

DOUBLE CASCADE MODEL OF FINANCIAL CRISES

T. R. HURD

*Mathematics & Statistics, McMaster University, 1280 Main St. West
Hamilton, Ontario, L8S4L8, Canada
hurd@mcmaster.ca*

DAVIDE CELLAI

*MACSI, Department of Mathematics and Statistics, University of Limerick
Limerick, V94 T9PX, Ireland
davide.cellai@gmail.com*

SERGEY MELNIK

*MACSI, Department of Mathematics and Statistics, University of Limerick
Limerick, V94 T9PX, Ireland
sergey.melnik@ul.ie*

QUENTIN H. SHAO

*Mathematics & Statistics, McMaster University, 1280 Main St. West
Hamilton, Ontario, L8S4L8, Canada
shaoq@math.mcmaster.ca*

Received (day,month,year)

Revised (day,month,year)

Abstract

The scope of financial systemic risk research encompasses a wide range of interbank channels and effects, including asset correlation shocks, default contagion, illiquidity contagion, and asset fire sales. This paper introduces a financial network model that combines the default and liquidity stress mechanisms into a “double cascade mapping”. The progress and eventual result of the crisis is obtained by iterating this mapping to its fixed point. Unlike simpler models, this model can therefore quantify how illiquidity or default of one bank influences the overall level of liquidity stress and default in the system. Large-network asymptotic cascade mapping formulas are derived that can be used for efficient network computations of the double cascade. Numerical experiments then demonstrate that these asymptotic formulas agree qualitatively with Monte Carlo results for large finite networks, and quantitatively except when the initial system is placed in an exceptional “knife-edge” configuration. The experiments clearly support the main conclusion that when banks respond to liquidity stress by hoarding liquidity, then in the absence of asset fire sales, the level of defaults in a financial network is negatively related to the strength of bank liquidity hoarding and the eventual level of stress in the network.

Keywords: Systemic risk, banking network, contagion, random graph, default, funding liquidity, liquidity hoarding.

MSC: 05C80, 91B30, 91B70, 91G40

1 Introduction

Since the banking crisis of 2007-2008, the types of shocks transmitted through interbank networks during a crisis now thought to be important include not only shocks arising from defaulting banks, but a number of additional phenomena, most notably asset shocks originating with forced sales by banks of illiquid assets and funding liquidity shocks as illiquid banks recall loans from other banks. A comprehensive treatment of these contagion effects and how they can be computed in stylized financial network specifications can be found in Hurd (2016).

A well-developed strand of literature on bank default cascade models, starting with Eisenberg & Noe (2001) and reviewed in Upper (2011), is based on a network of banks wherein insolvency of a given bank, defined as a bank whose net worth becomes non-positive, will generate shocks to the asset side of the balance sheet of each of its creditor banks. Under some circumstances, such “downstream” shocks can cause further insolvencies that may build up to create a global insolvency cascade. One such contribution is by Nier, Yang, Yorulmazer & Alentorn (2007) that uses Monte Carlo simulation to determine how the key network parameters for their model of defaults in a stylized random network of 25 banks can influence the total number of defaults in a nonlinear, indeed sometimes nonmonotonic, fashion. The paper of Gai & Kapadia (2010a) and its extension by Hurd & Gleeson (2011), adapt the Watts (2002) network model of information cascades to the context of financial systems, deriving both analytical and Monte Carlo results for the dependence of the default cascade on different structural parameters. May & Arinaminpathy (2010) present approximate analytical formulas for the Nier et al model and Gai-Kapadia model that can explain some of the main properties of the simulation results found in those papers. Amini, Cont & Minca (2012) and Amini, Cont & Minca (2016) develop a simple but general analytical criterion for resilience to default contagion, based on an asymptotic analysis of default cascades in heterogeneous networks.

More recently, after remarking on the observed “freezing” of interbank lending around the time of the Lehman collapse, papers by Gai & Kapadia (2010b), Gai, Haldane & Kapadia (2011) and Lee (2013) adopt variations of an idea that funding illiquidity^a of a bank can also be transmitted contagiously through interbank exposures. They argue that a bank whose liquid assets are insufficient to cover demands on its liabilities will be illiquid or “stressed”, and must reduce its interbank lending, thereby creating shocks to the liability side of its debtor banks’ balance sheets. Such “upstream” shocks can cause illiquidity stress in other banks that in some circumstances may build up to create a global illiquidity cascade.

A third channel, sometimes called market illiquidity or the asset fire sale effect, is identified by Cifuentes, Ferrucci & Shin (2005) who consider a network of banks

^aFunding illiquidity, being the insufficiency of liquid assets to cover a run on liabilities, is distinct from market illiquidity, where assets become difficult to sell due to an oversupply in the market. See Tirole (2010) and Brunnermeier & Pedersen (2009) for a detailed analysis of these concepts.

that hold a common asset. As stress in the system builds up, some banks are forced to sell large positions in the illiquid asset, leading to a downward price spiral coupled to worsening bank balance sheets across the network. This mechanism has been extended to the sharing of multiple assets in a paper by Caccioli, Shrestha, Moore & Farmer (2014).

These papers simplify by omitting economic details that obscure their focus on a single channel of systemic risk. They share the view that the end result of a crisis is a new equilibrium that is reached through a sequence of mechanistic steps where banks update and modify their balance sheets in response to price changes and shocks transmitted from other banks. Usually, properties of the system at step n of the cascade can be computed in terms of the properties at step $n - 1$, yielding a set of equations that can be called a *cascade mapping*. In some models, an analytic cascade mapping can be found that yields a detailed picture of systemic risk without the need for Monte Carlo simulation based computations. Even in such cases however, large scale benchmarking of the analytic approximation method against Monte Carlo simulations is undertaken to validate its use.

An important question now arises whether one can successfully integrate two or more such cascade mechanisms, incorporate “spillover” effects in a realistic, analytically tractable way, and discern new features of systemic risk. The purpose of the present paper is to do exactly this, by introducing a network model of a double illiquidity-insolvency cascade, and deriving an analytic cascade mapping that describes it. This double cascade model, shown schematically in Figure 1, can resolve issues not addressable within a single cascade model, for example, to quantify the effect of a bank’s behavioural response to liquidity stress on the level of eventual defaults in the entire system. In particular, one can show that a bank that reacts to stress at a time of crisis by shrinking its own interbank lending, thereby inflicting liquidity shocks to its debtor banks, will also protect itself from default due to interbank default shocks.

In the literature on how contagion channels “spill over” into each other, Diamond & Rajan (2005) explain the liquidity, solvency and behaviour of banks through the role they play in intermediating between lenders and borrowers. As discussed in Tirole (2010), the holding of very illiquid assets can cause solvent banks (those with positive equity) to default due to their inability to raise enough liquidity to meet short-term needs. Battiston, Gatti, Gallegati, Greenwald & Stiglitz (2012) model the robustness dynamics of a network of banks linked by their interbank exposures, highlighting the feedback effects of a “financial accelerator” mechanism on the level of systemic risk. Roukny, Bersini, Pirotte, Caldarelli & Battiston (2013) further develop this idea, investigating the systemic impact of network topologies and different levels of market illiquidity and applying this methodology to the Italian interbank money market from 1999 to 2011.

It should be made clear that in order to focus on pure contagion effects, the double cascade model and the above models simplify by explicitly ruling out certain other systemic mechanisms that have been explored in the economics and finance

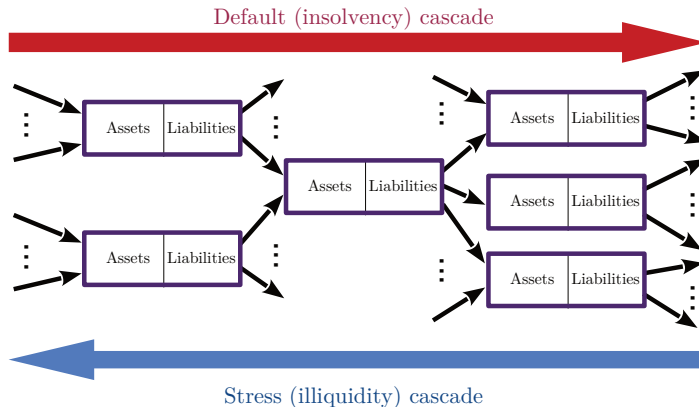


Figure 1: A schematic diagram of part of a financial network, showing banks and their balance sheets connected by directed edges representing interbank exposures. Default shocks are transmitted in the forward “downstream” direction, while liquidity or stress shocks are transmitted in the upstream direction.

literature. In the present paper, the most important economic effect that is ruled out is the asset “fire sale”, which means that for the duration of the crisis, changing values of fixed assets are ignored. Some of the above references discuss the impact of financial cascades on the non-financial economy and the consequent feedback into the financial markets. They find that fire sales of assets amplify any cascade that takes hold in the network. It will be an important future challenge to extend our double cascade framework to include the fire sale effect and determine whether our main conclusions are robust to such an extension.

Section 2 of the paper sets out the network framework and assumptions underlying the balance sheet structure of banks, the timing of the crisis and bank behaviour. These assumptions lead to stock flow consistent rules for the transmission of shocks through the double stress/default cascade, including the conditions for banks to become stressed or defaulted. Section 3 develops our main technical contribution, which yields an explicit analytic cascade mapping for default and stress probabilities valid in a large heterogeneous network of banks with random balance sheets and interbank exposures and connectivity given by a random “skeleton”. Section 4 provides a parallel development of default and stress probabilities for cascades on networks where it is assumed that the skeleton of interbank connections is known explicitly, but balance sheets and exposure sizes are random. Several representative financial experiments are reported in Section 5. First we summarize experiments that validate the accuracy of the cascade mapping by direct comparison of large network analytics to Monte Carlo simulation results. Secondly, we investigate the relationship between the level of stress and default, verifying our assertion that average default probability decreases as average stress probability increases. A final

experiment demonstrates that our analysis is still useful when pushed to a highly heterogeneous specification of the model that is consistent with known heuristics of financial networks, such as fat-tailed degree and exposure distributions, and the 2011 stress testing data on 90 large banks in the EU system. We observed that the stylized EU network is very resilient, and only extremely large shocks to the average default buffer size will trigger a cascade of defaulted and stressed banks.

One conclusion of this paper is that our analytic asymptotic results on default and stress probabilities are broadly consistent with results from Monte Carlo simulations on random networks. A second conclusion is that under the assumption of no asset fire sales, stress and default are inversely related: as banks respond to stress more vigorously, creating more network stress, they protect the network from default. It is left for future work to determine the extent to which this conclusion is valid when a fire sale mechanism provides an additional channel causing stressed but solvent banks to default.

2 Cascade Mechanisms

The assumptions we now make lead to a stylized network model admitting double cascades of illiquidity and insolvency that can be computed efficiently under a wide range of initial conditions. The picture to adopt is that of a system hit “out of the blue” by a large shock that triggers a systemic crisis that proceeds stepwise in time. Many realistic effects that would significantly complicate our model are left out. For example, banks respond mechanistically according to simple stylized rules rather than dynamically optimizing their behaviour. We rule out other channels of systemic risk, in particular asset fire sales. We take a static view, and ignore external cash flows, interest payments and price changes for the duration of the cascade. We make an assumption of zero-recovery on interbank assets on the default of a counterparty during the cascade. Our assumptions imply that insolvency (that is, when equity is nonpositive) is equivalent to default (that is, the bank can no longer pay some liability). This means that even when stressed, banks always pay recalled interbank liabilities in full, until the moment they become insolvent.

The network consists of a collection of N banks, labelled by numbers from the set $\mathcal{N} = \{1, 2, \dots, N\}$, each structured with a balance sheet as shown in Figure 2.

The subset of debtor banks of a bank $v \in \mathcal{N}$ is called the in-neighbourhood of v , denoted by \mathcal{N}_v^- , and the *in-degree* is the number of debtors $j_v = |\mathcal{N}_v^-|$. Similarly the bank’s creditor banks form a subset called the out-neighbourhood \mathcal{N}_v^+ of v whose size $k_v = |\mathcal{N}_v^+|$ is called the *out-degree*. The web of interbank counterparties is called the “skeleton”, and is identifiable as a directed graph, i.e. a collection of directed arrows, called edges, between pairs of nodes. Each debtor-creditor pair v, w with $w \in \mathcal{N}_v^+$ is denoted by an edge $\ell = (vw)$ pointing from v to w .^b The

^bThe convention that arrows point from debtors to creditors means that default shocks propagate in the downstream direction. Confusingly, much of the systemic risk literature uses the




Figure 2: The stylized balance sheet of a bank v with in-degree $j_v = 3$ and out-degree $k_v = 2$. Banks w_1, w_2, w_3 are debtors of v while w'_1, w'_2 are its creditors. The total exposure of v to w_1 is denoted with $\Omega_{w_1 v}$, and so on. The default buffer Δ of bank v is the difference between assets and liabilities, and the “stress buffer” Σ is the preferred asset class from which it pays creditors’ demands.

type (j, k) of a bank v is its in-degree $j = j_v$ and out-degree $k = k_v$, and we write $v \in \mathcal{N}_{jk}$. The *type* $(k, j) = (k_\ell, j_\ell)$ of an edge $\ell = (vw)$ is the out-degree $k_\ell = k_v$ of the debtor bank v and the in-degree $j_\ell = j_w$ of the creditor bank w , and we write $\ell \in \mathcal{E}_{kj}$.

The bank equity or net worth $\Delta_v = \max(A_v^F + A_v^{IB} + A_v^L - L_v^D - L_v^{IB}, 0)$ of bank v incorporates *limited liability*, which assumes that the bank must default and be forced to liquidate at the first moment it becomes insolvent, that is, its liabilities exceed its assets. Thus Δ_v can be interpreted as the bank’s *default buffer* that must always be kept positive to avoid default.

Banks are also concerned about the possibility of runs on their liabilities, and try to keep a positive buffer of liquid assets A^L (such as cash and government bonds), henceforth called the *stress buffer* Σ , which are the preferred asset class from which to pay creditors’ demands. A bank v with Σ_v fully depleted to zero will be called *stressed*. This does not mean the bank is *illiquid* in the sense of being unable to meet the creditors’ demands. Rather it means that the bank is experiencing a significant degree of stress in meeting these demands and must turn to other assets, first interbank assets, then fixed assets, to realize the needed cash. The no fire sale assumption means that as long as they are solvent, stressed banks always meet their payment obligations.

The interbank liabilities L^{IB} and assets A^{IB} decompose into bilateral interbank reverse convention that arrows point from creditors to debtors.

exposures. For any bank v and one of its creditors $w \in \mathcal{N}_v^+$, we denote by Ω_{vw} the total exposure of w to v . Then we have the constraints

$$A_v^{IB} = \sum_{w \in \mathcal{N}_v^-} \Omega_{wv}; \quad L_v^{IB} = \sum_{w \in \mathcal{N}_v^+} \Omega_{vw}. \quad (2.1)$$

Prior to the onset of the crisis, all banks are assumed to be in the normal state. Then, on day $n = 0$, a collection of banks, possibly all banks, are assumed to experience initial shocks. Two kinds of initial shocks are possible. First, an asset shock causes a drop in the mark-to-market value of the fixed asset portfolio, and reduces the default buffer. If the downward asset shock leads to $\Delta \leq 0$, the bank must default. The second kind of shock is a demand shock by the external depositors that reduces the stress buffer. If the demand shock leads to $\Sigma \leq 0$, the bank becomes stressed.

Timing and Bank Behaviour Assumptions:^c

1. Prior to the crisis, all banks are in the normal state, neither stressed nor insolvent. The crisis commences on day 0 after initial shocks trigger the default or stress of one or more banks;
2. Balance sheets are recomputed daily on a mark-to-market basis and banks respond daily on the basis of their newly computed balance sheets. All external cash flows, interest payments, and asset and liability price changes are ignored throughout the crisis.
3. An insolvent bank, characterized by $\Delta = 0$, is forced into liquidation by the regulator. At this moment, each of its creditor banks are obliged to write down its defaulted exposures to zero thereby experiencing a *solvency shock* that reduces its default buffer Δ .
4. A stressed bank, defined to be a non-defaulted bank with $\Sigma = 0$, reacts one time only, at the moment it becomes stressed, by reducing its interbank assets A^{IB} to $(1-\lambda)A^{IB}$ where λ taken to be a fixed constant across all banks during the crisis. It does so by recalling a fraction λ of its interbank loans, thereby transmitting a *stress shock* to the liabilities each of its debtor banks. Since nondefaulted banks are able to pay all of their liabilities, such recalled loans are repaid in full.
5. A newly defaulted bank also triggers maximal stress shocks (i.e. with $\lambda = 1$) to each of its debtor banks as its bankruptcy trustees recall all its interbank loans, reducing A^{IB} to 0;
6. Stress shocks reduce any bank's stress buffer Σ and stressed banks remain stressed until the end of the cascade or until they default.

^cTreating cascade steps as daily is simply an aid to understanding.

Remark 1. The stress parameter λ was introduced by Gai, Haldane & Kapadia (2011) to simplify banks' response to liquidity stress and to capture the effect of liquidity hoarding. We suppose that banks react sufficiently strongly to stress by choosing $\lambda \in [0.4, 1]$, to capture the picture that they hoard liquidity to preempt the need for further response later in the crisis. Unlike default, a bank is free to set its response to liquidity: both λ and the size of the stress buffer Σ are its own policy decisions. In a more realistic and complex model, λ would be endogenously determined for each bank, reflecting the cumulative demands on its liabilities.

The dynamics of the cascade that follow from these assumptions is stock flow consistent. This means that as long as both counterparties are solvent, every transaction involves four equal and opposite balance sheet adjustments: to both the assets and liabilities of both counterparties. The dynamics also has the property that it is completely determined by a reduced set of balance sheet data that consists of the collection of buffers Δ, Σ and interbank exposures Ω .

We shall now apply the cascade rules to a financial network found on day 0 of the crisis in an initial state described by the following elements: the interbank links which form a directed graph \mathcal{E} on the set of banks $v \in \mathcal{N}$, called the *skeleton*; the buffers Δ_v, Σ_v for nodes $v \in \mathcal{N}$; and the exposures Ω_{vw} for edges $\ell = (vw) \in \mathcal{E}$. After n cascade steps, we identify \mathcal{D}_n , the set of defaulted banks, \mathcal{S}_n , the set of undefaulted banks that are under stress after n steps, and $\mathcal{D}_n^c \cap \mathcal{S}_n^c$ which contains the remaining undefaulted, unstressed banks.^d In our model, banks do not recover from either default or stress during the crisis, so the sequences $\{\mathcal{D}_n\}_{n \in \mathbb{N}}$ and $\{\mathcal{D}_n \cup \mathcal{S}_n\}_{n \in \mathbb{N}}$ are non-decreasing.

To say that a bank v is defaulted at step n means that default shocks to step $n - 1$ exceed its default buffer:

$$\mathcal{D}_n := \begin{cases} \{\Delta_v = 0\} & \text{for } n = 0 \\ \left\{ v \mid \sum_{w \in \mathcal{N}_v^-} \Omega_{vw} \xi_{vw}^{(n-1)} \geq \Delta_v \right\} & \text{for } n \geq 1, \end{cases} \quad (2.2)$$

where the variables ξ indicate the fractional sizes of the various default shocks impacting v . Similarly, for any upstream neighbour $w \in \mathcal{N}_v^-$, we say that v is defaulted at step n *without regarding* w and write $v \in \mathcal{D}_n \textcircled{R} w$ under the conditions

$$\mathcal{D}_n \textcircled{R} w := \begin{cases} \{\Delta_v = 0\} \cap \mathcal{N}_w^+ & \text{for } n = 0 \\ \left\{ v \in \mathcal{N}_w^+ \mid \sum_{v' \in \mathcal{N}_v^- \setminus w} \Omega_{v'v} \xi_{v'v}^{(n-1)} \geq \Delta_v \right\} & \text{for } n \geq 1. \end{cases} \quad (2.3)$$

Finally, to say v is stressed at step n means both that it is not yet defaulted and

^dFor any set \mathcal{B} , \mathcal{B}^c denotes its complement. For an *event* defined by some condition P , for example $v \in \mathcal{D}_n$, the indicator function for that event is written $\mathbb{1}(P)$.

the stress shocks to step $n - 1$ exceed the stress buffer, i.e. $\mathcal{S}_n = \mathcal{D}_n^c \cap \hat{\mathcal{S}}_n$ where

$$\hat{\mathcal{S}}_n := \begin{cases} \{\Sigma_v = 0\} & \text{for } n = 0 \\ \left\{ v \mid \sum_{w \in \mathcal{N}_v^+} \Omega_{vw} \zeta_{vw}^{(n-1)} \geq \Sigma_v \right\} & \text{for } n \geq 1, \end{cases} \quad (2.4)$$

where the variables ζ indicate the fractional sizes of the stress shocks impacting v .

Accounting for the fact that at the moment when v becomes stressed it reduces its interbank exposures, one can see that for $n \geq 1$ the fractions $\xi^{(n)}$ are given recursively by

$$\xi_{vw}^{(n)} := \begin{cases} \xi_{vw}^{(n-1)} & \text{when } w \in \mathcal{D}_{n-1} \cup \mathcal{D}_n^c \\ 1 & \text{when } w \in \mathcal{D}_n \setminus \mathcal{D}_{n-1} \text{ and } v \in \hat{\mathcal{S}}_{n-1}^c \\ 1 - \lambda & \text{when } w \in \mathcal{D}_n \setminus \mathcal{D}_{n-1} \text{ and } v \in \hat{\mathcal{S}}_{n-1}, \end{cases} \quad (2.5)$$

with the initial values $\xi_{vw}^{(0)} = \mathbb{1}(w \in \mathcal{D}_0)$. Similarly, accounting for the assumption that defaulted creditors transmit a maximal stress impact, whereas stressed creditors only transmit a stress shock of a fraction λ of the interbank exposure, one has for $n \geq 0$

$$\zeta_{vw}^{(n)} := \begin{cases} 0 & \text{when } w \in \hat{\mathcal{S}}_n^c \cap (\mathcal{D}_n^c \textcircled{R} v) \\ \lambda & \text{when } w \in \hat{\mathcal{S}}_n \cap (\mathcal{D}_n^c \textcircled{R} v) \\ 1 & \text{when } w \in \mathcal{D}_n \textcircled{R} v. \end{cases} \quad (2.6)$$

The reader will be curious to see the \textcircled{R} condition in (2.6). The rationale is that the \textcircled{R} condition explicitly eliminates an apparent feedback where the stress shock to v from w at step n seems to depend on the default state of v at step $n - 2$. This feedback is apparent, not real, because any stress shock to a defaulted bank is inconsequential.

The above prescription completely characterizes the cascade mapping of the system to itself. This mapping arrives at the fixed point that represents the end of the crisis in at most $2N$ steps when acting on a network of N banks. The next two sections are devoted to the derivation of two probabilistic versions of this double cascade mapping: first for the case of networks with random skeletons, and second for the case of networks with random balance sheets on a fixed skeleton.

3 Networks with Random Skeletons

Our model of a random financial network has three layers of structure: the skeleton (a random directed graph $(\mathcal{N}, \mathcal{E})$); the buffer random variables Δ_v, Σ_v , $v \in \mathcal{N}$ and the random exposures $\Omega_\ell, \ell \in \mathcal{E}$. We give a complete specification of the distributional properties of these random variables.

The skeleton $(\mathcal{N}, \mathcal{E})$ is a random directed assortative configuration graph (ACG) of the type studied by Hurd (2016), generalizing the well-known undirected configuration random graph model introduced in Bender & Canfield (1978) and Bollobás (1980). The model is parametrized by N and the node and edge type distributions

$$P_{jk} = \mathbb{P}[v \in \mathcal{N}_{jk}], \quad Q_{kj} = \mathbb{P}[\ell \in \mathcal{E}_{kj}], \quad j, k \leq K \quad (3.1)$$

where for simplicity in the following we assume in and out degrees j, k are bounded by a constant K . P and Q can be considered as bivariate distributions which have marginals:

$$P_k^+ := \sum_j P_{jk}, \quad P_j^- := \sum_k P_{jk}, \quad Q_k^+ := \sum_j Q_{kj}, \quad Q_j^- := \sum_k Q_{kj}. \quad (3.2)$$

Assortativity is defined to be the Pearson correlation of Q , considered as a bivariate probability distribution. Several studies of real financial networks, notably Bech & Atalay (2010), highlight the fact and relevance of their observed negative assortativity.

Definition 1 (The ACG Construction). The node and edge type probability laws P, Q are consistent if:

$$\begin{aligned} z &:= \sum_j j P_j^- = \sum_k k P_k^+, \\ Q_k^+ &= k P_k^+ / z, \quad Q_j^- = j P_j^- / z, \quad \text{for all } j, k \leq K. \end{aligned} \quad (3.3)$$

Given consistent data N, P, Q , a random graph $(\mathcal{N}, \mathcal{E})$ with N nodes is sampled as follows:

1. Draw a sequence of N node-type pairs $X = ((j_1, k_1), \dots, (j_N, k_N))$ independently from P , and accept the draw if and only if it is feasible, i.e. $\sum_{v \in \mathcal{N}} j_v = \sum_{v \in \mathcal{N}} k_v$, and this defines the number of edges E that will result. Label the v th node with k_v *out-stubs* (each out-stub is a half-edge with an out-arrow, labelled by its degree k_v) and j_v *in-stubs*, labelled by their degree j_v . Define the partial sums

$$u_j^- = \sum_v \mathbb{1}(j_v = j), \quad u_k^+ = \sum_v \mathbb{1}(k_v = k), \quad u_{jk} = \sum_v \mathbb{1}(j_v = j, k_v = k), \quad (3.4)$$

the number $e_k^+ = k u_k^+$ of k -stubs (out-stubs of degree k) and the number of j -stubs (in-stubs of degree j), $e_j^- = j u_j^-$.

2. Conditioned on X , the result of Step 1, choose an arbitrary ordering ℓ^- and ℓ^+ of the E in-stubs and E out-stubs. The matching sequence, or “wiring”, W of edges is selected by choosing a pair of permutations $\sigma, \tilde{\sigma} \in S(E)$ of the set \mathcal{E} . This determines the edge sequence $\ell = (\ell^- = \sigma(\ell), \ell^+ = \tilde{\sigma}(\ell))$ labelled by $\ell \in \mathcal{E}$, to which is attached a probability weighting factor

$$\prod_{\ell \in \mathcal{E}} Q_{k_{\sigma(\ell)} j_{\tilde{\sigma}(\ell)}}. \quad (3.5)$$

The ACG simulation algorithm defines the required class of random graphs, but is infeasible because the acceptance condition in Step 1 is achieved only rarely, and drawing random permutations as in Step 2 is impractical. The paper Hurd (2016) proposes an approximate simulation algorithm that appears to work well in practice.

The non-negative default buffer random variables $\Delta_v, v \in \mathcal{N}$ have point masses at $x = 0$ that represent the bank initial default probability p_v^0 . We assume that the distribution functions of Δ_v depend only on the type (j, k) , and have the following form:

$$D_{jk}(x) = \mathbb{P}[\Delta_v \leq x \mid v \in \mathcal{N}_{jk}] ; \frac{d}{dx} D_{jk}(x) := p_{jk}^0 \delta_0(x) + d_{jk}(x) . \quad (3.6)$$

where $d_{jk}(x) \geq 0$ is a specified function with $\int_0^\infty d_{jk}(x) dx = 1 - p_{jk}^0$. Similarly, the stress buffer Σ_v has a point mass at $x = 0$ that represents this bank's initial stress probability q_v^0 and a distribution function that depends only on its node type (j, k) . Thus the stress buffer distribution functions of nodes $v \in \mathcal{N}_{jk}$ have the following form:

$$S_{jk}(x) = \mathbb{P}[\Sigma_v \leq x, v \notin \mathcal{D}_0 \mid v \in \mathcal{N}_{jk}] ; \frac{d}{dx} S_{jk}(x) := q_{jk}^0 \delta_0(x) + s_{jk}(x) . \quad (3.7)$$

where $s_{jk}(x) \geq 0$ is a specified function with $\int_0^\infty s_{jk}(x) dx = 1 - p_{jk}^0 - q_{jk}^0$. The exposure random variables $\Omega_\ell, \ell \in \mathcal{E}$ are positive (i.e. there is zero probability to have a zero weight) and have distributions that depend only on the edge type (k, j) . These can be specified by the distribution functions

$$W_{kj}(x) = \mathbb{P}[\Omega_\ell \leq x \mid \ell \in \mathcal{E}_{kj}] ; \frac{d}{dx} W_{kj}(x) = w_{kj}(x) . \quad (3.8)$$

Finally, conditional on the random skeleton $(\mathcal{N}, \mathcal{E})$, the collection of random variables $\{\Delta_v, \Sigma_v, \Omega_\ell\}$ is assumed to be mutually independent.

It was proven in Hurd (2016) that the probability measure on random networks just defined has a property called locally tree-like independence (LTI) extending the well-known locally tree-like property of configuration graphs that cycles of any fixed finite length occur with an asymptotically zero density as the number of nodes N goes to infinity for fixed P, Q . The probabilistic analysis to follow rests on this extended type of independence that holds as $N \rightarrow \infty$:

The locally tree-like independence (LTI) property: Consider the double cascade model defined by the collection of random variables $(\mathcal{N}, \mathcal{E}, \Delta, \Sigma, \Omega)$. Let $\mathcal{N}_1, \mathcal{N}_2 \subset \mathcal{N}$ be any two subsets that share exactly one node $\mathcal{N}_1 \cap \mathcal{N}_2 = \{v\}$ and let X_1, X_2 be any pair of random variables where for each $i = 1, 2$, X_i is determined by the information on \mathcal{N}_i . Then X_1 and X_2 are conditionally independent, conditioned on the information $\Delta_v, \Sigma_v, j_v, k_v$ located at the node v .

The cascade mapping we now propose is based on the observation that a probability such as $\mathbb{P}[v \in \mathcal{D}^n]$ that depends on the n th order neighbourhood of the node

v can be approximated iteratively, in terms of the probabilities of the states of the first order neighbouring nodes and edges of v at step $n - 1$. The accuracy of such a scheme depends on the degree of dependence between the states of these first order neighbours of v . In ideal situations, there is no dependence at all, and the cascade mapping is exact. In our present model, there are two sources of residual dependence amongst these neighbours, and our approximation amounts to neglecting this residual dependence.

Let us define

$$\begin{aligned}
p_{jk}^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \mid v \in \mathcal{N}_{jk}] , \\
q_{jk}^{(n)} &= \mathbb{P}[v \in \mathcal{S}_n \mid v \in \mathcal{N}_{jk}] , \\
\tilde{p}_{jk}^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \textcircled{R} w \mid v \in \mathcal{N}_{jk} \cap \mathcal{N}_w^+] , \\
\hat{q}_{jk}^{(n)} &= \mathbb{P}[v \in \hat{\mathcal{S}}_n \mid v \in \mathcal{N}_{jk}] .
\end{aligned} \tag{3.9}$$

where $p_{jk}^{(n)}$ is the probability that a (j, k) -node has defaulted by time n , $q_{jk}^{(n)}$ is the probability that a (j, k) -node is stressed at time n , $\tilde{p}_{jk}^{(n)}$ is the probability that a (j, k) -node has defaulted by time n , without regarding of one of its in-neighbours, and $\hat{q}_{jk}^{(n)}$ is the probability that a (j, k) -node satisfies the stress condition (2.4) at time n . Our aim is to compute $p_{jk}^{(n)}, q_{jk}^{(n)}, \tilde{p}_{jk}^{(n)}$, plus an auxiliary quantity

$$t_{kj}^{(n)} = \mathbb{P}\left[\xi_{wv}^{(n)} = 1 \mid (w, v) \in \mathcal{E}_{kj}\right], \tag{3.10}$$

recursively over n . It is also convenient to define

$$\begin{aligned}
p_k^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \mid k_v = k] = \sum_j p_{jk}^{(n)} P_{j|k} , \\
\tilde{p}_j^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \textcircled{R} w \mid j_v = j, v \in \mathcal{N}_w^+] = \sum_k \tilde{p}_{jk}^{(n)} P_{k|j} , \\
\hat{q}_j^{(n)} &= \mathbb{P}[v \in \hat{\mathcal{S}}_n \mid j_v = j] = \sum_k \hat{q}_{jk}^{(n)} P_{k|j} ,
\end{aligned} \tag{3.11}$$

where $P_{j|k} := \frac{P_{jk}}{P_k^+}, P_{k|j} := \frac{P_{jk}}{P_j^-}$.

Our computations will rely on two facts. The first is that if X, Y are two independent random variables with probability density functions (PDFs) $f_X(x) = F_X'(x), f_Y(y) = F_Y'(y)$, then

$$\begin{aligned}
\mathbb{P}[X \geq Y] &= \mathbb{E}[\mathbb{1}(X \geq Y)] = \int_{\mathbb{R}} \int_{\mathbb{R}} \mathbb{1}(X \geq Y) f_X(x) f_Y(y) dx dy \\
&= \int_{\mathbb{R}} F_Y(x) f_X(x) dx = \langle F_Y, f_X \rangle .
\end{aligned} \tag{3.12}$$

In general, the Hermitian inner product on \mathbb{R} is defined as $\langle f, g \rangle = \int_{-\infty}^{\infty} \bar{f}(x)g(x) dx$, but here, both operands are real functions and the conjugate operator disappears.

A second fact is that if X_1, X_2, \dots, X_n are n independent random variables with PDFs f_{X_i} , then the PDF of the sum $X = X_1 + X_2 + \dots + X_n$ is the convolution

$$f_X = f_{X_1} \otimes f_{X_2} \otimes \dots \otimes f_{X_n} = \otimes_{k=1}^n f_{X_k}, \quad (3.13)$$

where the convolution product of two functions is the function defined by $(f \otimes g)(x) = \int_{\mathbb{R}} f(y)g(x-y)dy$. For convolution powers, we write $\otimes_{k=1}^n f_X = f_X^{\otimes n}$.

The following cascade mapping provides a closed set of recursive formulas for $n \geq 1$ starting with $p_{jk}^{(0)} = \tilde{p}_{jk}^{(0)}, \hat{q}_{jk}^{(0)}, t_{kj}^{(0)} = p_k^{(0)}$ determined by the initial shock and stress probabilities.

Cascade Mapping: For any $n \geq 1$, suppose $p_{jk}^{(n-1)}, \hat{q}_{jk}^{(n-1)}, \tilde{p}_{jk}^{(n-1)}, t_{kj}^{(n-1)}$ are known. The cascade mapping gives the quantities $p_{jk}^{(n)}, \hat{q}_{jk}^{(n)}, \tilde{p}_{jk}^{(n)}, t_{jk}^{(n)}$ recursively by

$$p_{jk}^{(n)} = \left\langle D_{jk}, \left(g_j^{(n-1)} \right)^{\otimes j} \right\rangle, \quad (3.14)$$

$$\tilde{p}_{jk}^{(n)} = \left\langle D_{jk}, \left(g_j^{(n-1)} \right)^{\otimes j-1} \right\rangle, \quad (3.15)$$

$$\hat{q}_{jk}^{(n)} = \left\langle S_{jk}, \left(h_k^{(n-1)} \right)^{\otimes k} \right\rangle, \quad (3.16)$$

$$t_{kj}^{(n)} = t_{kj}^{(n-1)} + (p_k^{(n)} - p_k^{(n-1)})(1 - \hat{q}_j^{(n-1)}), \quad (3.17)$$

where D_{jk} and S_{jk} are defined in (3.6) and (3.7), respectively, and we use (3.11) to compute $p_{k'}^{(n-1)}, \hat{q}_{j'}^{(n-1)}, \tilde{p}_{j'}^{(n-1)}$. Moreover, the stress probabilities $q_{jk}^{(n)}$ are determined by

$$\begin{aligned} 1 - q_{jk}^{(n)} - p_{jk}^{(n)} &= \mathbb{P}[v \in \hat{\mathcal{S}}_n^c \cap \mathcal{D}_n^c \mid v \in \mathcal{N}_{jk}] \\ &= \left(1 - \hat{q}_{jk}^{(n)} \right) \left(1 - \left\langle D_{jk}, \left(\tilde{g}_j^{(n-1)} \right)^{\otimes j} \right\rangle \right). \end{aligned} \quad (3.18)$$

The probability distribution functions in these formulas are also computed recursively:

$$\begin{aligned} g_j^{(n-1)}(x) &= \sum_{k'} \left[(1 - p_{k'}^{(n-1)})\delta_0(x) + t_{k'j}^{(n-1)}w_{k'j}(x) \right. \\ &\quad \left. + (p_{k'}^{(n-1)} - t_{k'j}^{(n-1)}) \cdot \frac{1}{1-\lambda} w_{k'j}(x/(1-\lambda)) \right] \cdot Q_{k'|j}, \end{aligned} \quad (3.19)$$

$$\begin{aligned} h_k^{(n-1)}(x) &= \sum_{j'} \left[(1 - \hat{q}_{j'}^{(n-1)})(1 - \tilde{p}_{j'}^{(n-1)})\delta_0(x) + \tilde{p}_{j'}^{(n-1)}w_{kj'}(x) \right. \\ &\quad \left. + \hat{q}_{j'}^{(n-1)}(1 - \tilde{p}_{j'}^{(n-1)}) \cdot \frac{1}{\lambda} w_{kj'}(x/\lambda) \right] \cdot Q_{j'|k}, \end{aligned} \quad (3.20)$$

$$\tilde{g}_j^{(n-1)}(x) = \sum_{k'} \left[(1 - p_{k'}^{(n-1)})\delta_0(x) + p_{k'}^{(n-1)}w_{k'j}(x) \right] \cdot Q_{k'|j} \quad (3.21)$$

with $Q_{k|j} = \frac{Q_{kj}}{Q_j}, Q_{j|k} = \frac{Q_{kj}}{Q_k}$.

Justification: To justify (3.14) we need to suppose that the collection of random variables $\Omega_{wv}\xi_{wv}^{(n-1)}$ for different $w \in \mathcal{N}_v^-$ are mutually conditionally independent. However, there are two sources of dependence: the first is the usual breaking of the LTI property for finite N due to cycles in the skeleton, the second, specific to this model, is the dependence on whether v is in \mathcal{S}_{n-1} . Our approximation is to neglect both sources of dependence, allowing the use of (3.12) to give

$$\begin{aligned} p_{jk}^{(n)} &= \mathbb{P}[\Delta_v \leq \sum_{w \in \mathcal{N}_v^-} \Omega_{wv} \xi_{wv}^{(n-1)} \mid v \in \mathcal{N}_{jk}] \\ &= \left\langle D_{jk}, \left(g_j^{(n-1)} \right)^{\otimes j} \right\rangle. \end{aligned} \quad (3.22)$$

where

$$g_j^{(n-1)}(x) = \sum_{k'} Q_{k'|j} \frac{d}{dx} \mathbb{P}[\Omega_{wv} \xi_{wv}^{(n-1)} \leq x \mid v \in \mathcal{N}_{jk}, w \in \mathcal{N}_v^-, k_w = k'] . \quad (3.23)$$

Under the conditions $v \in \mathcal{N}_{jk}, w \in \mathcal{N}_v^-, k_w = k'$, the events $\{\xi_{wv}^{(n-1)} = 0\}, \{\xi_{wv}^{(n-1)} = 1\}, \{\xi_{wv}^{(n-1)} = 1 - \lambda\}$ have conditional probabilities $1 - p_{k'}^{(n-1)}, t_{k'j}^{(n-1)}, (p_{k'}^{(n-1)} - t_{k'j}^{(n-1)})$ respectively and the three events are conditionally independent of Ω_{wv} assuming the LTI property. These facts lead to (3.19).

To verify (3.15), we use (2.3) instead of (2.2) and follow these same steps. To verify (3.16), we use (2.4), (3.9) and (3.12) to give the formula

$$\hat{q}_{jk}^{(n)} = \mathbb{P}[\Sigma_v \leq \sum_{w \in \mathcal{N}_v^+} \Omega_{vw} \hat{\zeta}_{vw}^{(n)} \mid v \in \mathcal{N}_{jk}] = \left\langle S_{jk}, \left(h_k^{(n-1)}(x) \right)^{\otimes k} \right\rangle \quad (3.24)$$

where

$$h_k^{(n-1)}(x) = \sum_{j'} \frac{d}{dx} \mathbb{P}[\Omega_{vw} \hat{\zeta}_{vw}^{(n-1)} \leq x \mid v \in \mathcal{N}_{jk}, w \in \mathcal{N}_v^+, j_w = j'] Q_{j'|k} . \quad (3.25)$$

To verify (3.20), note that under the conditions $k_v = k, w \in \mathcal{N}_v^+, j_w = j'$, the events $\{\hat{\zeta}_{vw}^{(n-1)} = 0\}, \{\hat{\zeta}_{vw}^{(n-1)} = 1\}, \{\hat{\zeta}_{vw}^{(n-1)} = \lambda\}$ are equivalent to the events $\{w \in \hat{\mathcal{S}}_{n-1}^c \cap \mathcal{D}_{n-1}^c \boxtimes v\}, \{w \in \mathcal{D}_{n-1} \boxtimes v\}, \{w \in \mathcal{S}_{n-1}\}$ and hence have conditional probabilities $(1 - \hat{q}_{j'}^{(n-1)})(1 - \tilde{p}_{j'}^{(n-1)}), \tilde{p}_{j'}^{(n-1)}, \hat{q}_{j'}^{(n-1)}(1 - \tilde{p}_{j'}^{(n-1)})$ respectively. To verify (3.17), we apply the LTI property again to compute the conditional probability of the disjoint union $\{\xi_{wv}^{(n)} = 1\} = \{\xi_{wv}^{(n-1)} = 1\} \cup \{w \in \mathcal{D}_n \cap \mathcal{D}_{n-1}^c, v \in \hat{\mathcal{S}}_{n-1}^c\}$ defined by (2.5).

Finally, to verify (3.18), we note that $\hat{\mathcal{S}}_n^c \cap \mathcal{D}_n^c = \hat{\mathcal{S}}_n^c \cap \{\Delta_v > \sum_{w \in \mathcal{N}_v^-} \Omega_{wv} \mathbf{1}_{\{w \in \mathcal{D}_{n-1}\}}\}$. By LTI, these last two events are independent conditioned on $v \in \mathcal{N}_{jk}$, and the required formula results by following the steps taken to prove (3.14).

Remark 2. We see from this argument that the cascade mapping has two sources of error. The first, stemming from cycles in the skeleton that cause finite size corrections to the LTI property, is familiar, and known to behave like $O(N^{-1})$ on

configuration graphs as the network size N grows to infinity, and to be zero on finite tree graphs. The second source of error is less familiar and stems from dependence due to the assumption that the default shocks that impact v are diminished by a factor $1 - \lambda$ when v becomes stressed, that is, they depend on whether v is stressed or not.

The iterates of this cascade mapping converge as $n \rightarrow \infty$ to a fixed set of probabilities that represent the eventual state of the system at the end of the cascade. These final probabilities can be used to measure the overall impact of the crisis. For example, the expected number of eventually defaulted and stressed banks are

$$\text{Expected number of defaulted banks} = N \sum_{jk} P_{jk} p_{jk}^{(\infty)} = N \sum_k P_k^+ p_k^{(\infty)} \quad (3.26)$$

$$\text{Expected number of stressed banks} = N \sum_{jk} P_{jk} q_{jk}^{(\infty)}. \quad (3.27)$$

4 Networks with Fixed Skeletons

The goal of the present section is to derive approximate probabilistic formulas describing the double cascade on a network where the skeleton is actually known (deterministic) and finite, while the buffers and weights are random. This analysis will allow us to address systemic risk in tractable models of real observed financial networks, without the need for Monte Carlo simulations.

Let $A = A_{vv'}, v, v' \in \mathcal{N}$ be the nonsymmetric adjacency matrix of the fixed directed graph $(\mathcal{N}, \mathcal{E})$. We number the nodes in \mathcal{N} by $v = 1, 2, \dots, N$ and the links by $\ell = 1, 2, \dots, E$ where $E = \sum_{1 \leq v, v' \leq N} A_{vv'}$. The buffer random variables Δ_v, Σ_v at each node are assumed to have a mass $p_v^{(0)}, q_v^{(0)}$ at 0 (representing the initial default and stress probabilities) and continuous support with density functions $d(x), s_v(x)$ on the positive reals. The edge weights $\Omega_\ell, \ell \in \mathcal{E}$ have continuous support with densities $w_\ell(x)$ on the positive reals but no mass at 0. The random variables $\{\Delta_v, \Sigma_v, \Omega_\ell\}, v \in \mathcal{N}, \ell \in \mathcal{E}$ are assumed to be an independent collection.

The aim of this section is to use the LTI property as an approximation to derive formulas for the marginal likelihoods $p_v^{(\infty)}, q_v^{(\infty)}$ for the eventual default and stress of all individual nodes, as well as the possibility to compute formulas for more detailed systemic quantities. This approximation is not exact for the same two reasons as before: there is dependence between the default shocks hitting a bank v due to the stress response, and when there are cycles in the skeleton. In general, when the skeleton is a single random realization from a configuration graph ensemble, we expect the LTI approximation to get better with increasing N . We now present an approximate analysis, paralleling the previous section, of the sequence

of probabilities

$$\begin{aligned}
p_v^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n] , \\
q_v^{(n)} &= \mathbb{P}[v \in \mathcal{S}_n] , \\
\tilde{p}_{wv}^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \text{ (R) } w] , \\
\tilde{q}_v^{(n)} &= \mathbb{P}[v \in \hat{\mathcal{S}}_n] ,
\end{aligned} \tag{4.1}$$

for each node v or edge wv . For the same reason as before we need in addition to track

$$t_{wv}^{(n)} = \mathbb{P} \left[\xi_{wv}^{(n)} = 1 \mid v \in \mathcal{N}_w^+ \right] . \tag{4.2}$$

Inductively, we have

$$p_v^{(n)} = \left\langle D_v, \otimes_{v' \in \mathcal{N}_v^-} \left(g_{v'v}^{(n-1)} \right) \right\rangle , \tag{4.3}$$

$$\tilde{p}_{wv}^{(n)} = \left\langle D_v, \otimes_{v' \in \mathcal{N}_v^- \setminus w} \left(g_{v'v}^{(n-1)} \right) \right\rangle , \tag{4.4}$$

$$\tilde{q}_v^{(n)} = \left\langle S_v, \otimes_{v' \in \mathcal{N}_v^+} \left(h_{vv'}^{(n-1)} \right) \right\rangle , \tag{4.5}$$

$$t_{wv}^{(n)} = t_{wv}^{(n-1)} + (p_w^{(n)} - p_w^{(n-1)})(1 - \tilde{q}_v^{(n-1)}) , \tag{4.6}$$

$$q_v^{(n)} = 1 - p_v^{(n)} - \left(1 - \tilde{q}_v^{(n)} \right) \left\langle D_v, \otimes_{v' \in \mathcal{N}_v^-} \left(\tilde{g}_{v'v}^{(n-1)} \right) \right\rangle . \tag{4.7}$$

The PDFs can be computed as before under the LTI approximation:

$$\begin{aligned}
g_{wv}^{(n-1)}(x) &= (1 - p_w^{(n-1)})\delta_0(x) + t_{wv}^{(n-1)}w_{wv}(x) \\
&\quad + (p_w^{(n-1)} - t_{wv}^{(n-1)}) \cdot \frac{1}{1 - \lambda} w_{wv}(x/1 - \lambda) ,
\end{aligned} \tag{4.8}$$

$$\begin{aligned}
h_{vw}^{(n-1)}(x) &= (1 - \tilde{q}_w^{(n-1)})(1 - \tilde{p}_{vw}^{(n-1)})\delta_0(x) + \tilde{p}_{vw}^{(n-1)}w_{vw}(x) \\
&\quad + \tilde{q}_w^{(n-1)}(1 - \tilde{p}_{vw}^{(n-1)}) \cdot \frac{1}{\lambda} w_{vw}(x/\lambda) ,
\end{aligned} \tag{4.9}$$

$$\tilde{g}_{wv}^{(n-1)}(x) = (1 - p_w^{(n-1)})\delta_0(x) + p_{wv}^{(n-1)}w_{wv}(x) . \tag{4.10}$$

5 Numerical Experiments

In this section, we report briefly on numerical experiments that illustrate the methods developed in this paper. Firstly, we aim to convince the reader that the LTI method correctly computes the fixed point of the double cascade mapping in networks with large values of N by cross validating it using an independently coded Monte Carlo (MC) implementation. For efficiency, the LTI implementation uses a Fast Fourier Transform approach to compute the convolutions in (3.14)-(3.17). This technique was developed in Hurd & Gleeson (2013) and is sketched in Appendix A. Secondly, we will show how the LTI method can lead to answers to questions about the nature of systemic risk, particularly the intertwining of stress and default. Thirdly, we shall show how the method performs in a challenging stylized network specified to reflect the complex characteristics of a 2011 dataset on the network of 90 most systemically important banks in the European Union.

5.1 Experiment 1: Verifying the LTI Method

This experiment aims to verify that the LTI method performs as expected when applied to a stylized financial network whose specification is similar to that given in Gai & Kapadia (2010a). It consists of a random directed Poisson skeleton with $N = 20000$ nodes and mean degree $z = 10$, where each node v can be viewed as a bank with a default buffer $\Delta_v = 0.04$ and stress buffer $\Sigma_v = 0.035$. Unlike the deterministic interbank exposures used in Gai & Kapadia (2010a), the weight Ω_ℓ of an edge ℓ is taken from a log-normal distribution with mean $\mu_\ell = 0.2j_\ell^{-1}$, and standard deviation $0.383\mu_\ell$. Note that this specification makes the exposure size dependent on the lending bank. An initial shock is applied to the network that causes each bank to default with a 1% probability.

We compare the final fractions of defaulted and stressed banks as computed using MC simulation with 1000 realizations and the LTI analytic formulas in (3.26) and (3.27). It is simple to generate a directed Poisson random graph of size N with mean degree $z > 0$: one simply selects directed edges independently from all $N(N - 1)$ potential edges, each with probability $p = z/(N - 1)$. The resultant bi-degree distribution is a product of binomials, $P_{jk} = \mathbb{P}[v \in \mathcal{N}_{jk}] = \text{Bin}(N - 1, p, j) \times \text{Bin}(N - 1, p, k)$, which for large N nodes is approximately a product of Poisson(z) distributions.

Figure 3 plots the results as functions of the stress response parameter λ , with error bars that represent the 10th and 90th percentile of the MC result. It shows the expected agreement between MC and LTI analytics, with discrepancies that can be attributed to finite N effects present in the MC simulations.

One fundamental property of our model is clearly shown in this experiment: stress and default are negatively correlated. This fact can be explained by the stress response which enables banks to react to liquidity shocks before they default, by reducing their interbank exposures. This response creates yet more stress, but leads to a more resilient network. The “knife-edge” property of default cascades is also clearly shown: In the model parametrization we chose, a very small increase in λ dramatically alters the stability of the network. We also note that MC error bars are very large near the knife-edge.

5.2 Experiment 2: A Stylized Poisson Network

The next experiment focuses again on Poisson networks, with the aim to better understand the effects of various parameters on network resilience. In general, we continue to find confirmation that the LTI results accurately reflect observations from MC simulations.

5.2.1 Experiment 2A: Effects of Default and Stress Buffers

We consider how the parametrized financial network of Experiment 1 in a default-susceptible state with $\lambda = 0.5$ can be made resilient to random shocks by either




Figure 3: Experiment 1. The mean default and stress cascade sizes computed by MC simulations (symbols) and LTI analytics (lines). The skeleton is a Poisson directed network with $N = 20000$ nodes and mean degree $z = 10$. All banks have default buffers $\Delta_v = 0.04$ and stress buffers $\Sigma_v = 0.035$. An edge weight Ω_ℓ of an edge ℓ is taken from a log-normal distribution with mean $\mu_\ell = 0.2j_\ell^{-1}$ and standard deviation $0.383\mu_\ell$. Error bars indicate the 10^{th} and 90^{th} percentiles of the MC result.

increasing the default buffers or decreasing the stress buffers. Such changes, for example, can be prompted by financial regulators.

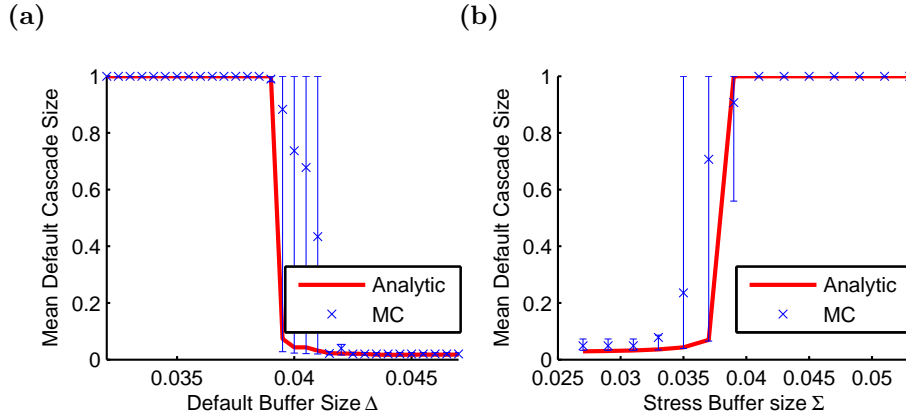


Figure 4: Experiment 2A. The mean default cascade size as a function of (a) default buffer Δ and (b) stress buffer Σ . The LTI analytic approximation (lines) correctly predicts MC simulations results (symbols). Here $\lambda = 0.5$ and other parameters are chosen as in Experiment 1 (Fig. 3).

In Fig. 4(a) we illustrate how the change in default buffers affects the default cascade size. We observe a very fast transition from 100% default cascade to almost no default as Δ increases over the interval $[0.04, 0.045]$. This knife-edge property is observable in both the LTI analytics and in the MC simulations. Note again that the MC error bars, representing the 10th and 90th percentiles, become very large near the knife-edge.

In Fig. 4(b), we examine the influence of the stress buffer on the default cascade size. If stress buffers are reduced, banks start to react to stress shocks more quickly, which can in turn dramatically reduce default cascade risk in the network. Figures 4(a) and (b) show that there may be several approaches to improving network resilience.

5.2.2 Experiment 2B: Effects of Graph Connectivity and Stress Response

Aside from mandating changes to the behaviour of banks during or prior to a crisis by imposing constraints on stress and default buffers, regulators can also influence the shape of the financial network as a whole. In Experiment 2B, we demonstrate how systemic risk is influenced by the skeleton itself. To this end, we calculate the sizes of default and stress cascades in a directed Poisson network as a function of the connectivity parameter z and the stress response λ . In our simple model specification, the mean degree z is the only parameter that controls the shape of the skeleton, whereas in a more realistic modelling approach the skeleton may have many more parameters.

In this experiment, we increased the model complexity by assuming each bank has a random default buffer taken from a log-normal distribution with mean 0.18 and standard deviation 0.18, and a stress buffer from an independent log-normal distribution with mean 0.12 and standard deviation 0.12. The edge weights Ω_ℓ come from a log-normal distribution with mean and standard deviation proportional to $(j_\ell k_\ell)^{-0.5}$, with the average edge weight on the entire network equal to 1. Once again we apply an initial shock so that each bank has 1% chance of defaulting initially.

In Figs. 5(a) and (b), we respectively show the mean sizes of default and stress cascades as functions of network mean degree z and the stress response λ . For clarity of the graphics, we do not show the MC simulations, as they agree, in a similar fashion as earlier, with the LTI analytics. Again, in these plots we notice the strong anti-correlation between stress and default probabilities as we vary z and λ . It is also interesting to observe that the final level of stress is not monotonic in the connectivity parameter z .

5.3 Experiment 3: An EU-Inspired Network with 90 Nodes

It is well known that Poisson random graphs are an inadequate description of real economic networks, so we are interested in having a rough picture of the systemic

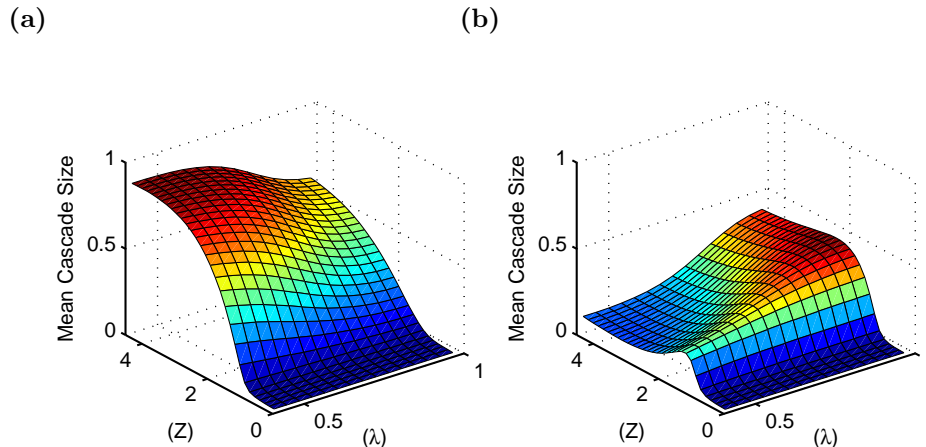


Figure 5: Experiment 2B. The sizes of the mean default cascade (a) and the mean stress cascade (b) on a directed Poisson network as functions of the network mean degree z and stress response parameter λ . Here the values of edge weights (interbank exposures), default and stress buffers are taken from log-normal distributions as specified in the text. Other parameters are the same as in Fig. 3.

risk of actual financial networks. In Experiment 3, we consider a single realization of a 90 node graph that aims to capture stylized features of the interbank network of the European Union. We computed cascade dynamics on this network using both the MC simulations and the LTI analytic method of Section 4 for networks with a fixed skeleton. As a preliminary step (not reported), we validated the consistency of the LTI analysis by verifying that the two methods agree as expected on a number of tree networks.

Numerous studies of real-world financial networks, notably Bech & Atalay (2010) and Cont, Moussa & Santos (2013), have observed their highly heterogeneous structure and concluded that in and out degrees have fat tailed distributions, as do exposure sizes, and presumably, buffers. Our schematic model of 90 EU banks was designed to capture these basic statistical features, and as well to fit aggregated statistics from data published on the 2011 ECB stress testing of systemically important banks in the European Union. We show the skeleton of our stylized network in Figure 6. The details of its construction and the specifications for buffer and exposure distributions are given in Appendix B.

In the first part of the experiment, we calculate the mean sizes of stress and default cascades that start from the default of one random bank in the EU network. Figure 7(a) shows the mean default and stress cascade sizes as functions of the stress response parameter λ . This graph shows no evidence of cascading (beyond the initially defaulted bank and the stress it causes to its immediate neighbors), demonstrating that the MC and analytic computations agree that the EU network




Figure 6: A representation of the skeleton of the 90 bank network of Experiment 3. We do not show the edge directions here to avoid cluttering the figure. The nodes are plotted with total degree increasing in the clockwise direction, with the minimally connected bank being the rightmost node.

in June 2011 was resilient to such a shock.




Figure 7: Experiment 3: sizes of default and stress cascades on a stylized EU interbank network starting from the default of a single randomly chosen bank. (a) Cascade sizes for various values of the stress response λ . The EU financial system at the time of the 2011 stress testing exercise appears to be resilient to single-bank shocks. (b) The same system as in (a) where a dire pre-shock crisis has reduced the bank default buffers to 10% of their original value, and stress buffers to a fraction of their value indicated on the x -axis.

To move the EU network to a knife-edge situation where a large scale double cascade can be triggered by the default of a single bank, we found it was necessary to imagine a dire crisis where prior to the default shock, the Core Tier 1 Capital of all institutions (i.e. their default buffers) has been decimated to 10% of their initial amount. Furthermore, each bank’s stress buffer is also reduced to a certain fraction of its original value. In Fig. 7(b) we take $\lambda = 0.7$ and show the cascade sizes versus the remaining fraction of stress buffers.

The MC simulations show that for smaller shocks the size of a stress cascade first dominate default cascades: the banking system becomes highly illiquid, but most banks are able to protect themselves from default. Beyond a certain critical value of the global shock, however, default cascades take over stress cascades, as banks are hit by a shock so high that they default immediately, without the possibility to protect themselves as they would if they got first in the stressed state. Our analytical approach is able to capture the essence of these mechanisms, identifying correctly the critical point where default cascades overtake stress cascades. The quantitative size of the cascades away from this critical point, however, is less well matched by the analytical model, probably because of the presence of short cycles in the skeleton. As cycles tend to slow down the cascades, the transitions are smoother in the MC. It is quite interesting that, in spite of our two approximations, our model is able to capture the general mechanisms: a peak in the number of stressed banks is a signal that can indicate quite precisely the point where systemic default takes over in a banking network.

6 Conclusions

Our double cascade model is a natural extension of a strand of systemic risk research that studies elementary models that build in either default or stress cascades, but not both. Only by combining the default and stress mechanisms into a single model can one measure such features as the intuitively obvious effect of banks using the stress response to reduce their risk of default.

Developing a feasible and reliable computation framework for a model as complex as ours poses challenges that have been met in the experiments described in Section 5. We have demonstrated how computations can be done by two complementary approaches: the Monte Carlo (MC) method and the locally tree-like independence (LTI) analytic method. Having these two approaches, each with its pros and cons, allows cross validation, increasing confidence in the results one obtains. MC simulation, with its natural flexibility, remains the workhorse computational framework for general systemic network risk. However, where it applies, the LTI method, given its relation to similar methods in other areas of network science, adds the possibility to understand the flow of the cascade in a different and sometimes better way than one can with MC simulations alone. For example, using LTI one can determine sensitivities to changing parameters through explicit differentiation.

In some situations, for example in simulating general assortative (P, Q) configuration models, the MC method is infeasible or leads to unacceptably long run times, whereas the LTI method can be computed without difficulty. Ultimately what is important is that both methods have complementary strengths and weaknesses, and when used in combination lead to robust and reliable conclusions about a wide range of network effects.

Our numerical experiments explore only a small sample of simple model specifications, leaving many promising financial networks to be investigated using our techniques. While the systemic importance of parameters such as network connectivity, mean buffer strength, and the size of the interbank sector have been studied previously, other parameters such as the stress response, the buffer and exposure variances, and graph assortativity, remain almost completely unexplored. The effect of market illiquidity and asset fire sales has been omitted from the present paper, but its impact on the cascade and consequently the greater economy merits careful investigation. Financial network databases, and the statistical methods for matching such data to the model, are still in an underdeveloped state, but are needed to tie down the wide range of parameters in our model. Planned future investigations of the double cascade model and its extensions will uncover and explain further interesting and unexpected systemic risk phenomena, and find uses by policy makers and regulators.

Appendix A Discrete Probability Distributions and the Fast Fourier Transform

Numerical implementation of these models follows the methods outlined in Hurd & Gleeson (2013). In this section, we analyze the case where all random variables $\{\Delta_v, \Sigma_v, \Omega_\ell\}$ take values in a specific finite discrete set $\mathcal{M} = \{0, 1, \dots, (M - 1)\}$ with a large value M . It is well known that in this situation the convolutions in (3.13) are slow to compute by direct integration, but can be performed exactly and efficiently by use of the discrete Fast Fourier Transform (FFT).

Let X, Y be two independent random variables with probability mass functions (PMF) p_X, p_Y taking values on the non-negative integers $\{0, 1, 2, \dots\}$. Then the random variable $X + Y$ also takes values on this set and has the probability mass function (PMF) $p_{X+Y} = p_X * p_Y$ where the convolution of two functions f, g is defined to be

$$(f * g)(n) = \sum_{m=0}^n f(m)g(n - m) . \tag{A.1}$$

Note that p_{X+Y} will not necessarily have support on the finite set \mathcal{M} if p_X, p_Y have support on \mathcal{M} . This discrepancy leads to the difficulty called “aliasing”.

We now consider the space $\mathbb{C}^{\mathcal{M}}$ of \mathbb{C} -valued functions on $\mathcal{M} = \{0, 1, \dots, M - 1\}$. The discrete Fourier transform, or fast Fourier transform (FFT), is the linear

mapping $\mathcal{F} : a = [a_0, \dots, a_{M-1}] \in \mathbb{C}^M \rightarrow \hat{a} = \mathcal{F}(a) \in \mathbb{C}^M$ defined by

$$\hat{a}_k = \sum_{l \in \mathcal{M}} \zeta_{kl} a_l, k \in \mathcal{M}. \quad (\text{A.2})$$

where the coefficient matrix $Z = (\zeta_{kl})$ has entries $\zeta_{kl} = e^{-2\pi i kl/M}$. The “inverse FFT” (IFFT), is given by the map $a \rightarrow \tilde{a} = \mathcal{G}(a)$ where

$$\tilde{a}_k = \frac{1}{M} \sum_{l \in \mathcal{M}} \bar{\zeta}_{kl} a_l, k \in \mathcal{M}. \quad (\text{A.3})$$

If we let \bar{a} denote the complex conjugate of a , we can define the Hermitian inner product between

$$\langle a, b \rangle := \sum_{m \in \mathcal{M}} \bar{a}_m b_m. \quad (\text{A.4})$$

We also define the convolution product of two vectors:

$$(a * b)(n) = \sum_{m \in \mathcal{M}} a(m) b(n - m \text{ modulo } M), \quad n \in \mathcal{M}. \quad (\text{A.5})$$

Note that this agrees with (A.1) if and only if the sum of the supports of a and b is in \mathcal{M} . Otherwise the difference is called an aliasing error: our numerical implementations reduce or eliminate aliasing errors by taking M sufficiently large.

The following identities hold for all $a, b \in \mathbb{C}^M$:

1. Inverse mappings: $a = \mathcal{G}(\mathcal{F}(a)) = \mathcal{F}(\mathcal{G}(a))$;
2. Conjugation: $\overline{\mathcal{G}(a)} = \frac{1}{M} \mathcal{F}(\bar{a})$;
3. Parseval Identity: $\langle a, b \rangle = M \langle \tilde{a}, \tilde{b} \rangle = \frac{1}{M} \langle \hat{a}, \hat{b} \rangle$;
4. Convolution Identities: $\tilde{a} \cdot \tilde{b} = \widehat{(a * b)}$, $\hat{a} \cdot \hat{b} = \widehat{(a * b)}$,

where $\cdot *$ denotes the component-wise product.

As an example to illustrate how the above formulas help, we observe that a typical formula (3.14) can be computed instead by the formula

$$p_{jk}^{(n)} = \frac{1}{M} \langle \mathcal{F}(D), \left(\hat{g}_j^{(n-1)} \right)^j \rangle, \quad (\text{A.6})$$

where $\hat{D} = \mathcal{F}(D)$, $\hat{g}_j^{(n-1)} = \mathcal{F}(g_j^{(n-1)})$ and the power is the component-wise vector multiplication. Such FFT-based formulas can be computed systematically, very efficiently, if the discrete probability distributions for Δ, Σ, Ω are initialized in terms of their Fourier transforms.

Appendix B EU Network Construction

In 2011, the European Banking Authority (EBA) made public a dataset^e on the interbank exposures of a selection of 90 medium to large European banks, as well as other information such as their Core Tier 1 Capital and Risk Weighted Assets. In this dataset, interbank exposures are aggregated by country on the liability side, which means we only know the aggregated amount each bank v has lent to all banks in each EU country c . In this Appendix we explain how we built the synthetic network in Experiment 3 that mimics stylized facts of real financial networks and uses this EBA data as a source of additional information about interbank liabilities and bank balance sheets.

Motivated by the ubiquitous relevance of networks with fat-tailed degree distributions, as reported in papers such as Cont, Moussa & Santos (2013); Bech & Atalay (2010), we built the skeleton based on the preferential attachment model of Bollobás, Borgs, Chayes & Riordan (2003). Given four parameters α , γ , δ_- and δ_+ and letting $\beta = 1 - \alpha - \gamma$, this model grows a random directed network from a finite initial “seed graph” using three rules:

1. With probability α , add a new vertex v together with an edge from v to an existing vertex w , where w is chosen according to $j_w + \delta_-$ (that is, with probability proportional to $j_w + \delta_-$).
2. With probability β , add an edge from an existing vertex v , chosen according to $k_v + \delta_+$, to an existing vertex w , chosen according to $j_w + \delta_-$.
3. With probability γ , add a new vertex v together with an edge from an existing vertex w to v , where w is chosen according to $k_w + \delta_+$.

This method is known to lead to fat-tailed degree distributions $P_j^- \sim j^{-\tau_-}$ and $P_k^+ \sim k^{-\tau_+}$ with Pareto exponents $\tau_- = 1 + \frac{1+\delta_-(\alpha+\gamma)}{\alpha+\beta}$, $\tau_+ = 1 + \frac{1+\delta_+(\alpha+\gamma)}{\gamma+\beta}$. The four parameters $\alpha = 0.169$, $\gamma = 0.169$, $\delta_- = \delta_+ = 4.417$ are determined by the following conditions: we assumed Pareto exponents $\tau_- = \tau_+ = 4$ (ensuring finiteness of certain moments), that $\alpha = \gamma$, and that the mean degree is $z = 10$ (which is a typical value observed in financial networks). To achieve this, we searched over a range of values of α , generating for each α 1000 samples of networks each consisting of $N = 1000$ nodes. For each realized network sample, we selected the subnetwork of the 90 most connected nodes and calculated its mean degree z . Finally we selected the α value which provided a mean degree z closest to 10. This set of parameters was then used to produce the skeleton of Experiment 3 by generating a single sample of a (directed) scale-free graph with $N = 1000$ nodes, and retaining its most connected subnetwork of 90 nodes.

In order to build bank balance sheets, we assumed an LTI compatible specification of the buffer and exposure random variables as log-normally distributed

^e<http://www.eba.europa.eu/risk-analysis-and-data/eu-wide-stress-testing/2011>

conditionally on the network topology:

$$\Delta_v = (k_v j_v)^{\beta_1} \exp[a_1 + b_1 X_v] , \quad (\text{B.1})$$

$$\Sigma_v = \frac{2}{3} (k_v j_v)^{\beta_1} \exp[a_1 + b_1 \tilde{X}_v] , \quad (\text{B.2})$$

$$\Omega_v = (k_v j_v)^{\beta_2} \exp[a_2 + b_2 X_\ell] . \quad (\text{B.3})$$

where the collection $\{X_v, \tilde{X}_v, X_\ell\}$ consists of independent standard normal random variables. To fix the parameter values, first we arbitrarily set $\beta_1 = 0.3$ and $\beta_2 = -0.2$, with the rationale that the default buffer should increase with bank connectivity, while a larger number of counterparties should imply lower average bilateral exposures. We used the reported Core Tier 1 Capital as a proxy for the default buffers, and thus we matched the first and second sample moments $E[\Delta_v]$ and $E[\Delta_v^2]$ using equation (B.1). Since we found no proxy in the data for the stress buffers Σ , we arbitrarily selected the same parameters as for Δ , but with a prefactor $2/3$. Finally, matching equation (B.3) with sample moments $E[A_v^{IB}]$ and $E[(A_v^{IB})^2]$, from the aggregated interbank exposure data, gives us enough equations to determine the full list of parameters: $\beta_1 = 0.3$, $a_1 = 8.03$, $b_1 = 0.9$, $\beta_2 = -0.2$, $a_2 = 8.75$, $b_2 = 1.16$.

Appendix C Acknowledgements

This project was germinated at the MITACS Canada International Focus Period devoted to Advances in Network Analysis and its Applications, organized by Evangelos Kranakis and held in Vancouver, Canada in July 2012. It was funded by awards from the Global Risk Institute for Financial Services in Toronto (TRH), from the Natural Sciences and Engineering Research Council of Canada (TRH), from the Irish Research Council co-funded by Marie Curie Actions under FP7 (INSPIRE fellowship, SM), from the FET-Proactive project PLEXMATH (DC), and from the Science Foundation Ireland (11/PI/1026, DC and SM)

We are indebted to Lionel Cassier of École Polytechnique France, and Huibin Cheng, Matheus Grasselli and Bernardo Costa Lima of McMaster University, Canada, who were active in developing an early version of our model. We are grateful to Grzegorz Halaj of the European Central Bank, for discussions about the EU financial network.

References

- H. Amini, R. Cont & A. Minca. Stress testing the resilience of financial networks. *International Journal of Theoretical and Applied Finance*, 15:1–20, 2012.
- H. Amini, R. Cont & A. Minca. Resilience to contagion in financial networks. *Mathematical Finance*, 26:329–365, 2016.

- S. Battiston, D. D. Gatti, M. Gallegati, B. Greenwald & J. E. Stiglitz. Liaisons dangereuses: Increasing connectivity, risk sharing, and systemic risk. *Journal of Economic Dynamics and Control*, 36(8):1121 – 1141, 2012.
- M. Bech & E. Atalay. The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22):5223–5246, 2010.
- E. A. Bender & E. R. Canfield. The asymptotic number of labeled graphs with given degree sequences. *Journal of Combinatorial Theory, Series A*, 24(3):296–307, 1978.
- B. Bollobás. A probabilistic proof of an asymptotic formula for the number of labelled regular graphs. *European Journal of Combinatorics*, 1:311, 1980.
- B. Bollobás, C. Borgs, J. Chayes & O. Riordan. Directed scale-free graphs. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA '03*, pages 132–139, Philadelphia, 2003. Society for Industrial and Applied Mathematics. ISBN 0-89871-538-5.
- M. K. Brunnermeier & L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.
- F. Caccioli, M. Shrestha, C. Moore & J. D. Farmer. Stability analysis of financial contagion due to overlapping portfolios. *Journal of Banking and Finance*, 46: 233–245, 2014.
- R. Cifuentes, G. Ferrucci & H. S. Shin. Liquidity risk and contagion. *Journal of the European Economic Association*, 5:556–566, 2005.
- R. Cont, A. Moussa & E. B. Santos. Network structure and systemic risk in banking systems. In J.-P. Fouque & J. A. Langsam, editors, *Handbook on Systemic Risk*, chapter 13, pages 327–368. Cambridge University Press, 2013.
- D. W. Diamond & R. G. Rajan. Liquidity shortages and banking crises. *Journal of Finance*, 60(2):615–647, 2005.
- L. Eisenberg & T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.
- P. Gai & S. Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A*, 466(2120):2401–2423, 2010a.
- P. Gai & S. Kapadia. Liquidity hoarding, network externalities, and interbank market collapse. *Proceedings of the Royal Society A*, 466:2401–2423, 2010b.
- P. Gai, A. Haldane & S. Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58:453–470, 2011.

- T. R. Hurd. The construction and properties of assortative configuration graphs. In R. Melnik, editor, *Recent Progress and Modern Challenges in Applied Mathematics, Modeling and Computational Science*, Fields Institute Communications. Springer-Verlag Berlin, October 2016. <https://arxiv.org/submit/1427020>.
- T. R. Hurd. *Contagion! Systemic Risk in Financial Networks*. SpringerBriefs in Quantitative Finance. Springer Verlag, Berlin Heidelberg New York, 2016. Available at <http://ms.mcmaster.ca/tom/tom.html>.
- T. R. Hurd & J. P. Gleeson. A framework for analyzing contagion in banking networks. arXiv:1110.4312 [q-fin.GN], October 2011.
- T. R. Hurd & J. P. Gleeson. On Watts cascade model with random link weights. *Journal of Complex Networks*, 1(1):25–43, 2013.
- S. H. Lee. Systemic liquidity shortages and interbank network structures. *Journal of Financial Stability*, 9(1):1–12, 2013.
- R. M. May & N. Arinaminpathy. Systemic risk: the dynamics of model banking systems. *Journal of The Royal Society Interface*, 7(46):823–838, 2010.
- E. Nier, J. Yang, T. Yorulmazer & A. Alentorn. Network models and financial stability. *Journal of Economic Dynamics and Control*, 31:2033–2060, 2007.
- T. Roukny, H. Bersini, H. Pirotte, G. Caldarelli & S. Battiston. Default cascades in complex networks: Topology and systemic risk. *Scientific reports*, 3, 2013.
- J. Tirole. Illiquidity and all its friends. BIS Working Papers 303, Bank for International Settlements, Mar. 2010. URL <http://ideas.repec.org/p/bis/biswps/303.html>.
- C. Upper. Simulation methods to assess the danger of contagion in interbank markets. *Journal of Financial Stability*, 7(3):111–125, 2011.
- D. J. Watts. A simple model of global cascades on random networks. *Proceedings of the National Academy of Sciences*, 99(9):5766–5771, 2002.