

Contextual Areas

Exit Spirals in Coupled Networked Markets

 Christoph Aymanns,^{a,*} Co-Pierre Georg,^b Benjamin Golub^c
^aQuantCo, Inc., Cambridge, Massachusetts 02139; ^bEDHEC Business School, 06200 Nice, France; ^cNorthwestern University, Evanston, Illinois 60208

*Corresponding author

Contact: christoph.aymanns@gmail.com (CA); cogeorg@gmail.com,  <https://orcid.org/0000-0002-6403-941X> (C-PG); ben.golub@gmail.com (BG)

Received: May 8, 2020

Revised: August 10, 2020

Accepted: June 7, 2022

Published Online in Articles in Advance: May 2, 2023

Area of Review: Financial Engineering

<https://doi.org/10.1287/opre.2023.2439>
Copyright: © 2023 INFORMS

Abstract. Strategic agents choose whether to be active in networked markets. The value of being active depends on the activity choices of specific counterparties. Several markets are *coupled* when agents' participation decisions are complements across markets. We model the problem of an analyst assessing the robustness of coupled networked markets during a crisis—an exogenous negative payoff shock—based only on partial information about the network structure. We give conditions under which *exit spirals* emerge—abrupt collapses of activity following shocks. Market coupling is a pervasive cause of fragility, creating exit spirals even between networks that are individually robust. The robustness of a coupled network system can be improved if one of two markets is replaced by a centralized one or if links become more correlated across markets.

History: C. Aymanns would like to acknowledge generous support from the LSE Systemic Risk Centre. C.-P. Georg gratefully acknowledges support from the Algorand Foundation. B. Golub gratefully acknowledges funding from the National Science Foundation [Grant SES-1629446 and SES-1847860].

Keywords: over-the-counter markets • networks • strategic complementarity

1. Introduction

Many important financial instruments are traded in *networked markets*: over-the-counter (OTC) markets in which agents trade with—or otherwise depend on—only a subset of possible counterparties. Because of search or other frictions, market participants cannot adjust their counterparties instantly, and so trading relationships are persistent. Examples of networked markets include the markets for corporate bonds (Hendershott and Menkveld 2014, Di Maggio et al. 2023), repo (Hüser et al. 2021), municipal bonds (Li and Schürhoff 2019), securitizations (Hollifield et al. 2017), and the overnight interbank market (Afonso et al. 2013). Such markets are large,¹ and their fragility has been discussed as an important component of the 2007–2008 global financial crisis. Gorton and Metrick (2012), for example, argued that a central aspect of the crisis was a system-wide run on short-term collateralized debt and in particular on certain nongovernment bond repo markets.²

Market participants' decisions about whether to be active in various markets are strategic; counterparties' activity decisions affect an agent's incentives to participate in a market. At the same time, there are linkages across markets. An agent's activities in different markets can be complementary, for example, when access to one market makes it easier to operate in another.³ Markets may also be linked because of fixed costs. Having established infrastructure for holding or trading

one type of asset makes the marginal costs of entering other markets smaller. A third channel comes from the benefits of engaging in multiple lines of business. A risk-averse agent will be more willing to intermeddle if the revenue stream from one market can be added to that from another market, increasing the risk-adjusted value of the combined revenue stream. Such spillovers have been studied to some extent in centralized markets (Brunnermeier and Pedersen 2009) but are not well understood in coupled networked markets. Indeed, the combination of within-market networked strategic interactions with the linkage of activity decisions across markets presents a rich and complex setting.

Understanding how networked markets respond to stress in the presence of complex payoff interactions is of fundamental importance. In this paper, we study the endogenous determination of activity in such markets. Motivated by our interest in fragility, we study the expectations of an analyst, for example, a policymaker, with only partial knowledge of the network structure, assessing the stability of the system. Our main results identify circumstances in which such analysts should be especially concerned about an *exit spiral*, a sudden collapse of activity in the system.

We model the level of activity of the market as the outcome of a strategic interaction in Section 2 by examining the Nash equilibria of a *networked market game*, where each agent decides whether to be active in various

markets. Activity may represent provision of liquidity, servicing payment obligations, or other actions that support counterparties' activity. An agent's benefit from being active depends on the activity decisions of its counterparties—also called network neighbors—in various markets, rather than in just a single market in isolation. We study a fairly rich class of payoffs that go beyond existing analyses. In our specification of payoffs, agents receive benefits that depend on their neighbors' activities (in a way that allows for flexible spillovers within and across markets) while facing a fixed cost if they are active in any market. Fixed costs are idiosyncratic and drawn from a distribution governed by a parameter (corresponding to exogenous economic conditions) that stochastically shifts costs up or down. After setting out a tractable framework for studying such markets, we state basic results on equilibrium existence in Section 2. Throughout, we make the most optimistic equilibrium selection possible by focusing on activity in the maximal equilibrium. In principle, an omniscient analyst with full knowledge of the network could use this information to fully assess the aggregate level of activity in the market as a whole as well as that of any individual neighborhood.

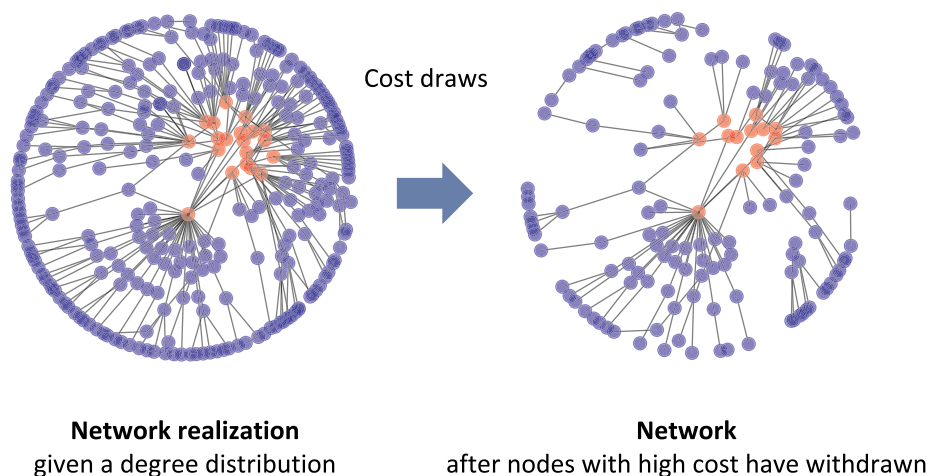
Our main results, however, do not assume an omniscient analyst. Instead, we take the perspective of a macroprudential stress test conducted before the game is played. An analyst specifies a crisis scenario and wants to know how the crisis affects the activity decisions of market participants. A fundamental challenge in such a stress test is that authorities have only limited information about the network structure in the crisis scenario.⁴ The timeline of our model is intended to accommodate this feature and is illustrated for one market in Figure 1. First, a random *crisis network* is realized

in each market, whose links specify the potential trading partners of the agents. The distribution of the crisis networks reflects the analyst's ex ante uncertainty, discussed above, and allows us to capture heterogeneities in the agents' network connectedness (number of counterparties). Next, the agents' costs are drawn from a distribution, with the aggregate cost parameter mentioned above determining the severity of the crisis scenario. For some agents, the fixed cost may be so large that they immediately withdraw, irrespective of any other agent's action.

In the realized crisis networks, the game we described earlier is played. The analyst conducting the macroprudential stress test wishes to assess the expected equilibrium activity in the maximal equilibrium from an ex ante perspective. Thus, our calculations focus on the expected activity measure, with expectations taken over the crisis network realization and the distribution of fixed costs. We examine the expected activity as a function of the severity of the aggregate cost shock in order to probe how sensitive activity is to the economic environment. In particular, we are interested in a situation when small shifts in costs can cause a collapse in activity.

Our first set of results on this are stated for a reasonably general class of network games of strategic complements. We provide a sufficient condition for there to be a discontinuous collapse of activity, an exit spiral. The sufficient condition is an upper bound on the benefit an agent derives from a single link. The condition entails that complementarities across links are crucial to supporting activity. We show how this condition obtains naturally in coupled markets but also in some environments with a single network, for example, a model of financial contagion.⁵ The main contribution of Section 3 is that it studies Nash equilibria with fairly

Figure 1. (Color online) An Illustration of the Timing of our Model in One Market



Notes: First a crisis network is realized. Second, agents' cost draws are realized, and the game unfolds, causing some agents to choose to withdraw from the market.

flexible payoffs and gives sufficient conditions for discontinuous collapses.⁶

Next, we seek to better understand the comparative statics of fragility and when coupling between markets creates fragility that would otherwise be absent. In Section 4, we thus focus on a class of situations in which the complementarity between markets is strong, whereas complementarities within a market are weak. More precisely, we introduce a setting of *simple complementarities*, where each agent finds it optimal to be active in all networks if and only if it has one active neighbor in each. The resulting equilibrium outcomes have a natural graph-theoretic interpretation. Building on results in graph theory, we can precisely describe the conditions for exit spirals. This setting turns out to be one in which coupling between markets is crucial to discontinuous collapses; exit spirals occur exactly when networked markets are coupled, whereas each network would be robust if it were operating on its own. In this setting, we can also give an intuition in terms of network structure for how coupling brings about exit spirals. We do this by introducing the notion of *fragile connectors*. A fragile connector is an agent whose connectedness to others is fragile—dependent on a few counterparties or even just a single counterparty—in one market and which provides support to many agents in other markets.

We use the explicit calculations available in the setting of simple complementarities to examine several follow-up questions about the conditions that foster fragility and policies that may ameliorate it. An important determinant of fragility is how correlated, or overlapping, links are across distinct markets. Our basic model takes neighborhoods across markets to be independent. We show that the fragility of coupled networked markets—their susceptibility to exit spirals—is reduced when markets have more links in common, that is, when an agent's counterparties in one market are more likely to be counterparties in the other. In addition, when one market is centralized with all agents being connected via an exchange, exit spirals do not arise. What is more, when one of the markets is centralized, the size of the shock needed to cause activity to dry up is larger than in the case of coupled noncentralized markets. Lastly, we study explicit numerical examples in canonical types of networks—coupled Erdős-Rényi networks and scale free (power law) networks—to illustrate the starkness of the discontinuity in the case of two coupled network markets, in particular in comparison with the continuous transition in the case where one of the markets is centralized.

Our closing discussion, in Section 5, examines two of our key assumptions: the existence of strategic complementarities and the exogenously fixed network structure. The discussion also surveys the related literature and our contributions relative to it. Although network models have been used extensively to describe interlinked financial obligations, local trading relationships,

and other over-the-counter markets, there has been much less study of the distinctive economic implications of *coupled* networked markets. The conclusions emphasized above—the existence of exit spirals, the benefits of centralizing at least one market, and the benefits of greater overlap between two networked markets—demonstrate some of the implications that are economically interesting and potentially relevant for policy analysis.

2. Model

In this section, we present a model of network games with coupled activity decisions and then illustrate it with some examples. This model serves as the setting of our most general results.

2.1. A Game: Participation in Coupled Markets

We will define a game in which agents strategically choose whether to be active in each of multiple markets. There is a set N of agents and a set \mathcal{M} of markets. Agents may trade bilaterally with one another in each market $\mu \in \mathcal{M}$. Agents interact in a given market only with a subset of other agents; this subset can depend on the market. The set of relationships among agents in market μ 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, that is, ordered pairs (i, j) with $i, j \in N$. We assume that there are no self-links.⁷ All markets share the same node set: $V(\mathcal{G}_\mu) = N$. We call a tuple $(\mathcal{G}_\mu)_{\mu \in \mathcal{M}}$ a *multilayer network*.

We write $i \xrightarrow{\mu} j$ for a link from i to j in market μ , which we interpret as any interaction that supports j 's activity. The μ above the arrow indicates the market in question; if the market is understood or its identity is unimportant, we omit μ . We offer a few examples. In a trading application, a link could represent i trading with j (fulfilling orders, offering quotes, etc. on request from j).⁸ In a debt contracts application, a link could represent a credit relationship whereby i makes (re)payments to j , which supports j 's ability to pay its creditors in turn.

It will be useful to define an agent's neighborhood (e.g., the set of its trading partners or debtors). Because there are multiple networked markets, we will keep track of multiple neighborhoods.

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

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

and can be interpreted as the set of agents supporting i . The in-degree of agent i in market μ is the size of the in-neighborhood $d_{i,\mu}^- = |K_{i,\mu}^-|$.

Analogously, we define the set of agents that are supported by 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 agents.

Agents play a complete-information game of strategic activity choice. Let $a_{i,\mu} \in \{0,1\}$ denote agent i 's decision of whether to be *active* in market μ , with the tuple $a_i = (a_{i,\mu})_{\mu \in \mathcal{M}}$ being agent i 's *action*. We use the notation \mathbf{a} for the profile of all agents' actions, that is, $\mathbf{a} = (a_i)_{i \in N}$. An agent's payoff will depend on the actions of other players and the realization of an exogenous, agent-specific cost c_i .

An agent's payoff is

$$u_i(\mathbf{a}) = b_i(\mathbf{a}; (\mathcal{G}_\mu)_{\mu \in \mathcal{M}}) - c_i \bigvee_{\mu \in \mathcal{M}} a_{i,\mu}, \quad (1)$$

where $\bigvee_{\mu \in \mathcal{M}} a_{i,\mu}$ is defined to be 1 if any of the $a_{i,\mu}$ is 1. Here, $b_i(\cdot)$ is a real-valued function describing the benefits of being active. The cost $c_i > 0$ is a random variable, independent of all other random realizations in the model, reflecting the fixed cost to be paid if the agent is active in any market. The distribution of costs is $F(c; x)$, where $x \in [0, 1]$ is an aggregate parameter that will shift all costs (a simple example is that x is subtracted from all costs).

We impose structure on the game via several assumptions.

Assumption 2.1 (Maintained Assumptions on Benefits and Costs). *We assume the following:*

1. If $x < x'$, then $F(c; x) \leq F(c; x')$, and $F(0; x) = 0$ for all x .
2. b_i is bounded: there is a real number \bar{b} so that $b_i(\mathbf{a}; (\mathcal{G}_\mu)_{\mu \in \mathcal{M}}) \leq \bar{b}$ for all values of the arguments. Moreover, whenever b_i is positive, it is strictly increasing in each $a_{i,\mu}$.
3. We have $\lim_{x \rightarrow 0} F(\bar{b}, x) = 0$.

Part (1) of the assumption says that as x increases, costs stochastically decrease. Part (2) of the assumption says that there is an upper bound \bar{b} on the benefits and that increasing one's action in any component has some (possibly very small) benefit.⁹ Part (3) says that as x decreases to 0, all costs shift to be above this upper bound.

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

$$\mathcal{R}_i(\mathbf{a}_{-i}) = \begin{cases} \mathbf{1} & \text{if } b_i(\mathbf{1}, \mathbf{a}_{-i}; (\mathcal{G}_\mu)_{\mu \in \mathcal{M}}) \geq c_i \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (2)$$

Here, $\mathbf{1}$ refers to the tuple of all ones, and similarly for $\mathbf{0}$. The notation $(\mathbf{1}, \mathbf{a}_{-i})$ refers to the action profile where i takes action 1 in every market and the others' actions are given by \mathbf{a}_{-i} . Note that (2) states that each agent i takes the same action in all markets μ in any best response. This holds because upon choosing to be active in one market the costs of being active in more markets are the same while the benefits are higher. As a result, for the purposes of studying equilibria and possible deviations,¹⁰ it will suffice to keep track of only a single action y_i for each agent across all layers. We will often call it simply i 's action and correspondingly refer

to a single-valued best response and single-valued equilibrium actions.

In view of the previous paragraph, we work with a vector of actions \mathbf{y} in which y_i is the action of node i in every layer.

Assumption 2.2 (Structural Assumptions on Benefit Function). *We assume the following:*

1. The value $b_i(\mathbf{y}; (\mathcal{G}_\mu)_{\mu \in \mathcal{M}})$ depends only on y_ℓ for those ℓ with a directed path to i using links in any market¹¹ in $(\mathcal{G}_\mu)_{\mu \in \mathcal{M}}$.
2. The function b_i is anonymous. Let φ be any measurable function $\varphi: N \rightarrow N$. Upon relabeling all agents according to φ , for all i the function $b_{\varphi(i)}$ is equal to b_i .
3. The benefit function $b_i(\mathbf{1}, \mathbf{y}_{-i}; (\mathcal{G}_\mu)_{\mu \in \mathcal{M}})$ is weakly increasing in the vector \mathbf{y}_{-i} of others' actions and is zero when all in-neighbors ℓ of i have $y_\ell = 0$.

Part (1) gives a meaning to the links in the network; payoffs of i can depend only on agents with a path of dependency links indirectly supporting i . Part (2) of the assumption says that labels do not matter; payoffs depend on network structure and actions, not on players' indices. Part (3) of the assumption is stronger than necessary but very convenient technically because it permits the use of monotone comparative statics techniques.

A key outcome of interest will be the fraction of active agents or, equivalently, the probability that an agent selected uniformly at random is active. For concreteness, we will focus on the maximum expected value of this outcome across pure-strategy equilibria of the game.¹² Denote the set of such equilibria Ξ . Letting I denote an agent sampled uniformly at random, we define

$$\mathcal{A}((\mathcal{G}_\mu)_{\mu \in \mathcal{M}}) = \max_{\sigma \in \Xi} \mathbb{E}[y_I | (\mathcal{G}_\mu)_{\mu \in \mathcal{M}}]. \quad (3)$$

Looking at activity in a maximal equilibrium, defined to be a σ that is a maximizer of the expectation in the above expression, corresponds to making the most optimistic equilibrium selection possible. The expectation is over draws of the costs for each node as well as the draw of the agent I . Because the game is super-modular, there is a unique equilibrium that is maximal, and it is maximal in a very strong sense; for each realization of cost draws, the action profile \mathbf{y} played in this equilibrium pointwise dominates the vector \mathbf{y}' corresponding to any other equilibrium. Moreover, natural iterative dynamics converge to this outcome; start with the profile where all agents are active and then have everyone iteratively best respond to previous actions. Formally, this corresponds to applying \mathcal{R} repeatedly to the vector $(\mathbf{1}, \dots, \mathbf{1})$. For more detail on super-modular games, maximal equilibria, and other natural learning dynamics reaching them, see, for example, Milgrom and Roberts (1990), Milgrom and Shannon (1994), and Adlakha and Johari (2013).

Recall also that we assumed that the game is one of complete information among the agents; this assumption is made for simplicity. However, it turns out that the iterative dynamic just described leads to the equilibrium in question even if agents have much less knowledge than in the complete-information benchmark. In particular, the myopic dynamic described above clearly does not require agents to know others' cost draws in order to best respond to their actions.

2.2. Examples

We now discuss two economic applications that fit into the environment described. These describe relevant environments in which economic surplus is generated from activity within the markets, motivating more concretely the interest in activity levels and exit spirals.

2.2.1. Trading Networks. Our model can be interpreted as describing the willingness of intermediaries to be active in trading networks. The agents N in this model are intermediaries. Each intermediary $i \in N$ is associated with a single *fundamental trader*—a customer to whom this intermediary exclusively provides market access—who may be a buyer or seller.¹³ There are some pre-existing relationships that are capable of facilitating over-the-counter transactions between intermediaries; these constitute the multilayer network $(\mathcal{G}_\mu)_{\mu \in \mathcal{M}}$.

Consider first a market μ in isolation. Fundamental buyers and sellers cannot interact directly but must trade through an intermediary. Each period, a number of random fundamental buyer-seller pairs are drawn uniformly at random, constituting a positive fraction of all possible pairs. A given fundamental seller can trade with its associated fundamental buyer only if there is a directed link between their respective counterparties. More precisely, if fundamental seller A is associated with intermediary i_A , who has a link from intermediary i_B , who is associated with fundamental buyer B , then A and B can trade. In that case, the four parties critical to a trade then share the surplus—for concreteness, equally, as in Goyal and Vega-Redondo (2007). Integrating over the uncertainty, the trading revenues to an intermediary are then proportional to its neighborhood size. When there are multiple markets, the same trading protocol plays out in each market separately, and revenues from trading accrue separately across markets.

There are several ways to model the economics of the coupling between markets. The simplest, and the one in our basic model, is that the two markets are coupled due to there being a common fixed cost paid for being active in at least one market. Markets could also be coupled via the benefits of being active. For instance, suppose an intermediary can choose to allocate a fixed capital budget across markets and that the benefit derived from a given market is proportional to the

allocated capital. Then, unless buyer-seller pairs across markets are perfectly correlated, a risk-averse intermediary will choose to diversify its activity across markets rather than focusing it on a single market, and will be better off for this diversification. Conversely, as an intermediary loses access to one market, it cannot compensate perfectly by reallocating capital to another market.

A final example of coupling comes from technological or institutional complementarities between trading in different markets. An interesting instance of this is the link between the financial market for repo (short-term debt secured by collateral such as bonds) and the market for the underlying collateral. Intermediaries use repo loans to fund intermediation activity in collateral; indeed, repo is a dominant source of short-term credit. At the same time, intermediaries use access to the market for collateral in order to support their own repo lending activities. Access to the collateral market is valuable because it allows an intermediary to liquidate the collateral securing repo loans in case of counterparty default, making repo trading less risky. For a deeper discussion of this institutional coupling, see Gorton and Metrick (2012) and Brunnermeier and Pedersen (2009).

In all of these cases, there are strategic complementarities across intermediaries in being active. Intermediaries' payoffs are increasing in the other intermediaries' activity because (a) intermediaries receive a payoff only if they provide access to trade for their fundamental buyers or sellers, and (b) this becomes more likely as other intermediaries become active.

Note that although the intermediation in this setting is (just for simplicity) restricted to operate in one step, there are clearly longer-distance strategic effects; if one intermediary's neighborhoods are too small to provide enough trading benefits to cover the costs of operation, that intermediary will withdraw and make other traders' neighborhoods smaller, thus reducing their trading revenues, and they may withdraw as well, and so on.

2.2.2. Contagion of Financial Distress. In this application, links reflect financial obligations. Agents are financial intermediaries who decide whether to be active, which means servicing their debts, or inactive, which means defaulting. Each network \mathcal{G}_μ represents a different type of financial obligation.¹⁴ The in-neighbors of intermediary i are intermediaries $j \xrightarrow{\mu} i$, who have an obligation to i via instrument μ . Conversely, the out-neighbors of intermediary i are intermediaries $i \xrightarrow{\mu} j$, to whom i has an obligation via instrument μ . We assume that within and across instruments, all of intermediary i 's creditors—that is, its out-neighbors—have equal seniority, and funds recouped in one market can be used to

pay obligations in any other market without restriction.¹⁵ Thus, in case of default of an intermediary, its creditors are paid *pro rata*. This means that an intermediary cannot default selectively in one market while remaining active in another, because doing so would violate the seniority constraint. Furthermore, as is common, we assume that defaulting intermediaries incur large bankruptcy costs.

In the following, we formalize this model of financial contagion. Let $L_{ji,\mu}$ denote the nominal value of obligations of intermediary j to intermediary i . Similar to Eisenberg and Noe (2001) and Acemoglu et al. (2015), let $\tau_{ji,\mu}(a_{j,\mu})$ be the payments by intermediary j to intermediary i in market μ if j chooses action $a_{j,\mu}$. We assume that for each intermediary i , nominal obligations match nominal assets, that is, $\sum_{j \in K_{i,\mu}^-} L_{ji,\mu} = \sum_{j \in K_{i,\mu}^+} L_{ij,\mu}$, recalling that our K notation denotes neighborhoods. We make the standard assumption that if intermediaries are active, they pay their obligations in full. If they choose to be inactive (i.e., default), they pay only a fraction of their obligations, that is, $L_{ji,\mu} = \tau_{ji,\mu}(1) > \tau_{ji,\mu}(0)$. The payoff of being active for intermediary i is then

$$u_i(\mathbf{a}; (\mathcal{G}_\mu)_\mu) = \sum_{\mu} \left[\sum_{j \in K_{i,\mu}^-} \tau_{ji,\mu}(a_{j,\mu}) - \sum_{j \in K_{i,\mu}^+} \tau_{ij,\mu}(a_{i,\mu}) - B(1 - a_{i,\mu}) \right] - c_i \bigvee_{\mu \in \mathcal{M}} a_{i,\mu},$$

where we assume that the bankruptcy cost B is sufficiently large to make voluntary default unattractive. As before, there is a fixed cost c_i that the intermediary incurs to be active. In this application, it can be thought of as an adverse shock to the intermediary. Note that an intermediary must default if its obligations and costs exceed the payments made to it and the seniority constraint implies that it must then default in all markets.¹⁶

Because intermediary i 's payoff increases in the activity of intermediaries who have an obligation to i , this game is super-modular. As in the general model in Section 2.1, because of the structure of the payoff function and the additional constraints, in any best response, intermediaries take the same action across markets.

It is worth remarking on how our work fits into financial contagion models. Such models have focused on the existence and properties of market clearing vectors (see, e.g., Eisenberg and Noe 2001 and Acemoglu et al. 2015) as well as the sizes or overall damage of cascades of default (see, e.g., Elliott et al. 2014, Acemoglu et al. 2015, and Glasserman and Young 2015). See Glasserman and Young (2016) for a detailed survey. In focusing on activity decisions, we are closer to the second strand.

3. Stress Tests in Random Networks

3.1. Macroprudential Stress Testing and Ex Ante Expected Activity

In our analysis, we focus on situations where the ex ante expected number of active agents is of interest. This is, for example, the case in macroprudential stress tests (see, e.g., Constâncio 2017 and Anderson et al. 2018) in which an analyst has imperfect knowledge about the network structure of financial obligations, anticipates a crisis scenario, and asks what will happen to agents' activity.

There are at least three reasons why analysts, even at central banks, do not have perfect knowledge of the structure of financial networks. First, and almost tautologically, only historical network data can be observed. An analyst undertaking a macroprudential stress test is interested in future scenarios, so the analyst must always extrapolate from historical data. Second, even as data collection efforts of regulators have increased following the global financial crisis, granular information about many financial instruments is not available.¹⁷ Third, many intermediaries are both nationally and internationally connected. However, granular data about international financial linkages between individual financial institutions are not usually available.¹⁸

Our model of the crisis network is thus captured by a distribution over multilayer networks, $G \in \Delta((\mathcal{G}_\mu)_{\mu \in \mathcal{M}})$.

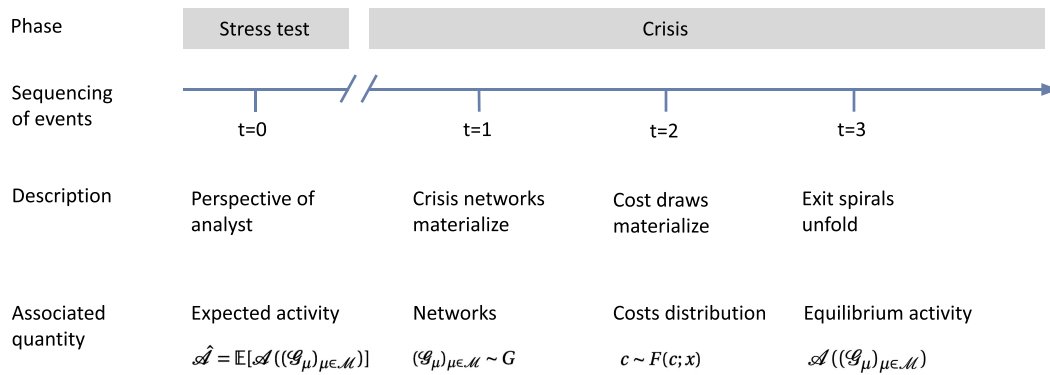
With this in mind, our model has a timeline that can be divided into two phases: a stress testing phase ($t = 0$) and a crisis phase ($t = 1, 2, 3$). Let us first consider the events of the crisis phase. Fix a cost distribution $F(c; x)$ and a distribution G over multilayer networks. At $t = 1$, a multilayer network $(\mathcal{G}_\mu)_{\mu \in \mathcal{M}}$ is drawn. At $t = 2$, the agents' costs are drawn. Finally, at $t = 3$, an equilibrium activity level, given the multilayer network and cost draws, is realized, subject to our selection of the maximum-activity equilibrium. In the stress testing phase, the analyst thus forms an expectation,

$$\hat{A} = \mathbb{E}[\mathcal{A}((\mathcal{G}_\mu)_{\mu \in \mathcal{M}})],$$

where the expectation is taken according to the distributions assessed by the analyst at $t = 0$, that is, $F(c; x)$ and G . In this application, the transformed cost parameter $1 - x$ can be interpreted as a measure of the size of the aggregate shock.¹⁹ In what follows, our object of interest will be the expected equilibrium activity measure \hat{A} and its dependence on the distributions $F(c; x)$ and G , reflecting different beliefs of the analyst over crisis scenarios. Figure 2 illustrates the sequencing of events and highlights where the primitives of the model are relevant.

A key question is whether—from the ex ante perspective of the analyst—the random crisis network is robust or fragile: whether activity degrades gradually as the crisis gets worse or whether it exhibits extreme sensitivity. We probe this by studying how sensitive

Figure 2. (Color online) Stress Testing and Crisis Timeline



the network is to the shock size $1 - x$. In particular, a configuration is said to be fragile if the total activity is highly sensitive to the value of $1 - x$.

3.2. Random Network Model

For the rest of this section, we study contagion in random networks, which we interpret as the crisis networks at $t = 1$, with the randomness modeling the uncertainty at the time of the stress test. We work for technical convenience with a continuum model, where the set of nodes is $N = [0, 1]$. For each market $\mu \in \mathcal{M}$, fix a joint distribution $(p_{jk,\mu})_{j,k}$ where $p_{jk,\mu}$ is the fraction of agents with in-degree j and out-degree k in market μ . (Here, j and k range over the nonnegative integers.) Assume that the distributions satisfy, for each μ , that $\sum_{j,k} j p_{jk,\mu} = \sum_{j,k} k p_{jk,\mu}$, so that in aggregate, the measure of in-links matches that of out-links. We use a model G that draws a multilayer graph on a continuum of agents with three key properties: (i) the realized degree distribution in market μ is $(p_{jk,\mu})_{j,k}$; (ii) any node ℓ 's probability of being in a given node's in-neighborhood in market μ is proportional to the out-degree of ℓ in market μ , corresponding to uniform matching conditional on degrees; and (iii) the neighborhoods of any countable set of nodes are independent conditional on the realized degrees. In Online Appendix B we present a continuum version of the Bollobás configuration model formalizing the construction of such a G .

We focus on degree distributions as the key parameters in the distribution G because this is the main dimension identified as important for contagions (see, e.g., Newman 2002); other structure could also be incorporated into the analyst's distribution over coupled networks. Models based on degree distributions make the study of contagions tractable. Indeed, we will be able to study conditions for fragility without parametric assumptions on the shape of the degree distribution.

We also impose the following.

Assumption 3.1. *The random networks \mathcal{G}_μ are mutually independent random variables with distribution G_μ .*

This implies that, for example, agent i 's in-degree in \mathcal{G}_μ is independent of its in-degree in $\mathcal{G}_{\mu'}$ for $\mu \neq \mu'$. This assumption is convenient for the sharpest statements of results, but not necessary. We relax this assumption in Section 4.7 in the context of our leading illustration.

3.3. General Sufficient Condition for Discontinuity

We now state sufficient conditions on the environment that guarantee that there is a discontinuity in the maximal equilibrium. We begin with general conditions and then discuss them and illustrate them in more specific cases.

Definition 3.1 (Exit Spiral). Fix the parameters of the environment (i.e., of the game and the network structure). We say an *exit spiral* occurs at x_{crit} , if:

1. There is an $A_{crit} > 0$ such that whenever $x > x_{crit}$, we have $\hat{\mathcal{A}} \geq A_{crit}$ asymptotically almost surely;
2. For $x < x_{crit}$, the activity measure satisfies $\hat{\mathcal{A}} = 0$ almost surely.

Let $\bar{b}(1)$ be the supremum of the functions b_i across all multilayer networks and action profiles a in which i has only one active neighbor in any network. Let e_μ^- be the expected in-degree of the origin of a random edge in \mathcal{G}_μ , and let

$$\bar{e} = \max_{\mu \in \mathcal{M}} e_\mu^-.$$

Proposition 3.1 (Sufficient Condition for Discontinuity). *Suppose that there is an x so that $\hat{\mathcal{A}} > 0$ (i.e., there is positive activity at some x). Suppose further that, uniformly over x , we have*

$$F(\bar{b}(1); x) \bar{e} < 1.$$

Then, there is an $x_{crit} > 0$ such that there is an exit spiral at x_{crit} .

A natural case in which the condition holds is when an agent derives no benefit from having only a single link in one network.

Corollary 3.1 (Sufficient Condition for Discontinuity). *Suppose that $\bar{b}(1) = 0$. Then, if there is an x such that $\hat{A} > 0$, then there is an exit spiral at some x_{crit} .*

This arises naturally in coupled markets if access to a market μ' is required to get positive net incremental value from relationships in a market μ . More generally, the sufficient condition in Proposition 3.1 holds if $\bar{b}(1)$ is small enough; note that the condition $F(\bar{b}(1); x)$ imposes a joint restriction on benefits and costs.

3.3.1. Proof Intuition. The idea behind the proof is to analyze a fixed-point condition that must hold at equilibrium and show that this condition is not consistent with equilibria having very low activity levels. As an extreme simplifying assumption, made only for exposition in this intuitive sketch, suppose that all nodes have the same in-degree in all layers.

Let $\rho(\alpha; x)$ be the probability that an agent has $u_i(1; \mathbf{1}_A) \geq u_i(0; \mathbf{1}_A)$ when all in-neighbors of any node²⁰ are playing $y_i = 1$ independently with probability α . This can be viewed as an “aggregate reaction function”; it describes the aggregate level of activity occurring in a best response to activity level α . We show in the proof that equilibrium requires

$$\rho(\alpha; x) = \alpha.$$

The fixed-point condition says that when neighbors are active independently with probability α , on average others have incentives to be active at the same rate.

Now, if the maximum equilibrium level of activity decreases to 0 continuously as x decreases, then there must be solutions of $\rho(\alpha; x) = \alpha$ with arbitrarily small α (and some values of x). But the proof shows that, under the assumed condition ($F(\bar{b}(1); x)\bar{e} < 1$), there cannot be such solutions, because $\rho(\alpha; x) < \alpha$ for α in a neighborhood of 0, uniformly over x . This, in turn, is proved by establishing that the derivative²¹ of ρ in its first entry is less than 1, uniformly across x . This is where the key sufficient condition $F(\bar{b}(1); x)\bar{e} < 1$ comes into play.

This gives some of the ideas of the proof. The actual proof is challenging for two reasons. First, some key steps in the above argument were heuristic, such as that we can treat neighbors of a node as independent at equilibrium, something that is far from obvious in general. Second, assuming uniform degrees reduced the problem to a one-dimensional one, but in reality, the problem is high-dimensional because there are many layers and each node has a “degree type,” with different degrees in each layer. The proof accounts for these subtleties and shows that, nevertheless, the basic idea sketched above can be carried through.

3.3.2. Discussion. The main value of Proposition 3.1 comes from its generality. Although we use the random network structure of the configuration model,

our assumptions in Section 2.1 on the game otherwise allow for a large class of payoffs. For example, we can have payoffs depend on the activity patterns of indirect neighbors, as in the debt example of Section 2.2.2. To our knowledge, despite active literature on phase transitions and, separately, on network games, rigorous results on phase transitions in large-network games are rare.

Although the result is motivated primarily by our study of multiple coupled markets, it is more general. The condition of the proposition, $F(\bar{b}(1); x)\bar{e} < 1$, applies whether the setting has a single market or many. The key aspect of the condition is a small enough benefit of individual links. Discontinuities occur when complementarities are crucial to generating incentives for activity, whether these complementarities are within or across market.²²

Strong complementarities across links can arise within a single market. For example, in the financial obligation example of Section 2.2.2, it may be unlikely that a node would function given just one actively repaying counterparty; multiple counterparties paying their obligations can be necessary to ensure an agent’s activity. Note that such a market would experience fragility essentially irrespective of the situation in other coupled markets. Even if all complementary markets were to perform well, within-market complementarities would cause fragility in such a case.

The presence of spillovers across markets generates a distinct force for fragility. Suppose that there are two markets, and each market individually is robust, unlike the example we just gave. Each individual market would not experience fragility if we fixed a certain positive level of activity in the other markets. Nevertheless, when multiple markets are coupled and subject to shocks, the system becomes fragile. This is, of course, also consistent with the proposition. Because $\bar{b}(1)$ is defined as an upper bound on benefits when a node has just one link across all markets, $\bar{b}(1)$ can be low entirely because of complementarity across, and not within, markets. Thus, coupling of markets can create fragility where none would be expected based on single-market considerations. The next section is devoted to fleshing out this point and exploring the mechanics and implications in more detail in a setting where spillovers across markets play the key role.

4. Coupled Networked Markets With Simple Complementarities

The previous section established a discontinuity result in some generality but left two main questions open. First, the setting there was abstract and yielded only a general, sufficient condition for discontinuity. It is valuable to have a more concrete specification of payoffs to examine exact conditions for discontinuity and the comparative statics of the resulting fragility. To this

end, we specialize to a game with *simple complementarities*. This game is designed to focus on the effects of coupling; links within a market are perfectly substitutable, whereas there is complementarity between an agent’s operations across markets. We will see that the analysis of our equilibrium activity measure is quite tractable and reduces to the graph-theoretic notion of a *mutual giant component*. This allows us to flesh out the point, anticipated at the end of the previous section, that coupling can create fragility when none would be present without coupling.

Second, at a more technical level, working with a continuum of agents in the previous section was a convenient idealization and technically necessary to obtain rigorous results in the generality we had there. However, because real networks contain finite numbers of agents, it is valuable to check that the discontinuities documented above do in fact occur with large but finite numbers of agents. To this end, we consider finite models in this section and study them for large enough populations, rather than working directly at the limit.²³

4.1. Model

We simplify the model introduced in Section 2.1 along three dimensions. First, we restrict attention to two coupled networked markets; we call them the market for asset R and asset C , for concreteness. It will be apparent how to generalize the definitions to an arbitrary number of markets. The choice of names is inspired by the case of the repo and collateral markets, which are complements due to their institutional coupling, discussed in Section 2.2.1 on trading games (but the force for coupling could also come from the other channels discussed there, such as risk-sharing considerations).

Second, as a stark version of the coupling between the two markets, we assume that the benefit function $b(\cdot)$ exhibits a strong form of complementarity; this assumption will make it possible to characterize the structure of equilibria and their fragility much more explicitly than we could in the previous section.

Assumption 4.1.

1. The function b is nonnegative and nondecreasing in \mathbf{a}_{-i} .
2. $b(\mathbf{a}_{-i}) = 0$ if and only if $S_{i,R} = 0$ or $S_{i,C} = 0$, where

$$S_{i,\mu} = \sum_{j \in K_{i,\mu}} a_{j,\mu}.$$

Part (i) of the assumption means that the game is super-modular. Thus, its maximal equilibrium can be found by starting from the maximum feasible actions $\mathbf{a} = (\mathbf{1}, \dots, \mathbf{1})$ and repeatedly applying the best response function, \mathcal{R} , until convergence (Milgrom and Roberts 1990). Part (ii) of the assumption means that an agent receives no benefit from being active if in one of the two markets it has no active neighbors but otherwise

derives a positive benefit. This will allow us to establish a simple link between the maximal equilibrium of the super-modular game and specific graph-theoretic concepts. In the context of coupled trading networks, this amounts to assuming that not having access to the other market prevents earning a positive return, whereas a single counterparty offers enough access to make activity valuable.²⁴

Third, we make the following assumption about the agents’ cost draws.

Assumption 4.2. *The agents’ costs are i.i.d with distribution $F(c, x)$, where, for some fixed $\epsilon > 0$,*

$$F(c, x) = \begin{cases} x & \text{for } c < \bar{c} + \epsilon \\ 1 & \text{otherwise.} \end{cases}$$

Thus, in the limit as $n \rightarrow \infty$, a fraction $1 - x$ of agents, drawn uniformly at random, will receive a cost draw that exceeds the maximum benefit of being active, and hence, these agents will choose to be inactive. Let W denote this set of agents. We will identify the realization of W with that of the cost draw c and will refer to the agents in W as the *shocked* agents and to $1 - x$ as the *size* of the shock. Under these assumptions, an agent’s best response to its in-neighbors’ actions simplifies to

$$\mathcal{R}_i(\mathbf{a}_{-i}) = \begin{cases} \mathbf{1} & \text{if } i \notin W \text{ and } S_{i,R}(\mathbf{a}_{-i}) \cdot S_{i,C}(\mathbf{a}_{-i}) > 0 \\ \mathbf{0} & \text{otherwise.} \end{cases} \quad (4)$$

As before, this allows us to reduce the analysis from now on to a one-dimensional description of an agent i ’s action, y_i , which, without loss of generality, is that agent’s action in all markets.

4.2. Equilibrium in Given Networks

In the following, we will provide a graph-theoretic description of the maximal equilibrium. Suppose that, at $x = 1$, the networks satisfy the following assumption.

Assumption 4.3. *All agents have at least one incoming edge in \mathcal{G}_R and \mathcal{G}_C .*

In this case, when $x = 1$, so that no firms are shocked, the maximal equilibrium is $\mathbf{y}^* = \mathbf{1}$. We can think of this as a *preshock* situation, in which no agents have received a cost shock; our assumption guarantees that no agents withdraw from the networks when $x = 1$. This result is stated and proved as Lemma A.2 in Appendix A.4.1 (using some terminology we introduce later in this section). Starting from such a baseline, we can consider equilibrium outcomes for more interesting realizations of the cost shock when $x < 1$. We refer to this regime as the *postshock* regime. In the following, we characterize how the postshock equilibrium depends on the structure of the networks \mathcal{G}_C and \mathcal{G}_R .

Let $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$ denote the networks after all of the edges corresponding to the agents in W (i.e., those

that have received a bad cost draw, $c > \bar{b}$) have been removed. To link the equilibrium outcome to the structure of $\mathcal{G}_C(W)$ and $\mathcal{G}_R(W)$, we first define a stable subset of nodes in a network \mathcal{G} .

Definition 4.1 (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 4.2 (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 all $\mu \in \{R, C\}$.

The existence and size of mutually stable subsets is closely related to the existence of a maximal equilibrium in our game. Indeed, the following proposition provides a precise connection between the two objects.

Proposition 4.1 (Maximal Equilibrium and Mutually Stable Subsets). *In the maximal equilibrium action profile conditional on the set W , denoted by \mathbf{y}^* , the set of active agents (those with $y_i^* = 1$) equals the maximal mutually stable subset of $(\mathcal{G}_R(W), \mathcal{G}_C(W))$.*

In other words, if we fix W and 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 action profile W induced in the maximal equilibrium. That this is an equilibrium follows from the form of \mathcal{R} . First, none of these agents have been shocked, because the shocked ones were removed from $(\mathcal{G}_R(W), \mathcal{G}_C(W))$. Second, all agents in such a set have, by definition of a mutually stable subset, at least one incoming link in both networks from other agents in the set, so by definition of \mathcal{R} , it is an equilibrium for all of them to be active. To complete the proof of Proposition 4.1, we must show that no equilibrium has a set of active agents 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 agents (by definition of \mathcal{R}) is mutually stable. Thus, it must be contained in the maximal mutually stable set.

4.3. Equilibrium in Random Networks

We now study the equilibrium in the random networks setting with the stress-testing motivation discussed in Section 3.1.

4.3.1. Network Model. We now define the network model we study in this section. As before, we will specify the distribution of crisis networks. The only difference relative to the continuum configuration model is that there is a finite number of nodes, allowing us to analyze the key questions without the idealization of the continuum limit.

Let $\mathbf{d}_\mu^+ = (d_{i,\mu}^+)_{i=1}^n$ and $\mathbf{d}_\mu^- = (d_{i,\mu}^-)_{i=1}^n$ be sequences of nonnegative integers representing the out-degree and in-degree, respectively, of all agents in market $\mu \in \{R, C\}$, whereas before, $n = |N|$. In Appendix A.4, we impose some technical conditions on these degree sequences that make random graphs generated from them well behaved. These assumptions ensure, for example, that the first and second moments of the degree distributions remain bounded in the limit $n \rightarrow \infty$ and that the sum of all out-degrees matches the sum of all in-degrees. 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 that 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 large networks, that is, ones with n sufficiently large. In this limit, our technical assumptions impose that these degree sequences are consistent with some joint distributions $(p_{jk,\mu})_{j,k}$ for $\mu \in \{R, C\}$, where $p_{jk,\mu}$ is the fraction of agents with in-degree j and out-degree k in network μ . We fix these degree distributions throughout the section. We assume that

Assumption 4.4. *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, agent i 's out-degree in \mathcal{G}_R is independent of its out-degree in \mathcal{G}_C . In Section 4.7, we relax this assumption. In addition, we assume the following.

Assumption 4.5. *For $\mu \in \mathcal{M}$ 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 4.3, and thus (without any shocks) all nodes are in a maximal stable set and indeed in a mutually stable set (see Lemma A.2 in Appendix A.4.1).

4.3.2. The Giant Out-Component. In the following, we will show that in the random network case as $n \rightarrow \infty$ the characterization of the maximal stable set in a given network is reducible to the study of the *giant out-component* in the network. Similarly, the maximal mutually stable set in a given multilayer network can be reduced to the *mutual giant out-component*. We now build up to the definition of these objects. We will begin by introducing the giant out-component for a single network and hence, will drop the network index where appropriate to simplify notation.

It is important to note that $\mathcal{G}(W)$ —the network with the shocked agents stripped of their edges—is equal in distribution to a draw from G 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 4.3 (Strongly Connected Subsets and Components). In a network $\mathcal{G} = (V, E)$, a subset $V' \subset V$ is *strongly connected* if, for any nonempty proper subset $V'' \subseteq V'$, there is an edge from V'' to $V' \setminus V''$.²⁵ A *strongly connected component* is a maximal strongly 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 A.4. We can now define the giant out-component.

Definition 4.4 (Giant Out-Component). Define $L_{\mu,n}$ to be any largest-cardinality strongly connected component $L'_{\mu,n}$ and all nodes reachable by following a directed path out from $L'_{\mu,n}$.²⁶ Then, the sequence of graphs is said to have a giant out-component if $L_{\mu,n}$ is well defined for all large enough n and for some constant $c > 0$ that depends only on $(p_{jk,\mu})_{j,k}$, $|L_{\mu,n}|$ tends, with probability one, to cn .

It can be shown that, under the technical assumptions made above, if the sequence has a giant out-component, then asymptotically $L_{\mu,n}$ is uniquely determined; see Cooper and Frieze (2004). Suppressing the n index, we denote the subgraph of \mathcal{G}_μ associated with it by $GC_o(\mathcal{G}_\mu)$.

The concept of a giant out-component generalizes to multilayer networks (Buldyrev et al. 2010). In particular, we make the following definition.

Definition 4.5 (Mutual Giant Out-Component). For large enough n , the *mutual giant out-component* $MGC_o(\mathcal{G}_R, \mathcal{G}_C)$ is defined to be the induced graph on the intersection of nodes in $GC_o(\mathcal{G}_R)$ and $GC_o(\mathcal{G}_C)$.

In other words, a node is in the mutual giant out-component if it is in the giant out-component in each network. Moreover, the fraction of nodes in a mutually stable subset that are outside the mutual giant out-component is negligible for large n (Buldyrev et al. 2010). Therefore, the size of the mutual giant out-component proves an asymptotically exact estimate of the size of the mutually stable set.

Thus, by Proposition 4.1, the limit fraction of nodes in the mutual giant out-component is the fraction of active agents in equilibrium as n goes to infinity, which we denote by \hat{A} , suppressing the subscript n . For further details, see Appendix A.4.2.

4.4. Exit Spirals

We can now present our result on exit spirals in coupled networked markets with complementarities. We need one more assumption.

Assumption 4.6. Fix $(p_{jk,\mu})_{j,k}$; select, uniformly at random, a fraction $1 - x$ of nodes and remove all their edges.²⁷ Let $c(x)$ be the size of the giant out-component in the

resulting graph. Giant component concavity holds if c is concave in x over the range where $c(x) > 0$.

Then we have the following.

Proposition 4.2. Under the assumptions in this section, in coupled networked markets with complementarities, there is an exit spiral at some x_{crit} , as defined in Definition 3.1.

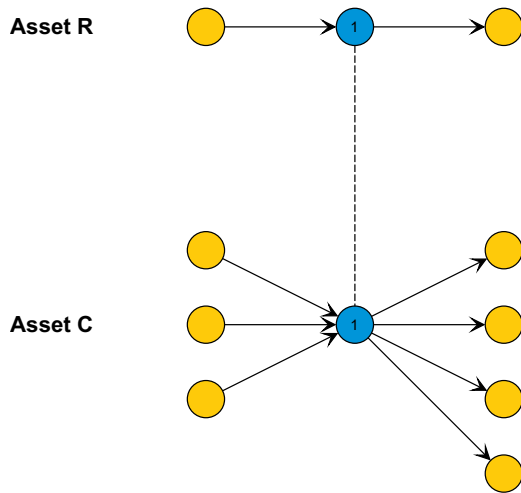
Proposition 4.2 states that there are two activity regimes, and the one that obtains depends on the fraction of agents that receive a bad cost draw. If the shock is sufficiently small, that is, less than some $1 - x_{\text{crit}}$ in size, agents in both markets remain active. However, if the shock size increases beyond its critical value by an arbitrarily small amount, activity in both markets vanishes; that is, an exit spiral occurs. The transition between the active and the “frozen” (i.e., zero-activity) market regime is discontinuous at the critical shock size; starting from an active market, the withdrawal of a very small measure of additional agents is amplified through the coupled structure of the markets to the extent that all activity disappears entirely. Formally, this is reflected in the fact that at $1 - x_{\text{crit}}$, activity goes from a strictly positive value to 0.²⁸

If shock size is increased, more agents become inactive, and some agents lose all their counterparties, being forced to withdraw as well. The question is when and why this should result in a discontinuous loss of activity.

Here, we can provide some intuition for how the interaction between the networked markets can amplify collapses. In subsequent discussions, we offer more technical discussions explaining when this possibility is realized. A key ingredient in an exit spiral is a type of node that we call a *fragile connector*, which we now define informally (see Figure 3 for an illustrative example). A fragile connector is an agent with few counterparties in the market for one asset (say, R) and many counterparties in the other market (say, C). The agent is fragile because it can easily lose access to its support in one market, which becomes more likely as the shock size is increased. Once this occurs, the withdrawal of the node causes a severe spillover effect on the other market, where it has a chance of destabilizing other fragile connectors of the “opposite” type. The proposition implies that there is indeed a critical shock size at which such a spiral unfolds.

The analysis of when such a cascade occurs at some shock size is closely related to our general discontinuity analysis. Indeed, as we now demonstrate, in a continuum configuration model, the result would follow from an application of Corollary 3.1. Assumption 4.5 implies that at $x = 1$ there exists a mutual giant out-component, and the fraction of agents outside it vanishes as $n \rightarrow \infty$ (see also Appendix A.4.2). Nodes in the mutual giant out-component form the set of active agents in the

Figure 3. (Color online) Illustrative Example of a Fragile Connector



Notes. Agent 1 is a stylized example of a fragile connector. It receives support in the market for asset R from only a single agent and is therefore susceptible to the withdrawal of this critical neighbor. At the same time, agent 1 is the sole provider of support to a number of agents in the market for asset C ; it acts as a connector in this. Thus, if agent 1 decided to be inactive, it would lead to a loss of access to support for a large number of agents in the market for asset C . The existence of such fragile connectors is crucial in the mechanism underlying Proposition 4.2.

maximal equilibrium. As a result, for $x = 1$, there is positive activity. Assumption 4.1 implies that $\bar{b}(1) = 0$. Agents derive no benefit from being active if there is no active neighbor in one of the two networked markets. It is straightforward to check that the rest of Assumption 2.1 holds as well. This means that Corollary 3.1 applies and exit spirals must obtain.

However, because one of the purposes of this section is to establish the analogous result in a finite-population model, we must undertake a more detailed analysis, which we do in the next subsection.

4.4.1. Graph-Theoretic Proof Intuition. We now sketch how Proposition 4.2 is proved by studying the size of the mutual giant out-component as we vary x (see Appendix A.5 for the detailed proof). First, recall that activity 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. As mentioned above, we can relate this to a notion from random graph theory called the *mutual giant out-component* (Buldyrev et al. 2010). The fraction of agents in a mutual giant out-component satisfies a fixed-point condition. A node's probability of being in the mutual giant out-component depends on its neighbors' probability of being in it. (The degree sequences \mathbf{d}_μ^+ and \mathbf{d}_μ^- (for both values of μ) and the size of the shock enter this equation.) The key insight is related to the intuition sketched in Section

3.3.1. If the mutual giant out-component is small (comprising a fraction c of nodes), then the probability that a node happens to link to nodes in it in two layers simultaneously is smaller (of order c^2). Thus, "small" solutions for the mutual giant out-component size are not sustainable, and activity must disappear discontinuously.

There is another perspective on the fixed-point equation analyzed in this argument. 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 shock determine the branching probabilities. The probability of nonextinction for such a process is determined by a system of coupled equations whose largest solution yields $\hat{A}(x)$, and these are precisely the equations discussed in the previous paragraph. If, starting from a randomly chosen agent, the branching process goes extinct, a node cannot form part of the giant out-component. The argument of the previous paragraph thus shows that under the conditions we have studied, small nonextinction probabilities are not sustainable. Appendix A.4 formalizes this further. Amini et al. (2016) and Elliott et al. (2014), as well as Buldyrev et al. (2010) use branching process ideas to characterize giant components, although none of these papers characterize outcomes for arbitrary degree distributions in coupled networks.

4.5. Avoiding Exit Spirals Through Market Centralization

One of the advantages of the explicit payoffs in this section is that we can characterize the threshold x_{crit} for an exit spiral and compare the robustness of the market to related situations of interest, such as those that might be brought about by policy interventions.

An important counterfactual is that in which one market, say C , is centralized. By that we mean that each agent active in the market can interact with all other agents, as it is the case in a securities exchange with a central clearinghouse. Formally, let $\bar{\mathcal{G}}_C$ denote the complete network, and assume for this subsection that $\mathcal{G}_C = \bar{\mathcal{G}}_C$. Then, the set of active agents in the maximal equilibrium defined in Section 2.1 corresponds to the maximal stable set in \mathcal{G}_R , because, since \mathcal{G}_C is complete, it is also a mutually stable set.

Following the same argument as in the previous section, it can be easily deduced that, asymptotically, the fraction of nodes in the giant out-component of network \mathcal{G}_R is the fraction of nodes in the maximal stable set in the random graph with n nodes. Thus, the size of the giant out-component c , as defined in Definition 4.4, is the equilibrium expected activity.

Recall that, in the shocked regime with $x < 1$, the degree distribution governing \mathcal{G}_R is thinned, and thus

as x changes c changes too. We can now state the main result for the case where one market is centralized:

Proposition 4.3. *Let \mathcal{G}_R be drawn as described in Section 4.3.1, and let $\mathcal{G}_C = \bar{\mathcal{G}}_C$ be a complete network. Then, there exists an x_{crit}^c such that*

- (a) • \hat{A} is continuous at x_{crit}^c ,
- $\hat{A} > 0$ for $x > x_{crit}^c$,
- $\hat{A} = 0$ for $x \leq x_{crit}^c$.

(b) *Furthermore, if $1 - x_{crit}$ is the shock size at which activity vanishes when $\mathcal{G}_R, \mathcal{G}_C$ are both random networks as defined in Section 4.3 and Assumption 4.6 holds, then $1 - x_{crit}^c > 1 - x_{crit}$.*

This result shows that, when one of the two markets is complete, for example, if it corresponds to a centralized exchange, the transition from the active to the frozen market regime is no longer abrupt but smooth. In addition, the transition always occurs at a larger shock size in the presence of a centralized exchange. Here, activity is less sensitive to the inactivity of a single agent and can only vary smoothly with the size of the shock. An exit spiral, with discontinuous withdrawal, is not possible. Comparing the two propositions emphasizes the stabilizing effect of a centralized exchange on activity in the presence of shocks.

The main difference between the case when both networks are genuinely over the counter and the centralized market case is the absence of fragile connectors. Because in the complete network for asset C all agents support each other, there can be no contagion through the complete network. Although an agent may be fragile in the market for asset R, its withdrawal cannot lead to a large number of subsequent withdrawals in the market for asset C. Therefore, the absence of fragile connectors removes the amplification effect that results from the complementarity of the markets for asset R and asset C. This leads to a smooth transition and an increased critical shock size.

The proof of Proposition 4.3(a) is a straightforward application of known results on the giant out-component in directed networks (see Newman 2002 and Cooper and Frieze 2004). These immediately yield the smooth transition of activity regimes in Proposition 4.3(a). As in our proof of Proposition 4.2, the technique behind the analysis of the size of the giant out-component as a function of x is expressing the probability that a random node is in the giant out-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 out-component and its behavior as we vary x . Intuitively, Proposition 4.3(b) obtains that because an exit spiral can never propagate through the complete network, all other choices of \mathcal{G}_C must lead to a smaller critical

shock size. For the full proof of Proposition 4.3, see Appendix A.5.

4.6. Numerical Examples

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 now illustrate the results in Propositions 4.2 and 4.3 by considering particular degree distributions: the binomial (Erdős-Rényi) and scale-free (power law) degree distributions.

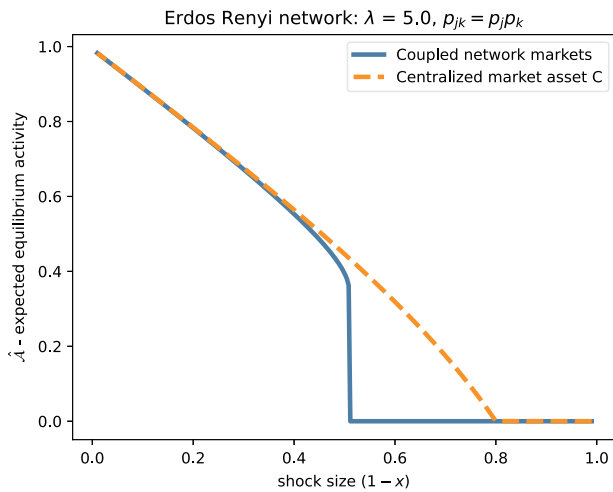
Erdős-Rényi. This is the simplest type of random network, corresponding to random matching of counterparties. This type of graph is obtained by letting each directed link exist with a given probability q . We hold the average in- and out-degree $\lambda = nq$ fixed as n varies. Here, because of the independence of in- and out-degrees, the joint degree distribution factorizes into $p_{jk} = p_j p_k$ with $p_j = p_k$ and $p_k = \binom{n-1}{k} q^k (1-q)^{n-k-1}$.

Scale Free. A more realistic random graph structure models the wide heterogeneity in degrees often seen in real networks. 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 activity measure of the maximal equilibrium \hat{A} numerically; the detailed calculations can be found in Online Appendix E. For each degree distribution, we also compute \hat{A} when the market for asset C is replaced by a complete network. In Figures 4 and 5, we present the results for the binomial and scale free degree distributions, respectively. The findings of Propositions 4.2 and 4.3 are apparent. First, in the case of two coupled networked markets ($\mathcal{G}_R, \mathcal{G}_C$), there is a discontinuous transition from the active to the frozen market regime. Second, when the market for asset C 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 the average degree is actually lower.

Figure 4. (Color online) Erdős-Rényi Network Example



Notes. Equilibrium expected activity \hat{A} as a function of the size of the shock $1 - x$ in an Erdős-Rényi network with average degree $\lambda = 5$.

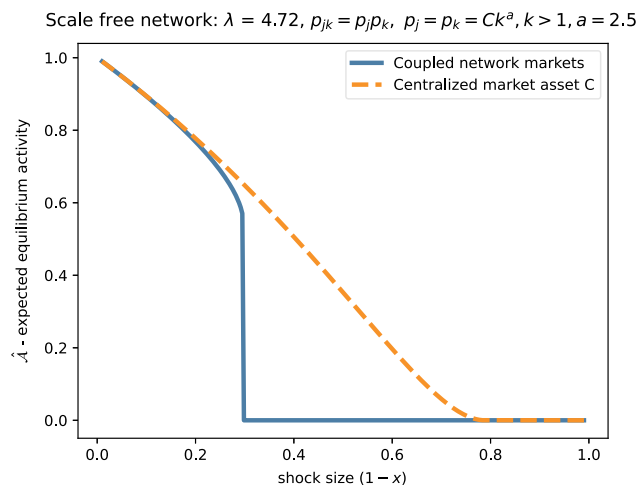
4.7. Correlations Between Networked Markets

Assumption 4.4 imposed that \mathcal{G}_R and \mathcal{G}_C are independent draws from their respective distributions. Under this assumption, an agent’s counterparties in the market for asset R are uncorrelated with its counterparties in the market for asset C . In many cases, however, the presence of an edge between two agents 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, there is a sequence $M(n)$ so that, because $n \rightarrow \infty$, both $|E(\mathcal{G}_C)|/M(n)$ and $|E(\mathcal{G}_R)|/M(n)$ tend to 1. Define the *overlap* measure of network similarity as

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

Figure 5. (Color online) Scale-Free Network Example



Notes. Equilibrium expected activity \hat{A} as a function of the size of the shock $1 - x$ in a scale free network.

If \mathcal{G}_R and \mathcal{G}_C are independent, because $n \rightarrow \infty$, 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 4.2 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 market for asset C 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 4.3 apply, whereas for $\omega = 0$ the results in Proposition 4.2 apply.

Varying the overlap parameter ω then interpolates between these two extremes. If the overlap is not too high, the sensitivity to shocks is still quite pronounced, as in the independent-networks case covered by Proposition 4.2. For higher values, the dependence of activity on x is much smoother. Numerical and heuristic calculations 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. The details can be found in Online Appendix F.

Thus, as the two networked markets become more similar, that is, as agents share more counterparties across markets, activity becomes more resilient to cost shocks. This seemingly contradicts the notion that a diversified set of counterparties protects against random shocks to an agent’s counterparties. However, the complementary nature of the networked markets makes diversification here harmful rather than helpful and leads instead to an amplification of cost shocks through the networked markets.

In Online Appendix F, we also compute \hat{A} as a function of the shock size explicitly for two levels of overlap $\omega = \{0.2, 0.8\}$. As expected for $\omega = 0.2$, we observe a discontinuous transition from the active to the frozen regime, whereas for $\omega = 0.8$, we observe a continuous transition. We compute \hat{A} both via our approximate method and numerically by iterating the best response functions. Note that the heuristic solution approximates the numerical solution quite well, although there are clear finite size effects for the numerical solution (we used $n = 2,000$).

5. Discussion

5.1. Strategic Complementarity

We have studied network games with strategic complementarities throughout. This was one of the main structural assumptions imposed on the game in Section 2.

Strategic complements is a restrictive assumption; it rules out, for instance, a situation where competition from a new market participant reduces another agent's intermediation rents more than it would increase that agent's new business.

However, the assumption of strategic complements is used only to ensure that certain technical conditions are satisfied in our analysis. For instance, it is used to show the existence of equilibria in the continuum configuration model and to establish certain monotonicities that underlie the techniques there. We are then able to formalize the kinds of "mean field" calculations that are often performed heuristically in related models (see, e.g., Newman 2002).

Once these technical properties are established, strategic complementarities play no further role in the analysis. Indeed, our main analysis of exit spirals studies an aggregate reaction function, as discussed in Section 3.3.1. It is here that the substantive conditions that govern exit spirals are introduced. These properties control the reaction function $\rho(\alpha; x)$ near $\alpha = 0$ (i.e., analyzes best responses when few other agents are adopting). None of this analysis relies on strategic complements or even monotonicity of the aggregate reaction function ρ . To be clear, these technical properties play an important role in our analysis. However, the core phenomena we analyze are more general, and carrying through the analysis discussed in Section 3.3.1 under weaker assumptions on the game presents an interesting direction for future work.

5.2. Network Formation

To keep our analysis tractable, we have assumed that the crisis networks are drawn independently at random across markets. This means that agents cannot choose with whom to connect but instead take their neighbors as given when deciding whether to be active or not. An important albeit difficult question is how outcomes would change if instead the network were (at least partly) formed endogenously. It is not obvious whether network formation mitigates or exacerbates exit spirals; there are forces in both directions, and which force dominates depends on the specifics of the network formation game as well as the subsequent activity game. A good way to build intuition about these forces is to ask when fragile connectors—the drivers of our results on exit spirals under simple complementarities—become more or less likely.

If endogenous network formation causes counterparties to become correlated across networks, overlap would increase. This would make fragile connectors less likely and thus mitigate exit spirals, as discussed in Section 4.7. Several authors have provided possible reasons for counterparties to be correlated (Erol and Vohra 2020, Farboodi 2021), including Elliott et al. (2021), who showed that risk shifting can cause an agent to link to predominantly similar agents in both networks. Overlap

could also increase if the cost of forming a link to an agent in one market is very small when a link to that agent already exists in another market. Finally, if agents know a potential neighbor's degree type, they might choose to avoid connecting to a (potential) fragile connector in the first place, thereby making them less prevalent in the network.

On the other hand, the potential harm caused by exit spirals constitutes an externality. In particular, fragile connectors do not internalize the damage that follows indirectly when they decide to be inactive. Therefore, fragile connectors will, from the perspective of a social planner, underinvest in creating and maintaining additional links to protect against shocks. Relatedly, risk-averse agents might prefer to link to different agents across networks such that if a neighbor withdraws in one market, the link in the other markets is not automatically lost as well. This would in turn reduce the overlap between networks, making exit spirals more likely (see again our discussion in Section 4.7).

Finally, the nature of exit spirals makes for a particularly stark contrast between the returns to investment for a planner and for individual agents. Take the case where an exit spiral exists at a value $x_{\text{crit}} > 0$, and consider an extended model where a social planner could, at a cost, mitigate the aggregate shock, effectively increasing x .²⁹ The social returns to this would be very high because aggregate activity has a very steep slope in x near the critical value of x_{crit} . Now suppose that, at the same configuration, each individual agent could instead invest to reduce the impact of the shock locally, for example, by creating more links. Holding other agents' behavior fixed, the system winds up at the critical status quo on the brink of collapse, and such extra investment would make no difference when all of the agent's neighbors' collapse in an exit spiral. Exit spirals are fundamentally an aggregate phenomenon, and individual agents are powerless to stop them when they happen, even if they can change their local environment. Avoiding or mitigating them requires a coordinated response or a planner's direct intervention.

This short discussion already highlights some of the intricacies of adding network formation and comparing a social planner with decentralized outcomes. We believe these issues raise some interesting economic questions for future work. Elliott et al. (2022) offered some analysis along these lines in a related setting motivated by a distinct supply chain application.

5.3. Relation to the Literature

Trading and Liquidity. Our model allows us to study the existence of exit spirals in networked markets. As discussed in Section 2.2.1, the decision modeled could be that of intermediaries choosing whether to be active in a trading game. This links our paper to the literature studying illiquidity spirals. The closest analog to our

model in this literature is the model of Brunnermeier and Pedersen (2009), who studied the coupling between repo and collateral. Our model of exit spirals in coupled networked markets explicitly takes into account the OTC network structure of these markets. This sets us apart from Brunnermeier and Pedersen (2009), who studied the feedback between market and funding liquidity in centralized markets, abstracting from network structure. Our model shows that the network structure alone, abstracting from haircut and pricing feedback, can lead to a significant amplification of exogenous shocks in coupled networked markets. Acharya et al. (2011) showed 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. Differently from these two papers, we show that, given the complementarity of two coupled networked markets, the networked nature of these markets is sufficient to generate a feedback between the two markets that amplifies an initial exogenous shock.

A related paper studying market freezes is Acemoglu et al. (2020), which showed that anticipated defaults can lead to ex ante credit freezes in interbank networks. Their model has a specific trading protocol, and market freezes obtain as subgame perfect equilibria of a borrowing and lending game. Although our paper also studies market freezes, our focus is on the role of complementarities and coupling and the question under what conditions a positive activity equilibrium disappears discontinuously, that is, when exit spirals are possible.

Financial Over-the-Counter Networks. Most empirical studies of OTC markets focus on financial exposure networks. A common finding of these studies is that OTC markets have nontrivial network structure, rather than being complete (see, for example, Di Maggio et al. 2023 on the interdealer corporate bond market and Li and Schürhoff 2019 on the market for municipal bonds). Craig and Von Peter (2014) showed, for example, that the German interbank market has a network structure that can best be described as a core-periphery network.³⁰ In the international context, Gabrieli and Georg (2014) showed that the Euroarea interbank market follows a core-periphery structure less closely, with relatively safe and large international banks connecting the different national core-periphery networks. The importance of relatively safe intermediaries in facilitating the functioning of an over-the-counter interbank market is also highlighted by Frei et al. (2022) in a rich model of trading in OTC markets. We take from this literature the point that the structure of over-the-counter trading networks is important. Our contribution to it is to highlight the importance of coupling between such

markets and to study qualitative conditions under which such coupling introduces new fragilities.

Contagion in Financial Networks. In Section 2.2.2, we discuss how our model can be applied to the case of financial contagion. 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, Zawadowski 2013, Elliott et al. 2014, Acemoglu et al. 2015, and Farboodi 2021, as well as Glassermann and Young 2016 for an extensive overview). Burkholz et al. (2016) studied cascading failures in a multiplex network. Firms in their model have a core and a subsidiary business unit that 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.

Multilayer Network Theory. There is a growing literature in applied mathematics and physics on coupled, or *multilayer*, networks. A seminal paper is Buldyrev et al. (2010), and since then there have been a variety of applications, including, for example, to the question of whether firms should spin off subsidiary units. In terms of the theory of coupled networks, our contribution is to study the equilibrium of network games defined on such networks and to formulate conditions on payoff and network structure that facilitate exit spirals. Our general results on coupled network games of strategic complements in Section 3.3 go well beyond the existing literature, accommodating arbitrary network games of strategic complements and highlighting the key role of complementarity across links as the main economic force.

Games on Networks. Our paper is also related to the networks literature in economic theory, especially to the literature on 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 such as (in our context) "everyone withdraws." Jackson and Yariv (2007) and Galeotti et al. (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. The main innovation relative to these papers 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 of the network than has been reported previously because of the discontinuities discussed above. Indeed, the study of large-scale phase transition phenomena has been challenging because of the difficulty of formulating equilibrium conditions capturing the key macroscopic outcomes (Jackson and Yariv 2007). Our results provide some methodological advances as well as concrete illustrations that there are interesting phenomena that can be studied using these advances.

6. Conclusion

We develop a general model of large coupled networked markets and study the maximal equilibria of the resulting super-modular activity game from the perspective of an analyst who is uncertain about the network. Networked markets can be coupled for a variety of reasons. In our canonical example, we consider a model where an agent derives benefits from other agents' activity—in any market—and incurs a fixed cost for being active in any market. But networked markets can also be coupled via the benefits of being active, that is, if activity in two markets provides a strictly larger benefit than the sum of the benefits of being active in each market individually. A particularly stark form of coupling arises if two markets are institutionally coupled, for example, when secured funding is necessary to create demand for collateral and where the availability of secured funding depends on the existence of a market for the underlying collateral.

Coupling creates fragility because it makes links highly complementary and thus prone to simultaneous collapse. Indeed, coupled markets are distinctively prone to fragility; we show that even without any complementarity between links in each constituent network, a coupled network system is prone to exit spirals. Replacing at least one networked market by a centralized exchange reduces the extent of exit spirals and increases the resilience against exogenous shocks. Increasing overlap—making the networked markets more similar—also increases resilience.

Our analysis highlights a number of pertinent issues for policymakers. First, we illustrate the potential fragility of activity in networked markets and show how it may be reduced by moving toward centralized exchanges. Second, our results highlight the importance of better measurement of the structure of these markets, in particular with respect to identifying fragile connectors.

Moving forward, a natural next step would be to extend the model to account for incomplete information among agents when they decide whether to be active. In this setting, the realization of the shock profile or the parameters entering agents' best response functions may

only be partially known to them. This would, for example, allow us to study how activity is affected by changes in agents' beliefs about the distribution of the exogenous shock. In turn, this would permit a study of the realistic phenomenon that market activity may cease entirely because of bad news, connecting to the literature on higher-order beliefs in games of incomplete information (Golub and Morris 2017, Morris and Yildiz 2019). Interactions between the fragilities we have identified and ones arising from information present a rich source of theoretical questions.

Acknowledgments

For very helpful comments, we thank our discussants Adam Copeland, Maryam Farboodi, Bernard Herskovic, and John Nash and also (in random order) Semyon Malamud, Andrei Shleifer, Yu An, Rama Cont, Darrell Duffie, Elizabeth Leicht, Ben Hambly, Robin Greenwood, Chaojun Wang, Martin Oehmke, Max Bruche, David Scharfstein, Selman Erol, Ana Babus, and John Campbell. We are 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, UPF, 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). Co-Pierre Georg gratefully acknowledges support by the Algorand Foundation.

Appendix A. Technical Details and Proofs

A.1. Equilibrium in the Continuum Configuration Model

We will, throughout, refer to each agent's action as the single-valued $y_i \in \{0, 1\}$; this is because in any best response, actions $a_{i,\mu}$ are identical as μ ranges over the set of layers \mathcal{M} , so there is no loss for the purpose of studying equilibria. We also introduce a piece of notation; write $u_i(y_i, \mathbf{v})$ for the payoff that i receives when i takes action y_i and the actions of all other players are set according to the vector \mathbf{v} . We write $\mathbf{1}_A$ for the vector (i.e., function) in \mathbb{R}^N such that $y_i = 1$ if $i \in A$ and $y_i = 0$ otherwise.

We now define a "best response" operator. For any measurable set of nodes A , let

$$\mathcal{B}(A) = \{i \in N : u_i(1; \mathbf{1}_A) \geq u_i(0; \mathbf{1}_A)\}.$$

This is the set of agents who weakly prefer to be active in response to everyone in A being active.

Now define the following family of sets iteratively:

$$A(0) = N = [0, 1]$$

$$\forall t \geq 0 \quad A(t+1) = \mathcal{B}(A(t)).$$

The set $A(t)$ can be interpreted as the set of tentatively active agents, determined by a certain algorithm. The algorithm starts by initializing everyone to be active, and at each stage, $t+1$, keeping agents active who have a best response of $y_i = 1$

given that every $\ell \in A(t)$ plays $y_\ell = 1$. We can then define

$$\alpha(t) = \text{measure}(A_t).$$

Note that although the $A(t)$ are random subsets of $[0, 1]$, their measures are deterministic, because in the continuum model there is no aggregate uncertainty.

Because the game is one of strategic complements, standard results (Milgrom and Shannon 1994 and Kleene’s fixed point theorem) imply that the $A(t)$ are decreasing sets, and so they converge to a limit $A(\infty) = \bigcap_t A(t)$ with measure $\alpha(\infty) = \lim_{t \rightarrow \infty} \alpha(t)$. Moreover, $\mathcal{B}(A(\infty)) = A(\infty)$, and $\mathbf{y} = \mathbf{1}_{A(\infty)}$ corresponds to the maximal equilibrium, in the sense that all other equilibria have strictly fewer active agents (up to measure zero sets).

A.2. General Discontinuity Result: Preliminaries and Proof

We will now be more explicit about the sequence $\alpha(t)$ described in the previous subsection. That is, we will consider the fraction of nodes made active at each successive stage of the process by considering the random in-neighborhood of any node and the activity outcomes in it.

Because of a standard sampling bias, the distribution of a node’s in-degree conditional on being an in-neighbor of some other node is not the unconditional distribution of in-degrees. We will need to account for this. However, because the layers are all independent, in the continuum matching process the probability of i matching with the same node in multiple layers is zero. Thus, each in-neighbor of i is an in-neighbor in exactly one layer. Let $q_{\delta^-, \mu}$ be the probability that an in-neighbor of i in layer μ has in-degrees $\delta^- = (d_{\mu}^-)_{\mu \in \mathcal{M}}$. For an explicit form of this distribution, see Section A.2.2.

In the iterative process of the previous section, define $\alpha_{\delta^-}(t)$ to be the probability that a random node with in-degree type δ^- is in $A(t)$. Then, define $\tilde{\alpha}(t; \mu)$ by

$$\tilde{\alpha}(t; \mu) = \sum_{\delta^-} q_{\delta^-, \mu} \alpha_{\delta^-}(t),$$

the probability that a randomly drawn in-neighbor in layer μ is in $A(t)$.

Definition A.1. Given a vector $\tilde{\alpha} = (\tilde{\alpha}(\mu))_{\mu \in \mathcal{M}}$, let $\rho_{\delta^-}(\tilde{\alpha}; x)$ be the probability that an agent of in-degree type δ^- has $u_i(1; \mathbf{1}_A) \geq u_i(0; \mathbf{1}_A)$ when all in-neighbors of any node ℓ in each layer μ are in A independently with probability $\tilde{\alpha}(\mu)$. In computing this probability, we also integrate over the random cost draw c_ℓ .

At this point, it will be useful to recall several implications of Assumption 2.2 that guarantee that the function $\rho_{\delta^-}(\tilde{\alpha}; x)$ is well defined. In the continuum configuration model, the fan-in of any node³¹ is a branching process with finitely many nodes. Part (1) of the assumption guarantees that the utility of the node depends only on the activity levels in that branching process, and thus the event $u_i(1; \mathbf{1}_A) \geq u_i(0; \mathbf{1}_A)$ is measurable with respect to a countable family of random variables. By part (2) of the assumption, the probability does not depend on the node’s identity.

It is straightforward to establish inductively that at each stage $t + 1$ of the best-response process of the previous subsection, for almost all i , it holds that neighbors’ memberships in $A(t)$ are independent random variables, because the extended in-neighborhoods of two neighbors of i are almost

surely nonoverlapping. (This is important to check because Definition A.1 presumes such independence.) Thus,

$$\alpha_{\delta^-}(t) = \rho_{\delta^-}(\tilde{\alpha}(t-1); x).$$

And then, averaging this equation weighted by the degree type of a random in-neighbor, we have the iteration

$$\tilde{\alpha}(t; \mu) = \sum_{\delta^-} q_{\delta^-, \mu} \rho_{\delta^-}(\tilde{\alpha}(t-1); x). \quad (\text{A.1})$$

To recover the actual share $\alpha(t)$ of active nodes at stage t of the process, we simply average values of ρ_{δ^-} according to the population distribution of in-degree types δ^- to get

$$\alpha(t) = \sum_{\delta^-} p_{\delta^-} \alpha_{\delta^-}(t) = \sum_{\delta^-} p_{\delta^-} \rho_{\delta^-}(\tilde{\alpha}(t-1); x). \quad (\text{A.2})$$

From this, it is clear that the sequence of vectors $(\tilde{\alpha}(t))_{t=1}^\infty$ is sufficient to determine the masses of active agents at each step and the (maximum) equilibrium active share $\alpha(\infty)$. This sequence, in turn, is determined by a function appearing in the iteration (A.1). More precisely, let $\mathbf{Q}: \mathbb{R}^{\mathcal{M}} \rightarrow \mathbb{R}^{\mathcal{M}}$ be defined by

$$\mathbf{Q}_\mu(\mathbf{z}; x) = \sum_{\delta^-} q_{\delta^-, \mu} \rho_{\delta^-}(\mathbf{z}; x). \quad (\text{A.3})$$

With this definition of \mathbf{Q} and the initialization $\tilde{\alpha}(0) = \mathbf{1}$, equation (A.1) implies that we can compute the sequence $(\tilde{\alpha}(t))_{t=1}^\infty$ using

$$\tilde{\alpha}(t) = \mathbf{Q}(\tilde{\alpha}(t-1); x).$$

Note that the function \mathbf{Q} is clearly continuous, and each of its entries is weakly increasing in each entry of the input under our maintained assumptions. Moreover, $\mathbf{Q}(\mathbf{1}) \leq \mathbf{1}$. Thus, $(\tilde{\alpha}(t))_{t=1}^\infty$ is a weakly decreasing sequence of real numbers, converging to the (elementwise) largest $\mathbf{z} \in \mathbb{R}^{\mathcal{M}}$ so that $\mathbf{Q}(\mathbf{z}; x) = \mathbf{z}$.

We summarize the key finding in this lemma:

Lemma A.1. Let \mathbf{Q} be defined by (A.3). There is a vector \mathbf{z}^* satisfying

$$\mathbf{Q}(\mathbf{z}; x) = \mathbf{z}$$

and such that for any \mathbf{z}' satisfying the displayed equation, we have $\mathbf{z}' \leq \mathbf{z}^*$ entry-wise. The activity measure at the largest equilibrium of the activity game is given by

$$\alpha(\infty) = \sum_{\delta^-} p_{\delta^-} \rho_{\delta^-}(\mathbf{z}^*; x).$$

Because \mathbf{Q} is a monotone mapping under the partial order on $\mathbb{R}^{\mathcal{M}}$ given by entry-wise \leq , Tarski’s fixed-point theorem guarantees the existence of \mathbf{z}^* as claimed. The other part of the result is just an application of (A.2). \square

A.2.1. Proof of Proposition 4.1. By the assumption made in the proposition statement, for large enough x , the function \mathbf{Q} has a strictly positive fixed point.

Define

$$m(x) = \max_{\mu \in \mathcal{M}} z^*(\mu; x).$$

We have noted that $m(x) > 0$ for some x . Also, by Assumption 2.1, for sufficiently small x , we have $m(x) = 0$.

We claim that $m(x)$ has a discontinuity as we vary x , which establishes the existence of an exit spiral. We must rule out

the existence of a sequence of $x_p \rightarrow x$ such that $m(x_p)$ is arbitrarily small and positive for large p .

To this end, recall that $\rho_{\delta^-}(\bar{\alpha}; x)$ is the probability that an agent of in-degree type δ^- has $u_i(1; \mathbf{1}_A) \geq u_i(0; \mathbf{1}_A)$ when all in-neighbors of any node ℓ in layer μ are in A independently with probability $\bar{\alpha}(\mu)$.

Also recall that

$$Q_\mu(z; x) = \sum_{\delta^-} q_{\delta^-, \mu} \rho_{\delta^-}(z; x).$$

Now we will argue that, for all $x \in (0, 1)$, we have

$$\begin{aligned} \frac{\partial Q_\mu(\mathbf{0}; x)}{\partial z_{\mu'}} &= \sum_{\delta^-} q_{\delta^-, \mu} \frac{\partial \rho_{\delta^-}(\mathbf{0}; x)}{\partial z_{\mu'}} && \text{definition of } Q_\mu \\ &\leq F(\bar{b}(1); x) \sum_{\delta^-} q_{\delta^-, \mu} \delta_{\mu'}^- && \text{see below} \\ &\leq F(\bar{b}(1); x) \bar{c} && \text{definition of } \bar{c}. \end{aligned}$$

In the middle step, we reason as follows. First, by Assumption 2.2(3), the derivative is identically zero when no neighbors of i are active, so the only contribution will be from events where at least one neighbor of i is active. Second, if $\|z\|$ is small, any events where two neighbors of i are active are negligible relative to events where exactly one neighbor of i is active. Because intersections of these independent events are negligible, we may simply add up over them. In the event that some single neighbor of i is active, the probability that this makes i want to be active is bounded by the probability that $c_i \leq \bar{b}(1)$, because the right-hand side bounds the benefit to i of neighbor activity; this probability is upper bounded by $F(\bar{b}(1); x)$. Adding up across all neighbors gives the expression, because there are $\delta_{\mu'}^-$ in-neighbors of i in layer μ .

Using the hypothesis that $F(\bar{b}(1); x) \bar{c} = c < 1$, this shows that for all small enough z , we have (using the Euclidean norm)

$$\|Q(z)\| \leq c\|z\| + O(\|z\|^2),$$

which means that it is impossible for Q to have a fixed point with $m(x)$ arbitrarily small, because that fixed point would have

$$\|Q(z)\| = \|z\|.$$

A.2.2. Details on In-Neighbor Degree Distribution. Recall that $(p_{jk, \mu})_{j,k}$ is the joint distribution of in- and out-degrees, with $p_{jk, \mu}$ being the probability that in-degree is j and out-degree is k in layer μ .

Define

$$p_{j, \mu}^- := \sum_k \frac{k}{\lambda} p_{jk, \mu},$$

where λ is a constant chosen so that $\sum_{j=1}^{\infty} p_j^- = 1$. This is the probability that the initial node of a randomly chosen link in layer μ in the continuum configuration model has in-degree j (see for example Cooper and Frieze (2004) and Newman (2010)). The idea is that that a node's probability of being the initial node is proportional to its out-degree k , because of the half-edge matching process of the configuration model.

Because the layers are all independent, in the continuum matching process the probability of matching with the same

node in multiple layers is zero. Thus, each in-neighbor is an in-neighbor in exactly one layer. The probability that an in-neighbor in layer μ has in-degree type $\delta^- = (d_{\mu'}^-)_{\mu' \in \mathcal{M}}$ is then

$$q_{\delta^-, \mu} = p_{d_{\mu, \mu}^-, \mu}^- \prod_{\mu' \neq \mu} p_{d_{\mu', \mu'}^-}^-.$$

A.3. The Configuration Model: Finite Networks

In the following, we introduce random networks that are drawn uniformly at random conditional on a degree distribution, formalizing the details of the model first introduced in Section 3.2. A standard device for generating and analyzing these graphs is the *configuration model*. All of the concepts introduced below apply equally to each market $\mu \in \mathcal{M}$. 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 agents $i \in N$, whereas before, $n = |N|$ is the cardinality of the set of agents. Note that all out-edges must have a corresponding in-edge; therefore $\sum_i d_{i,n}^+ = \sum_i d_{i,n}^-$. 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 A.2 (Bollobás Configuration Model; see, e.g., Amini et al. 2016). 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 $\tilde{G}(n, \mathbf{d}_n^+, \mathbf{d}_n^-)$ 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, we 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. (2016) 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 as follows

Assumption A.1. For each n , \mathbf{d}_n^+ and \mathbf{d}_n^- are sequences of non-negative integers such that $\sum_i d_{i,n}^+ = \sum_i d_{i,n}^-$ 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 that 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, whereas the maximum degree may go to infinity, the degree sequence does not become too dispersed:

Assumption A.2 (Proper Degree Sequences; Cooper and Frieze 2004). Let Δ_n denote the maximum degree. Then:

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

We call an infinite tuple that satisfies Assumptions A.1 and A.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 3.1) implies that $d_{iC,n}^+, d_{iR,n}^+, d_{iC,n}^-, d_{iR,n}^-$, $d_{iC,n}^+, d_{iR,n}^+$, 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.

A.4. Results on Coupled Networks with Simple Complementarities

A.4.1. General Facts

Lemma A.2. Suppose that Assumption 4.3 is satisfied. Then, in a maximal equilibrium, $y_i^* = 1$ for all i ; equivalently, the maximal stable set consists of all nodes.

Fix a \mathcal{G} , and define $\hat{\mathcal{G}}$ to be the same graph with all edges reversed. Note that 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 (because 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 $\hat{\mathcal{G}}$ 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 $\hat{\mathcal{G}}$. 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 are stable. This proves the result.

A.4.2. Equilibrium and the Mutual Giant Component.

Building on Section 4.3.2, we can now establish a connection between the equilibrium activity defined in Section 3.1 and certain asymptotic properties of the random graphs defined above. The size of the mutual giant out-component and the equilibrium activity measure are related as follows.

Lemma A.3. Let \mathbf{y}^* be an equilibrium for $\mathcal{G}_R, \mathcal{G}_C$ and with W being the set of shocked nodes. Then,

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

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 agents in equilibrium. It is only a lower bound because 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 that even the weak components in either of the two markets that are not part of the giant out-component have a negligible size as $n \rightarrow \infty$. Therefore, in the limit of large networks, the size of the mutual giant out-component is sufficient to compute the equilibrium activity. In the following, we will discuss how the mutual giant out-component can be found.

A.4.3. A Branching Process Approximation of Equilibrium Activity.

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.

A.4.3.1. 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.³²

Suppose one starts to explore the network from a randomly chosen link via a breadth-first search algorithm. How many agents 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 $(p_k^+)_{k=0}^\infty$.³³ 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 A.4. The probability f that the branching process defined by $(p_k^+)_{k=0}^\infty$ goes extinct is the smallest solution to $f = H(f)$.

Then, the size of the giant out-component is given by a simple corollary of Lemma 4 (see, for example, Newman 2010) that we summarize in the following lemma.

Lemma A.5. 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 .

A.4.3.2. Equilibrium Activity 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 A.4.3 with the markets for asset C and asset R 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 market for asset C.

Consider the following coupled branching process, first, with $x = 1$, that is, no shock. Choose a link at random in \mathcal{G}_R , and follow the node it goes into. Because \mathcal{G}_C and \mathcal{G}_R are independent by assumption, the agent 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:

$$\hat{p}_{jk,R}(s_C) := \underbrace{\sum_{l=j}^\infty \sum_{m=k}^\infty p_{lm,R}}_A \underbrace{\binom{l}{j} (1-s_C)^{l-j} s_C^j}_{B} \underbrace{\binom{m}{k} (1-s_C)^{m-k} s_C^k}_C. \tag{A.4}$$

The distribution $\hat{p}_{jk,R}(s_C)$ is the joint distribution of in- and out-degrees of the network for asset R 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 is present after thinning. Finally, C is the probability that k out of initially m out-links is present

after thinning. An equivalent argument can be made for the market for asset C. 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 because of the cost shock. We call x the size of the shock. The final size of the mutual giant out-component (and thus activity 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 agents. Let $\hat{A}(x)$ be the expected activity of the maximal equilibrium.

Lemma A.6. 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 R (C) network s_R^* (s_C^*) is the greatest solution to

$$\begin{aligned} s_R &= x g_R(f_R, s_C) = x \left(1 - \sum_{jk} \hat{p}_{jk,R}(s_C) f_R^k \right), \\ f_R &= H_R(f_R, s_C) = \sum_k \hat{p}_{k,R}^+(s_C) f_R^k, \\ s_C &= x g_C(f_C, s_R) = x \left(1 - \sum_{jk} \hat{p}_{jk,C}(s_R) f_C^k \right), \\ f_C &= H_C(f_C, s_R) = \sum_k \hat{p}_{k,C}^+(s_R) f_C^k. \end{aligned} \tag{A.5}$$

Activity is then

$$\hat{A}(x) := s^* = x g_R(x g_C(s^*)).$$

To see this, note that the expression $x g_R(x g_C(s^*))$ is monotonically increasing in s (see Lemma A.6) 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 activity, from above by x . This is simply because a fraction of $1 - x$ of agents withdraw from the markets because of the shock realization.

Suppose now that one of the markets is replaced 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 A.7. Let \mathcal{G}_R be a random network. Let $\bar{\mathcal{G}}_C$ be a complete network. Given a shock of size x , the size of the giant out-component in the R network s_R^* is the greatest solution to

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

Activity is then

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

Thus, if the C 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 R network taken on its own (we have discussed the study of its size above). To see this,

first note that if \bar{G}_C is complete, there will be no contagion through \bar{G}_C . All agents in \bar{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 A.6. The greatest fixed point exists by the same argument as in the proof of Lemma A.6.

A.5. Proofs of Random Network Results.

In the following, we will prove Propositions 4.2 and 4.3. Our main contribution is to provide conditions on the degree distributions of the random networks for which Propositions 4.2 and 4.3 hold. We will first provide a sketch of the proof of Proposition 4.3(a) because this is a standard result from the literature (see Cooper and Frieze 2004) and is useful for the subsequent proofs of 4.2 and 4.32.

For the proof of Proposition 4.3, we will use standard properties of a generic probability generating function (pgf) that we summarize in the following remark.

Remark A.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 $f'(1)$; that is, because $f'(1) \rightarrow 1$, we have that $s^* \rightarrow 1$.

We illustrate some graphical intuition for this proof in Figure 8 in Online Appendix D.

Proof of Proposition 4.3. Recall that for $x=1$ we have $H(z) = \sum_k p_k^+ z^k$ with $p_k^+ = \sum_{jk} p_{jk}/\lambda$. It can be shown (see, for example, Cooper and Frieze 2004 or Newman 2010) 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 A.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; that is, $g(1) = 0$. If $f < 1$, the size of the giant out-component is $g(f) > 0$; that is, 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

$$x_c = \frac{1}{H'(1)} = \frac{\lambda}{\sum_{jk} p_{jk} k}. \quad (\text{A.6})$$

Because f is continuous as the derivative $\hat{H}'(1)$ changes, it is also continuous in x , which determines $\hat{H}'(1)$. Note that Assumption 4.5 ensures that in the absence of a shock there exists a giant out-component of positive size. This concludes the proof. \square

Lemma A.8. 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 of Lemma A.8. The fact that $f(x) = H(f(x), x)$ is continuous follows from the proof of Proposition 4.3. To show that $f(x)$ is monotonically decreasing, we use the result from the proof of Proposition 4.3 that $x \in [0, x_c]: f = 1 \Rightarrow df/dx = 0$.

Now consider what happens when $x \in (x_c, 1]$ and $f < 1$. In this case, we derive for df/dx :

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

$$\frac{df}{dx} = \frac{dH}{df}(f, x) \frac{1}{x} \left(f + x \frac{df}{dx} - 1 \right),$$

$$\frac{df}{dx} = \frac{dH}{df}(f, x) \frac{1}{x} \left(1 - \frac{dH}{df}(f, x) \right)^{-1} (f - 1).$$

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, that is, for $x \in (x_c, 1]$ $\frac{dH}{df}(f, x) < 1$, where $H(f, x) = f < 1$; see Figure 8 in Online 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 because $\frac{dH}{df}(1, x) > 1$ 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 A.9. $g(f, x)$ is continuous and monotonically increasing in x for $x \in [0, 1]$.

Proof of Lemma A.9. The fact Lemma A.9 the fact that $g(f, x)$ is continuous follows directly from the proof of Proposition 4.3. $g(f, x)$ is monotonically increasing because

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

Let $F(s, x) := xg_R(xg_C(s))$. In order to prove Proposition 4.2, we first need to establish a couple of facts about $F(s, x)$, which we summarize in the Lemma A.10. We will use the index $\mu \in \{R, C\}$ whenever results apply to both R and C networks.

Lemma A.10. 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 of Lemma 10. For this proof, we invoke results from Lemmas A.8 and A.9. For $s \in (0, 1]$:

1. $F(s, x)$ is continuous; $g_\mu(s)$ is continuous as shown in Lemma A.9. $F(s, x)$ 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 A.9. $F(s, x)$ 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 because $g_\mu(s) \leq 1$. Furthermore, assuming a positive shock size, that is, $x < 1$, 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 $s = 1$ that scales with x , that is, 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 :

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

Because $\frac{dg_\mu}{ds}(s) > 0$, $\frac{d^2 g_\mu}{ds^2}(s) < 0$ (by Assumption 4.6) and $x > 0$, we must have $\frac{\partial^2 F}{\partial s^2} < 0$, that is, $F(s, x)$ concave.

5. $\lim_{s \rightarrow 0} F(s, x) \rightarrow 0$: Because 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 4.3. 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 4.3).

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

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

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, because 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 ($\partial F/\partial s = 0$); and (ii) for larger values of s , ($s > s_{c,\mu}$) $F(x, s)$ is strictly monotonically increasing and concave but bounded from above ($F(x, s) < 1$). □

Proof of Proposition 4.2. This proof invokes results from Lemma A.10 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 Figure 9 in Online Appendix D.

First, note that $s = 0$ is a trivial solution to $s = F(s, x)$ for all x because $g_\mu(0) = 0$. Furthermore, as shown in Lemma A.10, there exists a region for sufficiently small s in which $F(s, x)$ is constant and equal to zero. As seen in Lemma A.10, 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 (because 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 4.5 ensures that in the absence of a shock there exists a mutual giant out-component of positive size. Because we are investigating cascades following a small shock, we are interested only 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 because of 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 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 $s > s_{c,\mu}$ downward. 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 nontrivial 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 a 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 nontrivial solution $0 < s^* < 1$ exists.

Note that, because 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 nontrivial solution must always be greater than zero, that is, $s^* > 0$ 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 Proposition 4.3(b).

Proof of Proposition 4(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, that is, $g_R(r_c) = 0$. Also, $F(s, x_c) = x_c g_R(x_c g_C(s)) < r_c$ because $F(s, x_c) < x_c$ 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 4.2. This proves Proposition 4.3(b) by contradiction.

Endnotes

¹ The market for asset-backed securities had about \$652 billion outstanding at the end of the global financial crisis, whereas the municipal bond market had about \$4 trillion outstanding at the end of 2012. Copeland et al. (2014) estimated the sum of all repo outstanding on a typical day in July and August 2008 to be \$6.1 trillion.

² Similarly, Krishnamurthy et al. (2014) showed that markets for asset-backed commercial paper (ABCP) experienced a significant contraction.

³ For instance, Brunnermeier and Pedersen (2009) argued—in the context of a centralized market—that the market for repo (short-term loans secured by collateral such as bonds) and the market for the underlying collateral are coupled. The complementarity works as follows; when the bond market is less liquid, traders are less willing to take a bond as collateral (because they will have a hard time selling it in case of counterparty default), 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.

⁴ In particular, certain counterparty relationships may be effectively shut down because of bilateral exposure limits or other events; see,

for example, Perignon et al. (2018) and Di Maggio et al. (2023). Online Appendix C offers evidence from the 2008 financial crisis showing that active overnight lending links between large banks changed substantially during the crisis. In a large cross-country comparison of 13 jurisdictions, Anand et al. (2018) discussed how analysts make inferences about granular network data based on the partial data available to them.

⁵ For a recent survey of the growing literature, see Glassermann and Young (2016).

⁶ This stands in contrast to a physics literature, which typically considers more mechanical processes (Dodds and Watts 2004) and often examines them through approximate mean-field analyses.

⁷ That is, for all networks that we consider, $(i, i) \notin E(\mathcal{G})$.

⁸ Trading links are often reciprocal, but we allow for the possibility that they track the flow of one asset so that, for example, i always sells an asset to j but never buys it.

⁹ This assumption is very convenient technically, as we will see in the following few paragraphs.

¹⁰ When an agent deviates to a strategy with different actions across all layers, the agent can get at least as high a payoff from a deviation that involves only one action across all layers.

¹¹ That is, paths like $\ell \xrightarrow{\mu_1} \ell_1 \xrightarrow{\mu_2} \ell_2 \dots \xrightarrow{\mu_p} i$.

¹² Strategies in such equilibria correspond to maps $\sigma_i: \mathbb{R}^n \rightarrow \{0, 1\}$ describing each agent's entry decision for given cost draws; recall that the game is one of complete information, so these are known.

¹³ These customers are not part of the network but should be thought of as being attached to the nodes of the network.

¹⁴ For example, interbank loans (Allen and Gale 2000), certificates of deposit (Perignon et al. 2018), or derivatives. Donaldson et al. (2022) discussed netting in interbank networks. D'Errico and Roukny (2021) discussed when derivatives can be compressed. For our illustrative model, we abstract away from such issues and assume that obligations can be neither netted nor compressed.

¹⁵ Restrictions might arise through regulations such as the separation of investment and retail banking under the Glass-Steagall Act or limits to transfer funds between subsidiaries of a bank holding company. The benefit function could be modified to accommodate such restrictions.

¹⁶ It would be possible to include these constraints in the payoff function, but for clarity we prefer to include them as explicit constraints in the agents' decision problem.

¹⁷ See, for example, the discussion in Anand et al. (2018).

¹⁸ The Bank for International Settlements provides the BIS locational database, which provides country-level bilateral information, albeit for a subset of countries only.

¹⁹ Recall that a larger value of x corresponds to stochastically lower costs.

²⁰ Note that in general, nodes that are not immediate neighbors of i may matter for i 's payoffs.

²¹ Actually, this slope condition is weaker than the condition we impose in our result, and stronger results can be obtained by studying ρ (and its generalization, which appears in the proof) near 0.

²² Indeed, the proof shows that, were incentives from activity to come purely from the value of single links, this would not be enough to support activity under the condition $F(\bar{b}(1); x) \bar{c} < 1$.

²³ This analysis requires certain detailed characterizations from frontier results in random graph theory. Thus, we do not expect that an analogous finite- n analysis is available for the full generality of the results in the previous section, although this presents an important set of questions for further work.

²⁴ In the case of repo and collateral that we have discussed, this assumption corresponds to a strong institutional coupling whereby one cannot

profitably intermediate in the repo market without access to the collateral market, and one cannot obtain liquidity for collateral purchases without repo market access. Empirically, it would correspond to the property that conditional on entering, agents earn positive gross profits if and only if they have access to trading opportunities in both markets.

²⁵ Note that this implies that an edge exists in the other direction as well.

²⁶ Every nonempty proper subset has an edge to its complement or from it.

²⁷ If technical assumptions in the Appendix apply to the original graph, they also apply to the graph obtained from this process, called the "percolated" graph, so the results we use from Cooper and Frieze (2004) about giant out-components still hold.

²⁸ Although the result is stated in the $n \rightarrow \infty$ limit, numerical experiments show that exit spirals can make large networks susceptible to full collapse after the removal of a single node.

²⁹ This could be by directly increasing x or by making the network denser.

³⁰ The robustness of coupled core-periphery networks is studied in a previous version of this article, with analogs of the main findings presented here.

³¹ Defined to be all nodes that have a path to i using links in any layer.

³² The 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.

³³ Of course, this only corresponds to the number of agents 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", that is, 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) showed rigorously that this approximation is indeed valid.

³⁴ To produce this plot, we used the data from Gabrieli and Georg (2014), who studied 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 because banks are required to hold minimum average reserves over this period commensurate with their deposits. The overnight interbank market is the main mechanism of how these reserves are redistributed among banks.

³⁵ Consider 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.

³⁶ Note that, for confidentiality reasons, the CDF depicted must be censored at degree 50, but the figure is qualitatively unchanged without the censoring.

³⁷ Because the Target2 payment system became operational in the last country of the euro area only in May 2008, unfortunately, we do not have data for any date before then.

References

Acemoglu D, Ozdaglar A, Siderius J, Tahbaz-Salehi A (2020) Systemic credit freezes in financial lending networks. *Math. Finance Econom.* 15:185–232.

- Acemoglu D, Ozdaglar A, Tahbaz-Salehi A (2015) Systemic risk and stability in financial networks. *Amer. Econom. Rev.* 105(2):564–608.
- Acharya VV, Gale D, Yorulmazer T (2011) Rollover risk and market freezes. *J. Finance* 66(4):1177–1209.
- Adlakha S, Johari R (2013) Mean field equilibrium in dynamic games with strategic complementarities. *Oper. Res.* 61(4):971–989.
- Afonso G, Kovner A, Schoar A (2013) Trading partners in the interbank lending market. Unpublished manuscript, Federal Reserve Bank of New York.
- Allen F, Gale D (2000) Financial contagion. *J. Political Econ.* 108:1–33.
- Amini H, Cont R, Minca A (2016) Resilience to contagion in financial networks. *Math. Finance* 26:329–365.
- Anand K, van Lelyveld I, Ádám Banai, Friedrich S, Garratt R, Hałaj G, Figue J, Hansen I, Jaramillo SM, Lee H, Molina-Borboa JL, Nobili S, Rajan S, Salakhova D, Silva TC, Silvestri L, de Souza SRS (2018) The missing links: A global study on uncovering financial network structures from partial data. *J. Finance Stab.* 35:107–119.
- Anderson R, Baba C, Danielsson J, Das US, Kang H, Segoviano M (2018) Macroprudential stress tests and policies: Searching for robust and implementable frameworks. International Monetary Fund Working Paper, Washington, DC. <https://www.imf.org/en/Publications/WP/Issues/2018/09/11/Macroprudential-Stress-Tests-and-Policies-Searching-for-Robust-and-Implementable-Frameworks-46218>.
- Athreya KB, Jagers P (2012) *The IMA Volumes in Mathematics and Its Applications*, vol. 84 (Springer Science & Business Media, Berlin).
- Barabási AL, Albert R (1999) Emergence of scaling in random networks. *Science* 286(5439):509–512.
- Blume LE (1993) The statistical mechanics of strategic interaction. *Games Econom. Behav.* 5(3):387–424.
- Britton T, Janson S, Martin-Löf A, et al (2007) Graphs with specified degree distributions, simple epidemics, and local vaccination strategies. *Adv. in Appl. Probab.* 39(4):922–948.
- Brunnermeier MK, Pedersen LH (2009) Market liquidity and funding liquidity. *Rev. Finance Stud.* 22(6):2201–2238.
- Buldyrev SV, Parshani R, Pau G, Stanley HE, Havlin S (2010) Catastrophic cascade of failures in interdependent networks. *Nature* 464:1025–1028.
- Burkholz R, Leduc MV, Garas A, Schweitzer F (2016) Systemic risk in multiplex networks with asymmetric coupling and threshold feedback. *Phys. D* 323–324:64–72.
- Constâncio V (2017) Macroprudential stress-tests and tools for the non-bank sector. Prepared remarks at the ESRB Annual Conference, Frankfurt am Main, September 22, 2017, European Central Bank.
- Cooper C, Frieze A (2004) The size of the largest strongly connected component of a random digraph with a given degree sequence. *Combin. Probab. Comput.* 13(03):319–337.
- Copeland A, Martin A, Walker M (2014) Repo runs: Evidence from the tri-party repo market. *J. Finance* 69(6):2343–2380.
- Craig B, Von Peter G (2014) Interbank tiering and money center banks. *J. Finance Intermed.* 23(3):322–347.
- D’Errico M, Roukny T (2021) Compressing over-the-counter markets. *Oper. Res.* 69:1660–1679.
- Di Maggio M, Kermani A, Song Z (2023) The value of trading relationships in turbulent times. *J. Financial Econ.* Forthcoming.
- Dodds PS, Watts DJ (2004) Universal behavior in a generalized model of contagion. *Phys. Rev. Lett.* 92(21):218701.
- Donaldson JR, Piacentino G, Yu X (2022) Debt maturity in financial networks. Unpublished, Washington University in St. Louis, St. Louis, MO.
- Duffie D, Sun Y (2007) Existence of independent random matching. *Ann. Appl. Probab.* 17(1):386–419.
- Eisenberg L, Noe TH (2001) Systemic risk in financial systems. *Management Sci.* 47(2):236–249.
- Elliott M, Georg CP, Hazell J (2021) Systemic risk shifting in financial networks. *J. Econom. Theory* 191:105157.
- Elliott M, Golub B, Leduc MV (2022) Supply network formation and fragility. *Amer. Econom. Rev.* 112(8):2701–2747.
- Elliott ML, Golub B, Jackson MO (2014) Financial networks and contagion. *Amer. Econom. Rev.* 104(10):3115–3153.
- Ellison G (1993) Learning, local interaction, and coordination. *Econometrica* 61(5):1047–1071.
- Erol S, Vohra R (2020) Network formation and systemic risk. Preprint submitted June 1, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2546310.
- Farboodi M (2021) Intermediation and voluntary exposure to counterparty risk. NBER Working Paper No. 29467, NBER, Cambridge, MA.
- Frei C, Capponi A, Brunetti C (2022) Counterparty risk in over-the-counter markets. *J. Finance Quant. Anal.* 57:1058–1082.
- Gabrieli S, Georg CP (2014) A network view on interbank liquidity. Discussion Paper No. 44/2014, Deutsche Bundesbank.
- Galeotti A, Goyal S, Jackson MO, Vega-redondo F, Yariv L (2009) Network games. *Rev. Econ. Stud.* 77(1):218–244.
- Glasserman P, Young HP (2015) How likely is contagion in financial networks? *J. Banking Finance* 50:383–399.
- Glassermann P, Young HP (2016) Contagion in financial networks. *J. Econom. Lit.* 54(3):779–831.
- Golub B, Morris SE (2017) Expectations, networks, and conventions. Preprint submitted June 3, <http://dx.doi.org/10.2139/ssrn.2979086>.
- Gorton G, Metrick A (2012) Securitized banking and the run on repo. *J. Financial Econ.* 104(3):425–451.
- Goyal S, Vega-Redondo F (2007) Structural holes in social networks. *J. Econom. Theory* 137:460–492.
- Hendershott T, Menkveld AJ (2014) Price pressures. *J. Financial Econ.* 114(3):405–423.
- Hollifield B, Neklyudov A, Spatt C (2017) Bid-ask spreads, trading networks, and the pricing of securitizations. *Rev. Financial Stud.* 30:3048–3085.
- Hüser AC, Lepore C, Veraart L (2021) How do secured funding markets behave under stress? evidence from the gilt repo market. Staff Working Paper No. 910, Bank of England.
- Jackson MO, Yariv L (2007) Diffusion of behavior and equilibrium properties in network games. *Amer. Econom. Rev.* 97(2):92–98.
- Krishnamurthy A, Nagel S, Orlov D (2014) Sizing up repo. *J. Finance* 69(6):2381–2417.
- Li D, Schürhoff N (2019) Dealer networks. *J. Finance* 74:91–144.
- Milgrom P, Roberts J (1990) Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica* 58(6):1255–1277.
- Milgrom P, Shannon C (1994) Monotone comparative statics. *Econometrica* 62(1):157–180.
- 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.
- Morris S (2000) Contagion. *Rev. Econom. Stud.* 67(1):57–78.
- Morris S, Yildiz M (2019) Crises: Equilibrium shifts and large shocks. *Amer. Econom. Rev.* 109(8):2823–2854.
- Newman M (2010) *Networks: an introduction*. (Oxford University Press, Oxford, UK).
- Newman ME (2002) Spread of epidemic disease on networks. *Phys. Rev. E.* 66(1):016128.
- Perignon C, Thesmar D, Vuillemeys G (2018) Wholesale funding dry-ups. *J. Finance* 73:575–617.
- Son SW, Bizhani G, Christensen C, Grassberger P, Paczuski M (2012) Percolation theory on interdependent networks based on epidemic spreading. *EPL* 97(1):16006.
- Zawadowski A (2013) Entangled financial systems. *Rev. Financial Stud.* 26(1):1291–1323.

Christoph Aymanns's research is on developing computational and mathematical models of financial stability, with a focus on dynamic feedback loops and contagion on networks. He obtained his PhD from the Mathematical Institute at Oxford University. He was a postdoc at the London School of Economics and an assistant professor in the School of Finance at the University of St. Gallen. He currently works as a data scientist at QuantCo.

Co-Pierre Georg's research is located at the intersection of financial innovation and financial stability. He studies the role of networks in financial contagion and economic innovation processes. He received his PhD from the Friedrich Schiller University Jena in 2011 and was a Postdoctoral Researcher at the Universidad Carlos III de Madrid before joining the University of Cape Town as a

lecturer and associate professor within the School of Economics. In 2022, he joined the finance group at EDHEC Business School as an associate professor.

Benjamin Golub's research focuses on social and economic networks, particularly as they relate to social learning, peer effects, and financial and economic contagion. His work has examined the spread of rumors, the fragility of markets, and the optimal design of interventions in networks. Born in Ukraine, he received his PhD from the Stanford Graduate School of Business in 2012 and held appointments as a Junior Fellow at the Harvard Society of Fellows and as an assistant and associate professor at Harvard's Department of Economics before moving in 2021 to Northwestern, where he is currently a professor of economics and of computer science (by courtesy).

Copyright 2023, by INFORMS, all rights reserved. Copyright of Operations Research is the property of INFORMS: Institute for Operations Research and its content may not be copied or emailed to multiple sites or posted to a listserv without the copyright holder's express written permission. However, users may print, download, or email articles for individual use.