Illiquidity and Insolvency: a Double Cascade Model of Financial Crises

T. R. Hurd[1], Davide Cellai[2], Huibin Cheng[1],

Sergey Melnik[2], Quentin Shao[1]

[1] *Mathematics & Statistics, McMaster University, Canada*

[2] *MACSI, University of Limerick, Ireland*

October 15, 2013

Abstract

In the aftermath of the interbank market collapse of 2007-08, the traditional idea that systemic risk is primarily the risk of cascading bank defaults has evolved into the view that it involves both cascading bank defaults as well as funding liquidity shocks, and that both types of shocks impair the functioning of the remaining undefaulted banks.

In current models of systemic risk, these two facets, namely funding illiquidity and insolvency, are treated as two distinct and separate phenomena. The main goal of the double cascade model is to integrate these two facets as two faces of the same coin. In a default cascade, insolvency of a given bank will create a shock 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. The pivotal question concerning default cascades is: in a given financial network, what is the effect of the default of a single firm on the solvency of other banks in the system? In a stress cascade, illiquidity that hits a given bank will create a shock to the liability side of the balance sheet of each of its debtor banks. Under some circumstances, such "upstream" shocks can cause further illiquidity stresses that may build up to create a global illiquidity cascade. The pivotal question concerning stress cascades is: in a given financial network, what is the effect of the illiquidity of a single firm on the liquidity of other banks in the system?

Our paper will introduce a deliberately simplified model of insolvency and illiquidity in financial networks that can provide answers to the question of how illiquidity

or default of one bank can influence the overall level of liquidity stress and default in the network. To get there, a number of issues are addressed. First, this paper proposes a stylized model of individual bank balance sheets that builds in regulatory constraints, the most important types of interbank exposures, and collateralization of interbank exposures. Secondly, three different possible states of a bank, namely the normal state, the stressed state and the insolvent state, are identified with conditions on the bank's balance sheet. Thirdly, the paper models the behavioural response of a bank when it finds itself in the stressed or insolvent states. Importantly, a stressed bank will seek to shrink its balance sheet, by recalling short term interbank assets. This serves to protect the bank from the default of its counterparties, but creates stress in the network by forcing its debtor banks to raise cash, perhaps causing them to become stressed.

Versions of these proposed models can be designed to have a property we call "locally tree-like independence" that leads to large-network asymptotic cascade formulas. Details of numerical experiments are given that verify that these asymptotic formulas yield the expected quantitative agreement with Monte Carlo results for large finite networks. These experiments illustrate clearly our main conclusion that in financial networks, the average default probability is inversely related to strength of banks' stress response and therefore to the overall level of stress in the network.

Key words: Systemic risk, banking network, contagion, random graph, default, funding liquidity, credit risk, financial mathematics, repo

AMS Subject Classification: 05C80, 91B30, 91B70, 91G40

1 Introduction

In the aftermath of the interbank market collapse of 2007-08, financial systemic risk is increasingly seen to involve not only cascades of defaulting banks but cascading of illiquid banks as well.[1] In two seminal papers, Gai and Kapadia [2010] and Gai et al. [2011], two key facets of systemic risk, namely illiquidity and insolvency cascades, are treated as two distinct and separate phenomena that can arise in financial networks. The two models presented are in a sense "dual" to one another: insolvency shocks are transmitted from debtor to creditor while illiquidity shocks are transmitted in an analogous way but from creditor to debtor. The main goal of the double cascade model we present is to integrate these two facets of systemic risk as two faces of the same coin.

In Gai and Kapadia [2010], insolvency of a given bank, defined as a bank whose net worth becomes non-positive, will generate a shock to the asset side of the balance sheet of each of its creditor banks. Under some circumstances, such "downstream"

[1] A typical definition of systemic risk that focuses only on cascading bank defaults is given on the website of the Commodity Futures Trading Commission:
www.cftc.gov/consumerprotection/educationcenter/cftcglossary/glossary_s

shocks can cause further insolvencies that may build up to create a global insolvency cascade. In their model of insolvency contagion, the pivotal question is this: in a given financial network, what is the effect of the default of a single firm on the solvency of other banks in the system?

The paper Gai et al. [2011] adopts the idea that funding illiquidity of a bank[2] is when its liquid assets are insufficient to cover a possible run on its repo liabilities. They argue that an illiquid or stressed bank is most likely to tighten its balance sheet by reducing its interbank lending, thereby creating shocks to the liability side of its debtor banks' balance sheets. Under some circumstances, such "upstream" shocks can cause further illiquidity stress that may build up to create a global illiquidity cascade. In their model of illiquidity contagion, the pivotal question is this: in a given financial network, what is the effect of the funding illiquidity of a single firm on the liquidity of other banks in the system?

The purpose of the present paper is to construct a deliberately simplified model of systemic risk, integrating both sides, illiquidity and insolvency, of a single coin. Only in the double cascade model can one frame the important question: What effect does a bank's behavioural response to liquidity stress have on the probable level of eventual defaults in entire system? One would expect that a stressed bank that reacts to protect itself from eventual default by shrinking its own balance sheet will inflict liquidity shocks to its debtor banks. More significantly, one may also expect a system that on average reacts more strongly to stress will be more resilient to defaults. The average strength of this stress reaction, denoted λ, will be a key parameter that determines the eventual level of default and stress in the network.

Our paper can provide quantitative answers to this question. We find that in general, the overall level of default in the system is *negatively* related to the stress response parameter λ, and hence to the overall level of stress in the system. To arrive at models where this negative relationship can be quantified, we will need to address a number of issues. First, we need to create a stylized model of an individual bank balance sheet that builds in regulatory constraints, the most important types of interbank exposures, and the possibility of collateralization of interbank exposures. For this, we adopt the balance sheet framework outlined in Gai et al. [2011]. Secondly, three possible states of a bank, namely the normal state, the stressed state and the insolvent state, are defined by conditions on the bank's *default* and *stress buffers* derived from their balance sheet. Thirdly, we will model the behavioural response of a bank when it finds itself in the stressed or insolvent states. Importantly, a stressed bank will seek to shrink its balance sheet by recalling short term interbank assets. This reaction will reduce its exposure to defaulting counterparties, and therefore will reduce the overall level of default in the network. However, such a reaction also forces its debtor banks to raise cash, and if forced to raise too much too quickly, these debtor banks will themselves become stressed. Therefore it seems intuitively natural that the overall level of stress in the network goes up with the stress parameter λ and

[2]Funding illiquidity is distinct from market illiquidity, where assets become difficult to sell due to an oversupply in the market. See Brunnermeier and Pedersen [2009] for a detailed analysis of these concepts.

down with the overall level of default.

The double cascade picture we present is well justified in light of the new Basel III regulatory framework. At the heart of the new regulatory changes are three new accounting ratios. The leverage ratio (LR), the ratio of tier I capital (i.e. equity) to assets, must be kept above 3%. This condition can be directly translated into our assumption of a positive default buffer for each bank. The liquidity coverage ratio (LCR), a specific ratio of liquid assets and short term liabilities, must be kept above 100%. This regulation, intended to keep banks liquid during a short term (30 day) run on its repo and interbank liabilities, is analogous to our assumption of the positivity of a stress buffer. There is one key difference: by penalizing interbank liabilities where our model does not, the LCR constraint explicitly seeks to avoid the possibility of a stress cascade. It is reasonable to expect that banks that violate their LCR constraint will tighten their balance sheets: this is equivalent to our assumed stress response. Finally, the Net Stable Funding Ratio (NSFR) is an additional new regulatory constraint, similar to the LCR, but focusing on maintaining liquidity in the face of liability runs over a one year horizon.

The present paper continues a recent strand of literature, built on the economic interbank framework proposed by Eisenberg and Noe [2001], that focuses on deliberately simplified models of random financial networks. These papers share the aim of determining the key parameters that most impact the level of systemic risk in such networks. They also share a philosophy common in the physics literature that true understanding of complex phenomena should proceed via the numerical and analytic study of simple "toy models" chosen carefully to exhibit key features. Only after learning the true impact of various possible mechanisms in such toy models will it be worthwhile to investigate realistic models that try to represent the intricacies of observed financial systems. Nier et al. [2007] uses Monte Carlo methods to highlight how the key network parameters for a stylized network of 25 banks can influence the total number of defaults in a nonlinear, indeed sometimes nonmonotonic, fashion. The paper of Gai and Kapadia [2010], and its extension Hurd and Gleeson [2011], adapt the Watts [2002] model of information cascades to the context of financial systems, deriving both analytical and Monte Carlo results showing the dependence of the default cascade on different structural parameters. May and Arinaminpathy [2010] present analytical formulas for the NYYA and GK models based on a mean field approximation that can explain some of the main properties of the graphs found in those papers. Gai et al. [2011] provides a stylized model of liquidity shocks that is mathematically equivalent to the model of Gai and Kapadia [2010], but with cascades that flow in the reverse direction from creditor to debtor. An important homogeneity assumption made in these last two papers is that either a bank's interbank assets (in Gai and Kapadia [2010]) or its interbank liabilities (in Gai et al. [2011]) are constant across its counterparties. When combined into a double cascade model with shocks acting in opposing directions, such mathematically convenient, but over restrictive, assumptions would imply the trivial case of exposures that are constant across the entire network. Amini et al. [2012] develop a simple but general analytical criterion for

resilience to default contagion in random financial networks, based on an asymptotic analysis of default cascades in heterogeneous networks.

The key technical innovation of the present paper is introduce a model framework, generalizing the frameworks of Gai and Kapadia [2010], Hurd and Gleeson [2011] and Gai et al. [2011], that will allow us to study the intertwining of stress and default cascades, each acting in opposing directions. In particular, the present paper provides a full mathematical treatment of networks with random stress and default buffers and random interbank exposures, extending the techniques developed in Hurd and Gleeson [2013] to go beyond the type of over restrictive homogeneity assumptions made in earlier works. The main result of the paper, Theorem 1, provides an exact asymptotic analysis of default and stress probabilities at each step of the double cascade, in the limit as the number of banks N goes to ∞.

The remainder of this paper is structured as follows. Section 2 provides the basic network framework and assumptions underlying the balance sheet structure of banks. The timing of the crisis and the bank behaviour assumptions are introduced in Section 3. These assumptions lead to rules for the transmission of shocks through the double stress/default cascade. The precise dynamical rules of the double cascade, including the conditions for banks to become stressed or defaulted, are provided in Section 4. Section 5 develops our main theorem, which yields an explicit asymptotic analysis of default and stress probabilities in large heterogeneous networks with random connectivity, balance sheets and interbank exposures. A key technical assumption needed to prove this result is to assume a condition we call the locally tree-like independence (LTI) property. Section 6 provides a parallel development of default and stress probabilities for cascades on finite "real-world" networks, where it is assumed that the graph of interbank connections is known explicitly, but balance sheets and exposure sizes are still random. Several specific financial experiments are reported in Section 7. First we summarize experiments that verify the main theorem by direct comparison of large N analytics to Monte Carlo simulation results. Secondly, we investigate the relationship between the stress response parameter and the level of stress and default, verifying our statement that average stress probability increases and average default probability decreases as λ increases. A final experiment involves a detailed specification of the underlying random variables that is consistent with known heuristics of financial networks and the latest stress testing data on 89 large banks in the EU system. We observed that the network in this specification is highly resilient, and only by a very large shock to the average default buffer size will a high fraction of banks become either stressed or defaulted. The first main conclusion of the paper, discussed in Section 8, is that the analytic asymptotic results on default and stress probabilities that stem from the main theorem, when used carefully, are consistent with results from Monte Carlo simulations on finite random graphs. A second conclusion is that stress and default are inversely related: as banks respond to stress more vigorously, creating more network stress, they protect the network from default.

One critically important effect could be, but is not, addressed in this paper. Space

considerations led us to leave this aspect for future work. This is the impact of financial cascades on the non-financial economy, and the consequent feedback into the financial markets through "firesales" of assets. The effects of such market illiquidity on financial networks have been extensively studied, for example in Cifuentes et al. [2005] and Adrian and Shin [2010], and it is known that these effects will amplify any cascade after it takes hold in the network.

2 Balance Sheet Assumptions

In the context of our paper, a deliberately simplified model of a bank is a specification of a bank's balance sheet and its behavioural rules, at a moment in time. The characteristics of the bank's asset and liability classes should be simplified enough to allow feasible computations of network dynamics, and complex enough to capture the network effects under study. The deliberately simplified modelling assumptions we now specify combine elements of the insolvency model of Gai and Kapadia [2010] and the illiquidity model of Gai et al. [2011] and are designed to capture the dual effects of illiquidity and insolvency during a financial crisis.

The network of banks consists of a collection of N banks, each structured in a similar manner. We adopt a labelling system to identify banks and their interbank counterparties. The banks are labelled by numbers from the set $\mathcal{N} = \{1, 2, \ldots, N\}$. The set of debtor banks of a bank $v \in \mathcal{N}$ is called the in-neighbourhood of v, and is denoted by $\mathcal{N}_v^-$; similarly the bank's creditor banks form a set called the out-neighbourhood $\mathcal{N}_v^+$ of v. The basic graph theoretic measure of a bank v's size is the degree pair (j_v, k_v), where the *in-degree* is the number of debtors $j_v = |\mathcal{N}_v^-|$ and the *out-degree* is the number of creditors $k_v = |\mathcal{N}_v^+|$.

Balance Sheet Assumptions: The basic asset and liability classes have the following characteristics:

1. In addition to cash, the bank holds liquid collateral assets in the form of risk-free government bonds;
2. The retail loan book of the bank is composed of illiquid loans to the retail sector;
3. Interbank loans are always assumed to be overnight so we neglect the term structure of interbank lending;
4. Repo transactions, short for repurchase agreements, are overnight, and are collateralized by either collateral assets or rehypothecated reverse-repo assets;
5. Bank deposits are insured hence the possibility of a run by depositors can be ignored;
6. Bank equity is held only by investors that are external to the banking system;
7. Limited liability: no bank may have negative equity;
8. Prior to the onset of a crisis, balance sheet entries are book values reported quarterly.

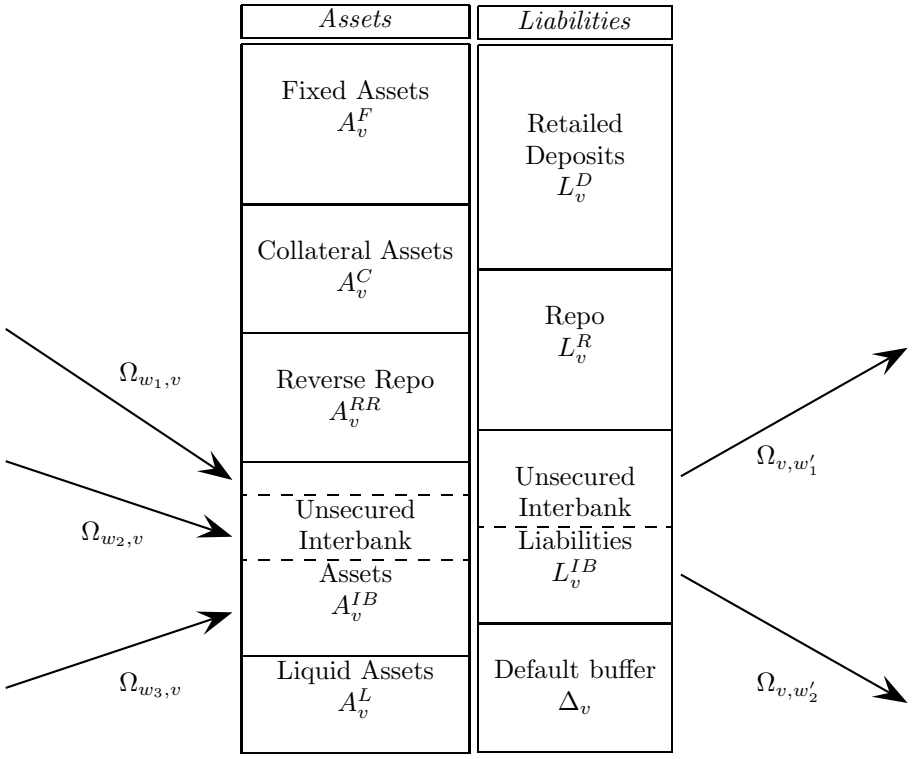

Figure 1: The stylized balance sheet of a bank v with in-degree $j_v = 3$ and out-degree $k_v = 2$. Banks w_1, w_2, w_3 are debtors of v while w_1', w_2' are its creditors. (Based on a figure from Gai et al. [2011].)

By these assumptions, the balance sheet of a bank $v \in \mathcal{N}$ has the structure shown in Figure 1. On the asset side of bank v's balance sheet we have: fixed assets A_v^F which denote illiquid retail loans; collateral assets A_v^C and reverse-repo assets A_v^{RR} which denote the assets that are allowable collateral for overnight repo borrowing; A_v^{IB} which are interbank loans; and cash A_v^L. On the bank's liability side, fixed deposits are denoted by L_v^D; repo liabilities by L_v^R; interbank liabilities by L_v^{IB}. The equity is defined to be

$$\Delta_v = A_v^F + A_v^C + A_v^{RR} + A_v^{IB} + A_v^L - L_v^D - L_v^R - L_v^{IB} \ . \tag{1}$$

when the right side is positive. When the right side is non-positive, the assumption of limited liability is taken to mean that the bank must default and Δ_v is defined to be 0. In what follows, we also refer to Δ_v as the bank's *default buffer*.

The interbank liabilities L^{IB} and assets A^{IB} decompose into bilateral interbank 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

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

The information of interbank counterparties is identifiable as a directed random graph, i.e. a set of nodes or vertices with 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 arrow pointing from v to w. The *type* (j_v, k_v) of a bank v is its in-degree j_v and out-degree k_v. The type (k_ℓ, j_ℓ) 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.[3]

Repo transactions are collateralized overnight loans, designed to be free of counterparty risk. At the *near date* (or inception date), the borrower (who is also the asset seller) and lender (who is also the asset buyer) enter into a swap of ownership of a collateral asset valued at \$1, either in A^C or A^{RR}, in exchange for \$ $(1-h)$. The *haircut* $h \in (0,1)$, which we assume in this paper to be constant over time and across banks, affords the lender a margin of safety: in case the borrower defaults, the lender's cost to dispose of the collateral is hoped to be less than h. At the *far date* (the end date of the contract, which we assume here to be the next day) the asset buyer agrees to resell the asset, at the pre agreed price \$ $(1-h)+r$, where the *repo rate* r is basically the overnight risk free interest rate. Repos are used by banks to keep current accounts positive, even during periods of stress and solvency shocks. The essence of the behaviour assumptions that will be made in the next section is that the bank's total capacity to raise cash by repo funding must exceed its total repo liabilities. The *stress buffer* Σ measures this capacity:

$$\Sigma = A^L + (1-h)[A^C + A^{RR}] - L^R \ . \tag{2}$$

Banks with a non-positive stress buffer are called *stressed* and trigger stress shocks with a magnitude proportional to the *stress response parameter* $\lambda \leq 1$.

In the next section, we will make additional assumptions about how insolvency and illiquidity cascade through the network. It turns out that the mechanical cascade rules will depend on the interbank exposures Ω together with the reduced set of balance sheet variables (Δ, Σ), rather than the full vector $X = (A^F, A^C, A^{RR}, A^{IB}, A^L, L^D, L^R, L^{IB})$. In what follows we will focus on the variables (Δ, Σ) and ignore the impact of cascades on the banks' full balance sheets and consequently on the greater economy.

3 Systemic Crisis Assumptions

Prior to the onset of the crisis, all banks are assumed to be in the normal state. Then, on day $n = 0$, a random collection of banks are assumed to experience initial shocks that deplete either their default buffer, causing $\Delta = 0$ or their stress buffer, causing $\Sigma = 0$.

We suppose that following this initial moment of the crisis, the financial regulator forces banks to recompute their balance sheets daily, using mark-to-market valuation. The initial stochastic state of the network is then described by the following random

[3]The convention that arrows point from debtors to creditors means that default shocks propagate in the downstream direction. Unfortunately, in some of the systemic risk literature, this convention is reversed.

elements: the interbank links form a directed random graph $\mathcal{E}$ on the set of banks $v \in \mathcal{N}$; the buffers Δ_v, Σ_v for $v \in \mathcal{N}$ are random; and the random variables Ω_{vw} for edges $\ell = (vw) \in \mathcal{E}$ each represent the interbank exposure of a bank $w \in \mathcal{N}_v^+$ to v. We now make assumptions about the day-by-day cascade dynamics on a general network. Subsequent sections will provide a probabilistic analysis of the cascade in both finite and infinite networks.

Crisis Timing Assumptions: Prior to the crisis, all banks are in the normal state, neither stressed nor insolvent.

1. The crisis commences on day 0 triggered by the default or illiquidity of one or more banks;
2. Balance sheets are recomputed daily on a mark-to-market basis;
3. Banks respond daily on the basis of their newly computed balance sheets;
4. All external cash flows, interest payments, and asset price changes are ignored throughout the crisis.

The daily response of banks will also be governed by simple rules:

Bank Behaviour Assumptions: On each day of the crisis, banks respond as follows:

1. An insolvent bank, characterized by $\Delta = 0$, is forced by the regulator into receivership. At this moment, each of its creditor banks are obliged to write down their defaulted exposures to zero thereby experiencing a *solvency shock*.
2. Solvency shocks reduce a bank's default buffer.
3. An illiquid or stressed bank, defined to be a non-defaulted bank with $\Sigma = 0$, reacts by shrinking its assets. Specifically, it reduces its interbank assets A^{IB} to $(1 - \lambda)A^{IB}$. It does so by terminating a constant fraction λ of its interbank loans. In doing so, it transmits a *stress shock* to the liabilities each of its debtor banks. λ is taken to be a constant across all banks during the crisis.
4. 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;
5. Stress shocks reduce a bank's stress buffer Σ.

This framework admits many variations and extensions. In the type of model specification we consider in this paper, we ignore the possibility of market illiquidity and changes in the underlying asset valuations during the short crisis period. In this case, the above rules apply mechanistically starting on day 0. The crisis evolves in a deterministic manner starting from the initial random state characterized by the quintuple $(\mathcal{N}, \mathcal{E}, \Delta, \Sigma, \Omega)$.

4 Double Cascade Dynamics

We now specify the precise dynamics of the double cascade on a financial network that on day $n = 0$ of the crisis consists of a random directed graph $\mathcal{E} \subset \mathcal{N} \times \mathcal{N}$, random buffers Δ_v and Σ_v for each bank (node) $v \in \mathcal{N}$ and random exposure sizes Ω_ℓ for each edge $\ell \in \mathcal{E}$.

The set $\mathcal{D}_n$ contains all the defaulted banks after n cascade steps, the set $\mathcal{S}_n$ comprises the undefaulted banks that are under stress after n steps, and $\mathcal{U}_n = \mathcal{D}_n^c \cap \mathcal{S}_n^c$ contains the remaining undefaulted, unstressed banks.[4] 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. We provide an inductive characterization of the sequence of random sets $\mathcal{D}_n$ and $\mathcal{S}_n$. We use the notation that an *event* defined by some condition P is written $\{P\}$, for example $\{v \in \mathcal{D}_n\}$, and the indicator random variable for that event is written $\mathbb{1}_{\{P\}}$.

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

$$\mathcal{D}_n := \begin{cases} \{\Delta_v = 0\} & \text{for } n = 0 \\ \left\{ v \mid \sum_{w\in\mathcal{N}_v^-} \Omega_{wv}\xi_{wv}^{(n-1)} \geq \Delta_v \right\} & \text{for } n \geq 1 \end{cases} \tag{3}$$

where the random variables ξ indicate the fractional sizes of the various default shocks impacting v. Similarly, to say that a bank v is stressed at step n means both that it is not yet defaulted and the stress shocks to step $n-1$ exceed the stress buffer, i.e. $\mathcal{S}_n = \mathcal{D}_n^c \cap \tilde{\mathcal{S}}_n$ where

$$\tilde{\mathcal{S}}_n := \begin{cases} \{\Sigma_v = 0\} & \text{for } n = 0 \\ \left\{ v \mid \sum_{w\in\mathcal{N}_v^+} \Omega_{vw}\zeta_{vw}^{(n-1)} \geq \Sigma_v \right\} & \text{for } n \geq 1 \end{cases} \tag{4}$$

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

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

$$\xi_{wv}^{(n-1)} := \begin{cases} 0 & \text{when } w \in \mathcal{D}_{n-1}^c \\ 1 & \text{when } w \in \Delta\mathcal{D}_m \text{ and } v \in \tilde{\mathcal{S}}_m^c, m = 0, \ldots, n-1 \\ 1-\lambda & \text{when } w \in \Delta\mathcal{D}_m \text{ and } v \in \tilde{\mathcal{S}}_m, m = 0, \ldots, n-1 \end{cases} \tag{5}$$

In formulas such as this, we use notation

$$\Delta\mathcal{D}_j = \mathcal{D}_j \setminus \mathcal{D}_{j-1},\ \mathcal{D}_{-1} = \emptyset; \qquad \Delta\mathcal{S}_j = \mathcal{S}_j \setminus \mathcal{S}_{j-1},\ \mathcal{S}_{-1} = \emptyset\ . \tag{6}$$

[4] For any set $\mathcal{B}$, $\mathcal{B}^c$ denotes its complement.

Note that in this definition, "fictitious" default shocks with fraction $\xi = 1-\lambda$ continue to impact after the default of v. Similarly, accounting for the assumption that defaulted creditors have a maximal impact on the bank's stress buffer, whereas stressed creditors only shock its stress buffer by a portion λ of the interbank exposure, one has for $n \geq 1$

$$\zeta_{vw}^{(n-1)} := \begin{cases} 0 & \text{when } w \in \tilde{\mathcal{S}}_{n-1}^c \cap \mathcal{D}_{n-1}^c \\ \lambda & \text{when } w \in \tilde{\mathcal{S}}_{n-1} \cap \mathcal{D}_{n-1}^c \\ 1 & \text{when } w \in \mathcal{D}_{n-1} \end{cases} . \tag{7}$$

At this point a subtle feedback effect arises that if not resolved will create difficulties in the subsequent probabilistic analysis. Similar to a situation discussed in Hurd and Gleeson [2013], there is dependence between the random variables Ω_{vw} and $\zeta_{vw}^{(n-1)}$ that appear in (4). This is because the condition $w \in \mathcal{D}_{n-1}$ in the definition of ζ_{vw} is not independent of Ω_{vw}. The difficulty can be resolved by considering properties that hold "without regarding" (WOR) a fixed bank. To this end, we will say "bank v is defaulted at step n, without regarding a bank $w \in \mathcal{N}_v^-$", and write $\{v \in \mathcal{D}_n \not\circledR w\}$, if the default condition is true without including an in-link from w to v. That is, for any $w \in \mathcal{N}$

$$\mathcal{D}_n \not\circledR w := \begin{cases} \{\Delta_v = 0\} \cap \mathcal{N}_w^+ & \text{for } n = 0 \\ \left\{ v \in \mathcal{N}_w^+ \mid \sum_{v' \in \mathcal{N}_v^- \setminus w} \Omega_{v'v} \xi_{v'v}^{(n-1)} \geq \Delta_v \right\} & \text{for } n \geq 1 \end{cases} \tag{8}$$

Using this definition, we can equivalently express the stress sets as $\mathcal{S}_n = \mathcal{D}_n^c \cap \hat{\mathcal{S}}_n$ where

$$\hat{\mathcal{S}}_n := \begin{cases} \{\Sigma_v = 0\} & \text{for } n = 0 \\ \left\{ v \mid \sum_{w \in \mathcal{N}_v^+} \Omega_{vw} \hat{\zeta}_{vw}^{(n-1)} \geq \Sigma_v \right\} & \text{for } n \geq 1 \end{cases} \tag{9}$$

with

$$\hat{\zeta}_{vw}^{(n-1)} := \begin{cases} 0 & \text{when } w \in \hat{\mathcal{S}}_{n-1}^c \cap \mathcal{D}_{n-1}^c \not\circledR v \\ \lambda & \text{when } w \in \hat{\mathcal{S}}_{n-1} \cap \mathcal{D}_{n-1}^c \not\circledR v \\ 1 & \text{when } w \in \mathcal{D}_{n-1} \not\circledR v \end{cases} \tag{10}$$

The reason is that for $v \in \mathcal{D}_n^c$, the $\not\circledR v$ condition is redundant and thus $\mathcal{D}_n^c \cap \tilde{\mathcal{S}}_m = \mathcal{D}_n^c \cap \hat{\mathcal{S}}_m$ for any $m \leq n$.

We need to check that the definition (5) of ξ remains essentially unaffected if $\tilde{\mathcal{S}}$ is replaced by $\hat{\mathcal{S}}$ on the right hand side, i.e. that ξ can change only in a way that leaves the default set unchanged. This is true because the sets $\hat{\mathcal{S}}_m \cap \mathcal{D}_m^c$ and $\tilde{\mathcal{S}}_m \cap \mathcal{D}_m$ are equal, while the sets $\hat{\mathcal{S}}_m \cap \mathcal{D}_m$ cannot be larger than the sets $\tilde{\mathcal{S}}_m \cap \mathcal{D}_m$. Thus the fraction ξ_{wv} won't change if v hasn't yet defaulted, and may change but cannot decrease if v has already defaulted.

The next two sections are devoted to the probabilistic analysis of the double cascade in two distinct settings: the case of infinite networks and the case of finite real-world networks.

5 Infinite Networks

Our model of a financial network has three layers of random structures: the skeleton graph (a random directed graph $(\mathcal{N}, \mathcal{E})$); the buffer random variables Δ_v, Σ_v, $v \in \mathcal{N}$ and the exposure weights $\Omega_\ell, \ell \in \mathcal{E}$.

The skeleton graph $(\mathcal{N}, \mathcal{E})$ is a random assortative directed graph as characterized by the following construction developed in Hurd and Gleeson [2011] of an assortative extension of the well-known configuration random graph model of Bollobãs [2001]:

Definition 1. *The infinite directed configuration graph model should be considered as the limit of any infinite sequence $(\Omega_N, \mathcal{F}_N, \mathbb{P})$ of random graphs of expected size N, and asymptotically consistent with compatible probability laws P, Q for the degree types of nodes and edges. More specifically:*

1. *For each $N = 1, 2, \ldots$ we have a random directed graph such that: the number of nodes of degree type (j, k) for $j, k \in \mathcal{K} = \{0, 1, \ldots, K\}$ is a random integer d^N_{jk} with expected value NP_{jk}; the number of edges of degree type (k, j) is a random integer e^N_{kj} with expected value NQ_{kj}.*
2. *For each feasible realization of the integer random variables $\{d^N_{jk}, e^N_{kj}\}$, the assortative configuration graph construction of Hurd and Gleeson [2011] is applied;*
3. *The normalized random sequences converge in probability to P and Q uniformly in j, k:*
$$d^N_{jk}/N = P_{jk} + o(1); \quad e^N_{kj}/N = Q_{kj} + o(1)$$
where $X^N = Y + o(1)$ for a sequence of random variables X^N means for each $\epsilon > 0$, $\mathbb{P}[|X^N - Y| > \epsilon]$ has the limit 0 as $N \to \infty$.
4. *For each $j, k \in \mathcal{K}$, $P_{jk} := \mathbb{P}[\mathcal{N}_{jk}]$ is the asymptotic probability of a random node having type (j, k) . This distribution has marginals $P^+_k := \sum_j P_{jk}, P^-_j := \sum_k P_{jk}$ and mean in and out degree $z = \sum_j jP^-_j = \sum_k kP^+_k$.*
5. *For each $j, k \in \mathcal{K}$, $Q_{kj} := \mathbb{P}[\mathcal{E}_{kj}]$ is the asymptotic probability of a type (k, j) edge. This distribution has marginals $Q^+_k := \sum_j Q_{kj}, Q^-_j := \sum_k Q_{kj}$ that are subject to the feasibility constraints $Q^+_k = kP^+_k/z, Q^-_j = jP^-_j/z$.*
6. *To simplify the analysis that follows, we fix a finite set of possible degrees $\mathcal{K} = \{k_1, k_2, \ldots, k_K\}$ and assume that $P^-_j, P^+_k > 0$ for all $j, k \in \mathcal{K}$.*

The non-negative random variables Δ_v have point masses at $x = 0$ that represent their initial default probability p^0_v. We assume that the distribution functions of Δ_v depend only on the type (j, k), and have the following form:

$$D_{jk}(x) = \mathbb{P}[\Delta_v \leq x | v \in \mathcal{N}_{jk}] \; ; \frac{d}{dx} D_{jk}(x) := p^0_{jk}\delta_0(x) + d_{jk}(x) \; . \tag{11}$$

where $d_{jk}(x) \geq 0$ is a specified function with $\int_0^\infty d_{jk}(x)dx = 1 - p_{jk}^0$. Similarly, 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:

$$S_{jk}(x) = \mathbb{P}[\Sigma_v \leq x, v \notin \mathcal{D}_0 | v \in \mathcal{N}_{jk}] \; ; \frac{d}{dx} S_{jk}(x) := q_{jk}^0 \delta_0(x) + s_{jk}(x) \; . \tag{12}$$

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 edge weight random variables Ω_ℓ 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

$$W_{kj}(x) = \mathbb{P}[\Omega_\ell \leq x | \ell \in \mathcal{E}_{kj}]; \; \frac{d}{dx} W_{kj}(x) = w_{kj}(x) \tag{13}$$

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

It was argued in Hurd and Gleeson [2011] that the above construction of a probability measure on random networks implies a property called locally tree-like independence (LTI) extending the locally tree-like property of random graphs that cycles of any fixed finite length occur in an infinite configuration graph only with zero probability. The probabilistic analysis to follow rests on this extended type of independence:

The locally tree-like independence (LTI) property: Consider a double cascade model on $\mathcal{N}$ defined by a 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, conditioned on information located at the node v, that is conditioned on Δ_v, Σ_v, j_v or k_v, X_1 and X_2 are independent.[5]

We are now in a position to derive exact formulas for the probabilistic cascade dynamics in infinite networks. We first define (here $\mathcal{U}_n = \mathcal{N} \setminus (\mathcal{S}_n \cup \mathcal{D}_n)$ are the undefaulted and unstressed nodes):

$$\begin{aligned} p_{jk}^{(n)} &= \mathbb{P}\left[v \in \mathcal{D}_n | v \in \mathcal{N}_{jk}\right] \; , \\ q_{jk}^{(n)} &= \mathbb{P}\left[v \in \mathcal{S}_n | v \in \mathcal{N}_{jk}\right] \; , \\ u_{jk}^{(n)} &= \mathbb{P}\left[v \in \mathcal{U}_n | v \in \mathcal{N}_{jk}\right] \; . \end{aligned} \tag{14}$$

[5]More precisely, to any subset of nodes $\mathcal{N}' \subset \mathcal{N}$ we associate the sigma-algebra $\mathcal{G}'$ generated by the balance sheets and degrees of nodes in $\mathcal{N}'$ and edges in $\mathcal{N}' \times \mathcal{N}'$. Let the sigma-algebras corresponding to $\mathcal{N}_1, \mathcal{N}_2, \{v\}$ be denoted $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_v$. Then $\mathcal{G}_1$ and $\mathcal{G}_2$ are statistically independent, conditioned on $\mathcal{G}_v$.

as well as additional probabilities

$$\begin{aligned}\hat{p}_{jk}^{(n)} &= \mathbb{P}\left[v \in \mathcal{D}_n \not\Vdash w | v \in \mathcal{N}_{jk} \cap \mathcal{N}_w^+\right] \ , \\ \hat{q}_{jk}^{(n)} &= \mathbb{P}\left[v \in \hat{\mathcal{S}}_n | v \in \mathcal{N}_{jk}\right] \ , \\ \Delta\hat{q}_{jk}^{(n)} &= \mathbb{P}\left[v \in \Delta\hat{\mathcal{S}}_n | v \in \mathcal{N}_{jk}\right] \ .\end{aligned} \tag{15}$$

Recall that $\Delta\hat{\mathcal{S}}_n = \hat{\mathcal{S}}_n \setminus \hat{\mathcal{S}}_{n-1}, n \geq 1$ and $\Delta\hat{\mathcal{S}}_0 = \hat{\mathcal{S}}_0$. Finally, we define conditional probabilities for $0 \leq m \leq n$:

$$p_{jk}^{(n,m)} = \begin{cases} \mathbb{P}\left[v \in \mathcal{D}_n | v \in \mathcal{N}_{jk} \cap \Delta\hat{\mathcal{S}}_m\right] & m = 0, \ldots, n-1 \ ; \\ \mathbb{P}\left[v \in \mathcal{D}_n | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}^c\right] & m = n \ , \end{cases} \tag{16}$$

$$\hat{p}_{jk}^{(n,m)} = \begin{cases} \mathbb{P}\left[v \in \mathcal{D}_n \not\Vdash w | v \in \mathcal{N}_{jk} \cap \Delta\hat{\mathcal{S}}_m \cap \mathcal{N}_w^+\right] & m = 0, \ldots, n-1 \ ; \\ \mathbb{P}\left[v \in \mathcal{D}_n \not\Vdash w | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}^c \cap \mathcal{N}_w^+\right] & m = n \ . \end{cases} \tag{17}$$

In terms of these quantities, we have the relations:

$$p_{jk}^{(n)} = \sum_{m=0}^{n-1} p_{jk}^{(n,m)} \Delta\hat{q}_{jk}^{(m)} + p_{jk}^{(n,n)}(1 - \hat{q}_{jk}^{(n-1)}) \ , \tag{18}$$

$$\hat{p}_{jk}^{(n)} = \sum_{m=0}^{n-1} \hat{p}_{jk}^{(n,m)} \Delta\hat{q}_{jk}^{(m)} + \hat{p}_{jk}^{(n,n)}(1 - \hat{q}_{jk}^{(n-1)}) \ . \tag{19}$$

$$u_{jk}^{(n)} = (1 - p_{jk}^{(n,n)})(1 - \hat{q}_{jk}^{(n)}) \ , \tag{20}$$

$$q_{jk}^{(n)} = 1 - u_{jk}^{(n)} - p_{jk}^{(n)} \ . \tag{21}$$

In (20), we use the fact

$$\mathbb{P}\left[v \in \mathcal{D}_n^c | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_n^c\right] = \mathbb{P}\left[v \in \mathcal{D}_n^c | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}^c\right] = (1 - p_{jk}^{(n,n)}) \ .$$

Inductively over n, it will be sufficient to compute the quantities $p_{jk}^{(n)}, \hat{q}_{jk}^{(n)}, p_{jk}^{(n,m)}, \hat{p}_{jk}^{(n,m)}$ for $m = 0, 1, \ldots, n$. Implementing these computations will require two important 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

$$\begin{aligned}\mathbb{P}\left[X \geq Y\right] = \mathbb{E}\left[\mathbb{1}_{\{X \geq Y\}}\right] &= \int_{\mathbb{R}} \int_{\mathbb{R}} \mathbb{1}_{\{X \geq Y\}} f_X(x) f_Y(y) \, \mathrm{d}x \mathrm{d}y \\ &= \int_{\mathbb{R}} F_Y(x) f_X(x) \, \mathrm{d}x = \left\langle F_Y, f_X \right\rangle \ .\end{aligned} \tag{22}$$

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

A second result from probability tells us that if $X_1, X_2, \cdots, X_n$ are n independent random variables with PDFs f_{X_i}, then the PDF of the sum $X = X_1 + X_2 + \cdots + X_n$ is the convolution

$$f_X = f_{X_1} * f_{X_2} * \cdots * f_{X_n} = \circledast_{k=1}^{n} f_{X_k} \tag{23}$$

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

The following induction over $n \geq 1$ is our main theorem.

Theorem 1. *For any $n \geq 1$, suppose $p_{jk}^{(m)}, \hat{q}_{jk}^{(m)}, p_{jk}^{(n-1,m)}, \hat{p}_{jk}^{(n-1,m)}$ are known for $m = 0, 1, \ldots, n-1$. Define the quantities*

$$p_k^{(m)} = \mathbb{P}[v \in \mathcal{D}_m | k_v = k] = \sum_j p_{jk}^{(m)} \frac{P_{jk}}{P_k^+} , \tag{24}$$

$$q_j^{(n-1)} = \mathbb{P}[v \in \mathcal{S}_{n-1} | j_v = j] = \sum_k q_{jk}^{(n-1)} \frac{P_{jk}}{P_j^-} , \tag{25}$$

$$u_j^{(n-1)} = \mathbb{P}[v \in \mathcal{U}_{n-1} | j_v = j] = \sum_k u_{jk}^{(n-1)} \frac{P_{jk}}{P_j^-} , \tag{26}$$

$$\hat{p}_j^{(n-1)} = \mathbb{P}[v \in \mathcal{D}_{n-1} \circledR w | j_v = j, v \in \mathcal{N}_w^+] = \sum_k \hat{p}_{jk}^{(n-1)} \frac{P_{jk}}{P_j^-} \tag{27}$$

as well as PDFs

$$\begin{aligned} g_j^{(n-1,m)}(x) = & \sum_{k'} \Big[(1 - p_{k'}^{(n-1)}) \delta_0(x) + p_{k'}^{(m-1)} w_{k'j}(x) \\ & + (p_{k'}^{(n-1)} - p_{k'}^{(m-1)}) \cdot \frac{1}{1-\lambda} w_{k'j}(x/(1-\lambda)) \Big] \cdot \frac{Q_{k'j}}{Q_j^-} , \end{aligned} \tag{28}$$

$$\begin{aligned} h_k^{(n-1)}(x) = & \sum_{j'} \Big[(1 - \hat{p}_{j'}^{(n-1)} - q_{j'}^{(n-1)}) \delta_0(x) + \hat{p}_{j'}^{(n-1)} w_{kj'}(x) \\ & + q_{j'}^{(n-1)} \cdot \frac{1}{\lambda} w_{kj'}(x/\lambda) \Big] \cdot \frac{Q_{kj'}}{Q_k^+} \end{aligned} \tag{29}$$

In terms of these quantities, we can compute the quantities $p_{jk}^{(n)}, q_{jk}^{(n)}$ by equations (18), (21) *where we compute $p_{jk}^{(n,m)}, \hat{p}_{jk}^{(n,m)}, \hat{q}_{jk}^{(m)}$ for $m = 0, 1, \ldots, n$ by the following equations*

$$p_{jk}^{(n,m)} = \left\langle D_{jk}, \left(g_j^{(n-1,m)}(x) \right)^{\circledast j} \right\rangle , \tag{30}$$

$$\hat{p}_{jk}^{(n,m)} = \left\langle D_{jk}, \left(g_j^{(n-1,m)}(x) \right)^{\circledast j-1} \right\rangle , \tag{31}$$

$$\hat{q}_{jk}^{(n)} = \left\langle S_{jk}, \left(h_k^{(n-1)} \right)^{\circledast k} \right\rangle . \tag{32}$$

Proof: To verify (30) for $m \leq n$, we use (3), (16) and (22) to give the formula

$$\begin{aligned} p_{jk}^{(n,m)} &= \begin{cases} \mathbb{P}[\Delta_v \leq \sum_{w \in \mathcal{N}_v^-} \Omega_{wv} \xi_{wv}^{(n)} | v \in \mathcal{N}_{jk} \cap \Delta \hat{\mathcal{S}}_m] & m < n \\ \mathbb{P}[\Delta_v \leq \sum_{w \in \mathcal{N}_v^-} \Omega_{wv} \xi_{wv}^{(n)} | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}^c] & m = n \end{cases} , \\ &= \left\langle D_{jk}, \left(g_j^{(n-1,m)}(x) \right)^{\circledast j} \right\rangle \end{aligned}$$

where

$$g_j^{(n-1,m)}(x) = \begin{cases} \frac{d}{dx} \mathbb{P}[\Omega_{wv} \xi_{wv}^{(n)} \leq x | v \in \mathcal{N}_{jk} \cap \Delta \hat{\mathcal{S}}_m, w \in \mathcal{N}_v^-] & m < n \\ \frac{d}{dx} \mathbb{P}[\Omega_{wv} \xi_{wv}^{(n)} \leq x | v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}^c, w \in \mathcal{N}_v^-] & m = n \end{cases} .$$

To verify (28) note that when $m < n$ we can write

$$g_j^{(n-1,m)}(x) = \sum_{k'} \frac{d}{dx} \mathbb{P}[\Omega_{wv} \xi_{wv}^{(n)} \leq x | v \in \mathcal{N}_{jk} \cap \Delta \hat{\mathcal{S}}_m, w \in \mathcal{N}_v^-, k_w = k'] \mathbb{P}[k_\ell = k' | j_\ell = j] .$$

Under the conditions $v \in \mathcal{N}_{jk} \cap \Delta \hat{\mathcal{S}}_m, w \in \mathcal{N}_v^-, k_w = k'$, the events $\{\xi_{wv}^{(n-1)} = 0\}$, $\{\xi_{wv}^{(n-1)} = 1\}$,$\{\xi_{wv}^{(n-1)} = 1 - \lambda\}$ are equivalent to the events $\{w \notin \mathcal{D}_{n-1}\}$, $\{w \in \mathcal{D}_{m-1}\}$, $\{w \in \mathcal{D}_{n-1} \setminus \mathcal{D}_{m-1}\}$ and hence have conditional probabilities $1 - p_{k'}^{(n-1)}$, $p_{k'}^{(m-1)}$, $p_{k'}^{(n-1)} - p_{k'}^{(m-1)}$ respectively. These latter three events are conditionally independent of Ω_{wv} by the LTI property. Also by the LTI property, the collection of random variables $\Omega_{wv} \xi_{wv}^{(n)}$ for different $w \in \mathcal{N}_v^-$ are mutually conditionally independent, justifying the use of formula (22). Similarly, when $m = n$, and under the conditions $v \in \mathcal{N}_{jk} \cap \hat{\mathcal{S}}_{n-1}$, $w \in \mathcal{N}_v^-, k_w = k'$ the events $\{\xi_{wv}^{(n-1)} = 0\}$, $\{\xi_{wv}^{(n-1)} = 1\}$ are equivalent to the events $\{w \notin \mathcal{D}_{n-1}\}$, $\{w \in \mathcal{D}_{m-1}\}$ while $\{\xi_{wv}^{(n-1)} = 1 - \lambda\}$ cannot occur. Hence the events have conditional probabilities $1 - p_{k'}^{(n-1)}$, $p_{k'}^{(n-1)}$ and 0 respectively, and the mixture CDFs $g_{jk}^{(n-1,m)}(x)$ for all $m \leq n$ are given by (28).

To verify (30), we use (8) instead of (3) and follow these same steps. To verify (32), we use (9), (15) and (22) to give the formula

$$\hat{q}_{jk}^{(n)} = \mathbb{P}[\Sigma_v \leq \sum_{w \in \mathcal{N}_v^+} \Omega_{vw} \zeta_{vw}^{(n)} | v \in \mathcal{N}_{jk}] = \left\langle S_{jk}, \left(h_k^{(n-1,m)}(x) \right)^{\circledast k} \right\rangle \tag{33}$$

where

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

To verify (29), note that under the conditions $v \in \mathcal{N}_{jk}, w \in \mathcal{N}_v^+, j_w = j'$, 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 $\{w \in \mathcal{U}_{n-1}\}$, $\{w \in \mathcal{D}_{m-1} \circledR v\}, \{w \in \mathcal{S}_{n-1}\}$ and hence have conditional probabilities $1 - \hat{p}_{j'}^{(n-1)} - q_{j'}^{(n-1)}, \hat{p}_{j'}^{(n-1)}, q_{j'}^{(n-1)}$ respectively.

□

Remark 2. *This theorem is in a non-Markovian form, as indicated by the sum over m in equations* (18) *and* (19)*. This sum will pose a difficulty in situations where we need to compute a very large number of cascade steps before converging to the fixed point. It is possible to reformulate the theorem in a Markovian form that avoids such a sum over m.*

6 Real-World Networks

The goal of the present section is to derive approximate probabilistic formulas describing the double cascade on a real-world network where the skeleton graph 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 $\mathcal{E}$. We number the nodes in $\mathcal{N}$ by $v = 1, 2, \ldots, N$ and the links by $\ell = 1, 2, \ldots, L$ where $L = \sum_{1 \leq v, v' \leq N} A_{vv'}$. 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 approximate formulas for the marginal likelihoods $p_v^{(\infty)}, q_v^{(\infty)}$ for the eventual default and stress of individual nodes, as well as the possibility to compute formulas for more detailed systemic quantities. This approximation is not exact whenever there are cycles in the skeleton graph. In general, when the skeleton graph is a single random realization from a configuration graph ensemble, we expect the LTI approximation to get better with increasing N. The LTI property will be exactly true in the special case of skeleton graphs that are trees: This fact will be used in Section 7.3 to verify the consistency of our numerical implementations.

We now present an approximate analysis, paralleling the $N = \infty$ analysis of the previous section, of the sequence of probabilities

$$\begin{aligned} p_v^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n] \ , \\ q_v^{(n)} &= \mathbb{P}[v \in \mathcal{S}_n] \ , \\ u_v^{(n)} &= \mathbb{P}[v \in \mathcal{U}_n] \end{aligned}$$

for each node v. For the same reasons as before we need in addition to track

$$\begin{aligned} \hat{q}_v^{(n)} &= \mathbb{P}[v \in \hat{\mathcal{S}}_n] \ , \\ p_v^{(n,m)} &= \begin{cases} \mathbb{P}\left[v \in \mathcal{D}_n | v \in \Delta\hat{\mathcal{S}}_m\right] & m = 0, \ldots, n-1 \ ; \\ \mathbb{P}\left[v \in \mathcal{D}_n | v \in \hat{\mathcal{S}}_{n-1}^c\right] & m = n \end{cases} \end{aligned}$$

and the WOR probabilities

$$\begin{aligned}
\hat{p}_{wv}^{(n)} &= \mathbb{P}[v \in \mathcal{D}_n \circledR w] \ , \\
\hat{p}_{wv}^{(n,m)} &= \begin{cases} \mathbb{P}\left[v \in \mathcal{D}_n \circledR w | v \in \Delta\hat{\mathcal{S}}_m \cap \mathcal{N}_w^+\right] & m = 0, \ldots, n-1; \\ \mathbb{P}\left[v \in \mathcal{D}_n \circledR w | v \in \hat{\mathcal{S}}_{n-1}^c \cap \mathcal{N}_w^+\right] & m = n \end{cases}
\end{aligned}$$

for each in-link of v. Note that

$$p_v^{(n)} = \sum_{m=0}^{n-1} p_v^{(n,m)} \Delta\hat{q}_v^{(m)} + p_v^{(n,n)}(1 - \hat{q}_v^{(n-1)}) \ , \tag{34}$$

$$\hat{p}_{wv}^{(n)} = \sum_{m=0}^{n-1} \hat{p}_{wv}^{(n,m)} \Delta\hat{q}_v^{(m)} + \hat{p}_{wv}^{(n,n)}(1 - \hat{q}_v^{(n-1)}) \ , \tag{35}$$

$$u_v^{(n)} = (1 - p_v^{(n,n)})(1 - \hat{q}_v^{(n)}) \ , \tag{36}$$

$$q_v^{(n)} = 1 - u_v^{(n)} - p_v^{(n)} \ . \tag{37}$$

Inductively, we have

$$p_v^{(n,m)} = \left\langle D_v, \circledast_{v' \in \mathcal{N}_v^-} \left(g_{v'v}^{(n-1,m)} \right) \right\rangle \ , \tag{38}$$

$$\hat{p}_{wv}^{(n,m)} = \left\langle D_v, \circledast_{v' \in \mathcal{N}_v^- \setminus w} \left(g_{v'v}^{(n-1,m)} \right) \right\rangle \ , \tag{39}$$

$$\hat{q}_v^{(n)} = \left\langle S_v, \circledast_{v' \in \mathcal{N}_v^+} \left(h_{vv'}^{(n-1)} \right) \right\rangle \ , \tag{40}$$

$$\tag{41}$$

where

$$g_{wv}^{(n,m)}(x) = \begin{cases} \frac{d}{dx}\mathbb{P}[\Omega_{wv}\xi_{wv}^{(n)} \leq x | v \in \Delta\hat{\mathcal{S}}_m \cap \mathcal{N}_w^+] & m = 0, \ldots, n-1 \ ; \\ \frac{d}{dx}\mathbb{P}[\Omega_{wv}\xi_{wv}^{(n)} \leq x | v \in \Delta\hat{\mathcal{S}}_{n-1}^c \cap \mathcal{N}_w^+] & m = n \end{cases} \tag{42}$$

$$h_{vw}^{(n)}(x) = \frac{d}{dx}\mathbb{P}[\Omega_{vw}\zeta_{vw}^{(n)} \leq x | v \in \mathcal{N}_w^-] \ . \tag{43}$$

These PDFs can be computed using the LTI approximation by following the logic of the proof of Theorem 1:

$$\begin{aligned}
g_{wv}^{(n,m)}(x) &= (1 - p_w^{(n)})\delta_0(x) + p_w^{(m-1)} w_{wv}(x) \\
&\quad + (p_w^{(n)} - p_w^{(m-1)}) \cdot \frac{1}{1-\lambda} w_{wv}(x/1-\lambda) \ ,
\end{aligned} \tag{44}$$

$$h_{vw}^{(n)}(x) = (1 - \hat{p}_{vw}^{(n-1)} - q_w^{(n-1)})\delta_0(x) + \hat{p}_{vw}^{(n-1)} w_{vw}(x) + q_w^{(n-1)} \cdot \frac{1}{\lambda} w_{vw}(x/\lambda) \tag{45}$$

7 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 double cascade in stylized networks with large values of N. Secondly, we will show how the LTI method can lead to answers to questions about the nature of systemic risk. 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.

To crossvalidate the LTI method, we developed a Monte Carlo (MC) implementation of the double cascade model, and compared the final fraction of defaulted and stressed nodes generated using the LTI and MC computations. Before we present the results of the comparison, we mention two modifications that are necessary to implement these methods. First, implementation of the LTI method uses an FFT method sketched in Appendix A, which requires that the edge weight and buffer random variables take values on a common discrete lattice $\{m\delta\}_{m=0,1,\ldots,M-1}$. The second modification is that the MC implementation generated finite configuration graphs with $N = 20000$ nodes, with the specified P and Q matrices, rather than the infinite configuration graphs assumed in the LTI method.

7.1 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 and Kapadia [2010]. It consists of a random directed Poisson skeleton graph with mean degree[6] $z = 10$, where each node v can be viewed as a bank with a default buffer $\Delta_v = 0.04$ and stress buffer $\Sigma_v = 0.02$. Unlike the deterministic weights used in Gai and Kapadia [2010], the edge weight Ω_ℓ of an edge ℓ, representing the exposure between two banks, is taken from a log normal distribution with mean $\mu_\ell = 0.2 j_\ell^{-1}$, 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 fraction of defaulted bank and stressed banks as computed using MC simulation with 1000 realizations and the LTI analytic formulas. Directed Poisson random graphs are particularly amenable to study by Monte Carlo: To generate a random graph of size N with mean degree $z > 0$ from this class, 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 independent binomials, $\mathbb{P}[v \in \mathcal{N}_{jk}] = \mathrm{Bin}(N-1, p, j) \times \mathrm{Bin}(N-1, p, k)$, which for large N nodes is approximately a product of independent Poisson(z) distributions.

[6]The mean degree is the average number of "in" edges per node, which equals the average number of "out" edges per node.

Figure 2 plots the results as functions of the stress response parameter λ, with error bars that represent the 10^{th} and 90^{th} 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 reaction 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.

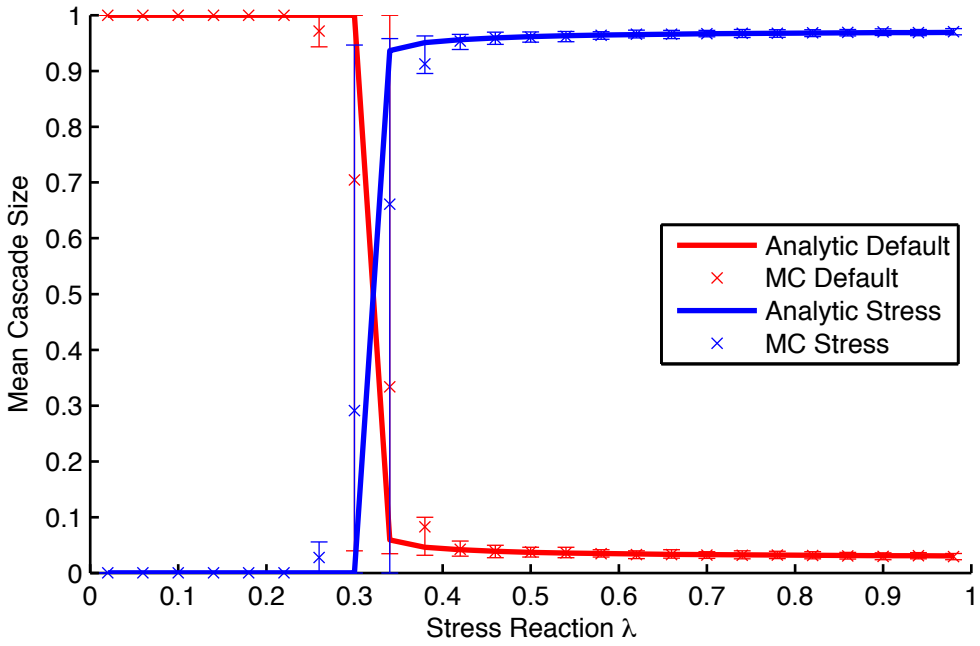

Figure 2: A comparison of the mean default (light/red) and stress (dark/blue) cascade size in Experiment 1 as computed by Monte Carlo (crosses) and LTI analytics (solid lines). Error bars indicate the 10^{th} and 90^{th} percentiles of the MC result.

7.2 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.

7.2.1 Experiment 2A: Effects of Default and Stress Buffers

We consider how the parametrized financial network of Experiment 1 in a default-susceptible state with $\lambda = 0.25$ can be made resilient to random shocks by varying either the default buffers away from 0.04 or the stress buffers away from 0.02.

Figure 3(a) imagines what would happen if regulators had required all banks to have higher default buffers, without any change in their stress behaviour. We observe a very fast transition to a stable network 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 10^{th} and 90^{th} percentiles, become very large near the knife-edge.

If banks with $\Delta = 0.04$ reduce their stress buffers below 0.02, they will react more quickly to stress shocks: This can also reduce default cascade risk in the network. Figure 3(b) shows that such a change dramatically reduces the average default cascade size in the network. Taken together, these two plots show that there may be many different approaches to dealing with network resilience.

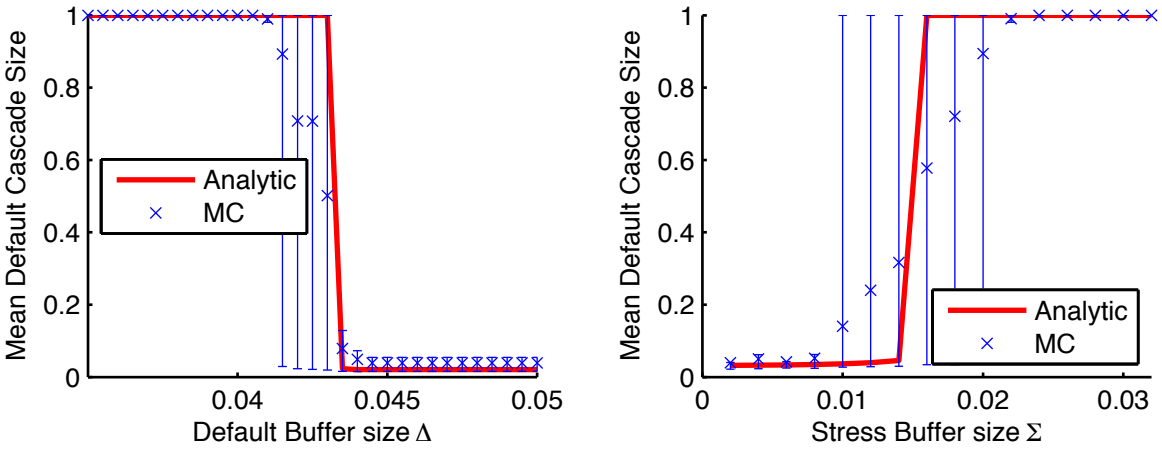

Figure 3: The results of Experiment 2A, showing the effects on the default resilience of a Poisson network when the default buffer (left) and stress buffer (right) are varied away from their benchmark values. Here $\lambda = 0.25$ and other parameters are chosen as in Experiment 1.

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

Aside from mandating changes to the behaviour of FIs during or prior to a crisis, regulators can also influence the shape of the financial network as a whole. Experiment 2B aims to demonstrate the importance of the skeleton graph itself, so we

observed the systemic risk in a directed Poisson network as a function of the connectivity parameter z and λ, the stress response. In our simple model specification, the mean degree z is the only parameter that controls the shape of the skeleton graph, whereas in a more realistic modelling approach the skeleton graph may have many more parameters.

In this experiment, we increased the model complexity by assuming each node to have 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 lognormal 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 FI has 1% chance of defaulting initially.

Figure 4(a) shows a surface plot of the mean default cascade size in the network as a function of z and λ. Figure 4(b) shows the mean stress cascade size of the network. For clarity of the graphics, we show LTI analytics only: the Monte Carlo results not shown continue to agree with LTI analytics. Again, in these plots we notice the strong anti-correlation between stress and default probabilities, and the effect of increasing the stress response. It is also interesting to observe that the final level of stress is not monotonic in the connectivity parameter z.

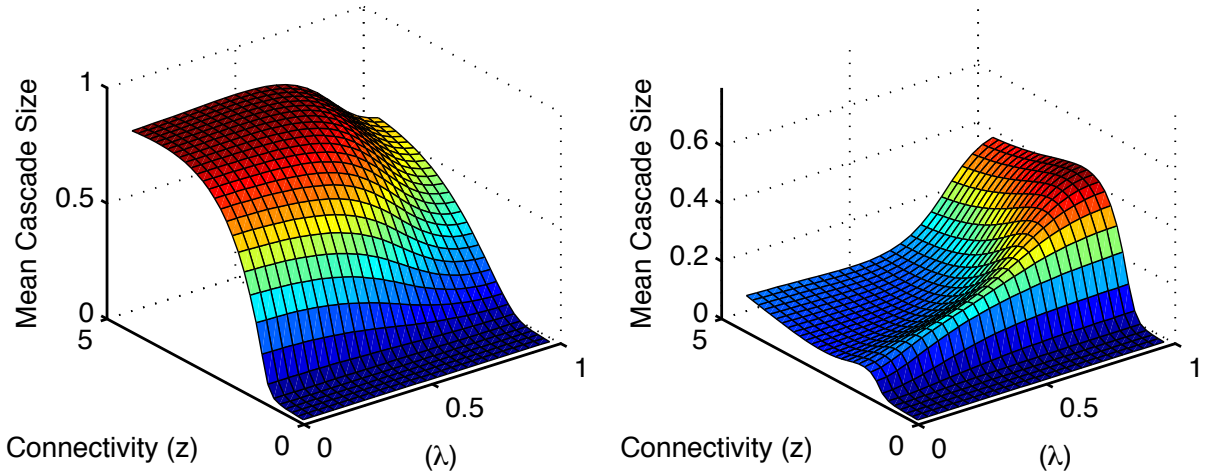

Figure 4: The results of Experiment 2B, showing the mean default (left) and stress (right) cascade sizes on a multitude of Poisson networks parametrized by connectivity z and stress response λ. Here buffers and exposures are all random.

7.3 Experiment 3: A Real-World Network with 90 Nodes

While the above experiments on hypothetical financial networks demonstrate the range of options available in our framework, we are of course very interested in having

a rough picture of the systemic risk of actual financial networks. In Experiment 3, we use the "real-world" method of Section 6 to compute the cascade dynamics on a single realization of an 90 node graph that aims to capture stylized features of the European Union network in 2011.

As pointed out in Section 6, the LTI property is exactly true for real-world tree networks. We used this fact to verify that the real-world network software used in Experiment 3 is correct on a number of tree networks. While we do not show the results here, such tests provide a strong independent validation of the real-world method.

Numerous studies of real-world financial networks, notably Bech and Atalay [2010] and Cont et al. [2010], have concluded that in and out degrees have fat tailed distributions, as do the exposure sizes, and presumably the buffers. We addressed these "stylized facts" by specifying a schematic model that intends to capture some key statistical features of data published on the 2011 ECB stress testing of 90 systemically important banks in the European Union. Our stylized model is specified as follows.

1. The skeleton graph, shown in Figure 5, is the subgraph consisting of the 90 most connected nodes of a single realization with $N = 1000$ nodes of a scale-free directed random graph drawn from the family of preferential attachment models introduced in Bollobãs et al. [2003]. Following that paper's notation, the parameters defining the skeleton graph are given by

$$\alpha = 0.169, \beta = 0.662, \gamma = 0.169$$

   and lead to fat-tailed marginal in and out degree distributions with Pareto exponents $\gamma_+, \gamma_- = 4$. Self loops are removed. The mean connectivity of the subgraph containing the 90 most connected nodes is 10.

2. The buffer random variables Δ_v and Σ_v are defined by

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

   The exposure random variables Ω_ℓ are defined by

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

   The parameters $\beta_1 = 0.3, a_1 = 8.03, b_1 = 0.9, \beta_2 = -0.2, a_2 = 8.75, b_2 = 1.16$ are determined by matching moments of the aggregated interbank exposure data[7].

3. The collection $\{X_v, \tilde{X}_v, X_\ell\}$ consists of independent standard normal random variables.

Figure 6(a) presents the mean default and stress cascade sizes that result in a numerical experiment that shocked one bank at random to default. The real-world

[7]Data are available at
http://www.eba.europa.eu/EU-wide-stress-testing/2011/2011-EU-wide-stress-test-results.aspx.

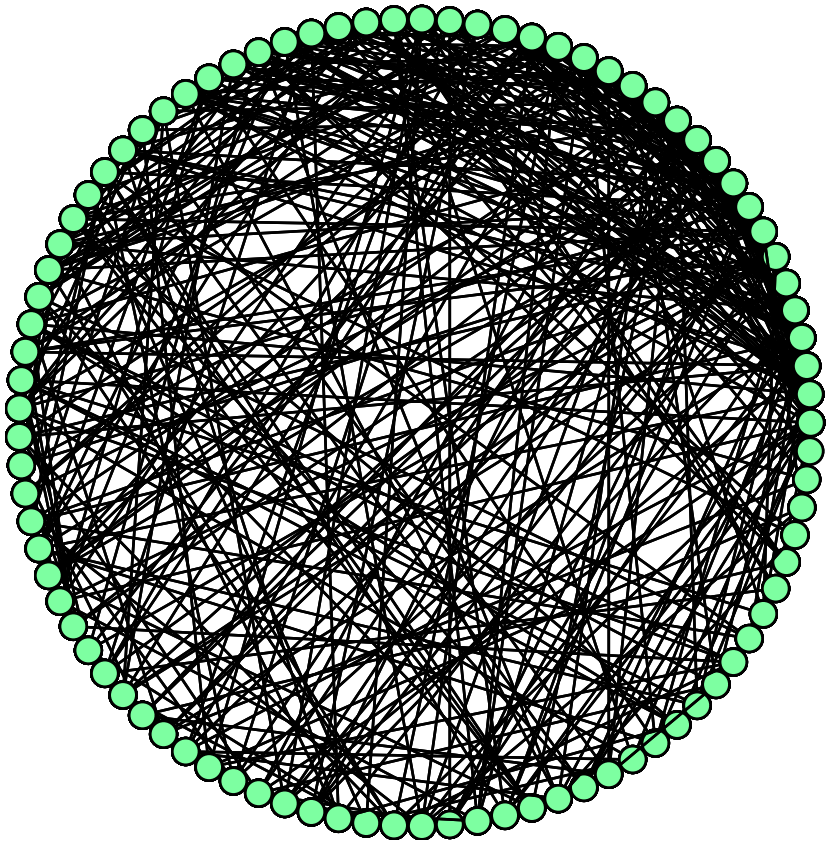

Figure 5: A representation of the undirected skeleton graph of the 90 bank network of Experiment 3. The nodes are plotted with total degree decreasing in the counterclockwise direction, with the maximally connected bank being the rightmost node.

analytical computations used follow the method of Section 6, and comparable MC computations were also performed. This graph shows that the MC and analytic computations agree that the EU network in 2011 was resilient to such a shock.

To move this network to a knife-edge situation where a large scale double cascade is triggered by a single bank default, we found it was necessary to imagine a dire crisis where prior to the default shock, the T1 capital of all institutions has been decimated to 1/10th of their initial amount. Figure 6(b) shows the behaviour of the resultant cascade as a function of the stress response parameter λ. This graph also shows limitations of the LTI approximation. In this specification of the network, while the LTI and MC methods agree that there is an intermediate size cascade, they disagree strongly on the actual mean default cascade size. This discrepancy is an example of a general tendency we observed: LTI and MC tend to agree best when the network is far from the critical surface where the cascade is intermediate in size. In Figure 6(b), both LTI and MC agree that the double cascade is intermediate in size, that is, that

the probability of banks becoming either stressed or defaulted is not close to either 0 or 1. However, they disagree strongly on the precise size of the cascade.

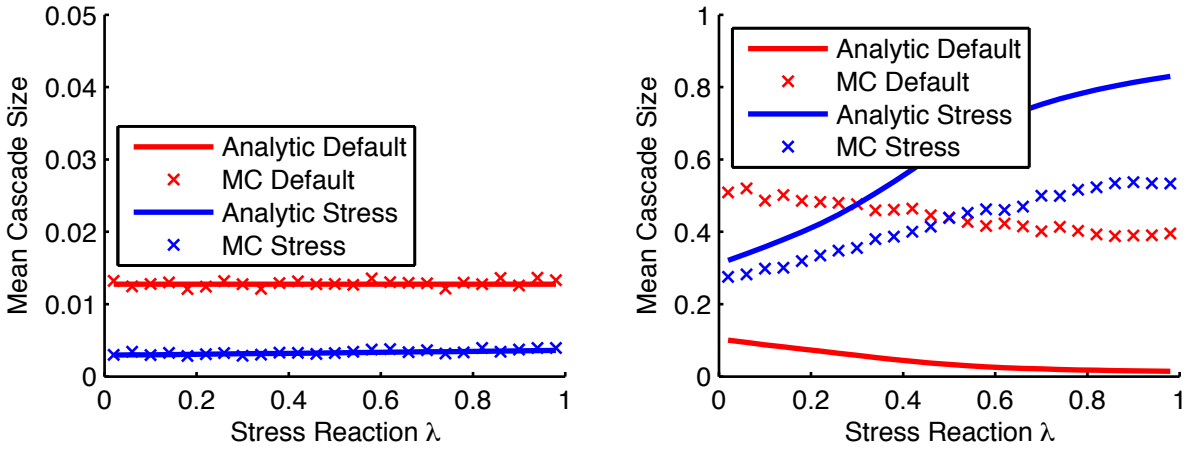

Figure 6: The results of Experiment 3 that show the effects of changing the stress response λ in a stylized finite network. The left plot corresponds to parameters that represent the EU financial system at the time of the 2011 stress testing exercise. The right plot shows the same system where a dire pre-shock crisis has decimated the banks' default buffers.

Finally, in real-world networks, we can get a picture of which banks are most susceptible to the default of the single bank. Figure 7 shows the eventual stress and default probabilities bank by bank, in increasing order of the LTI estimated default probability. We see that LTI analytics and MC agree quite well on the ordering of the banks, on the eventual probability of stress, but not well on the eventual probability of default.

8 Conclusions

The double cascade model we present in this paper is a natural extension of the previous systemic risk research that studies deliberately simplified models which build in either default or stress cascades, but not both. Only by combining the default and stress mechanisms into a single model, a non-trivial accomplishment given their fundamentally opposing natures, can one determine quantitatively 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 our double cascade model is a difficult undertaking. 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. In general, having these two approaches allows us to independently cross validate both methods.

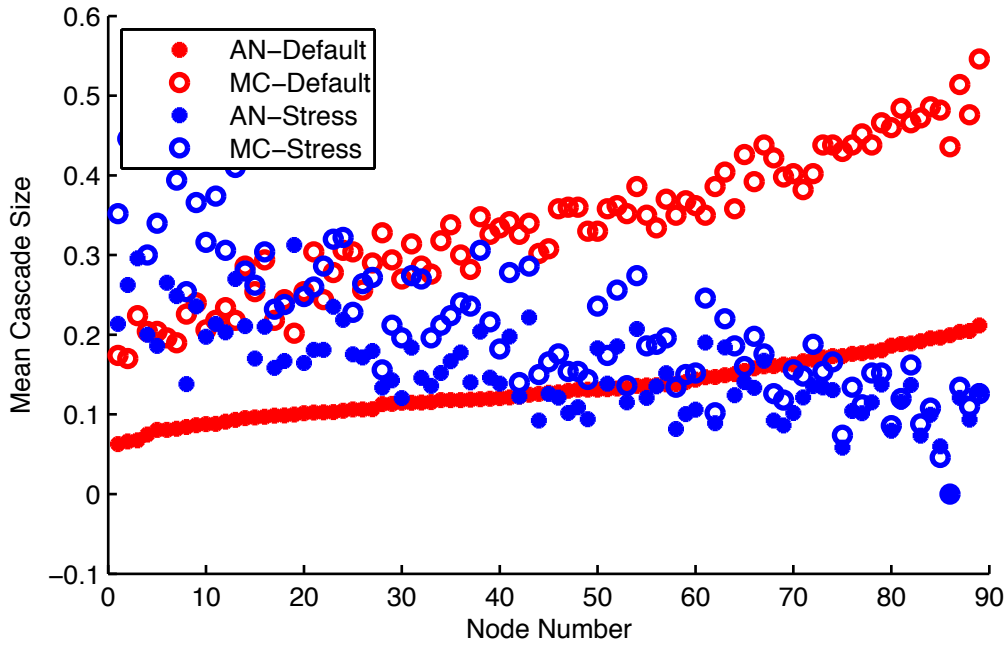

Figure 7: The eventual default and stress probabilities of individual banks after a severe crisis in the stylized EU financial system. The banks are numbered in increasing order of their default probability as computed by the real-world LTI method.

This validation capability leads to a dramatic increase in confidence in the results one obtains.

As to the pros and cons of MC versus LTI, we mention some of the key issues. To counter the natural flexibility of MC methods, the LTI method, where it applies, adds the possibility to better understand the flow of the cascade. For example, using LTI one can determine sensitivities to changing parameters through explicit differentiation. A con for the MC method is that simulating general assortative (P, Q) configuration models has not been well-studied in the literature, while the LTI method handles this generality without difficulty. On the other hand, relevant random graphs such as preferential attachment models have a straightforward MC simulation algorithm, but are not LTI. Another pro for MC is that the LTI approximation is uncontrolled for finite N configuration graphs, meaning we can only learn how accurate it is by comparing to MC results. Through experience, we are learning rules of thumb for when LTI gives acceptable results, for example when N is large and the cascade is far from critical. Apart from the graph generation step, MC is usually easier to program than LTI. But in some situations where the MC method leads to unacceptably long run times, LTI can be computed in seconds. Ultimately what is important is that both methods have complementary strengths and weaknesses, and

when used in combination we can arrive at robust and reliable conclusions about a wide range of network effects.

Future theoretical work must proceed by improving and extending both the MC and LTI methods. For example, we need better MC algorithms for simulating different random graph models and we would like to find computable analytical methods for networks that are not LTI.

Many promising specifications of financial networks remain 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. We hope we have demonstrated that the double cascade model developed in this paper has the potential to realistically represent observed financial networks, and that further investigations of such networks will uncover interesting and unexpected systemic phenomena.

9 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 (T.R.H.), from the Natural Sciences and Engineering Research Council of Canada (T.R.H.), from the Irish Research Council co-funded by Marie Curie Actions under FP7 (INSPIRE fellowship, S.M.), from the FET-Proactive project PLEXMATH (D.C.), and from the Science Foundation Ireland (11/PI/1026, D.C. and S.M.)

We are deeply indebted to Lionel Cassier of École Polytechnique France, Matheus Grasselli and Bernardo Costa Lima, both 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 lengthy discussions about the EU financial network, and to participants of the INET Workshop "Interlinkages and Systemic Risk" that took place on July 4-5, 2013 in Ancona, Italy and where this paper was presented.

A Discrete Probability Distributions and the Fast Fourier Transform

Numerical implementation of these models follows the methods outlined in Hurd and Gleeson [2013]. In this section, we analyze the case where the random variables $\{\Delta_v, \Sigma_v, \Omega_\ell\}$ all take values in a specific finite discrete set $\mathcal{M} = \{0, 1, \dots, (M-1)\}$ with a large value M. In such a situation, the convolutions in (23) 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

$$(f * g)(n) = \sum_{m=0}^{n} f(m)g(n-m) \tag{46}$$

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

$$\hat{a}_k = \sum_{l \in \mathcal{M}} \zeta_{kl} a_l \ , k \in \mathcal{M} \ .$$

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 $a \to \tilde{a} = \mathcal{G}(a)$ where

$$\tilde{a}_k = \frac{1}{M} \sum_{l \in \mathcal{M}} \bar{\zeta}_{kl} a_l \ , k \in \mathcal{M} \ .$$

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

$$\langle a, b \rangle := \sum_{m \in \mathcal{M}} \bar{a}_m b_m$$

We also define the convolution product of two vectors:

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

Note that this agrees with (46) 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 $a, b \in \mathbb{C}^M$: (i) Inverse mappings: $a = \mathcal{G}(\mathcal{F}(a)) = \mathcal{F}(\mathcal{G}(a))$; (ii) Conjugation: $\overline{\mathcal{G}(a)} = \frac{1}{M}\mathcal{F}(\bar{a})$; (iii) Parseval Identity:

$\langle a, b\rangle = M\langle \tilde{a}, \tilde{b}\rangle = \frac{1}{M}\langle \hat{a}, \hat{b}\rangle$; (iv) Convolution Identities: $\tilde{a} \cdot * \tilde{b} = \widetilde{(a * b)},\ \hat{a} \cdot * \hat{b} = \widehat{(a * b)}$, where $\cdot *$ denotes the component-wise product.

As an example to illustrate how the above formulas help, we observe that a typical formula (30) can be computed instead by the formula

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

where $\hat{D} = \mathcal{F}(D), \hat{g}_j^{(n-1,m)} = \mathcal{F}(g_j^{(n-1,m)})$ 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.

References

T. Adrian and H. S. Shin. Liquidity and Leverage. *Journal of Financial Intermediation*, 19(3):418–437, July 2010.

H. Amini, R. Cont, and A. Minca. Stress Testing the Resilience of Financial Networks. *International Journal of Theoretical and Applied Finance*, 15:1–20, 2012.

M. Bech and E. Atalay. The topology of the federal funds market. *Physica A: Statistical Mechanics and its Applications*, 389(22):5223–5246, 2010.

B. Bollobãs. *Random Graphs*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2nd edition, 2001.

B. Bollobãs, C. Borgs, J. T. Chayes, and O. Riordan. Directed scale-free graphs. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete Algorithms*, SIAM, p. 132–139, 2003.

M. K. Brunnermeier and L. H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 22(6):2201–2238, 2009.

R. Cifuentes, G. Ferrucci, and H. S. Shin. Liquidity risk and contagion. *Journal of the European Economic Association*, 5:556–566, 2005.

R. Cont, A. Moussa, and E. B. Santos. Network Structure and Systemic Risk in Banking Systems. *SSRN eLibrary*, 2010.

L. Eisenberg and T. H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001.

P. Gai and S. Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A*, 466(2120):2401–2423, 2010.

P. Gai, A. Haldane, and S. Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58:453–470, 2011.

T. R. Hurd and J. P. Gleeson. A framework for analyzing contagion in banking networks. arXiv:1110.4312 [q-fin.GN], October 2011.

T. R. Hurd and J. P. Gleeson. On Watts cascade model with random link weights. *Journal of Complex Networks*, 1(1):25–43, 2013.

R. M. May and 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, and A. Alentorn. Network Models and Financial Stability. *J. Econ. Dyn. Control*, 31:2033–2060, 2007.

S. L. Schwarcz. Systemic risk. *Georgetown Law Journal*, 97(1), 2008.

D. J. Watts. A simple model of global cascades on random networks. *PNAS*, 99(9): 5766–5771, 2002.