, 1]$ and $f < 1$. In this case we derive for df/dx :

$$\begin{aligned} \frac{df}{dx} &= \frac{1}{\lambda} \underbrace{\sum_{jk} jk p_{jk} (1 - x + xf)^{k-1}}_{dH/df|_{1/x}} \left(f + x \frac{df}{dx} - 1 \right), \\ \frac{df}{dx} &= \frac{dH}{df}(f, x) \frac{1}{x} \left(f + x \frac{df}{dx} - 1 \right), \\ \frac{df}{dx} &= \frac{dH}{df}(f, x) \frac{1}{x} \left(1 - \frac{dH}{df}(f, x) \right)^{-1} (f - 1). \end{aligned}$$

Note that for supercritical x the derivative of $H(f, x)$ with respect to f evaluated at the intersection with the diagonal is less than one, i.e. for $x \in (x_c, 1]$ $\frac{dH}{df}(f, x) < 1$, where $H(f, x) = f < 1$; see Fig. 11 in Appendix D for a graphical intuition. This can be seen as follows. Clearly for there to exist a solution $f < 1$ to $f = H(f, x)$, $H(f, x)$ must cross the diagonal. But since $\frac{dH}{df}(1, x) > 1$ for $x \in (x_c, 1]$ and $H(1, x) = 1$, $H(f, x)$ must cross the diagonal from below when approaching the intersection from the right. This implies that $\frac{dH}{df}(f, x) < 1$ at the intersection. This together with $0 < x, f < 1$ and $\frac{dH}{df}(f, x) > 0$ implies that $\frac{df}{dx} < 0$. \square

Lemma 10. $g(f, x)$ is continuous and monotonically increasing in x for $x \in [0, 1]$.

Proof. Lemma 10. The fact that $g(f, x)$ is continuous follows directly from the proof of Propo-

sition 4. $g(f, x)$ is monotonically increasing since

$$\frac{dg}{dx} = - \sum_{jk} p_{jk} k (f(x)x + 1 - x)^{k-1} \underbrace{\left(\frac{df}{dx} x + f(x) - 1 \right)}_{\leq 0} \geq 0.$$

□

Let $F(s, x) := xg_R(xg_C(s))$. In order to prove Proposition 5 we first need to establish a couple of facts about $F(s, x)$ which we summarize in the following lemma. We will use the index $\mu \in \{R, C\}$ whenever results apply to both repo and collateral networks.

Lemma 11. For $s \in (0, 1]$

1. $F(s, x)$ is continuous in s ,
2. $F(s, x)$ is monotonically increasing in s ,
3. $F(s, x)$ is bounded from above: $F(s, x) \leq x$,
4. $F(s, x)$ is concave in s ,
5. $\lim_{s \rightarrow 0} F(s, x) \rightarrow 0$,
6. $\lim_{s \rightarrow 0} \frac{\partial F(s, x)}{\partial s} \rightarrow 0$.

Proof. Lemma 11. For this proof we invoke results from Lemmas 9 and 10. For $s \in (0, 1]$:

1. $F(s, x)$ is continuous: $g_\mu(s)$ is continuous as shown in Lemma 10. $F(s, x)$ is a function of $g_\mu(s)$ and therefore also continuous in s .
2. $F(s, x)$ is monotonically increasing: $g_\mu(s)$ is monotonically increasing as shown in Lemma 10. $F(s, x)$ is therefore also monotonically increasing in s .
3. $F(s, x)$ is bounded from above - $F(s, x) < 1$: Clearly $g_\mu(s)$ is bounded from above since $g_\mu(s) \leq 1$. Furthermore, assuming a positive shock size, i.e. $x < 1$, we have $F(s, x) = xg_R(xg_C(s)) < 1$. Also note that the above implies that $F(s, x)$ has a maximum at $s = 1$ which scales with x , i.e. as x is decreased the maximum of $F(x, s)$ decreases by at least the same amount.
4. $F(s, x)$ is concave in s :

$$\frac{\partial^2 F}{\partial s^2}(s, x) = x^2 \left(x \frac{d^2 g_R}{ds^2}(s) \left(\frac{dg_C}{ds}(s) \right)^2 + \frac{dg_R}{ds}(s) \frac{d^2 g_C}{ds^2}(s) \right)$$

Since $\frac{dg_\mu}{ds}(s) > 0$, $\frac{d^2g_\mu}{ds^2}(s) < 0$ (by Assumption 5.3) and $x > 0$ we must have $\frac{\partial^2 F}{\partial s^2} < 0$, i.e. $F(s, x)$ concave.

5. $\lim_{s \rightarrow 0} F(s, x) \rightarrow 0$: Since for $s < s_{c,\mu}$ we have $f_\mu(s) = 1$ and $g_\mu(s) = 0$, where $s_{c,\mu}$ is the threshold for network μ at which the giant out-component vanishes as given in Proposition 4. In other words, there exists a critical $s_{c,\mu}$ at which the giant out-component in one of the intermediation networks vanishes (recall that we assume that $\lambda < \infty$, hence there always exists this critical $s_{c,\mu}$ by Proposition 4).

6. $\lim_{s \rightarrow 0} \frac{\partial F}{\partial s}(s, x) \rightarrow 0$:

$$\lim_{s \rightarrow 0} \frac{\partial F}{\partial s}(s, x) = x^2 \frac{dg_R}{dv}(v) \frac{dg_C}{ds}(s) \rightarrow 0.$$

Since for $s < s_{c,\mu}$ we have $f_\mu(s) = 1$ and $g_\mu(s) = 0$. Hence for $s < s_{c,\mu}$ we have $\frac{dg_\mu}{ds}(s) = 0$. In other words, since there exists a critical $s_{c,\mu}$ at which the giant component vanishes in one of the intermediation networks, there is a region for values of $s < s_{c,\mu}$ in which $F(x, s)$ is flat.

These observations show that, under the assumptions made here, $F(x, s)$ can be decomposed into two regions: (i) for small values of s ($s < s_{c,\mu}$) $F(x, s)$ vanishes ($F(x, s) = 0$) and is flat ($\partial F/\partial s = 0$). (ii) for larger values of s ($s > s_{c,\mu}$) $F(x, s)$ is strictly monotonically increasing and concave but bounded from above ($F(x, s) < 1$). \square

Proof. Proposition 5 A. This proof invokes results from Lemma 11 and relies in particular on our observations of the shape of $F(x, s)$ in the interval $s \in [0, 1]$. We illustrate the graphical intuition for this proof in Fig. 12 in Appendix D.

First note that $s = 0$ is a trivial solution to $s = F(s, x)$ for all x since $g_\mu(0) = 0$. Furthermore as shown in Lemma 11 there exists a region for sufficiently small s in which $F(s, x)$ is constant and equal to zero. As seen in Lemma 11, for all $s > s_{c,\mu}$ the function $F(s, x)$ is strictly increasing and concave provided $g_\mu(s)$ is concave. The fact that $F(x, s)$ is constant and flat close to $s = 0$ implies that in at least some of the interval $s \in [0, 1]$, $F(x, s)$ must lie below the diagonal. If for $s > s_{c,\mu}$ the function $F(x, s)$ increases sufficiently fast to cross the diagonal there will exist two solutions in addition to the trivial solution (since for $x < 1$ $F(x, s) < 1$ and hence cannot remain above the diagonal for the entire interval $s \in [0, 1]$).

Note that Assumption 5.1 ensures that in the absence of an exit shock there exists a mutual giant out-component of positive size. Since we are investigating cascades following a small exit shock we are only interested in the largest fixed point s^* of the map $s_n = F(s_{n-1}, x)$ with $s_0 = x$. This fixed point will be stable due to the concavity of $F(s, x)$ and because at s^* the slope of $F(x, s)$ is $\partial F/\partial s(s^*, x) < 1$.

Now consider how the largest fixed point s^* changes when the initial exit shock $1 - x$ is increased. Clearly, when x goes down, s^* goes down as well. This is because for a smaller value of x the curve $F(s, x)$ will have a smaller maximum value. This pushes the entire segment of the curve of $F(x, s)$ for $s > s_{c,\mu}$ downwards. Therefore $F(s, x)$ will intersect the diagonal at a smaller value. When both x and s^* decrease further the curve $F(s, x)$ will ultimately become tangent to the diagonal. This will correspond to some critical value x_c . At this point the largest solution s^* merges with the second largest on the diagonal.

If x is decreased further ($x < x_c$) both non trivial solutions vanish and only the trivial solution at $s = 0$ remains. In summary, if there exists some fixed point of $F(x, s)$, s^* , and an exit shock of a critical size $1 - x_c$ such that $F(x, s)$ is tangent to the diagonal ($\frac{\partial F}{\partial s}(s^*, x_c) = 1$), then there will be a region below x_c where only the trivial solution exists ($s^* = 0$) and a region above x_c where a non trivial solution $0 < s^* < 1$ exists.

Note that, since there exists some value $s_{c,\mu} > 0$ at which the derivative $\partial F / \partial s(s, x)$ vanishes, $F(x, s)$ must lie below the diagonal close to $s = 0$. Therefore, the non trivial solution must always be greater than zero, i.e. $s^* > 0$ for $x \geq x_c$. Therefore

$$\lim_{\epsilon \rightarrow 0} F(s^*, x_c - \epsilon) = 0 \neq F(s^*, x_c) > 0.$$

Hence $F(s, x)$ is discontinuous in x at $x = x_c$. From the above it also follows that, if there exists no $0 < s^* < 1$ such that at some $x = x_c > 0$, $\frac{\partial F}{\partial s}(s^*, x_c) = 1$, then only the trivial solution can exist and $F(s^*, x) = 0 \forall x < 1$. In this case a minimal disturbance of the network leads always to a complete collapse of the network. □

Now let us turn to the Proposition 5 B.

Proof. Proposition 5 B. Let's write $r_c = r_c(\mathcal{G}_R, \bar{\mathcal{G}}_C)$ and $x_c = x_c(\mathcal{G}_R, \mathcal{G}_C)$. Suppose we have $1 - x_c \geq 1 - r_c$ ($x_c \leq r_c$). Note that by definition at r_c , the size of the giant component in the repo network vanishes, i.e. $g_R(r_c) = 0$. Also, $F(s, x_c) = x_c g_R(x_c g_C(s)) < r_c$ since $F(s, x_c) < x_c$ for $s < 1$ and $x_c \leq r_c$ by assumption. However, at a fixed point we must have that $F(s, x_c) = s$. Thus for any solution s , we have that $s < r_c$ and hence $x_c g_C(s) < r_c$. But we must have that $g_R(s) = 0$ for all $s < r_c$. This implies that at the fixed point $s^* = 0$. However this contradicts $F(s^*, x_c) > 0$ which is required by Proposition 5 A. This proves Proposition 5 B by contradiction. □

D Additional Figures




Figure 11: Graphical intuition for proof of Proposition 4. We are interested in fixed points $f^* = H(f^*, x)$ with $f^* < 1$. We plot $H(f, x)$ for two choices of x . Note that the value of x determines the slope of $H(f, x)$ at $f = 1$. The dashed green line corresponds to the case when x is such that $dH(1, x)/df > 1$ while the continuous red line corresponds to the case when x is such that $dH(1, x)/df = 1$. Due to the convexity of $H(f, x)$ in f , $dH(1, x)/df \leq 1$ implies that there will be no fixed point apart from $f^* = 1$ in the interval $[0, 1]$. Thus $dH(1, x_c)/df = 1$ determines a critical value of x at which $f^* < 1$ merges with $f^* = 1$.




Figure 12: Graphical intuition for proof of Proposition 5. We are interested in the greatest fixed point $y^* = F(y^*, x)$ with $y^* > 0$. We plot $F(y, x)$ for two choices of x . Note that the value of x determines the slope of $F(y, x)$ at y^* . The dashed green line corresponds to the case when x is such that $dF(y^*, x)/dy > 1$ while the continuous red line corresponds to the case when x is such that $dF(y^*, x)/df = 1$. As x is decreased $F(1, x)$ and y^* decrease. At some critical x_c the curve $F(y, x)$ will become tangent to the diagonal. If x_c is decreased any further, $y^* > 0$ disappears and only the trivial fixed point $y^* = 0$ remains.

E Calculations for example random networks

E.1 Erdos-Rényi network

Let q denote the probability that a randomly chosen intermediary is connected to another intermediary by an outgoing or incoming link. Here, due to the independence of in- and out-degrees the joint degree distribution factorizes into $p_{jk} = p_j p_k$ with $p_j = p_k$ and

$$p_k = \binom{n-1}{k} q^k (1-q)^{n-k-1}.$$

When we hold the average in- and out-degree $\lambda = nq$ fixed and take the limit $n \rightarrow \infty$ the generating function for the out-degree distribution of a random node becomes

$$G(z) = e^{\lambda(z-1)},$$

Note that for the Erdős-Rényi network the generating function for the out-degree of a random node is equal to the generating function of the out-degree of the terminal node reached by following a random link (Newman, 2002). Thus, we have $G(z) = H(z)$. As shown in appendix C.3, after an exit shock removing a fraction $1 - x$ of nodes, the generating functions become

$$\hat{G}(z, x) = \hat{H}(z, x) = G(1 - x + zx) = H(1 - x + zx) = e^{\lambda x(z-1)},$$

As before, we compute equilibrium liquidity as the size of the giant out-component of the repo network: $\mathcal{L}^*(x) = s^*$. In Figure 6 we solve for s^* numerically.

E.2 Scale free networks

Now let's consider the case where \mathcal{G}_C and \mathcal{G}_R are directed networks with the same power law in- and out-degree distributions (also known as scale free networks). Networks with this degree distribution can be formed for example through a preferential attachment process as outlined in Barabási and Albert (1999). As for the Erdős-Rényi networks we assume that the in- and out-degrees are independent, such that $p_{jk} = p_j p_k$. We also assume that $p_j = p_k = C_\mu k^{-\alpha}$ for $\alpha \in (2, 3]$ and $k > 1$. The constant that normalizes the degree distribution is $C = 1/(\zeta(\alpha) - 1)$, where $\zeta(\cdot)$ is the Riemann zeta function. Also define the generating functions with their usual meanings

$$G(z) = C \sum_{k>1} k^{-\alpha} z^k = C(\text{Li}_\alpha(z) - z),$$

where $\text{Li}_s(z)$ is the polylogarithmic function defined by:

$$\text{Li}_s(z) = \sum_{k=1}^{\infty} \frac{z^k}{k^s},$$

where s is complex number and z is a complex number with $|z| < 1$, which is clearly valid here. In the following we will only consider real s and z . We also have

$$H(z) = \frac{1}{\lambda} \sum_j p_j j \sum_k p_k z^k = G(z).$$

As before we have that $\hat{G}(z, x) = G(1 - x + zx)$ and $\hat{H}(z, x) = H(1 - x + zx)$. We can make the substitution $w = 1 - x + zx$, i.e. $z = (w + x - 1)/x$. Then to find the extinction probability of the branching process we must solve

$$(w + x - 1)/x = H(w) = C(\text{Li}_\alpha(w) - w)$$

Again we compute equilibrium liquidity as the size of the giant out-component of the repo network: $\mathcal{L}^*(x) = s^*$. In Figure 7 we solve for s^* numerically.

F Overlap between repo and collateral networks

It is useful to write the equations (7) slightly differently. In particular let us introduce

$$\begin{aligned} \tilde{s}_R &= g_R(f_R, x - x(1 - \tilde{s}_C)), \\ f_R &= H_R(f_R, x - x(1 - \tilde{s}_C)), \\ \tilde{s}_C &= g_C(f_C, x - x(1 - \tilde{s}_R)), \\ f_C &= H_C(f_C, x - x(1 - \tilde{s}_R)), \end{aligned}$$

where $s_R = x\tilde{s}_R$ and $s_C = x\tilde{s}_C$. Clearly $x - x(1 - \tilde{s}_C)$ is simply the fraction of nodes remaining after the initial shock $1 - x$ minus the number of nodes that are not in the giant component of \mathcal{G}_C but remain in the network after the initial shock $1 - x$. We can make a crude, but simple, approximation to the effect of overlap as follows. Only intermediaries which do not lie outside the giant component in \mathcal{G}_R can withdraw upon their withdrawal in \mathcal{G}_C . The fraction of nodes that are in the giant component in \mathcal{G}_R but not in the giant component of \mathcal{G}_C is approximately




Figure 13: Equilibrium liquidity s^* as a function of the fraction of intermediaries $1 - x$ that withdraw from the repo and collateral markets following an exit shock in an Erdős-Rényi network with different levels of network overlap.

$(1 - \tilde{s}_C)(1 - \omega)$. Thus we obtain

$$\begin{aligned}\tilde{s}_R &= g_R(f_R, x - x(1 - \tilde{s}_C)(1 - \omega)), \\ f_R &= H_R(f_R, x - x(1 - \tilde{s}_C)(1 - \omega)), \\ \tilde{s}_C &= g_C(f_C, x - x(1 - \tilde{s}_R)(1 - \omega)), \\ f_C &= H_C(f_C, x - x(1 - \tilde{s}_R)(1 - \omega)),\end{aligned}$$

Note that this formulation reduces to the centralized market benchmark for $\omega = 1$ and the usual two network case for $\omega = 0$.

Recall from section E that the generating functions for the Erdős-Rényi network are given by

$$\hat{G}(z, x) = \hat{H}(z, x) = G(1 - x + zx) = H(1 - x + zx) = e^{\lambda x(z-1)},$$

Then it can be shown that

$$\begin{aligned}\tilde{s}_R &= 1 - e^{-\lambda_R x(1 - \tilde{s}_C)(1 - \omega)\tilde{s}_R}, \\ \tilde{s}_C &= 1 - e^{-\lambda_C x(1 - \tilde{s}_R)(1 - \omega)\tilde{s}_C}.\end{aligned}$$

If we take $\lambda_R = \lambda_C$, due to the symmetry of the expressions above we must have $\tilde{s}_R = \tilde{s}_C$, hence we can reduce the above to a single equation

$$s = 1 - e^{-\lambda x(1-s)(1-\omega)s}. \quad (9)$$

We know that there exists a regime for ω for which we observe a continuous transition at the critical exit shock (e.g. $\omega = 1$) as well as a regime with a discontinuous transition (e.g. $\omega = 0$). The critical value of ω at which the transition switches from continuous to discontinuous is often referred to as the tri-critical point. We can follow the standard procedure to determine the tri-critical point at which the transition becomes discontinuous, cf. [Son et al. \(2012\)](#). Let us first define the deviation measure

$$h(s) = s - (1 - e^{-\lambda x(1-s)(1-\omega)s}).$$

Suppose we are in a regime of ω in which the transition is continuous. Close to the critical exogenous shock we have $\epsilon = s \approx 0$ and we can expand around $h(0)$ to approximate $h(\epsilon)$, i.e.

$$h(\epsilon) = h'(0)\epsilon + \frac{1}{2}h''(0)\epsilon^2 + \frac{1}{6}h'''(0)\epsilon^3 + O(\epsilon^4).$$

Suppose for now that the first and second derivatives are non zero. At a solution of Eq. (9) we

must have $h(\epsilon) = 0$. If we ignore higher order terms and solving for ϵ we obtain

$$\epsilon \approx \frac{2h'(0)}{h''(0)},$$

At the critical point $\epsilon = 0$. Thus, provided $h''(0) \neq 0$, at the critical point we must have $h'(0) = 0$. It can be shown that $d\epsilon/dx$ does not diverge at the critical point in this case. Now suppose that $h''(0) = 0$. When solving for ϵ we now need to include higher order terms. Thus

$$\epsilon \approx \sqrt{\frac{6h'(0)}{h'''(0)}},$$

By applying the chain rule we find that $d\epsilon/dx = \partial\epsilon/\partial h'(0)\partial h'(0)/\partial x + R$, where R corresponds to the remaining terms of the derivative. Note that $\partial\epsilon/\partial h'(0) \propto 1/\sqrt{h'(0)}$. Thus, when $h'(0) = h''(0) = 0$, the derivative $d\epsilon/dx$ diverges and a discontinuous transition emerges. Solving for the value of ω at which the first and second derivatives go to zero, we obtain that $\omega_c = 2/3$. Thus, for coupled Erdős-Rényi networks there exists a discontinuous transition as long as approximately one third of the links differ between the two networks.

G Interpolating between scale-free and core-periphery networks

Many financial networks can be characterized as core-periphery networks. In Section 4.2, we showed how illiquidity spirals unfold in stylized core-periphery networks. We showed that for such perfect core-periphery networks, a discontinuous transition as observed in Section 5 does not occur. However, for scale-free networks a discontinuous transition is observed. Arguably, real financial networks are neither perfectly core-periphery, nor perfectly scale-free but somewhere in between these extremes.

In the following, we propose a very simple approach to interpolate between scale-free and core-periphery networks. We begin by constructing a scale-free network as outlined in Appendix C with tail exponent $\alpha = 2.5$ and N nodes using the configuration model. We then designate the N_C nodes with the highest degree as the core. With probability p_C we connect to core nodes that are not yet connected. Similarly, with probability $1 - p_P$, we remove an existing link between two periphery nodes. Clearly for $p_C = 0$ and $p_P = 1$ we leave the scale-free network unchanged. For $p_C = 1$ and $p_P = 0$ we obtain a perfect core-periphery network. We repeat this procedure to generate two coupled but independent networks.

Figure 14 shows that the transition is smoothed out (as opposed to being discontinuous) as the core becomes more connected. If the core remains unchanged, but peripheral links are

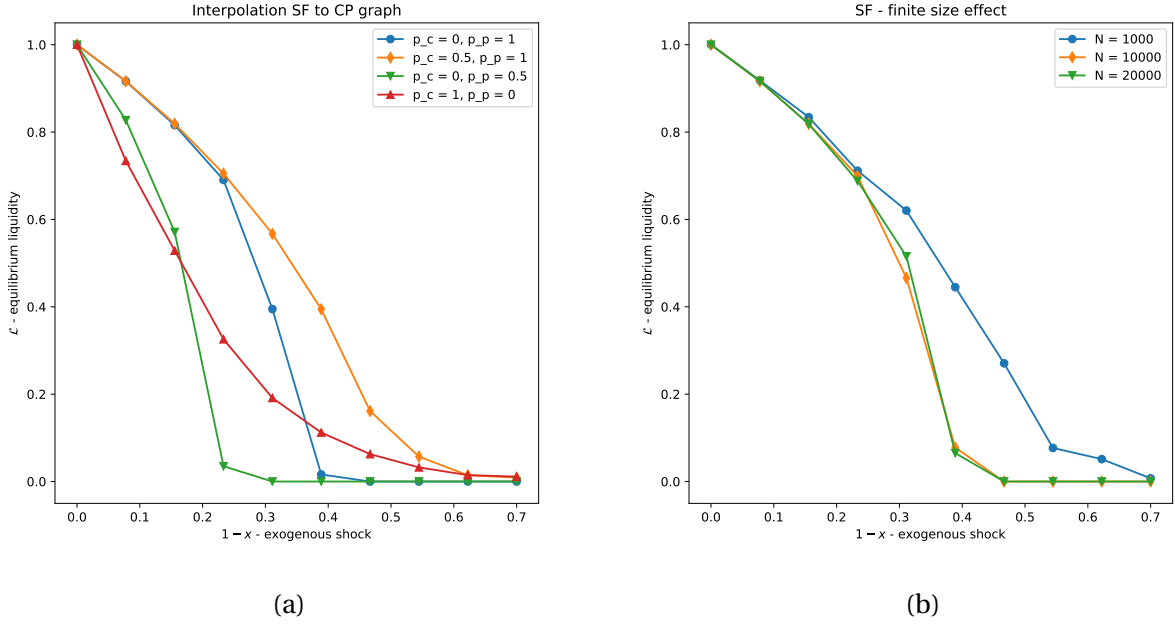


Figure 14: (a) Interpolating between scale-free and core-periphery networks. $N = 10000$, $N_C = 1000$. (b) Finite size effects for scale-free networks.

removed, the transition is less smooth and liquidity evaporates quicker for smaller shocks. We conjecture that there will be some critical p_C and p_P at which the discontinuous transition disappears. This can be found via a grid search over these parameters. Also note that the discontinuous transition is smoothed when the network is smaller, see Figure 14.

In Figure 15 we study how the equilibrium liquidity measure depends on the size of an exit shock for different core-periphery networks. We vary $p_c \in [0, 0.02]$, i.e. we slightly increase the number of links within the core. This has a sizeable impact on the resilience to a shock. The high sensitivity of the system to the existence of additional links within the core is an important insight from our analysis with implications for policy makers tasked with safeguarding financial stability: a market freeze does not have to be complete to leave a system of coupled core-periphery networks much more vulnerable to exit shocks. It matters where previously existing links are cut.

H Cascade sizes in small networks

Suppose we are interested in understanding how fragile a particular network of relatively small size $N \approx 100$ is to the removal of a single node. The fragility of the network depends on the structure of the network. So far, we have always studied networks generated according to some




Figure 15: Interpolation between scale-free and core-periphery networks (a) shows how the equilibrium liquidity measure depends on the size of the exit shocks for core-periphery networks with varying probability p_c of two previously unconnected core nodes becoming connected. (b) shows the maximum slope of the curves shown in (a) for every value of p_c .

canonical model, such as core-periphery networks, Erdős-Rényi networks or scale-free networks. In the set up we have been studying small shocks typically resulted in relatively small cascades. For some networks we then show that a discontinuous transition occurs for sufficiently large shocks.

It is worth noting that, ex-ante, we do not know whether a particular financial network is in a regime where small shocks lead to small changes in liquidity or in a fragile regime where small shocks can have drastic consequences. To understand this statement, note that any network with N nodes can be interpreted as the remains of a larger network with $M > N$ nodes following an exogenous shock of size $1 - x$.

In the following, we consider the size of cascades caused by the removal of a single node from two “critical” coupled scale-free networks. The critical coupled scale-free networks are generated by initializing two coupled scale-free network with $N = 100$ nodes in the usual way. Then, we randomly remove a fraction $1 - x = 0.35$ of the nodes and iterate the best response algorithm until the best responses have converged. This leaves us with two coupled networks that are close to the discontinuous transition of scale-free networks as $N \rightarrow \infty$. Ex-ante we have no reason to believe that this critical network is less likely than other network configurations.

We then remove a single node at random and study the size of the cascade. Figure 16 shows a histogram of cascade sizes. The distribution is bimodal. In the majority of cases, the removal of a single node does not lead to any cascade at all. However, for a significant fraction of cases, the removal of a single node is catastrophic and the resulting cascade leads to a complete




Figure 16: (a) Small network cascade sizes. (b) Small network cascade sizes vs. fragile set centrality.

evaporation of liquidity.

What determines whether the removal of a node is catastrophic? One way of studying this question is by studying the “fragile set” of a particular node i . The fragile set of node i is the set of nodes whose best response to the withdrawal of node i is to withdraw from both markets. Intuitively, the extent of the cascade following removal of node i increases with the size of a node’s fragile set and interconnectedness of nodes in the fragile set. Figure 16 shows the the fraction of surviving nodes conditional on the sum over the eigenvector centralities of the nodes in the fragile set of node i . To produce the plot we aggregate the results of 500 runs into 5 bins based on the sum of the eigenvector centralities of the nodes in the fragile set. Intuitively, eigenvector centrality is a measure of a node’s influence in a network. The more influence the nodes in i ’s fragile set have, the larger is the size of the ensuing cascade. These results are particularly relevant for supervisory authorities tasked with safeguarding financial stability since they capture two aspects of systemic risk: the probability of a systemic event and it’s extent.

I A simple trading model with intermediation and complementary assets

In Section 2 we introduce a model of intermediaries' willingness to trade in two coupled over-the-counter markets. We intentionally abstract from the end customers of repo and collateral assets (e.g. investors) and focus instead on the actions of the intermediaries. We then interpret the activity of intermediaries as a proxy for liquidity in these two markets. Here, we show how our framework can be embedded into a simple but more complete trading model with intermediation. That is, we outline the structure of trading in an OTC market with intermediaries and end customers. Based on this structure, we propose a model of liquidity provision in two coupled OTC markets that provides micro-foundations for some of the assumptions made in Section 2. This allows us to illustrate why our proxy for liquidity can indeed be related to the liquidity accessible to end customers.

To set the stage for the model, we first describe a OTC market for a particular asset, not coupled to another market. As in the main text, there is a set N of intermediaries connected by a graph \mathcal{G} . For simplicity, in this section, we assume the graph is undirected, i.e. that for any edge $(i, j) \in E$ there is also an edge $(j, i) \in E$.

A directed link $(i, j) \in E(\mathcal{G})$ represents that i can buy the asset from j . In addition to the trading relationships that intermediaries maintain among themselves, each intermediary $i \in N$ has relationships with a set of end customers M_i . For simplicity, assume that the sets of intermediaries' end customers are disjoint and of equal cardinalities.

Consider a static trading setting. Two customers, s and b , are selected uniformly at random from $\cup_i M_i$. The seller s is endowed with a unit of the asset, which it values at v_s . The buyer b has a valuation of $v_b > v_s$ for the one unit. Focus on the case where the two customers do not share the same intermediary, i.e. $b \in M_i$ and $s \in M_j$ with $i \neq j$.⁵² The only way to realize the potential surplus $v_b - v_s$ is to trade via a path in \mathcal{G} : the asset can only be transacted along links in the network. As in the main text, each intermediary has an activity $a_i \in \{0, 1\}$.

For a given action profile \mathbf{a} , provided there is a path from i to j in the network of trading relationships \mathcal{G} consisting of active intermediaries, b and s will be able to trade and the total surplus of $v_b - v_s$ will be realized.⁵³ The existence of such a path, and hence the likelihood that i and j can trade, is thus determined by the component structure of \mathcal{G} . Indeed, an asset can be transacted between customers of i and j in either direction precisely when they are both in the

⁵²As $|N|$ grows and $|M_i|$ remains fixed, this case is by far the most likely.

⁵³The distribution of this surplus among buyer, seller and intermediaries will depend on the bargaining protocol as well as the network topology, see for example Choi, Galeotti, and Goyal (2017) or Manea (2015).

same stable component of \mathcal{G} .

Trading in OTC repo and collateral markets

Having established the basic link between component structure and the liquidity available to customers, we next turn to intermediary behavior in a coupled network model. We show how the intermediaries' best response functions outlined in Section 2 can be obtained from an intermediary profit function that is still stylized, but includes more financial detail than the simplified form in our main text.

There is an OTC market for repo, as well as one for collateral, represented by the networks of intermediaries \mathcal{G}_R and \mathcal{G}_C , respectively. Now the fundamental demands driving trade are realized in both networks. Let s_R and b_R be the ultimate lender and borrower of repo, respectively. Let s_C and b_C be the seller and buyer of collateral, respectively. (We identify these with the corresponding intermediaries, in view of the discussion above.) We study a static setting: the sale and intermediation of collateral, as well as the issuance and intermediation of the repo loan, occur simultaneously in a single period.⁵⁴ We further assume that, also within the same period, a repo loan issued by an intermediary may default due to an exogenous shock. Define the following indicator random variables:

$$\begin{aligned} X_i^C &= \mathbb{1}\{i \text{ is on a directed path from } b_C \text{ to } s_C \text{ in } \mathcal{G}_C\}, \\ X_i^R &= \mathbb{1}\{i \text{ is on a directed path from } s_R \text{ to } b_R \text{ in } \mathcal{G}_R\}, \\ D_i &= \mathbb{1}\{\text{repo loan by } i \text{ has defaulted}\}. \end{aligned}$$

All variables in the following are intermediary-dependent, but the index i is suppressed. The intermediary has two binary decision variables, $a^R, a^C \in \{0, 1\}$, which have the same interpretation as in Section 2.

We now write down a profit function that captures the essential aspects of intermediation of collateral. If the intermediary is on a directed path from the seller to the buyer of collateral, then it derives some (expected) intermediation spread if it successfully intermediates the transaction of the asset (say, a bond). However, this requires that the intermediary have access to repo financing; that is its source of cash (since it does not hold cash reserves). It can obtain this financing if it is on a directed path of active intermediaries from the ultimate source of repo financing to its ultimate "sink." In that case it pays interest on the loan it uses to finance its collateral purchase. Regardless of whether it successfully intermediates, the intermediary, if it

⁵⁴In doing so, we follow classic models such as [Goyal and Vega-Redondo \(2007\)](#). Of course, dynamics are critical to intermediation, and this examined in work such as [Manea \(2015\)](#) and [Condorelli et al. \(2016\)](#). To focus on the distinctive aspects of the coupled network setting, we abstract from the complexities of dynamics.

chooses to be active, pays a cost (e.g., for the staff to operate the trading). Putting values to these costs and benefits, the profit of an intermediary from being active in the collateral market can be written as:

$$\Pi^C = [(vX^C - r)X^R a^R - c] a^C.$$

Here c is the cost of being active. The expected revenue from collateral intermediation conditional on being on an intermediation path is v , and it is multiplied by X^C , which is the indicator variable of being able to intermediate the collateral.⁵⁵ The rate r is the interest rate that must be paid for repo financing during that period. The variables v , X^C and X^R depend on other intermediaries' activity decisions (in the case of v , because the expected profit depends on how many other intermediaries are available to intermediate).

We note several features of this profit function. First, if the intermediary is not active in the collateral market ($a^C = 0$), the profit is zero. Second, if the intermediary is not active in the repo market ($a^R = 0$), the profit is $-c a^C$. Third, in order for intermediating collateral to be profitable, the profit from intermediating the collateral asset vX^C must be larger than the cost of repo financing r . Fourth, notice that the multiplicative term $X^C X^R$ captures the complementary nature of repo and collateral.

Now we consider intermediation of repo loans. Here, the story is parallel to the one above. An external demand for a loan is realized at intermediary b_R and a unit of cash is available at an external repo source, at intermediary s_R . An intermediary can obtain intermediation rents, but suffers a liquidation cost if a counterparty defaults and the intermediary lacks access to the collateral market. The profit of an intermediary from being active in the repo market is

$$\Pi^R = [wX^R - c_D X^R D (1 - X^C a^C) - c] a^R,$$

where w is the repo intermediation spread, c is a cost of being active in the repo market, and c_D is a default cost. The intermediary earns revenue wX^R from intermediation if it is on a repo intermediation chain. If a repo loan that it provided defaults and if it cannot liquidate the collateral in the collateral market, the intermediary incurs a default cost c_D . To liquidate collateral, the intermediary has to be active in the collateral market and in a trading chain to a collateral buyer.

Now assume for simplicity that the random variables X^C , X^R and D are independent *given other intermediaries' activity decisions* (which we suppress as arguments in our notation). Further, let us introduce the following notation: $\bar{X}^C := \mathbb{E}[X^C]$, $\bar{X}^R = \mathbb{E}[X^R]$ and $\bar{D} = \mathbb{E}[D]$. The ex-

⁵⁵The share of revenue an intermediary can expect depends on the exact protocol of trade, etc., but we average over this in setting the value v .

pected profits are then:

$$\begin{aligned}\mathbb{E}[\Pi^C] &= \left[\left(v\bar{X}^C - r \right) \bar{X}^R a^R - c \right] a^C, \\ \mathbb{E}[\Pi^R] &= \left[w\bar{X}^R - c_D \bar{X}^R \bar{D} \left(1 - \bar{X}^C a^C \right) - c \right] a^C.\end{aligned}$$

Note that when the intermediary has no incoming links from active intermediaries in the repo (collateral) market, we must have $\bar{X}^R = 0$ ($\bar{X}^C = 0$). This is because the intermediary can only lie on a trading chain in a particular market if it has an incoming link from active intermediaries.

To understand when we can recover the best responses of Section 2, let us now consider the intermediary's best response for different values of \bar{X}^C and \bar{X}^R . Clearly, when $\bar{X}^C = 0$ and $\bar{X}^R = 0$, the intermediary's best response is to be inactive, i.e. $a_C = a_R = 0$. Also, when $\bar{X}^C > 0$ and $\bar{X}^R = 0$, the intermediary's best response is to be inactive. If $\bar{X}^C = 0$ and $\bar{X}^R > 0$ the intermediary's best response is to be inactive if

$$w\bar{X}^R - c_D \bar{X}^R \bar{D} - c < 0.$$

This condition ensures that the default costs are sufficiently large such that the intermediary is not willing to intermediate repo if it cannot liquidate the collateral. Finally, when $\bar{X}^C > 0$ and $\bar{X}^R > 0$, the intermediary's best response is to be active, i.e. $a_C = a_R = 1$, if the following conditions hold:

$$\begin{aligned}\left(v\bar{X}^C - r \right) \bar{X}^R a^R - c &> 0, \\ w\bar{X}^R - c_D \bar{X}^R \bar{D} \left(1 - \bar{X}^C a^C \right) - c &> 0.\end{aligned}$$

and furthermore $w\bar{X}^R - c_D \bar{X}^R \bar{D} - c < 0$. The first condition states that the revenue from collateral intermediation minus the cost of repo financing is larger than the cost of being active. The second condition states that the revenue from being an intermediary in the repo market minus the expected cost of being unable to liquidate the collateral in case the borrower defaults is larger than the cost of being active.⁵⁶ If the three equations hold (for all nodes), then the best responses are, indeed, as outlined in Section 2.

As for the case of a single OTC market, the intermediaries' activity decision is directly related to the liquidity available to the end customers who wish to trade collateral and repo. For example, two randomly selected buyers and sellers of collateral will only be able to trade if there is an intermediation chain of intermediaries that operate in both markets simultaneously. In a random network model of the type studied in Section 5, this will boil down to the probability that the seller and buyers' intermediaries are in the mutual giant component. Thus, the number

⁵⁶If any of the inequalities is reversed, the intermediary will shut down in one, and hence in both markets.

of active intermediaries can serve as a proxy for the liquidity available for end customers.