Journal of Banking and Finance 146 (2023) 106684

Contents lists available at ScienceDirect

Journal of Banking and Finance

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

Scenario-free analysis of financial stability with interacting contagion channels

Garbrand Wiersema[a,b,*], Alissa M. Kleinnijenhuis[a,b,d], Thom Wetzer[a,c,e], J. Doyne Farmer[a,b,c,f]

[a] Institute for New Economic Thinking at the Oxford Martin School, University of Oxford, Oxford OX2 6ED, UK
[b] Mathematical Institute, University of Oxford, Oxford OX1 3LP, UK
[c] Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, OX2, 6ED, UK
[d] Stanford Institute for Economic Policy Research, Stanford University, Stanford, CA 94305, US
[e] Faculty of Law, University of Oxford, Oxford, OX1 3UL, UK
[f] Santa Fe Institute, Santa Fe, NM 87501, USA

ARTICLE INFO

Article history:
Received 13 September 2020
Accepted 11 September 2022
Available online 16 September 2022

JEL classification:
C63
G01
G17
G18
G21
G23

Keywords:
Financial Stability
Systemic Risk
Interacting Contagion Channels
Financial Contagion
Multiplex Networks
Stress Test
Solvency-Liquidity Nexus

ABSTRACT

Financial stress tests that capture multiple interactions between contagion channels are conditional on specific, subjectively-imposed stress scenarios. Eigenvalue-based approaches, in contrast, provide a scenario-independent measure of systemic stability, but so far only handle a single contagion mechanism. We develop an eigenvalue-based approach that brings the best of both worlds, enabling the analysis of multiple interacting contagion channels without the need to impose a subjective stress scenario. Our model captures the solvency-liquidity nexus, which allows us to demonstrate that the instability due to interacting channels can far exceed that of the sum of the individual channels acting in isolation. The framework we develop is flexible and allows for calibration to the microstructure and contagion channels of real financial systems. Building on this framework, we derive an analytic stability criterion in the limit of a large number of institutions that gives the instability threshold as a function of the relative size and intensity of contagion channels. This analytical formula requires comparatively little data to elucidate the mechanisms that drive instability in real financial systems and thus complements the insights gained from traditional stress tests.

1. Introduction

One of the revelations of the financial crisis was the importance of *systemic risk*.[1] Systemic risk is transmitted between institution to spread and amplify across the financial system through mechanisms referred to as *contagion channels* (Allen and Gale, 2000). Risk control measures that are prudent for a single institution acting on its own may be counterproductive when many institutions act in unison.[2] This problem is complicated by the fact that the financial system is heterogeneous, with different types of actors, such as banks, pension funds, hedge funds, money market funds and insurance companies, and many types of interactions.[3] This makes the financial system a prime example of a complex system and raises the key challenge for policymakers to capture the complex microstructure of the system in models of financial stability (Arinaminpathy et al., 2012; Aymanns et al., 2018).

The microstructural models currently used by policymakers to evaluate the stability of financial systems are premised on stress scenarios consisting of hypothetical exogenous shocks that could potentially threaten the stability of the system.[4] This approach has the obvious drawback that the specification of such scenar-

* Corresponding author.
E-mail address: garbrandwiersema@outlook.com (G. Wiersema).

[1] See e.g. Cont et al. (2010); Gai and Kapadia (2010); Fouque and Langsam (2013); Glasserman and Young (2016).

[2] See e.g. Adrian and Shin (2010); Thurner et al. (2012); Adrian and Shin (2014); Aymanns and Farmer (2015); Aymanns et al. (2016).

[3] See e.g. Allen and Babus (2009); Gai et al. (2011); Arinaminpathy et al. (2012); Caccioli et al. (2012); Cont et al. (2013); Cont and Schaanning (2017); Aymanns et al. (2018); Farmer et al. (2020).

[4] See e.g. Burrows et al. (2012); IMF (2014); Kok and Montagna (2016); Budnik et al. (2019).

https://doi.org/10.1016/j.jbankfin.2022.106684
0378-4266/

ios is inherently subjective, giving rise to debates about their realism and relevance to current market conditions (Borio et al., 2014; Aymanns et al., 2018). Moreover, scenarios are by their very nature not comprehensive – the financial system might be stable in one set of scenarios and collapse in other unforeseen or misspecified scenarios. This challenge can be partially overcome by analyzing *ensembles* of shock scenarios (see e.g. Elsinger et al. (2006); Montagna et al. (2021)), but because the space of potential scenarios is effectively infinite and the probability of specific scenarios is unknown, the problem is never entirely overcome.

An alternative method that does not suffer from these shortcomings explicitly models the financial network as a dynamical system, so that its stability can be analyzed in terms of its eigenvalues. This approach has been used for studying the effect of contagion channels in isolation (Caccioli et al., 2014; Bardoscia et al., 2017; Cont and Schaanning, 2019). However, financial systems have multiple interacting contagion channels, and studies have shown that the interaction of multiple channels can dramatically amplify instability compared to channels operating in isolation.[5] Up until now there has been no general method for treating interacting channels as a dynamical system. This causes a tenuous state of affairs for policymakers: To take into account interacting contagion channels, they are forced to rely on subjectively imposed stress scenarios.

Our key contribution in this paper is to offer a novel approach that combines the best of both worlds: a systematic method to analyze a financial network with multiple interacting contagion channels as a dynamical system, which significantly complements our ability to understand and monitor the stability of the financial system. This novel approach is realized by expressing contagion channels in terms of shocks to the liquidity and solvency of institutions, which allows us to reduce the multiple layers of the financial system's contagion network to a simple two-layer system and study the nexus of the interactions between liquidity and solvency. Using this method, we compute the linear stability of a financial system exposed to small shocks in a general setting. This makes it possible to estimate the stability of the financial system without having to impose a specific, subjective risk scenario. In contrast to methods such as those used by the EBA and the FED (EBA, 2018; FED, 2018), for example, this has the potential to yield accurate estimates of financial stability that are robust under a wide range of stress scenarios.

The paper proceeds as follows. Section 2 introduces our framework, which we refer to as the "shock transmission matrix", explains our modelling of contagion dynamics, and derives our main results. Section 3 applies our framework to randomly generated financial systems. The purpose of this exercise is to elucidate the interactions between solvency- and liquidity-mediated contagion channels and to demonstrate that neglecting these interactions may lead to a potentially substantial overestimation of stability. Section 4 concludes with a discussion of the implications of our findings for financial stability policy and financial stress testing practices.

2. Capturing the solvency-liquidity nexus

Financial contagion can take many forms, many of which have been extensively studied (see e.g. Allen and Gale (2000); Eisenberg and Noe (2001); Gorton and Metrick (2012)). In this paper, we analyze four principal contagion channels of the financial system, which we call funding contagion, overlapping portfolio contagion, counterparty risk contagion, and deleveraging contagion.

Funding contagion occurs when a borrowing institution depends on short-term loans to provide liquidity and runs the risk that the lender might withdraw its loans (Diamond and Dybvig, 1983; Acharya and Skeie, 2011; Caccioli et al., 2013). *Overlapping portfolio contagion* can materialize when two institutions hold common securities. If either institution sells securities this drives prices down, lowering the securities' value.[6] *Counterparty risk* occurs when a lender runs the risk that a borrower might default.[7] Finally, *deleveraging contagion* takes place when an institution uses borrowed funds to purchase assets.[8] Borrowing creates debt and the ratio of debt to equity is called the *leverage* λ. As part of good risk-management practices, it is common for financial institutions to target a particular leverage to control risk. If the value of assets drops, the debt burden remains constant but the equity value decreases, so leverage increases. This forces a leverage-targeting institution to pay off debt to maintain its leverage target, an action that drains the institution's liquidity.

The culmination of a severe financial crisis is usually the default of one or more institutions (Brunnermeier, 2009; Roukny et al., 2013). A default can be forced by insolvency or illiquidity. *Insolvency* occurs when asset values drop to the point where equity becomes negative – that is, when the value of an institution's liabilities exceeds that of its assets (Amini et al., 2016). Default due to *illiquidity*, on the other hand, occurs when an institution is unable to meet its payment obligations (Cont and Schaanning, 2017). Insolvency and liquidity can be related, but are analytically distinct: an institution can default due to a liquidity shock even when it is solvent, and vice versa. During financial crises, liquidity tends to be the more direct threat; an institution may survive temporary insolvency by maintaining liquidity and regaining its solvency at a later date, but for our purposes we neglect this possibility here. In normal economic times, a solvent institution is expected to borrow to avert a liquidity shortage. In times of economic crisis, however, this may not be possible because lending markets malfunction due to uncertainty about asset values, escalating collateral requirements, liquidity hoarding and capital flight, etc. (Gorton and Metrick, 2012).

We can analyze the stability of the financial system in terms of its resilience to shocks, which we can classify either as liquidity shocks or valuation shocks, depending on the type of default they threaten to cause. For the purposes of this paper, we define a liquidity shock as an unexpected outflux of liquid assets and a valuation shock as a drop in the (expected) value of an institution's assets.[9]

The key insight is that the four contagion channels we distinguish here can be described in terms of the propagation of liquidity and valuation shocks and the conversion of one type of shock into the other:

- *Propagation of liquidity shocks by funding contagion:* If institution i depends on a short-term loan from institution j, if j suddenly withdraws the loan to meet a liquidity shock it receives, then this causes a liquidity shock to i.

[5] See e.g. Caccioli et al. (2013); Poledna et al. (2015); Kok and Montagna (2016); Detering et al. (2021).

[6] See e.g. Adrian and Shin (2010); Caccioli et al. (2013, 2014, 2015); Duarte and Eisenbach (2018); Cont and Schaanning (2017, 2019).

[7] See e.g. Eisenberg and Noe (2001); Furfine (2003); Gai and Kapadia (2010); Battiston et al. (2012); Elliott et al. (2014); Acemoglu et al. (2015); Bardoscia et al. (2015, 2017).

[8] See e.g. Fostel and Geanakoplos (2008); Brunnermeier and Pedersen (2009); Adrian and Shin (2010); Geanakoplos (2010); Adrian and Shin (2014); Aymanns et al. (2016).

[9] We consider *expected* inflows and outflows of liquid assets as part of regular day-to-day liquidity management, and therefore do not classify such flows as a liquidity shock. For simplicity, we assume that shocks are non-negative. In principle, the framework could also capture negative shocks (i.e. liquidity and asset value *gains*), but this would cause the framework to lose some of the convenient properties guaranteed by the Perron Frobenius theorem.

- *Propagation of valuation shocks by counterparty risk contagion:* If a valuation shock causes institution i's probability of default to rise, the risk-adjusted value of its debt to institution j falls, causing a valuation shock to j.
- *Conversion of liquidity shocks to valuation shocks by overlapping portfolio contagion:* If institution i suffers a liquidity shock it may be forced to sell securities to raise liquidity. This depresses their price. If institution j also has a position in these securities it experiences a valuation shock.
- *Conversion of valuation shocks to liquidity shocks by deleveraging:* If a valuation shock decreases institution i's equity, its leverage rises. To return to its target leverage, the institution must raise cash to pay off its debt, essentially triggering a liquidity shock to itself (we do not consider slower mechanisms to raise equity-capital, such as issuing new shares or retaining earnings).

Note that we use the term "contagion channel" to refer to a specific mechanism that propagates or converts a financial shock, and not as a reference to the shock itself.

In the remainder of this section, we show how to describe the collective dynamics of these four interacting contagion channels in a scenario-independent framework. This allows us to characterize the financial system's resilience to a wide range of liquidity and valuation shocks based on the corresponding largest eigenvalue.

2.1. The shock transmission matrix

The interactions of the four contagion channels can be captured in a single matrix A which we call the *shock transmission matrix*, as shown in Figure 1a. Assume discrete dynamics and let $x^l_{t,i}$ denote the liquidity shock suffered by institution i at time t. The N-dimensional vector $\boldsymbol{x}^l_t$ gives the liquidity shocks to all institutions, where N is the number of financial institutions. Similarly, $x^v_{t,i}$ denotes the valuation shock to institution i at time t and $\boldsymbol{x}^v_t$ the N-dimensional vector of valuation shocks to all institutions. The combined shock vector $\boldsymbol{x}_t$ of length $2N$ is

$$\boldsymbol{x}_t = \begin{bmatrix} \boldsymbol{x}^l_t \\ \boldsymbol{x}^v_t \end{bmatrix}. \tag{1}$$

The shock transmission matrix A is the $2N \times 2N$ matrix that acts on the shock vector $\boldsymbol{x}_t$ according to

$$\boldsymbol{x}_{t+1} = A\boldsymbol{x}_t. \tag{2}$$

Given the distinction between the top and bottom half of $\boldsymbol{x}_t$, we decompose the shock transmission matrix into its four quadrants,

$$A = \begin{bmatrix} A^{ll} & A^{vl} \\ A^{lv} & A^{vv} \end{bmatrix}, \tag{3}$$

where each of the components A^{ll}, A^{lv}, A^{vl} and A^{vv} are $N \times N$ matrices, so that Eq. (2) can be written in the form

$$\boldsymbol{x}_{t+1} = \begin{bmatrix} \boldsymbol{x}^l_{t+1} \\ \boldsymbol{x}^v_{t+1} \end{bmatrix} = \begin{bmatrix} A^{ll}\boldsymbol{x}^l_t + A^{vl}\boldsymbol{x}^l_t \\ A^{lv}\boldsymbol{x}^l_t + A^{vv}\boldsymbol{x}^v_t \end{bmatrix}. \tag{4}$$

Eq. (4) makes explicit how the diagonal quadrant A^{ll} describes the propagation of liquidity shocks and A^{vv} the propagation of valuation shocks. The off-diagonal quadrant A^{lv} gives the conversion of liquidity to valuation shocks and A^{vl} the conversion of valuation to liquidity shocks. Figure 1b shows the corresponding contagion channels.

The shock transmission matrix A is the adjacency matrix of a weighted, directed, duplex network, where the nodes are institutions and the edges represent the transmission of shocks. Each institution is represented by a node in each layer. As shown in Figure 2, the top layer describes the propagation of liquidity shocks

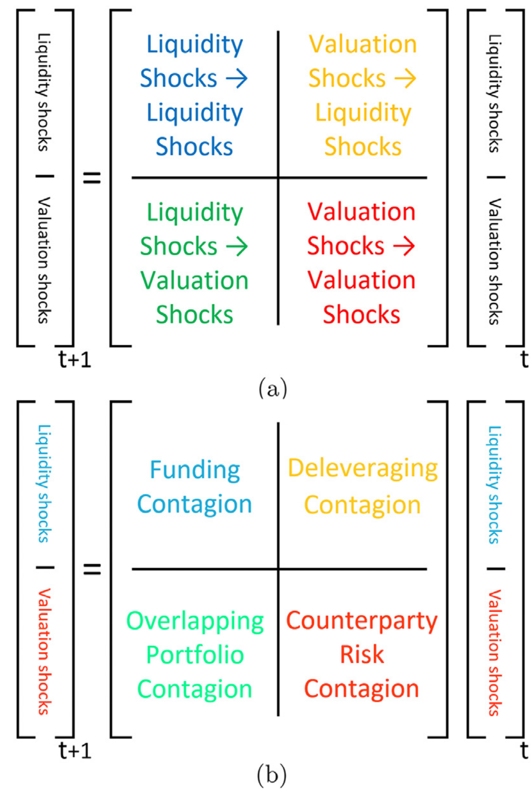

Fig. 1. Decomposition of shock dynamics. The vector of shocks to the financial system can be described as a concatenation of the vector of liquidity shocks and the vector of valuation shocks to each institution. The shock transmission matrix maps the complete vector of shocks in one period to the vector of shocks in the next period. It can be decomposed into its four quadrants as shown in the figure, corresponding to the propagation and conversion of both shock types. Note the correspondence of the quadrants in (a) and (b): Funding contagion propagates liquidity shocks, counterparty risk propagates valuation shocks, overlapping portfolio contagion converts liquidity shocks to valuation shocks and deleveraging converts valuation shocks to liquidity shocks.

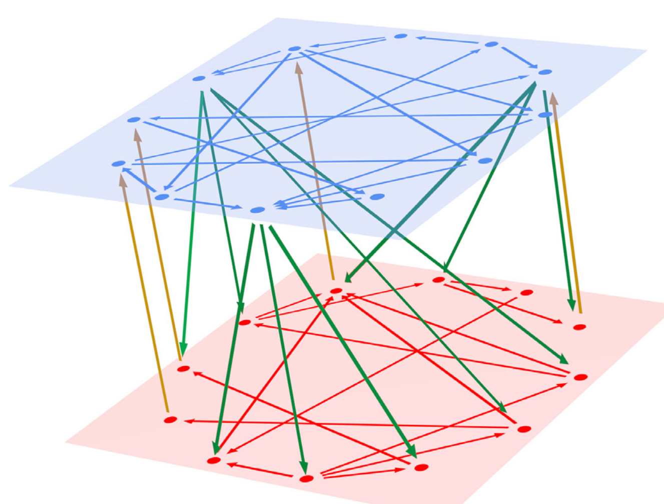

Fig. 2. The duplex network underlying the shock transmission matrix *A*. The nodes represent institutions and the edges represent the shock transmission between institutions. In this two-layer system, the top layer represents the liquidity shock network (in blue) and the bottom layer the valuation shock network (in red). The green and yellow arrows represent interactions between these two networks; the green arrow represents conversions of shocks from liquidity to valuation, whereas yellow arrows represent the conversion of shocks from valuation to liquidity. The shock transmission matrix is a weighted adjacency matrix that describes both layers and their interactions simultaneously.

by funding contagion and is referred to as the *liquidity shock network*. The bottom layer describes the propagation of valuation shocks by counterparty risk contagion and is referred to as the *valuation shock network*. The edges between the two layers describe liquidity shocks transitioning to valuation shocks and vice versa, according to overlapping portfolio and deleveraging contagion respectively. Because we can express all four contagion mechanisms in this two-layer system, in contrast to earlier methods[10], this method does not require a separate layer for each contagion mechanism.

The shock transmission matrix can be used to study the system's stability and resilience to shocks. Because all its elements are non-negative, the Perron-Frobenius theorem guarantees that the matrix has a non-negative real eigenvalue greater than or equal to (the absolute values of) the matrix' other eigenvalues. This largest eigenvalue describes the systemic properties of the financial system (Caccioli et al., 2014; Bardoscia et al., 2017; Cont and Schaanning, 2019): If the largest eigenvalue is greater than one, shocks that are not orthogonal to the corresponding eigenvector are amplified without bound and the system is *unstable*; if the eigenvalue is smaller than one, shocks are damped and we refer to it as *stable*. While no system is resilient to arbitrarily large shocks, an unstable system under this definition is not even resilient to small shocks, as it amplifies them without bound over time.

In Eq. (2), each contagion mechanism manifests itself in a single time step t. This setup implicitly assumes that all four contagion mechanisms act equally fast, which we know is not necessarily true in reality. However, this simplifying assumption does not affect our results: As we show in the Supplementary Materials, the set of conditions under which the largest eigenvalue is equal to one is independent of this assumption.

2.2. Institutions' responses to shocks

To study how a system's stability depends on its composition in terms of different types of financial institutions, we classify institutions based on the contagion they transmit. The response to a financial shock generally depends on its magnitude, but here we focus on shocks that are sufficiently small that the dynamics are approximately linear. This can be extended to deal with larger shocks by dynamically updating the shock transmission matrix as the shocks propagate.

When a liquidity shock hits, an institution may have multiple options available to respond. We assume that each institution has a *pecking order* that specifies the sequence in which it uses these options (Kok and Montagna, 2016; Hałaj, 2018). For example, once an institution has fully sold its position in a given security, it may move on to selling another, less liquid, security. The assumption of a liquidity pecking order underpins the design of regulatory measures like the Liquidity Coverage Ratio and Net Stable Funding Ratio requirements (BIS, 2013; 2014). We focus on shocks that are sufficiently small for us to assume that the liquidation option at the top of any institution's pecking order is not exhausted. In reality, the pecking order is institution-specific. Our methodology assumes that every institution has a pecking order, but it is in principle agnostic to what that pecking order looks like.

Following the approach of Kok and Montagna (2016) and Hałaj (2018), we assume that institutions adopt the pecking order that minimizes liquidation costs. As a result, any institution that holds sufficient cash on its balance sheet can absorb liquidity shocks without causing any contagion. Such institutions do not transmit any shocks in response to the receipt of a liquidity shock (i.e. its column in the left half of the shock transmission matrix is zero). We refer to these institutions as liquidity sinks. Institutions that can easily access cash, for example by borrowing on the interbank market (Rochet and Tirole, 1996) or accessing central bank credit (Bagehot, 1873), can also act as liquidity sinks (note, however, that borrowing cash is not an option when the liquidity shock arises because the institution needs to pay off its debts to decrease its leverage to return to its leverage target). If an institution holds insufficient cash but has made short-term loans, it can raise cash by not rolling over these loans. Finally, an institution can liquidate securities; we assume that this is done in descending order of liquidity (the most liquid securities are liquidated first, to minimise price impact). In sum: cash sits at the top of the pecking order, followed by withdrawal of short-term loans, and finally by liquidation of securities in descending order of liquidity.

Following a similar logic to that applying to liquidity sinks, we also define valuation sinks. For our purposes, a valuation sink is an institution without leverage. Because it has no creditors to transmit contagion to (it cannot go bankrupt) and cannot deleverage (it has no debt), it *absorbs* valuation shocks (i.e. its column in the right half of the shock transmission matrix is zero). An example of a valuation sink is a defined contribution pension fund which has no debt to the financial system.

We assume that each leveraged institution has a *leverage ceiling*, which reflects the maximum risk an institution is willing or allowed to take. The ceiling may be set by regulation, or it may be implicitly imposed by haircuts on collateralized loans. Because the haircut requires the collateral value to exceed the value of the loan, the borrower must finance the excess collateral with its own funds. This limits the amount of (collateralized) debt the institution can finance given its equity. If an institution operates sufficiently close to its ceiling that a valuation shock would force it to deleverage, then we say that it is *leverage targeting*.[11] In contrast, if the leverage is sufficiently below the ceiling (e.g. due to a leverage buffer, as proposed in recent regulation (Goodhart, 2013; FSB, 2017)) we say that it is *passively leveraged*. The shocks in our model are sufficiently small not to push passively leveraged institutions towards their ceiling to the point where they have to transition towards becoming leverage targeting.

2.3. Contagion equations

We now derive simple representative formulas for each contagion channel.

- *Funding contagion:* Suppose institution i extends a short-term loan of size S_{ij} to institution j, which is part of its short-term loan portfolio of size S_i. On receiving a liquidity shock x_i^l, assume institution i proportionately reduces the size of its short-term loans to each institution j to absorb the entire shock. This means that the liquidity shock that is transmitted to institution j is $A_{ji}^{ll} x_i^l$, where

$$A_{ji}^{ll} = \frac{S_{ij}}{S_i}. \tag{5}$$

- *Overlapping portfolio contagion:* Suppose that institution i holds n_{si} shares of security s, which makes up the top of i's pecking order[12], and that i experiences a liquidity shock x_i^l that causes

[10] See e.g. Caccioli et al. (2013); Kok and Montagna (2016); Poledna et al. (2015); Hüser et al. (2018); Bardoscia et al. (2018).

[11] See e.g. Adrian and Shin (2010); Duarte and Eisenbach (2018); Greenwood et al. (2015); Cont and Schaanning (2017); Bookstaber (2017).

[12] We assume for simplicity that institution i has a single security at the top of its pecking order. Hence, although the institution may hold positions in various securities, i only sells shares in security s (until the position is exhausted) to raise liquidity (but the model can allow for multiple securities in the top pecking order layer by assuming that i liquidates these securities proportionally to its position in each security).

i to sell $\Delta n_{si} = x^l_i / p_s$ shares, where p_s is the price of security s. Assume a price impact function of the form

$$\frac{\Delta p_s}{p_s} = \mu_s \frac{\Delta n_{si}}{n_s}, \tag{6}$$

where n_s is the total number of shares of security s in circulation, and the price impact factor μ_s is a nondimensional constant of order one that is inversely proportional to the liquidity of security s. Setting $\mu_s = 1$ implies that selling n_s shares drives the price to zero. Under the assumption of linearity, $\mu_s = 1$ is an upper bound because the price cannot be negative. The resulting valuation shock to any institution j that holds n_{sj} shares of security s is $\Delta p_s n_{sj} = \mu_s x^l_i n_{sj}/n_s$, which implies

$$A^{lv}_{ji} = \mu_s \frac{n_{sj}}{n_s}. \tag{7}$$

Note that the diagonal component A^{lv}_{ii} is nonzero. We assume for simplicity that institutions do not short securities, so we always have $n_{sj} \geq 0$.[13]

- *Counterparty risk contagion:* Assume passively leveraged institution i has equity E_i and total debt D_i, so that its leverage is $\lambda_i = D_i/E_i$. When institution i experiences a valuation shock, its probability of default rises and the risk-adjusted value of its debt falls (Bardoscia et al., 2017). Institutions with more equity can withstand larger valuation shocks without becoming insolvent. Therefore, we assume that the fractional drop in the value of the debt is proportional to the fractional loss in equity x^v_i/E_i. If institution i owes debt D_{ij} to institution j, then the valuation shock transmitted to institution j is $\delta_i x^v_i / E_i D_{ij}$, so

$$A^{vv}_{ji} = \delta_i \frac{1}{E_i} D_{ij} = \delta_i \lambda_i \frac{D_{ij}}{D_i}, \tag{8}$$

where the risk adjustment factor δ_i is a nondimensional constant of order one. Choosing $\delta_i = 1$ implies that a shock of size E_i (which causes bankruptcy) causes the full value of the debt to be lost and passed onto i's creditors as a valuation shock. Under the assumption of linearity, $\delta_i = 1$ is an upper bound as the loss cannot exceed the value of the debt. D includes short-term as well as long-term debt, so in general $D_{ij} \geq S_{ji}$.
- *Deleveraging contagion.* Suppose leverage targeting institution i maintains a leverage target λ_i. If it receives a valuation shock x^v_i it must pay off debt to return to its target. The amount by which it must reduce debt is $\lambda_i x^v_i$, so

$$A^{lv}_{ii} = \lambda_i. \tag{9}$$

We assume that institution i's leverage targeting prevents the institution from transmitting counterparty risk contagion to its creditors. This is because the institution averts the risk associated with increased leverage by paying off its debts to keep its leverage constant.

We want to stress that the parameters in the equations vary over time, and hence the shock transmission matrix is defined with respect to a specific time t. For simplicity we omit the time subscripts to the matrix and its entries.

The four contagion equations are summarized in Table 1. This set of contagion mechanisms is not exhaustive; for example, information contagion is not included (Aharony and Swary, 1996; Acharya and Yorulmazer, 2008). Furthermore, in times of crisis, institutions sometimes hoard liquidity in response to liquidity shocks (Acharya and Skeie, 2011; Heider et al., 2009). Liquidity hoarding can be included in the funding (5) and overlapping portfolio (7) contagion equations by adding a hoarding term that captures the additional liquidity an institution hoards proportionally to the received liquidity shock. We make the simplifying assumption that liquidity hoarding is absent. We have chosen the four forms we study here because they are all important, but we restrict ourselves to only four contagion channels for simplicity. Our basic methodology applies to any contagion channels and does not depend on the details of the interaction terms.

Table 1
Contagion Equations

Contagion Mechanism	Contagion Equation	Description
Funding Contagion	$A^{ll}_{ji} = \frac{S_{ij}}{S_i}$	Short-term lending withdrawal
Counterparty Risk Contagion	$A^{vv}_{ji} = \delta_i \lambda_i \frac{D_{ij}}{D_i}$	Probability of default increases due to lower valuations
Overlapping Portfolio Contagion	$A^{lv}_{ji} = \mu_s \frac{n_{sj}}{n_s}$	Price-impact of selling securities
Leverage Targeting Contagion	$A^{vl}_{ii} = \lambda_i$	Delevering requires raising liquidity

According to the Perron-Frobenius theorem, the largest eigenvalue of the shock transmission matrix is bounded by its smallest and largest column sums. The sum of a column's entries gives the size of the aggregate shock the institution transmits relative to a received liquidity or valuation shock (depending on whether the columns is in the left or right half of the matrix). When no column-sum exceeds one, no institution ever transmits an aggregate shock that exceeds the received shock, so there is no shock amplification and the system is stable. Conversely, when all column-sums exceed one, shocks are always amplified and the system is unstable.

The sum of a column corresponding to an institution's transmission of funding contagion (5) is equal to $\sum_j S_{ij}/S_i = 1$ and the sum of a column corresponding to overlapping portfolio contagion (7) is given by $\sum_j \mu_s n_{sj}/n_s = \mu_s \leq 1$. Hence, the aggregate shock transmitted in response to a liquidity shock is never amplified (but note that the addition of a liquidity hoarding term could change this). Furthermore, the sum of a counterparty risk contagion column is equal to $\sum_j \delta_i \lambda_i D_{ij}/D_i = \delta_i \lambda_i \leq \lambda_i$ and the sum of a deleveraging column is given by its only non-zero element λ_i. Therefore, under the assumption of no liquidity hoarding, *leverages exceeding one are the only source of shock amplification in the system.* This acts through the counterparty risk and deleveraging channels.

2.4. Stylized example

We illustrate the approach outlined above using a simple example of a self-contained financial system that includes all four contagion mechanisms as well as both liquidity and valuation sinks. Consider four institutions, as summarized in Figure 3.

- Pension fund h has no debt and a cash surplus, making it both a valuation and a liquidity sink. It makes long-term loans L_{hi}, L_{hj} and L_{hk} to institutions i, j and k and has a position n_{sh} in security s.
- Bank i is passively leveraged. It has a position n_{si} in security s, which sits at the top of its pecking order, and debt $D_{ih} = L_{hi}$ and $D_{ij} = S_{ji} + L_{ji}$ to institutions h and j.
- Bank j targets leverage λ_j. It makes short and long-term loans S_{ji}, L_{ji} and S_{jk}, L_{jk} to institutions i and k and has debt $D_{jh} = L_{hj}$ to institution h. The short-term loans sit at the top of its pecking order.

[13] The framework can accommodate short positions by simply allowing n_{sj} to be negative, but some convenient properties of the matrix would no longer be guaranteed by the Perron Frobenius theorem.

Balance sheet pension fund h	
Assets	Liabilities
Cash: surplus Position in security s Long-term loans to institutions i, j and k	 Equity

(a)

Balance sheet bank i	
Assets	Liabilities
Cash: required minimum Position in security s	Debt to h and j Equity

(b)

Balance sheet bank j	
Assets	Liabilities
Cash: required minimum Short and long-term loans to i and k	Debt to h Equity

(c)

Balance sheet bank k	
Assets	Liabilities
Cash: surplus Position in security s	Debt to h and j Equity

(d)

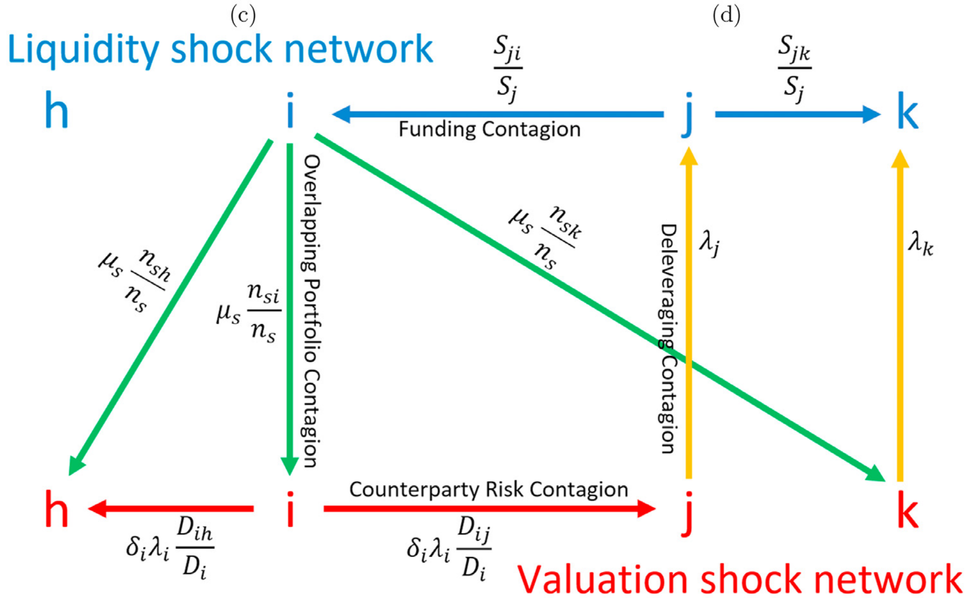

(e)

Fig. 3. A stylized example illustrating the interaction of multiple channels of contagion. Consider four institutions, h, i, j and k, whose balance sheets are given at the top of the figure. Each institution's liquidity is represented by a node in the liquidity shock network, while each institution's solvency is represented by a node in the valuation shocks network. The label of an edge indicates the type of contagion transmitted (each node transmits a single type of contagion so only one of a node's out-edges is labeled) and the expression next to the edge represents the size of the interaction. The out-edges of a node in the liquidity shock network give the node's response to liquidity shocks and are dictated by the asset that sits at the top of the institution's pecking order. Similarly, out-edges in the valuation shock network give the response to valuation shocks as given by the institution's leverage strategy: unleveraged (institution h), passively leveraged (institution i), or leverage targeting (institutions j and k).

- Bank k has a cash surplus, making it a liquidity sink, and maintains a leverage target λ_k. It has a position n_{sk} in security s, short-term debt $D_{kj} = S_{jk} + L_{jk}$ and long-term debt $D_{kh} = L_{hk}$.

The shock transmission matrix (and dynamic of the shock vector) of this system is:

$$\begin{array}{c} \\ h^l \\ i^l \\ j^l \\ k^l \\ h^v \\ i^v \\ j^v \\ k^v \end{array} \begin{array}{c} \begin{array}{cccccccc} h^l & i^l & j^l & k^l & h^v & i^v & j^v & k^v \end{array} \\ \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \frac{S_{ji}}{S_j} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \lambda_j & 0 \\ 0 & 0 & \frac{S_{jk}}{S_j} & 0 & 0 & 0 & 0 & \lambda_k \\ 0 & \mu_s \frac{n_{sh}}{n_s} & 0 & 0 & 0 & \delta_i \lambda_i \frac{D_{ih}}{D_i} & 0 & 0 \\ 0 & \mu_s \frac{n_{si}}{n_s} & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & \delta_i \lambda_i \frac{D_{ij}}{D_i} & 0 & 0 \\ 0 & \mu_s \frac{n_{sk}}{n_s} & 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix} \end{array} \begin{bmatrix} x^l_{t,h} \\ x^l_{t,i} \\ x^l_{t,j} \\ x^l_{t,k} \\ x^v_{t,h} \\ x^v_{t,i} \\ x^v_{t,j} \\ x^v_{t,k} \end{bmatrix} = \begin{bmatrix} x^l_{t+1,h} \\ x^l_{t+1,i} \\ x^l_{t+1,j} \\ x^l_{t+1,k} \\ x^v_{t+1,h} \\ x^v_{t+1,i} \\ x^v_{t+1,j} \\ x^v_{t+1,k} \end{bmatrix}, \tag{10}$$

where $S_j = S_{ji} + S_{jk}$, $D_i = D_{ih} + D_{ij}$ and $n_s = n_{sh} + n_{si} + n_{sk}$.

To simplify the discussion we set the price-impact and risk adjustment factors to their upper bounds $\mu_s = \delta_i = 1$. For this system, the largest eigenvalue of the shock transmission matrix is equal to

$$\nu = \left(\frac{S_{ji} n_{si} \lambda_i D_{ij} \lambda_j}{S_j n_s D_i} \right)^{1/4}, \tag{11}$$

and so the largest eigenvalue in this example is the product of each of the four contagion mechanisms, i.e.

$$\nu^4 = \frac{S_{ji}}{S_j} \times \frac{n_{si}}{n_s} \times \lambda_i \frac{D_{ij}}{D_i} \times \lambda_j.$$

The factors S_{ji}/S_j, n_{si}/n_s and D_{ij}/D_i are all less than or equal to one, and so exert a stabilizing force competing against the potentially destabilizing forces of the leverages λ_i and λ_j. If any of the four channels of contagion is removed the largest eigenvalue becomes zero and the system becomes unconditionally stable. The possibility for instability is caused by the interaction of all four channels, so analyzing each channel in isolation, as often done, gives a misleading result.

Consider some plausible numbers: if $S_{ji}/S_j = D_{ij}/D_i = 1/3$, $n_{si}/n_s = 1/4$ and $\lambda_i = \lambda_j = 6$, then $\nu = 1$ and the system is at its margin of stability. For any system, we define the *critical leverage* $\hat{\lambda}$ as the maximum leverage that all *leveraged* institutions can attain simultaneously without rendering the system unstable (i.e. $\nu \leq 1$, where all unleveraged institutions are assumed to remain so). For the system outlined in the stylized example, the critical leverage is thus given by $\hat{\lambda} = 6$.

The Stabilizing Role of Sinks. Sinks play an essential role in stabilizing the financial system. All else equal, if more shocks are transmitted to sinks, then less shocks are transmitted to potentially destabilizing institutions.[14] The ability of a sink to absorb shocks increases when it holds more short-term lending, more securities and more of other institutions' debt. This can be seen for the example by expanding the denominators in Eq. (11),

$$\nu = \left(\frac{S_{ji}}{S_{ji} + S_{jk}} \times \frac{n_{si}}{n_{sh} + n_{si} + n_{sk}} \times \lambda_i \frac{D_{ij}}{D_{ij} + D_{ih}} \times \lambda_j \right)^{1/4}. \qquad (12)$$

The sinks' short-term debt S_{jk}, securities holdings n_{sh} and n_{sk}, and lending D_{ih} all appear only in the denominator and therefore exert a stabilizing force on the system. Hence, *sinks stabilize the system by absorbing shocks that would otherwise have been transmitted to other institutions*.

2.5. Application to stress testing

This framework can be calibrated using granular data to accurately represent the microstructure of the financial system and it can be used to study any channels of contagion (he description can be made at any level of granularity, down to individual contracts). Because the shock transmission matrix' largest eigenvalue quantifies stability independently from any specific stress scenario, it provides an objective, robust measure of stability that allows for comparison across time, jurisdictions, policy interventions, and so on.

The eigenvectors associated with the largest eigenvalue also provide useful diagnostic information.[15] The right eigenvector provides a measure of institutions' in-degree centralities that takes the whole system into account and the left eigenvector provides a measure of institutions' out-degree centralities that takes the whole system into account (Newman, 2018). The larger entries of the right eigenvector flag the institutions that are likely to receive the largest shocks and those of the left eigenvector flag the institutions that play the biggest role in transmitting shocks. This can potentially help guide policymakers in identifying systemically important institutions and designing stress tests and interventions.

Scenario-Dependent Stability.

Although the framework here assumes infinitesimal shocks, it is also useful for understanding large shocks. In reality institutions do not respond to a shock instantaneously, but rather take a series of actions that are initially close to the dynamics captured by our framework but diverge from these over time, so as the financial system responds to a large shock its stability changes. As the shock plays out, institutions may exhaust the top layers of their pecking orders, market liquidities may fall, and risk adjustment factors may rise. By investigating how this affects (linear) stability, we may gain insight into how large shocks affect stability. The properties of the financial system may also change due to central bank or government intervention in response to the shock, or because of adaptive processes that take place as part of the ongoing evolution of the financial system (Farmer et al., 2021)).

The evolution of the stability of the financial system in response to a large shock embodies the *scenario-dependent* component of the system's shock-dynamics, which is not captured in its linear stability. As the framework can be calibrated to any state of the system, by continuously re-calibrating the shock transmission matrix as the system evolves in response to a shock, we may understand how the instantaneous stability changes, track it over time, and distinguish between the scenario-dependent and independent components of the system's stability. This can be done for both the empirically observed or forecast (by means of some simulation model) evolution of real financial systems.[16] By investigating the sensitivity of a system's stability to specific stress scenarios using the method developed here, policymakers may gain insight into the factors that generate financial instability in any given scenario.

The effect of the liquidity pecking order in particular warrants further investigation, because the stability of financial systems strongly depends on which assets are at the top of institutions' pecking orders. For example, when each institution has a single-layer pecking order (i.e. it liquidates a vertical slice across all its liquid assets), the shock transmission matrix remains constant as long as institutions do not default or change their leverage strategies (under reasonable assumptions about μ_s and δ_i). Liquidation-cost minimizing pecking orders, on the other hand, may be highly sensitive to the choice of stress scenario, because liquidities may change and layers of the pecking order may be exhausted. Therefore, investigating what pecking orders institutions use during crises may yield valuable insights into the predictability of financial stability.

3. Stability overestimation due to ignoring interactions

In the previous sections we derived the shock transmission matrix and explained how it may be used to assess the stability of financial systems with interacting contagion channels. In this section, to underscore the importance of using the shock transmission matrix to complement existing stress tests we demonstrate the dangers of ignoring the interactions between contagion channels. In the stylized example in section 2.4, we showed that ignoring any of the four contagion channels may overestimate stability. In this section we demonstrate this in a more general setting. Because most contagion literature studies the counterparty default (risk) channel alone[17], we focus on the overestimation due to only considering counterparty risk contagion and ignoring all other channels. This *stability of pure counterparty risk contagion* is given by the largest eigenvalue of the counterparty risk component of the shock transmission matrix, i.e its bottom-right quadrant. We first discuss two extremes for which the overestimation follows intuitively and then consider intermediate cases.

In the case where all institutions in the financial system are passively leveraged or no institution has tradeable securities at the top of their pecking order, the system's shock transmission matrix is block-triangular. Due to the properties of block-matrix determinants, the largest eigenvalue of this system is given by the largest eigenvalue of the diagonal quadrants, which correspond to counterparty risk and funding contagion. Such a system is stable when pure counterparty risk contagion is stable (as funding contagion is never amplified under the assumption of no liquidity hoarding) so that there is little potential for overestimating stability.

[14] In any system, the eigenvalues are determined by the shock transmission between strongly connected nodes. Sinks are by definition not strongly connected, so the more shocks transmission to sinks, the less shock transmission between strongly connected nodes and the lower the largest eigenvalue.

[15] The Perron Frobenius theorem guarantees that the right eigenvector is non-negative and, assuming that the network has a single strongly connected component, that it is unique (i.e. the coresponding eigenspace of the largest eigenvalue ν is one-dimensional). The same is true for the left eigenvector.

[16] Note that the purpose of our framework is to identify instabilities as they emerge, rather than forecast the nature and magnitude of the crisis that may ensue. The framework may provide insight into the simulated evolution of a financial system, but the simulation itself would require another model (or an extension of the one developed here).

[17] See e.g. Eisenberg and Noe (2001); Furfine (2003); Gai and Kapadia (2010); Battiston et al. (2012); Elliott et al. (2014); Acemoglu et al. (2015); Bardoscia et al. (2015, 2017).

In the opposite situation where all institutions are leverage targeting, the counterparty risk quadrant is zero. Hence, considering pure counterparty risk contagion would lead to the conclusion that the system is unconditionally stable, regardless of the true stability. To understand how ignoring interactions between contagion channels leads to an overestimation of stability in systems with both passively leveraged and leverage targeting institutions, we must understand how network composition affects stability. As a tractable example, in the sections that follow we apply the framework developed in section 2 to randomly generated financial systems.

3.1. Application to randomly generated financial systems

To gain insight into the dynamics of interacting contagion channels we study randomly generated financial systems. To do this we populate the balance sheets of N institutions with randomly generated securities and loans. We do this in such a way that the types of institutions introduced in section 2.2 have the following proportions:

1. A fraction ϕ_l of institutions have sufficient cash to absorb shocks. We call ϕ_l the *fraction of liquidity sinks*.
2. A fraction ϕ_v of institutions have no leverage. We call ϕ_v the *fraction of valuation sinks*.
3. A fraction F of institutions provide short-term loans. We call F the *fraction of short-term lenders*.
4. A fraction Λ of *leveraged* institutions are leverage targeting. We call Λ the *fraction of leverage targeters*.

These proportions constrain the random assignment of loans. For each security s out of a possible number N^w of distinct securities, we divide the total number of outstanding shares n_s into N^s blocks of n_s/N^s shares and assign each block to a randomly chosen institution. We do this with uniform probability and with replacement. Similarly, each institution makes N^d loans, each to a randomly chosen *leveraged* institution; for simplicity, all loans an institution *receives* are set equal in size. Any institution that was designated as leveraged but ended up not receiving any loans is allocated a single loan from a randomly chosen institution.

For any institution i, let N_i^s denote the number of blocks of security s received, N_i^d the total number of loans received, and N_{ji}^d the number of loans received from institution j. The leverages of the $N^v = (1-\phi_v)N$ leveraged institutions are set to the critical leverage $\lambda_i = \hat{\lambda}$, which fixes their debts relative to their equities. Once we choose the N^w distinct securities' market capitalizations $C_s = p_s n_s$, the requirement that any institution's assets (LHS) must equal the sum of its equity and debt (RHS) provides N constraints that determine the N institutions' equity E_i,

$$\sum_{s=1}^{N^w} C_s \frac{N_i^s}{N^s} + \sum_{j=1}^{N^v} D_j \frac{N_{ij}^d}{N_j^d} = E_i(\lambda_i + 1). \tag{13}$$

This allows us to generate a random financial system with any desired values of the parameters ϕ_l, ϕ_v, F and Λ.

Financial systems tend to have sparse, heterogenous topologies (Boss et al., 2004; Cont et al., 2013) that can frequently be characterized as core-periphery structures (Craig and Von Peter, 2014; Fricke and Lux, 2015). Similarly, our randomly generated systems here include many sources of heterogeneity and have a core-periphery structure: Since an institution can either be a liquidity sink or not be a liquidity sink, provide short-term loans or not provide short term loans, and be unleveraged, passively leveraged or have a leverage target, this implies that there are $2 \times 2 \times 3 = 12$ different *types* of institutions. The short-term lending network that results from our method of random construction has a core of institutions that both provide and receive short-term loans. There are also three distinct peripheries - one of institutions that only provide short-term loans, one of institutions that only receive short-term loans, and one of institutions that do not partake in the short-term lending network at all. The long-term lending network has a similar topology but with different institutions at its core and peripheries.

Furthermore, because loans and securities are chosen *with replacement*, institutions can receive multiple loans from the same institution and hold multiple shares in the same security. Consequently, the weights of the edges vary across institutions in all networks. Finally, because the institutions' assets are randomly determined, the endogenously-determined balance sheet sizes vary across institutions. By varying the system parameters ϕ_l, ϕ_v, F, Λ, N, N^d, and N^s, we can control the level of heterogeneity present in the system.

3.2. Mean-field approximation

In the limit where $N, N^d/N, N^s/N \to \infty$, the randomly generated financial systems reduce to a mean-field model. In fact, as we show below, the mean field model remains a reasonable approximation for much smaller, sparser systems as well. Rather than studying the stability of the generated systems explicitly we focus on the mean-field approximation, which is more insightful because it provides an analytic stability condition.

In Supplementary Materials, we explain that the dynamics of this mean-field model are uniquely defined by the transmission of the aggregate liquidity shock $x_t^l = \sum_i x_{t,i}^l$ and aggregate valuation shock $x_t^v = \sum_i x_{t,i}^v$. This allows us to reduce the system's full shock transmission matrix to a 2×2 matrix that describes the dynamics of the aggregate shocks,

$$\begin{bmatrix} x_{t+1}^l \\ x_{t+1}^v \end{bmatrix} = \left[\begin{array}{c||c} (1-\phi_l)F & \lambda(1-\phi_l)\Lambda \\ \hline \mu(1-\phi_v)(1-F) & \delta\lambda(1-\phi_v)(1-\Lambda) \end{array}\right] \begin{bmatrix} x_t^l \\ x_t^v \end{bmatrix}, \tag{14}$$

where μ is the price impact factor of the most liquid security (in which all institutions have a position when $N^s/N \to \infty$), and we have set $\lambda_i = \lambda$ for all leveraged institutions and $\delta_i = \delta$ for passively leveraged institutions.

We compute the characteristic equation of the matrix in Eq. (14) and solve for its largest eigenvalue $\nu = 1$ to find the *mean-field critical leverage*,

$$\hat{\lambda} = \frac{1-(1-\phi_l)F}{\mu\Lambda(1-\phi_l)(1-F) + \delta(1-\Lambda)(1-(1-\phi_l)F)}(1-\phi_v)^{-1}. \tag{15}$$

Eq. (15) demonstrates that financial stability is the result of a balance between the destabilizing force of leverage and the stabilizing force of sinks (note that the mean-field critical leverage (15) is an increasing function of both the liquidity sinks fraction ϕ_l and valuation sinks fraction ϕ_v). We first discuss the accuracy of the mean-field model before we use it to understand how ignoring interactions between contagion channels overestimates financial stability.

Accuracy of the Mean-Field Model. The mean-field model was derived in the limit of a dense network with an infinite number of institutions. In Figure 4, we compare the mean-field critical leverage predicted by Eq. (15) to that of randomly generated systems with varying sizes and densities and show that the mean-field critical leverage is a good approximation not only for large, dense systems, but also for fairly small, sparse systems.

To choose plausible estimates of the system parameters we do a rough calibration to the European financial system. To estimate the parameters ϕ_v, Λ, F and ϕ_l, we calculate the fraction of financial assets in the Eurosystem held by various sectors, as shown at

Sector	Assets
Pension Funds (PF)	2.6[a]
Insurance Corporations (IC)	8.0[b]
Monetary Financial Insts. (MFI)	30.6[c]
Investment Funds (IF)	12.4[d]
Financial Vehicle Corps. (FVC)	1.9[e]

(a)

Fraction	Calculation	Value
ϕ_v	$\frac{PF+IC}{PF+IC+MFI+IF+FVC}$	0.2
F	$\frac{MFI}{PF+IC+MFI+IF+FVC}$	0.5
ϕ_l	$\frac{PF+IC+MFI}{PF+IC+MFI+IF+FVC}$	0.75
Λ	$\frac{MFI}{MFIs+IFs+FVCs}$	0.75

(b)

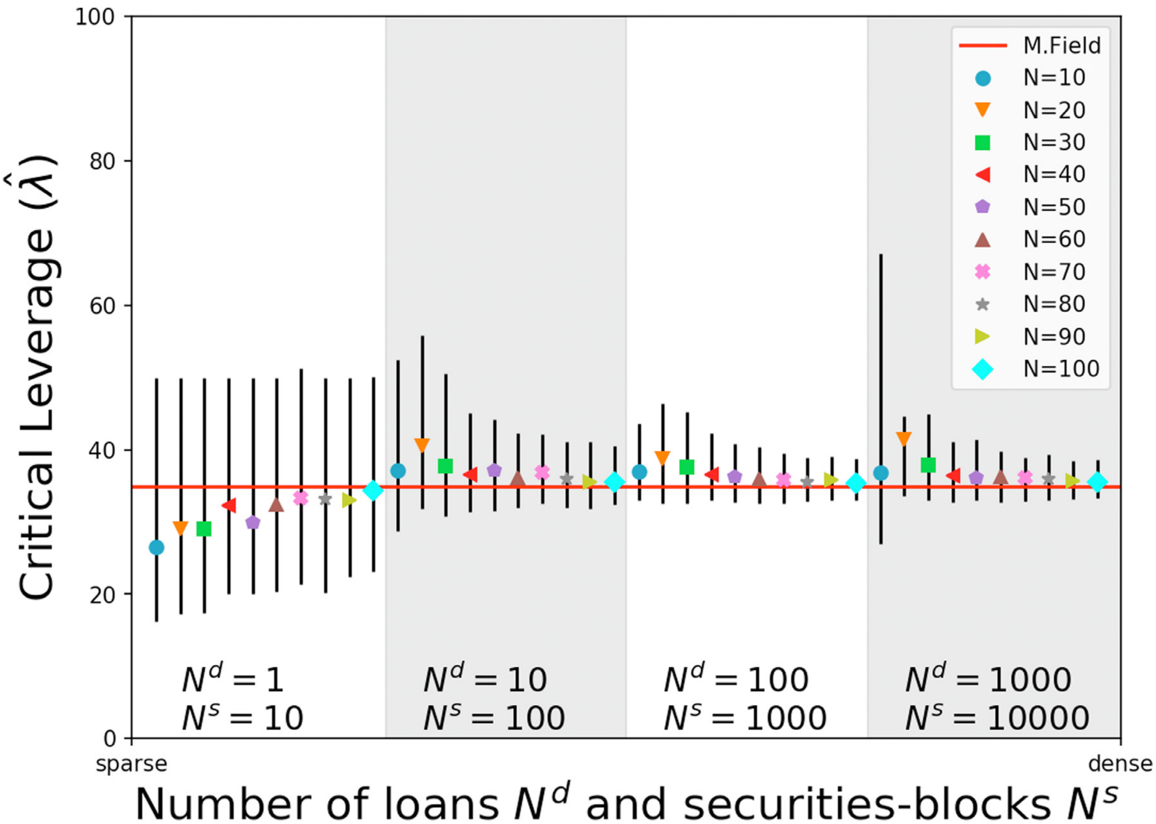

(c)

Fig. 4. Comparison of the mean field model to randomly generated financial systems. The plot compares the critical leverage (horizontal line) predicted by the mean field model, Eq. (15), to the critical leverages of randomly generated financial systems. For various combinations of the number of loans N^d, the number of securities-blocks N^s and the number of institutions N, we generate 500 random systems and plot the median (using various markers to indicate the number of institutions) and the 15^{th} to 85^{th} percentile interval (vertical bars) of the distribution of critical leverages. The simulated values converge to the predictions of the mean-field model as N^d, N^s, N increase. [a]Source: http://sdw.ecb.europa.eu/reports.do?node=1000005664; accessed November 3[rd] 2018 [b]Source: http://sdw.ecb.europa.eu/reports.do?node=1000005659; accessed November 3[rd] 2018 [c]Source: http://sdw.ecb.europa.eu/reports.do?node=1000005718; accessed November 3[rd] 2018 [d]Source: http://sdw.ecb.europa.eu/reports.do?node=1000003516; accessed November 3[rd] 2018 [e]Source: http://sdw.ecb.europa.eu/reports.do?node=1000003621; accessed November 3[rd] 2018

the top of Figure 4.[18] We include $N^w = 10$ distinct securities whose markets caps are calibrated to the ten largest stocks on the Euronext exchange.[19] The Eurosystem's 119 most *significant institutions* as designated by the ECB have leverages ranging between 10 and 20 (ECB, 2019). Given that the Eurosystem appears to be stable, we choose values of $\delta = \mu = 0.1$, which gives a critical leverage of $\hat{\lambda} = 35$. We also think these are reasonable values for other reasons.[20]

We explore different combinations of the parameters N, N^d and N^s, generating 500 realizations for each set of parameter values. In

[18] The aggregate sector assets are listed in Table (rounded to one decimal) and the parameter calculations are listed in Table (rounded to a multiple of $1/4^{th}$ or $1/5^{th}$ so the number of each type is an integer for any multiple of $N = 10$). We assume that pension funds and insurance corporations are unleveraged, monetary financial institutions (which are mostly banks) are leverage targeting, and investment funds and financial vehicle corporations are passively leveraged. Monetary financial institutions are assumed to be the only providers of short-term loans. Finally, we assume favorable market liquidity conditions such that all institutions but investment funds and financial vehicle corporations are liquidity sinks. Although these assumptions imply correlations between institutions' types, the are assigned independently from one another. Note that these are rough parameter estimates. Accurate calibration requires further (empirical) investigation.

[19] Source: https://www.statista.com/statistics/546298/euronext-market-capitalization-leading-companies/; accessed January 22[nd], 2019.

[20] The fact that a large fraction of the value of a loan can usually be recovered suggests that δ is substantially below one. Similarly, since institutions liquidate assets in order of liquidity, the price impact factor μ_s of any security at the top of institutions' pecking orders is likely to be substantially less than one.

Figure 4 we plot the median critical leverage as well as the 15^{th} to 85^{th} percentile intervals against the prediction of Eq. (15).[21] The figure shows that the mean-field model gives a good approximation for systems with at least $N \geq 30$ institutions and at least $N_d \geq 10$ loans made by each institution.

Although this is not a large effect, note that the smallest, sparsest systems in Figure 4 are the least stable. This is in contrast to Bardoscia et al. (2017), who find the sparsest systems are most stable. Our model approximately reduces to theirs in the limit where all institutions are passively leveraged ($\Lambda \to 0$). However, Bardoscia et al. (2017) require the sparsest systems to be acyclic, which is the most stable configuration, as the largest eigenvalue is zero by definition. We do not impose this requirement, which is why we get the opposite result.

Our model does not address important features of the financial system, such as hedging with derivatives or short positions. Figure in the appendix shows that introducing additional sources of heterogeneity increases the variation in critical leverage. Nonetheless, the mean field model does a good job of qualitatively capturing some of the key features of financial stability. We want to stress, however, that when fine-grained data is available, it is far preferable to use the full model developed in section 2, which uses weaker assumptions and contains fewer approximations.

3.3. The misclassification region

We now demonstrate that stability is almost always overestimated, and sometimes dramatically so, when ignoring the interactions between contagion channels. For simplicity, because this is only a qualitative demonstration, we set the price-impact and risk-adjustment factors in the mean-field critical leverage (15) equal to their upper bounds $\mu = \delta = 1$. This simplifies Eq. (15) to

$$\hat{\lambda} = \left(\frac{1-(1-\phi_l)F}{1-(1-\phi_l)F-\phi_l\Lambda} \right)(1-\phi_v)^{-1}. \tag{16}$$

Eq. (16) is a product of two terms; the first captures the intensity of the feedback loop between solvency and liquidity, and the second captures the way in which valuation sinks counterbalance the destabilizing force of leverage.

The stability of pure counterparty risk is determined by the counterparty risk quadrant alone. We find the critical leverage of pure counterparty risk contagion in the mean-field model by solving for the leverage for which the counterparty risk quadrant in Eq. (14) is equal to one,

$$\hat{\lambda} = \frac{1}{1-\Lambda}(1-\phi_v)^{-1}. \tag{17}$$

This is equivalent to setting ϕ_l or F in Eq. (16) equal to one (so there is no feedback loop between solvency and liquidity).

The counterparty risk critical leverage (17) may be shown to severely overestimate the true mean-field critical leverage when the interaction of other contagion channels is taken into account. To take a simple case, when $\phi_l = 0$, i.e. when there are no liquidity sinks, the true critical leverage equals

$$\hat{\lambda} = (1-\phi_v)^{-1}. \tag{18}$$

In Figure 5 we plot the counterparty risk critical leverage (17) and the true critical leverage (18). As the fraction of leveraged institutions increases, the discrepancy between the counterparty risk critical leverage and true critical leverage becomes arbitrily large, and hence so does the overestimation of stability due to ignoring the interactions between contagion channels. To show that this happens for any value of $\phi_l < 1$, in Figure of the appendix we plot the critical leverage for various values of ϕ_l. For $\phi_l = .75$, $F = .5$ and $\Lambda = .75$, for example, we find that the counterparty risk critical leverage overestimates the true critical leverage by 45%, but this soars to 300% when $\phi_l \to 0$, as may be the case when liquidity dries up during financial crises.

Detering et al. (2021) also show in a scenario-independent setting (which also only depends on aggregate system parameters) that stability is overestimated when ignoring the interaction between contagion channels. However, they only include counterparty default and overlapping portfolio contagion and do not capture liquidity in their model. By capturing the solvency-liquidity nexus (and all its contagion channels) in its entirety, we demonstrate that this overestimation is determined by the intensity of the feedback loop between solvency and liquidity.[22] The manifestation of this overestimation in all but unrealistic scenarios makes clear that taking account of interacting contagion channels, as our approach outlined in Section 2 proposes, is critical when evaluating financial stability.

4. Concluding remarks

Financial instability is caused by the endogenous amplification of shocks.[23] We are the first to introduce a scenario-independent measure of the stability of the solvency-liquidity nexus that takes into account the interactions of an arbitrary number of financial contagion channels. By describing the interactions of liquidity and valuation shocks, our method captures the most important contagion mechanisms and their interactions in a duplex network consisting of a liquidity and a valuation shock layer. The largest eigenvalue of the system provides a robust measure of the system's stability that is complementary to the insights provided by existing methods because it does not rely on subjectively imposed stress scenarios. Furthermore, the associated eigenvectors provide detailed insights that are valuable for important policy considerations, such as the identification of the most systemically important institutions.

With appropriate microdata this method can be calibrated against real financial systems. To do this it is necessary to estimate which institutions absorb liquidity shocks, to identify their leverage strategies and pecking orders, and measure the price-impact and risk-adjustment factors μ_s and δ_i. While this is not a trivial task, it is feasible with the right data. We hope that our model will help provide an incentive for central banks to collect this data. The analysis presented here relies on the assumption that shocks are small enough that it is only necessary to consider the top level of the pecking order. However, as we have outlined, the method may also be used to monitor the stability of a financial system as it evolves in response to larger shocks. This will be investigated in a follow-up paper.

The framework presented here has the advantage over black-box simulations, such as Kok and Montagna (2016); Cont and Schaanning (2017) or Farmer et al. (2020), that it provides insight into the mechanisms that cause contagion. In fact, it can be used in conjunction with such simulations to provide a deeper understanding of their results. Furthermore, our analysis here complements asymptotic graph techniques (see e.g. Gai and Kapadia (2010); Amini et al. (2016); Detering et al. (2021)), which study the final state to which the system converges in response to a shock. An

[21] Note that the percentile bars visualize the width of the observed distributions. They are not error bars – these are negligible given the large number of samples and are therefore not plotted.

[22] Comparison of equations 17 and 18 shows that the overestimation is caused by the first term on the right hand side of Eq. 16, as the second term appears identically in all three equations.

[23] See e.g. Danielsson and Shin (2003); BIS (2009); Krishnamurthy (2010); Acemoglu et al. (2012); Anderson et al. (2018).

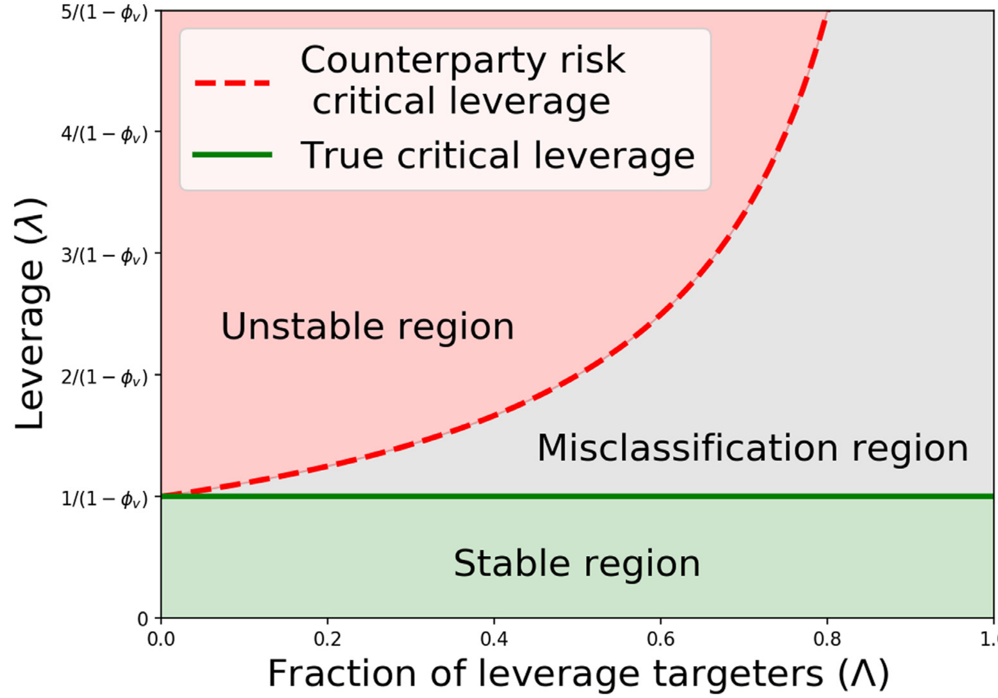

Fig. 5. Overestimation of financial stability by omitting the interaction of channels of contagion. The solid line plots the critical leverage of the mean-field model for the fraction of liquidity sinks $\phi_l = 0$, and the dashed curve plots the counterparty risk critical leverage as a function of the fraction of leverage targeting institutions Λ. The leverages are expressed in units of $(1-\phi_v)^{-1}$ to reduce the number of free parameters. The region between the true critical leverage and the counterparty risk critical leverage is the *misclassification region* and consists of destabilizing leverages that *seem* stable when considering only pure counterparty risk contagion.

important advantage of the instantaneous stability quantified by the shock transmission matrix is that it does not suffer from the compounding inaccuracies that any iterative model is inevitably exposed to.

We derive an analytic stability criterion in the limit of a large number of institutions which demonstrates that financial stability results from the balance between stabilizing and destabilizing forces. Although the stability criterion is simple, with only a few parameters, it is powerful enough to derive a wide range of insights about the stability of financial systems. It shows that the shock-amplifying forces of the levels of leverage common in real financial systems must be offset by damping to maintain stability. Understanding the conditions under which institutions absorb financial shocks is crucial to policymakers. Despite this, the absorption of shocks by sinks and the damping of financial shocks have received little attention in the literature so far. The only previous work that we are aware of that stresses this point is Aymanns et al. (2016), who study the balance between shock-amplification due to deleveraging banks and shock-damping by fundamental value investors. We study the interaction between shock-amplifying and shock-dampening forces in a much more general network setting. The stability criterion that we derive demonstrates the fundamental importance of the balance between stabilizing and destabilizing forces, highlighting that this interplay deserves further investigation.

Building on the work of others who have observed that interactions between contagion channels amplify instabilities in particular settings (Caccioli et al., 2013; Kok and Montagna, 2016; Poledna et al., 2015; Detering et al., 2021), the analytic stability criterion that we develop here elucidates the mechanism responsible for this amplification in general. We show that a feedback loop between liquidity and valuation shocks always exists and that when the interactions between contagion channels are ignored, this feedback loop is overlooked and stability is overestimated, sometimes dramatically so. Because most studies focus on a single type of shocks (see e.g. Eisenberg and Noe (2001); Caccioli et al. (2014); Bardoscia et al. (2017)), financial instabilities may be structurally underestimated. Hence, comprehensive measures of the financial stability implications of interacting contagion channels like the framework developed here are highly complementary to other existing methods.

Author contributions

Garbrand Wiersema is the lead author; he designed the theoretical framework, carried out the numerical analysis, and wrote the supplementary materials. All authors contributed to all aspects of the article.

Declaration of competing interest

The authors declare no competing interests.

Acknowledgments

The authors thank Rama Cont and Christoph Reisinger for valuable comments and suggestions. This work was supported by Baillie Gifford, the Institute for New Eco- nomic Thinking at the Oxford Martin School, the Oxford-Man Institute of Quantita- tive Finance, the Clarendon Fund at the University of Oxford, the Rebuilding Macroe- conomics grant from the Economics and Social Reserach Council, and a grant from Robert Thornton. The support of the Economic and Social Research Council (ESRC) is gratefully acknowledged, via the Rebuilding Macroeconomics Network (Grant Ref: ES/R00787X/1).

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.jbankfin.2022.106684

References

Acemoglu, D., Carvalho, V.M., Ozdaglar, A., Tahbaz-Salehi, A., 2012. The network origins of aggregate fluctuations. Econometrica 80 (5), 1977–2016.

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi, A., 2015. Systemic risk and stability in financial networks. Am. Econ. Rev. 105 (2), 564–608.

Acharya, V.V., Skeie, D., 2011. A model of liquidity hoarding and term premia in inter-bank markets. J. Monet. Econ. 58 (5), 436–447.

Acharya, V.V., Yorulmazer, T., 2008. Information contagion and bank herding. J. Money Credit Bank. 40 (1), 215–231.

Adrian, T., Shin, H.S., 2010. Liquidity and leverage. J. Financ. Intermediat. 19 (3), 418–437.

Adrian, T., Shin, H.S., 2014. Procyclical leverage and value-at-risk. Rev. Financ. Stud. 27 (2), 373–403.

Aharony, J., Swary, I., 1996. Additional evidence on the information-based contagion effects of bank failures. J. Bank. Finance 20 (1), 57–69.

Allen, F., Babus, A., 2009. Networks in finance. The Network Challenge: strategy, profit, and risk in an interlinked world 367.

Allen, F., Gale, D., 2000. Financial contagion. J. Polit. Econ. 108 (1), 1–33.

Amini, H., Cont, R., Minca, A., 2016. Resilience to contagion in financial networks. Math. Finance 26 (2), 329–365.

Anderson, R., Dannielson, J., Baba, C., Das, U.S., Kang, H., Segoviano, M., 2018. Macroprudential Stress Tests and Policies: Searching for Robust and Implementable Frameworks. Technical Report. IMF Working Paper No. 18/197.

Arinaminpathy, N., Kapadia, S., May, R.M., 2012. Size and complexity in model financial systems. Proc. Natl. Acad. Sci. 109 (45), 18338–18343.

Aymanns, C., Caccioli, F., Farmer, J.D., Tan, V.W., 2016. Taming the basel leverage cycle. J. Financ. Stab. 27, 263–277.

Aymanns, C., Farmer, J.D., 2015. The dynamics of the leverage cycle. J. Econ. Dyn. Control 50, 155–179.

Aymanns, C., Farmer, J.D., Kleinnijenhuis, A.M., Wetzer, T., 2018. Models of financial stability and their application in stress tests. Handbook Comput. Econ. 4, 329–391.

Bagehot, W., 1873. Lombard Street: A description of the money market. London: HS King.

Bardoscia, M., Battiston, S., Caccioli, F., Caldarelli, G., 2015. Debtrank: a microscopic foundation for shock propagation. PloS One 10 (6), e0130406.

Bardoscia, M., Battiston, S., Caccioli, F., Caldarelli, G., 2017. Pathways towards instability in financial networks. Nature Commun. 8, 14416.

Bardoscia, M., Bianconi, G., Ferrara, G., 2018. Multiplex network analysis of the UK OTC derivatives market. Technical Report. Bank of England Working Paper No. 726.

Battiston, S., Puliga, M., Kaushik, R., Tasca, P., Caldarelli, G., 2012. Debtrank: too central to fail? Financial networks, the fed and systemic risk. Sci. Rep. 2, 541.

BIS, 2009. Strengthening the resilience of the banking sector. Technical Report. Bank of International Settlements.

BIS, 2013. Basel III: the Liquidity Coverage Ratio and liquidity risk monitoring tools. Technical Report. Bank of International Settlements.

BIS, 2014. Basel III: the net stable funding ratio. Technical Report. Bank of International Settlements.

Bookstaber, R., 2017. Agent-based models for financial crises. Ann. Rev. Financ. Econ. 9, 85–100.

Borio, C., Drehmann, M., Tsatsaronis, K., 2014. Stress-testing macro stress testing: does it live up to expectations? J. Financ. Stab. 12, 3–15.

Boss, M., Elsinger, H., Summer, M., Thurner 4, S., 2004. Network topology of the interbank market. Quant. Finance 4 (6), 677–684.

Brunnermeier, M.K., 2009. Deciphering the liquidity and credit crunch 2007-2008. J. Econ. Perspect. 23 (1), 77–100.

Brunnermeier, M.K., Pedersen, L.H., 2009. Market liquidity and funding liquidity. Rev. Financ. Stud. 22 (6), 2201–2238.

Budnik, K.B., Balatti, M., Covi, G., Dimitrov, I., Groß, J., Hansen, I., Kleemann, M., Reichenbachas, T., Sanna, F., Sarychev, A., et al., 2019. Macroprudential stress test of the euro area banking system. Technical Report. ECB Occasional Paper Series.

Burrows, O., Learmonth, D., McKeown, J., Williams, R., 2012. Ramsi: a top-down stress-testing model developed at the bank of england. Bank England Q. Bull. Q3.

Caccioli, F., Catanach, T.A., Farmer, J.D., 2012. Heterogeneity, correlations and financial contagion. Adv. Complex Syst. 15 (supp02), 1250058.

Caccioli, F., Farmer, J. D., Foti, N., Rockmore, D., 2013. How interbank lending amplifies overlapping portfolio contagion: a case study of the austrian banking network.

Caccioli, F., Farmer, J.D., Foti, N., Rockmore, D., 2015. Overlapping portfolios, contagion, and financial stability. J. Econ. Dyn. Control 51, 50–63.

Caccioli, F., Shrestha, M., Moore, C., Farmer, J.D., 2014. Stability analysis of financial contagion due to overlapping portfolios. J. Bank. Finance 46, 233–245.

Cont, R., Moussa, A., et al., 2010. Network structure and systemic risk in banking systems. Edson Bastos e, Network Structure and Systemic Risk in Banking Systems (December 1, 2010).

Cont, R., Moussa, A., et al., 2013. Network structure and systemic risk in banking systems. In: Fouque, J.-P., Langsam, J.A. (Eds.), Handbook on Systemic Risk. Cambridge University Press, Cambridge, pp. 327–336.

Cont, R., Schaanning, E., 2017. Fire sales, indirect contagion and systemic stress testing. Technical Report. Norges Bank Working Paper No. 2017/2.

Cont, R., Schaanning, E., 2019. Monitoring indirect contagion. J. Bank. Finance 104, 85–102.

Craig, B., Von Peter, G., 2014. Interbank tiering and money center banks. J. Financ. Intermediat. 23 (3), 322–347.

Danielsson, J., Shin, H.S., 2003. Endogenous risk. Modern Risk Manag.: History 297–316.

Detering, N., Meyer-Brandis, T., Panagiotou, K., Ritter, D., 2021. An integrated model for fire sales and default contagion. Math. Financ. Econ. 15 (1), 59–101.

Diamond, D.W., Dybvig, P.H., 1983. Bank runs, deposit insurance, and liquidity. J. Polit. Econ. 91 (3), 401–419.

Duarte, F., Eisenbach, T.M., 2018. Fire-sale spillovers and systemic risk. FRB of New York Staff Report (645).

EBA, 2018. Adverse macro-financial scenario for the 2018 EU-wide banking sector stress test. Technical Report. European Banking Authority.

ECB, 2019. Eurosystem Supervisory Data. https://www.bankingsupervision.europa.eu/banking/statistics. [Online; accessed 10-April-2019].

Eisenberg, L., Noe, T.H., 2001. Systemic risk in financial systems. Manag. Sci. 47 (2), 236–249.

Elliott, M., Golub, B., Jackson, M.O., 2014. Financial networks and contagion. Am. Econ. Rev. 104 (10), 3115–3153.

Elsinger, H., Lehar, A., Summer, M., 2006. Risk assessment for banking systems. Manag. Sci. 52 (9), 1301–1314.

Farmer, J.D., Kleinnijenhuis, A.M., Nahai-Williamson, P., Wetzer, T., 2020. Foundations of system-wide financial stress testing with heterogeneous institutions. Bank of England Staff Working Paper No. 861.

Farmer, J.D., Kleinnijenhuis, A.M., Wetzer, T., 2021. Stress testing the financial macrocosm. Forthcoming Handbook Financ. Stress Test..

FED, 2018. Supervisory Scenarios for Annual Stress Tests Required under the Dodd-Frank Act Stress Testing Rules and the Capital Plan Rule 2018. Technical Report. Board of Governors of the Federal Reserve System.

Fostel, A., Geanakoplos, J., 2008. Leverage cycles and the anxious economy. Am. Econ. Rev. 98 (4), 1211–1244.

Fouque, J.-P., Langsam, J.A., 2013. Handbook on systemic risk. Cambridge University Press.

Fricke, D., Lux, T., 2015. Core–periphery structure in the overnight money market: evidence from the e-mid trading platform. Comput. Econ. 45 (3), 359–395.

FSB, 2017. Basel III: Finalising post-crisis reforms. Technical Report. Financial Stability Board.

Furfine, C.H., 2003. Interbank exposures: Quantifying the risk of contagion. J. Money Credit Bank. 111–128.

Gai, P., Haldane, A., Kapadia, S., 2011. Complexity, concentration and contagion. J. Monet. Econ. 58 (5), 453–470.

Gai, P., Kapadia, S., 2010. Contagion in financial networks. In: Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences. The Royal Society. rspa20090410

Geanakoplos, J., 2010. The leverage cycle. NBER Macroeconomics Annual 24 (1), 1–66.

Glasserman, P., Young, H.P., 2016. Contagion in financial networks. J. Econ. Lit. 54 (3), 779–831.

Goodhart, C., 2013. Ratio controls need reconsideration. J. Financ. Stab. 9 (3), 445–450.

Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. J. Financ. Econ. 104 (3), 425–451.

Greenwood, R., Landier, A., Thesmar, D., 2015. Vulnerable banks. J. Financ. Econ. 115 (3), 471–485.

Hałaj, G., 2018. System-wide implications of funding risk. Physica A: 503, 1151–1181.

Heider, F., Hoerova, M., Holthausen, C., 2009. Liquidity hoarding and interbank market spreads: the role of counterparty risk. Technical Report. ECB Working Paper Series.

Hüser, A.-C., Hałaj, G., Kok, C., Perales, C., van der Kraaij, A., 2018. The systemic implications of bail-in: a multi-layered network approach. J. Financ. Stab. 38, 81–97.

IMF, 2014. Canada Financial System Stability Assessment. Technical Report. International Monetary Fund.

Kok, C., Montagna, M., 2016. Multi-layered interbank model for assessing systemic Risk. Technical Report. ECB Working Paper No. 1944.

Krishnamurthy, A., 2010. Amplification mechanisms in liquidity crises. Am. Econ. J.: Macroecon. 2 (3), 1–30.

Montagna, M., Torri, G., Covi, G., 2021. Staff working paper no. 906 on the origin of systemic risk.

Newman, M., 2018. Networks. Oxford university press.

Poledna, S., Molina-Borboa, J.L., Martínez-Jaramillo, S., Van Der Leij, M., Thurner, S., 2015. The multi-layer network nature of systemic risk and its implications for the costs of financial crises. J. Financ. Stab. 20, 70–81.

Rochet, J.-C., Tirole, J., 1996. Interbank lending and systemic risk. J. Money Credit Bank. 28 (4), 733–762.

Roukny, T., Bersini, H., Pirotte, H., Caldarelli, G., Battiston, S., 2013. Default cascades in complex networks: topology and systemic risk. Sci. Rep. 3, 2759.

Thurner, S., Farmer, J.D., Geanakoplos, J., 2012. Leverage causes fat tails and clustered volatility. Quant. Finance 12 (5), 695–707.