Algorithmic Finance 5 (2016) 95–110
DOI:10.3233/AF-160166
IOS Press

Sensitivity and computational complexity in financial networks

Brett Hemenway*,[1] and Sanjeev Khanna[2]
Department of Computer Science, University of Pennsylvania, Philadelphia, PA, USA

Abstract. Determining the causes of instability and contagion in financial networks is necessary to inform policy and avoid future financial collapse.

In the American Economic Review, Elliott, Golub and Jackson proposed a simple model for capturing the dynamics of complex financial networks. In Elliott, Golub and Jackson's model, the institutions in the network are connected by linear dependencies (cross-holdings) and if any institution's value drops below a critical threshold, its value suffers an additional failure cost. This work shows that even in this simple model there are fundamental barriers to understanding the risks that are inherent in a network.

First, if institutions are not required to maintain a minimum amount of self-holdings, any change in investments by a single institution can have an arbitrarily magnified influence on the net worth of the institutions in the system. This implies that if institutions have small self-holdings, then estimating the market value of an institution requires almost perfect information about *every cross-holding in the system*.

Second, even if a regulator has complete information about all cross-holdings in the system, it may be computationally intractable to estimate the number of failures that could be caused by a small shock to the system.

Keywords: Financial contagion, computational complexity, network analysis, network stability, sensitivity

1. Introduction

The recent financial crisis and subsequent bailout have highlighted the need for a better understanding of the dynamics of financial networks. Indeed, the complexity of modern financial networks has been blamed for our collective failure to recognize the presence of serious risks in these systems. In this work, we show that even in extremely simple financial networks understanding the risk present in the system is computationally intractable in general – even in the presence of perfect information about all participants in the system. We hope that these insights will help regulators and policymakers to better understand the dynamics of financial networks.

Understanding an individual's risk in a financial network is a difficult task, because an institution's ability to fulfill its outgoing financial obligations is not a local property, *i.e.,* it cannot be understood by examining a single individual in isolation. An institution's ability to make its outgoing payments may depend on whether its incoming payments are made by its debtors, which in turn may depend on whether the incoming obligations are made to those institutions, etc. Thus each institution, acting without a global view of the network, cannot effectively understand its risk. Nevertheless, one might hope that a regulator, with a global view of the network could better understand the opportunities and risks inherent in the system.

Even with a global view, the situation remains relatively complex. One complicating factor is the

*Corresponding author: Brett Hemenway, Department of Computer Science, University of Pennsylvania, Levine Hall 3330 Walnut Street, Philadelphia, PA, USA. Tel.: +1 215 746 1738; Fax: +1 215 898 0587; E-mail: fbrett@cis.upenn.edu.

[1]Supported in part by a grant from the Jacobs-Levy Foundation.

[2]Supported in part by National Science Foundation grants CCF-1552909 and CCF-1617851, and a grant from the Jacobs-Levy foundation. E-mail: Sanjeev@cis.upenn.edu.

2158-5571/16/$35.00 © 2016 – IOS Press and the authors. All rights reserved

existence of cycles in the financial network, for example institution A may have obligations to institution B which has obligations to institution C, which in turn has obligations to institution A. The existence of these cyclical interdependencies has been put forward as one of the primary sources of complexity in financial networks. In this work, we will examine network dynamics in both cyclic (Theorem 1) and acyclic networks (Theorem 3).

One of the driving forces in the study of financial networks is their ability to magnify risk: if an institution, A, defaults on its obligations to institution B, this may cause institution B to default on its obligations to institution C etc. The spread of risk through a financial network is known as *financial contagion* and has been carefully modeled and studied (Allen and Gale, 2000; Eisenberg and Noe, 2001; Gouriéroux, Héam, and Monfort, 2012; Acemoglu, Ozdaglar, and Tahbaz-Salehi, 2015; Glasserman and Young, 2014; Elliott, Golub, and Jackson, 2014).

This work focuses on quantifying the ability of financial networks to amplify and conceal risk.

2. Our contributions

In this work, we study two questions related to the stability of financial networks. First, we look at how sensitive market valuations can be to small changes in network structure. Second, we examine the computational complexity of determining how far a given network is from a massive failure. Throughout this work we use the network model put forward by Elliott (Golub, and Jackson, 2014). In this model, financial institutions own shares of underlying assets, and the institutions are connected by cross-holdings, which are modeled as linear dependencies. If an institution's market value drops below a certain critical threshold, its value suffers a further discontinuous shock, modeling the effects of a loss of investor confidence, or the failure to pay everyday operating costs. See (Elliott, Golub, and Jackson, 2014) for an in-depth discussion of the interpretation and real-world validity of this model. The formal mathematical model is described in detail in Section 4.

Our first result (Theorem 2), shows that financial networks can be highly sensitive to small changes in their link structure. Concretely, we show that if a single institution changes a single cross-holding by ϵ, the market values of institutions in the system can change by as much as $\epsilon/2r$, where $r \leq 1$ is the minimum amount of self-holdings of the institutions in the network. The minimum self-holdings, r, is a measure of *integration* of the network, where $r = 0$ corresponds to a fully integrated network, and $r = 1$ corresponds to a network with no integration. This result shows that if each institution retains only, say, 5% self-ownership, a change in a single holding by ϵ can result in a 20ϵ change in an institution's market value. This amplification is directly caused by cycles in the network, and in acyclic networks this type of magnification cannot occur (see Lemma 1). Our bounds are essentially tight, and Theorems 1 and 2, show that the true sensitivity is in fact $\Theta(\epsilon/r)$.

This sensitivity magnification has many consequences. First, it means that in order to estimate market values of institutions within the system, all cross-holdings need to be known to an extremely high degree of precision. Investors or regulators who wish to calculate market values can be extremely far off if even a single cross-holding in the network remains unknown to them. Second, because small changes in a single institution's investments can have large effects on market values throughout the system, this indicates a potential for extreme instability in the system as a whole; the small portfolio changes in one institution can have a drastically magnified effect market values of other institutions, and so small changes by a myopic institution could topple even seemingly stable institutions. Third, this extreme sensitivity means that calculating market values in a privacy-preserving manner can be extremely difficult (Narayan, Papadimitriou, and Haeberlen, 2014). This is the flip-side of the first point, in order to calculate market values for institutions in the system, *all* the interbank holdings need to be known with high precision, which means that revealing market prices has the potential to reveal extremely detailed information about each institution's interbank holdings.

Our second result addresses the question of how well a regulator can assess the stability of a network. Suppose a regulator or oversight agency is presented with a network in which every bank is solvent, and suppose the regulator believes that the underlying assets cannot drop in value by more than some fixed amount d. What is the maximum number of failures that can be caused by a drop of this size in asset values? We emphasize that in this scenario, the regulator has complete information about the entire structure of the financial network, and the only uncertainty is in which specific assets may decline in value. We show calculating the number of institutions that can fail in a network is NP-Hard. In particular, we show

that is as hard as calculating the maximum balanced clique in a bipartite graph (Theorem 3). The maximum balanced bipartite subgraph (BCBS) problem is NP-hard (Gary and Johnson, 1979; D.S. Johnson, 1987), and there is evidence that it is even hard to approximate. It is known that if 3-SAT is not in $DTIME\left(2^{n^{3/4+\epsilon}}\right)$ for some ϵ, then there is no polynomial time algorithm for calculating the maximum balanced clique to within a factor of $2^{(\log n)^{\delta}}$ for some $\delta > 0$. Feige and Kogan (2004) go further, conjecturing that there is no polynomial time algorithm to approximate the maximum balanced bipartite clique to within a factor of n^{δ} for some $\delta > 0$. This would imply that there is no polynomial time algorithm that can even *estimate* the maximum number of failures caused by a drop in asset values of a given magnitude to within a factor of n^{δ} in a network with $2n$ institutions. In particular, this means that there are financial networks where stress testing (which is inherently computationally feasible) cannot hope to even approximate the magnitude of collapse that could be caused by some bounded drop in asset values.

Unlike our first result, which crucially relies on cycles in the network, this result holds even in acyclic networks: even when there are no cycles it is computationally intractable to estimate the maximum number of failures that could be caused by a bounded drop in asset values. This complexity arises not from cycles, but from the nonlinear dynamics that occur when a bank drops below its critical threshold value.

3. Previous work

Many different models have been proposed to study financial networks, and specifically models of stability and contagion.

Allen and Gale (1998) considered a model consisting of depositors and banks. Depositors deposit their money in the banks, and the banks must choose between making *short-term* investments or *long-term* investments. This model has three time periods, $t = 0, 1, 2$. At $t = 0$, investments are made, at $t = 1$ short-term investments pay off, and depositors choose whether to withdraw their money, and at time $t = 2$ long-term investments pay off. The banks' investment strategies then depend on the probability that depositors withdraw their money at time $t = 1$. In a follow-up work Allen and Gale (2000) introduced a network component, where banks can exchange deposits with each other in an effort to mitigate risk, and they showed simple contagion effects in this model.

Acemoglu, Ozdaglar, and Tahbaz-Salehi (2015) build on Allen and Gale's 3 time-step model. At time $t = 0$ banks can make a short-term investment, a long-term investment or loan money to other banks. Long term investments yield a fixed return at $t = 2$. Long term investments that are liquidated at $t = 1$ receive a return that is randomly distributed between two values. Banks whose investments returned the lower of the two values were said to have received a *shock*. Acemoglu, Ozdaglar and Tahbaz-Salehi considered how two extremal types of networks serve to propagate these shocks. They considered the ring (where each bank only has debts to its two neighbors) and the complete network where each bank's debts are spread to all other banks, as well as all convex combinations of these two. They showed that if the magnitude of the shocks are small then the complete network is more stable and resilient than the ring, and if the magnitudes of the shock are large enough then both the ring and complete networks are the *least* stable networks, thus the complete network exhibits a phase transition, moving from the most stable to the least stable network as the magnitude of the shocks increase.

Eisenberg and Noe (2001) developed a very simple and appealing network model where each bank has cash reserves and fixed debts to other banks. Eisenberg and Noe's work focused on showing that (under some basic restrictions on the network) there is always a unique clearing vector (indicating how much of its debts each bank pays to its creditors), and they gave a linear program and simple iterative algorithm for calculating this clearing vector and hence the equilibrium valuation of each bank.

Gai and Kapadia (2010) considered a modification of Eisenberg and Noe's network model, forcing all incoming edges to have the same weight, but allowed banks to have additional illiquid assets. Gai and Kapadia then considered the question of how a single bank failure propagates through a network. For this analysis, they considered different models for generating the underlying graph topology, and plotted contagion effects for different network models characterized by their degree distribution. They found that a single large shock could have devastating effects on the network, but that this was highly dependent on where in the network the shock hit.

Gouriéroux, Héam, and Monfort (2012) considered a model that allows interbank investment via shares (like (Elliott, Golub, and Jackson, 2014))

and lending or insurance (like (Eisenberg and Noe, 2001)). Unlike (Elliott, Golub, and Jackson, 2014), they do not introduce discontinuous failure costs. Gourieroux, Héam and Monfort extend Eisenberg and Noe's uniqueness results to show that (under mild constraints on the network) this extended model has a unique equilibrium value for all institutions. They then examined the effects of exogenous shocks on the network (*i.e.*, drops in asset values) using synthetic data and data obtained from the French banking sector.

Morris (2000) studied contagion in local interaction systems. In that model, each node represents a player, and each player engages in a local game with each of its neighbors. Morris focused on the case where each player has a binary strategy space, {0, 1}, and there is a global 2×2 payoff matrix, such that for each edge in the graph, the two neigboring players receive a payoff according to this global payoff matrix. In that model, a strategy is said to be "contagious," if it can spread from a finite set of players to an infinite set of players by the best-response dynamics of the underlying local game.

The notion of failure cascades and contagion have also been studied in the computer science literature by Blume et al. (2011), Blume et al. considered general cascades in graphs where the edges were unweighted and a node was said to fail if some critical threshold of its neighbors failed. By choosing all edges to have equal weight, and choosing each institution's failure threshold carefully, the failure model of (Elliott, Golub, and Jackson, 2014) can be made to overlap with this general network failure model.

In this work, we use the model of Elliott (Golub, and Jackson, 2014). In this model, institutions can own shares in each other, or in "primitive assets" that have intrinsic value outside of the network. The model is explained in detail in the next section. Elliott, Golub and Jackson introduced this model to help analyze and understand contagion effects in networks. Because a cascade of collapse requires an initial failure, Elliott, Golub and Jackson began by showing that the weakest institution can never be made strictly more stable by any fair trade between the institutions. Next they examined the contagion dynamics in a wide class of networks parametrized by "integration" and "diversification." Integration increases as institutions in the network increase their inter-network holdings, *i.e.*, integration increases as the percentage of each institution owned by shareholders external to the network decreases. As integration increases, the institutions' fates are more closely tied together. Diversification measures how risk is spread within the network. Diversification increases as institutions increase their *number* of cross-holdings. Neither integration nor diversification have strictly positive or negative effects on network stability, but instead have slightly more complex non-monotonic effects.

The notion of computational complexity has been studied in the context of financial products by Arora et al. who showed that banks can create derivatives that are computational intractable to price accurately (Arora et al., 2010, 2011). The result of Arora et al. crucially relies on the information asymmetry between the institution the seller (who creates the derivatives) and the buyer who only sees their resulting composition. Braverman and Pasricha (2014) show that even in the full information setting pricing compound options is PSPACE complete.

4. Model

We use the model put forward by Elliott (Golub, and Jackson, 2014). In this model there are n financial institutions, these can be viewed as countries, banks or private firms, and m underlying assets, that can be viewed as any object or project with intrinsic value. The financial institutions own shares of the underlying assets which impart value into the system. The values of the institutions themselves are interconnected via a network (modeled as a weighted, directed graph). The interdependencies (cross-holdings) between the institutions are modeled as simple linear dependencies. These linear cross-holdings can model simple equity stakes (one institution owning shares in another) or they can be viewed as an approximation of more complicated debt contracts between the institutions (see (Elliott, Golub, and Jackson, 2014, Section 2.5) for a more thorough discussion of the validity and generality of this model).

Although institutions invest in one another, all value in the system originates from the underlying assets. The price of asset k is denoted by p_k, and we use $D_{ik} \geq 0$ to denote the percentage of asset k owned by institution i. The $n \times m$ matrix of ownership is denoted by $\mathbf{D} = (D_{ik})$.

We define $\mathbf{C} = (C_{ij})$ to be the $n \times n$ matrix indicating the cross-holdings of institutions. Thus institution i owns a C_{ij} fraction of institution j. It will be be useful to view the network of cross-holdings as a directed graph with n nodes representing the financial institutions, and an edge from institution j to i of weight C_{ij} whenever $C_{ij} > 0$. Following (Elliott, Golub, and

Jackson, 2014), we set $C_{ii} = 0$ for all i. Now, $\sum_i C_{ij}$ is the fraction of institution j that is owned by institutions external to j. The remainder, the amount of self-ownership, is denoted by $\hat{C}_{jj} \stackrel{\text{def}}{=} 1 - \sum_i C_{ij}$. The matrix $\hat{\mathbf{C}}$ will be a diagonal matrix with $\hat{C}_{ii}$ on the diagonal.

As noted by Brioschi (Buzzacchi), this type of model introduces two types of valuations, the *equity valuation* ($\vec{V}$) and the *market valuation* ($\vec{v}$). The equity valuation of institution i is denoted by

$$V_i = \underbrace{\sum_k D_{ik} p_k}_{\text{Value of assets held by } i} + \underbrace{\sum_j C_{ij} V_j}_{\text{Values of institutions held by } i} \quad (1)$$

In matrix notation, this becomes $\vec{V} = \mathbf{D}\vec{p} + \mathbf{C}\vec{V}$ which implies $\vec{V} = (\mathbf{I} - \mathbf{C})^{-1}\mathbf{D}\vec{p}$. The matrix $\mathbf{I} - \mathbf{C}$ is guaranteed to be invertible because we assume that $\hat{C}_{jj} > 0$, so the column sums of $\mathbf{C}$ are all strictly less than one (see Lemma 2). In fact, the matrix $\mathbf{I} - \mathbf{C}$, is an M-Matrix (Poole and Boullion, 1974), and so $(\mathbf{I} - \mathbf{C})^{-1}$ is an inverse M-Matrix, about which many properties are known (Willoughby, 1977; Johnson, 1982).

This equity valuation significantly overvalues the institutions. In particular, we can see that $\|\vec{V}\|_1 \geq \|\vec{p}\|_1$, so the total value of the institutions in the system will (in general) be much larger than the total value of the underlying assets. This occurs because each asset counts towards the equity value of the institution that owns it and also to the institutions that have an equity stake in the asset's owner. The network's inflation of equity values is well-known and validated both theoretically and empirically (French and Poterba, 1991; Fedenia, Hodder, and Triantis, 1994).

To find an institution's *market value*, we must scale the institution's equity value by the percent stake it has in itself, thus the market value of institution i is $v_i = \hat{C}_{ii} V_i$, so the market values are the solution to the system

$$\vec{v} = \hat{\mathbf{C}}\vec{V} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{D}\vec{p} \quad (2)$$

The matrix $\mathbf{C}$ is column sub-stochastic because column i sums to $1 - \hat{C}_{ii}$. The system can also be viewed as a flow, where at each time step money flows between banks according to the link structure of the network (see Appendix A).

5. Sensitivity

Our first result concerns concerns the sensitivity of valuations to small changes in the structure of the network. Suppose a single institution shifts its holdings by a small quantity, ϵ, how much can this small change affect the market valuations in the network? This question is motivated by questions about network stability, and the possibility of privacy-preserving oversight. If small changes in network holdings can lead to large changes in the market values of the institutions, this indicates a fundamental instability in the financial network. Additionally, if small changes in interbank holdings can lead to large changes in market values, then any attempt at financial oversight must incorporate all the interbank holdings to a high degree of accuracy in order to predict market values.

A high sensitivity also has implications towards privately computing network statistics. Flood et al. (2013) proposed using tools from differential privacy (Dwork, 2006; Dwork et al., 2006) to provide a means of computing global network characteristics while preserving the privacy of each individual institution's holdings. A high sensitivity implies a worse trade-off between privacy and accuracy when calculating network statistics.

We begin our sensitivity analysis with a simple observation: if the total value of the underlying assets is $\|\vec{p}\|$, then an ϵ change in holdings can easily change the market valuations by $\epsilon\|\vec{p}\|$ (see Fig.1). We show that if the network is acyclic, then this is the largest change possible, but if there are cycles in the network, the sensitivity can be much larger.

Throughout this section, we use r (for "reserve") to measure the fraction of each institution held by investors *outside* the system.[1] We define $r = \min_i \hat{C}_{ii}$. Using the terminology of (Elliott, Golub, and Jackson, 2014), r is just a concrete metric of the *integration* of the network, and integration increases as $r \to 0$. Another interpretation of reserve is the discrepancy between equity valuation ($\vec{V}$) and market valuation ($\vec{v}$). Since $\vec{v} = \hat{C}\vec{V}$, we have that $r = \min_i \frac{v_i}{V_i}$.

5.1. Sensitivity in acyclic networks

We begin by noting that in the acyclic case, there is a strong bound on each institution's equity valuation,

[1] The reserve, or self-holdings, can be viewed as the amount of an institution that is not sold, or is held by private shareholders, who retain complete ownership of themselves. These private shareholders buy shares of institutions in the network, but no entity in the network owns shares of the private shareholders.

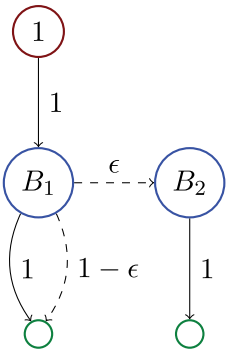

Fig. 1. An ϵ change results in an $\epsilon\|\vec{p}\|$ change in market values. Banks are in blue, external shareholders are in green, and the asset is shown in red. In this example $\mathbf{D}\vec{p} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$, and $\mathbf{C} = \mathbf{0}$. Thus $\vec{v} = \begin{bmatrix} 1 \\ 0 \end{bmatrix}$. Changing $\mathbf{C}$ to $\tilde{\mathbf{C}} = \begin{bmatrix} 0 & \epsilon \\ 0 & 0 \end{bmatrix}$ leads to a valuation of $\tilde{\vec{v}} = \begin{bmatrix} 1-\epsilon \\ \epsilon \end{bmatrix}$, thus the market valuation of B_2 changes by $\epsilon\|\vec{p}\|$.

i.e., the equity valuation cannot be too much larger than the market valuation.

Lemma 1. If the banking network has no cycles, then every institution's equity valuation is at most $\|\vec{p}\|_1$ where $\vec{p}$ is the vector of asset values.

Proof. Organize the financial network into layers, so that each institution only owns shares of institutions at lower layers. Thus institutions at level 1 do not have any cross-holdings, they only own the underlying assets. This means that the incoming edges to level one carry a total weight of at most $\|\vec{p}\|_1$. By adding fictitious institutions that pass on all of their incoming wealth, we can also ensure that every institution only owns shares in the entities at the preceding level, *i.e.*, institutions at level i only own shares in institutions (or fictitious institutions) at level $i-1$.

Now each institution's equity value is the sum of the values on all incoming edges. Since outgoing edges carry a value that is a percentage of equity value, the sum of the values on each institution's outgoing edges is at most the sum of the values on its incoming edges. (For real institutions the outgoing sum will be strictly less because the reserve rate $r > 0$, but for the fictitious institutions it can be exactly equal.)

Now, the sum of the values coming into layer 1 is at most the sum of the assets, $\|\vec{p}\|_1$. Thus the outgoing edges from level 1 to level 2 carry a total weight of at most $\|\vec{p}\|_1$. Proceeding inductively through the levels, we see that the sum of the values on the incoming edges at level i is at most $\|\vec{p}\|_1$. Thus the equity value of all the institutions on level i is at most $\|\vec{p}\|_1$ and in particular the equity value of any given institution is at most $\|\vec{p}\|_1$.

Corollary 1. If the banking network is acyclic, and one edge changes by at most ϵ, then no institution's market value can change by more than $\epsilon\|\vec{p}\|_1$.

Proof. Since each institution's equity value is at most $\|\vec{p}\|_1$, an ϵ change in any edge corresponds to an absolute change of at most $\epsilon\|\vec{p}\|_1$.

5.2. Sensitivity in general networks

In this section, we explore how much the market valuations can change when one bank changes its holdings by a small amount *in the presence of cycles in the network graph*. We begin by showing an upper bound on the change in market valuations that depends on the minimum self-ownership ($\hat{C}_{ii}$) of the institutions. For our upper bound, we do not require changes to occur in the holdings of a single bank. Instead, we allow any change in network structure, as long as the total (ℓ_1) change is bounded by ϵ. Formally, this means that we have two network matrices $\mathbf{C}$ and $\tilde{\mathbf{C}}$ such that $\|\mathbf{C} - \tilde{\mathbf{C}}\| \leq \epsilon$, and we would like to bound how much the market valuations can change between these two situations. Because we are showing an upper bound, allowing more general perturbations only strengthens our result.

Theorem 1. If $\|\mathbf{C} - \tilde{\mathbf{C}}\| < \epsilon$, then $\|\vec{v} - \tilde{\vec{v}}\| < \frac{\epsilon}{r}\|\mathbf{D}\vec{p}\|$, where $r = \min_i(\tilde{\hat{C}}_{ii}, \hat{C}_{ii})$ is the minimum reserve or "self-holdings" of the financial institutions. In addition, it is always true that $\|\vec{v} - \tilde{\vec{v}}\| \leq 2\|\mathbf{D}\vec{p}\|$, thus

$$\frac{\|\vec{v} - \tilde{\vec{v}}\|}{\|\mathbf{D}\vec{p}\|} \leq \min\left(\frac{\epsilon}{r}, 2\right)$$

Proof. Let $\tilde{\mathbf{C}} = \mathbf{C} + \mathbf{E}$, and $\tilde{\hat{\mathbf{C}}} = \hat{\mathbf{C}} + \hat{\mathbf{E}}$. By hypothesis $\|\mathbf{E}\| + \|\hat{\mathbf{E}}\| < \epsilon$. Then we have

$$\begin{aligned}
&\tilde{\hat{\mathbf{C}}}(\mathbf{I} - \tilde{\mathbf{C}})^{-1} - \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1} \\
&\quad = \left[(\hat{\mathbf{C}} + \hat{\mathbf{E}}) - \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}(\mathbf{I} - \tilde{\mathbf{C}})\right](\mathbf{I} - \tilde{\mathbf{C}})^{-1} \\
&\quad = \left[(\hat{\mathbf{C}} + \hat{\mathbf{E}}) - \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}(\mathbf{I} - \mathbf{C} - \mathbf{E})\right](\mathbf{I} - \tilde{\mathbf{C}})^{-1} \\
&\quad = \left[(\hat{\mathbf{C}} + \hat{\mathbf{E}}) - \hat{\mathbf{C}}(\mathbf{I} - (\mathbf{I} - \mathbf{C})^{-1}\mathbf{E})\right](\mathbf{I} - \tilde{\mathbf{C}})^{-1} \\
&\quad = \left[\hat{\mathbf{E}} + \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{E}\right](\mathbf{I} - \tilde{\mathbf{C}})^{-1}
\end{aligned}$$

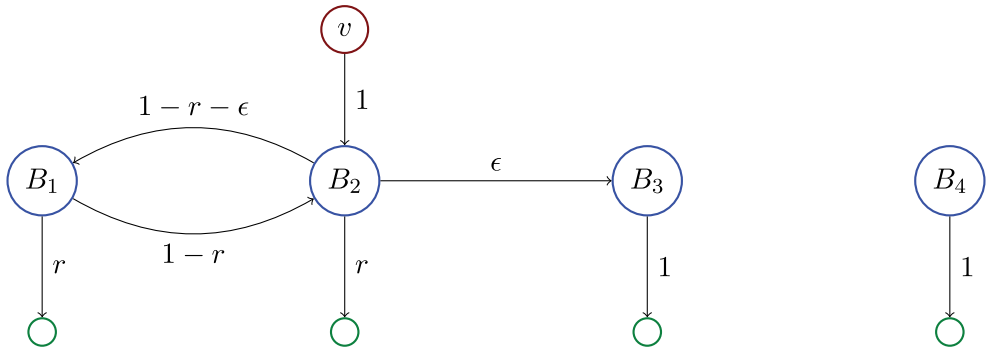

Fig. 2. The initial configuration, banks are in blue, external shareholders are in green, and the asset is shown in red.

Now, we notice that

$$\left\|(\mathbf{I}-\tilde{\mathbf{C}})^{-1}\right\| = \left\|\sum_{k=0}^{\infty}\tilde{\mathbf{C}}^k\right\| \leq \sum_{k=0}^{\infty}\left\|\tilde{\mathbf{C}}^k\right\| \leq \sum_{k=0}^{\infty}(1-r)^k = \frac{1}{r}$$

Because money is never created or destroyed, we have

$$\left\|\hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\vec{v}\right\|_1 = \|\vec{v}\|_1$$

Thus we have

$$\begin{aligned}&\left\|\hat{\tilde{\mathbf{C}}}(\mathbf{I}-\tilde{\mathbf{C}})^{-1} - \hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\right\| \\ &= \left\|\left[\hat{\mathbf{E}} + \hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\mathbf{E}\right](\mathbf{I}-\tilde{\mathbf{C}})^{-1}\right\| \\ &\leq \|\hat{\mathbf{E}}\| \cdot \left\|(\mathbf{I}-\tilde{\mathbf{C}})^{-1}\right\| \\ &\quad + \left\|\hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\right\| \cdot \|\mathbf{E}\| \cdot \left\|(\mathbf{I}-\tilde{\mathbf{C}})^{-1}\right\|\end{aligned}$$

Thus we immediately get the bound

$$\left\|\hat{\tilde{\mathbf{C}}}(\mathbf{I}-\tilde{\mathbf{C}})^{-1} - \hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\right\| \leq \frac{\epsilon}{r}$$

Since $\left\|\hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\right\| \leq 1$, we also have the trivial bound

$$\left\|\hat{\tilde{\mathbf{C}}}(\mathbf{I}-\tilde{\mathbf{C}})^{-1} - \hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\right\| \leq 2$$

The multiplicative bound of ϵ/r in Theorem 1 is much weaker than the bound of ϵ in the acyclic case (Corollary 1), as the minimum self-holdings approaches 0, this difference tends towards infinity. This discrepancy is not a limitation of our proof, but arises as an artifact of the effect that holdings cycles can have on the equity (and market) valuations of the institutions in the network. In Theorem 2 we show that there exist networks where changing a single institution's holdings by ϵ results in a change of $\frac{\epsilon}{2r+\frac{1-r}{2}\epsilon}\|\mathbf{D}\vec{p}\|$ in one of the institution's market values, where r is the minimum self-holdings in the network.

Theorem 2. There exist networks where $\|\mathbf{C}-\tilde{\mathbf{C}}\| \leq \epsilon$, and $\|\vec{v}-\tilde{\vec{v}}\| \geq \frac{\epsilon}{2r+\frac{1-r}{2}}\|\mathbf{D}\vec{p}\|$, where

$$\vec{v} = \hat{\mathbf{C}}(\mathbf{I}-\mathbf{C})^{-1}\mathbf{D}\vec{p}$$

and $r = \min_i(\tilde{\hat{C}}_{ii}, \hat{C}_{ii})$ is the minimum "self-holdings" of the financial institutions.

Proof. We exhibit an initial network in Fig. 2 and its perturbation in Fig. 3.

In Fig. 2, the equity values for the banks satisfy

$$\begin{aligned}B_1 &= (1-r-\epsilon)B_2 \\ B_2 &= v + (1-r)B_1 \\ B_3 &= \epsilon B_2 \\ B_4 &= 0\end{aligned}$$

So

$$\begin{aligned}B_2 &= v + (1-r)B_1 \\ &= v + (1-r)(1-r-\epsilon)B_2 \\ &= v + (1-r-\epsilon-r+r^2+r\epsilon)B_2\end{aligned}$$

Rearranging gives

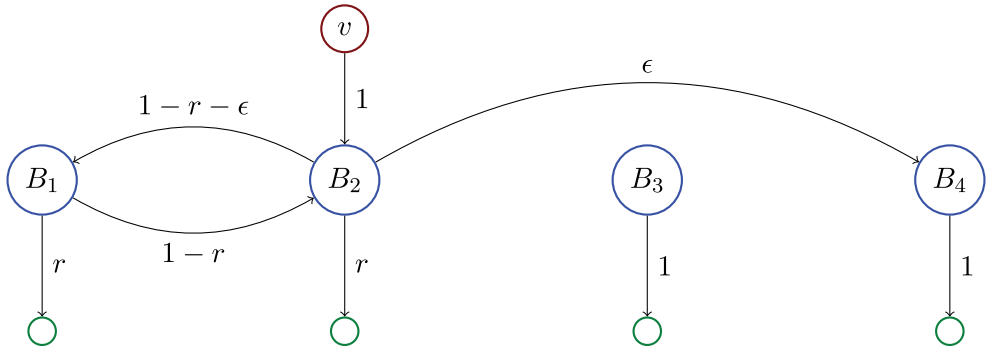

Fig. 3. The perturbed configuration, where one link of weight ϵ has been moved from B_3 to B_4.

$$
\begin{aligned}
B_2 &= v + (1 - r - \epsilon - r + r^2 + r\epsilon)B_2 \\
&\Downarrow \\
v &= (2r + \epsilon - r^2 - r\epsilon)B_2 \\
&\Downarrow \\
B_2 &= \frac{v}{2r + \epsilon - r^2 - r\epsilon} \\
&= \frac{v}{r(2-r) + (1-r)\epsilon} \\
&\geq \frac{1}{2}\left(\frac{v}{r + \frac{1-r}{2}\epsilon}\right)
\end{aligned}
$$

Thus the market valuation of B_3 is ϵB_2 which is at least $\frac{1}{2}\left(\frac{\epsilon v}{r+\frac{1-r}{2}\epsilon}\right)$.

If the link from $B_2 \to B_3$ were moved to $B_2 \to B_4$ (as in Fig. 3) then B_3's value drops to zero and B_4's value increases to ϵB_2.

Thus the change in ℓ_1 norm of the market valuations between the two situations is at least

$$\frac{\epsilon v}{r + \frac{1-r}{2}\epsilon}$$

Writing this in matrix notation, we have

$$\mathbf{C} = \begin{bmatrix} 0 & 1-r-\epsilon & 0 & 0 \\ 1-r & 0 & 0 & 0 \\ 0 & \epsilon & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\check{\mathbf{C}} = \begin{bmatrix} 0 & 1-r-\epsilon & 0 & 0 \\ 1-r & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & \epsilon & 0 & 0 \end{bmatrix}$$

$$\hat{\mathbf{C}} = \hat{\check{\mathbf{C}}} = \begin{bmatrix} r & 0 & 0 & 0 \\ 0 & r & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}$$

Note that increasing one link by ϵ and decreasing another by ϵ is actually a change in 2ϵ in the $\|\mathbf{C} - \hat{\mathbf{C}}\|$. Letting $\epsilon' = \frac{\epsilon}{2}$, we have a change of ϵ in $\|\mathbf{C} - \hat{\mathbf{C}}\|$ yields a change of at most $\frac{\epsilon}{2r+\frac{1-r}{2}\epsilon}$. Notice that as $r \to 0$, this is approaches 2, *i.e.*, the resulting valuation is *as far as possible* in terms of the ℓ_1 norm of the market values of the institutions.

The type of equity amplification necessary for the lower bound in Theorem 2 cannot happen in an acyclic network (see Corollary 1).

Now, we compare the upper and lower bounds on sensitivity. We have an upper bound of $\max(2, \frac{\epsilon}{r})$ and a lower bound of $\frac{\epsilon}{2r+\frac{1-r}{2}\epsilon}$. The lower bound implies the following.

- For any $\epsilon, \delta > 0$, there exists an r with $0 < r < 1$ and a network with minimum reserve r such that a change in edge weights of ϵ can lead to a change in market valuations of at least $(1 - \delta)2$. Thus the upper bound cannot be decreased below 2.
- For any r, δ, with $0 < r < 1$, and $0 < \delta$ there exists an ϵ with $0 < \epsilon < 1$ and a network with minimum reserve r such that a change in edge weights of ϵ can lead to a change in market valuations of at least $(1 - \delta)\frac{\epsilon}{2r}$, thus the upper bound cannot be decreased below $\frac{\epsilon}{2r}$.

Together, these show that the upper bound cannot be decreased below $\max(2, \frac{\epsilon}{2r})$, so our upper bound of $\max(2, \frac{\epsilon}{r})$ is essentially tight. Another interpretation of the lower bound is that for any $n > 0$, there exists

an $\epsilon, r > 0$ and a network such that an ϵ change in one edge weight can cause a multiplicative change of $n\epsilon$ in market malues of the institutions.

6. Bank failures

6.1. Losses caused by failure

The model of (Elliott, Golub, and Jackson, 2014) includes a notion of "failure," whereby institutions whose market value drops below a certain critical threshold suffer a further (discontinuous) loss in market value. These discontinuous penalties capture the notion that if an institution cannot pay its operating costs, it may see a further drop in revenues. Similarly, if confidence in the institution is shaken, and its debt rating is downgraded, it may see spike in the cost of capital, and hence see a further drop in value.

These discontinuous penalties are operationalized by a threshold value $\underline{v}_i$, such that if institution i's market value, v_i, drops below $\underline{v}_i$ then it incurs a failure cost and its market value drops by an additional $\beta_i(\vec{p})$.

Defining $I_{v_i<\underline{v}_i}$ to be the indicator variable which is 1 if $v_i < \underline{v}_i$ and 0 otherwise, and $b_i(\vec{v}, \vec{p}) = \beta_i(\vec{p})I_{v_i<\underline{v}_i}$, the market value of the institutions satisfies the equation

$$\vec{v} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}(\mathbf{D}\vec{p} - \vec{b}(\vec{v})) \tag{3}$$

Compare Equation 3 to the linear system given by Equation 2. The introduction of non-linear terms into the model adds significant complexity to the dynamics of the system, and can lead to failure cascades. One of the primary goals of (Elliott, Golub, and Jackson, 2014) was to characterize what network features affect the likelihood and severity of failure cascades.

6.2. Overview of our complexity results

Suppose a regulator has complete information about a financial network, including all cross-holdings and the prices of all underlying assets. Further, suppose that at equilibrium, this network has no failures. If the regulator believes that tomorrow, asset prices may drop by some fixed amount d, (*i.e.*, the sum total of asset prices may drop by d, but the exact drop of each asset price is unknown), what is the maximum number failures that could occur as a result of this drop? In other words, tomorrow, when the new equilibrium is calculated based on the new, lower, asset prices, what is the maximum number of banks that could have failed?

Our primary result is that the introduction of discontinuous failure costs increases the computational complexity of calculating basic network dynamics. If the complete network, as well as the price of all underlying assets is known, then the number of failures can be computed efficiently (Elliott, Golub, and Jackson, 2014, Section 3.2.3). If, on the other hand, there is some uncertainty in the prices (or future prices) of the underlying assets, calculating the maximum number of failures that could occur in the network is computationally intractable.

Thus we address the following question: given a stable network (where no banks have failed), if the total prices of the assets drop by some small amount, what is the maximum number of failures that occur at equilibrium? One potential complication is that in networks with discontinuous failure penalties, there may not be a unique equilibrium, and hence the number of failures may not be uniquely defined. To address this, for any fixed set of asset prices, we use standard practice and only consider the "best-case" equilibrium (the one with fewest failures). On the other hand, we consider the worst drop in asset prices (*i.e.*, the drop in asset prices that causes the most failures in its best-case equilibrium). Thus the "maximum" number of failures means the maximum over all bounded drops in asset values, of the minimum number failures that could occur at this each fixed drop in asset values (*i.e.*, the minimum over all equilibria at these new, lower values). We discuss the issue of multiple equilibria in more detail in Section 6.3.

In real-world financial networks, regulators, as well as the institutions themselves routinely perform "stress-tests" to assess the robustness of the network to financial shocks. Our results indicate that there are financial networks where no (computationally feasible) stress test, will ever be able to even approximate the maximum number of failures that could occur from some small shock to the system.

Note, however, that our results are not probabilistic, in the sense that we do not model the *probability* of a specific drop in asset prices, instead we just impose a bound on the magnitude of the total drop in asset values. Without assigning probability distributions to the asset prices, we cannot assess the probability that such a cascade could occur, only that it is feasible. Thus our question is whether there exists a small specific drop in asset prices that could cause a huge failure cascade – whether this drop is *probable* is outside the scope of our model. Thus, if a stress-test

could accurately model the probability distribution of asset prices, it may be able to assess the *likely* number of failures caused by a shock to the system, but any computationally feasible stress test will not be able to identify whether a small (but possibly unlikely) shock to the system could cause a far greater failure cascade.

6.3. Multiple equilibria

When discontinuous failure penalties ($\vec{b}(\vec{v})$) are introduced into the system, then there may be multiple equilibrium values for the institutions in the system, *i.e.*, the market values may not be uniquely defined (Elliott, Golub, and Jackson, 2014, Section 2.6 and Appendix A). As noted (Elliott, Golub, and Jackson, 2014), there are two distinctly different ways a network can have multiple equilibria. The first type falls into the standard theory of bank runs (Diamond and Dybvig, 1983), and this type of multiplicity can occur even in an acyclic network. For example suppose there is a network consisting of single institution holding a single asset of value p. Let v denote the value of the institution and $\underline{v}$ denote its failure threshold, and β its failure penalty. If $p > \underline{v} > p - \beta$, then a valuation of $v = p$, and $v = p - \beta$ are both consistent with Equation 3. The second type of multiplicity is caused by cycles in the network. See Fig. 4 for a simple (cyclic) network that has multiple equilibria.

In a network with multiple equilibria, we can talk about the "best-case" equilibrium (the one with the fewest failures) and the "worst-case" equilibrium (the one with the most failures). As in (Elliott, Golub, and Jackson, 2014), we focus on the best-case equilibrium, thus when we refer to the "the number of failures" we mean the number of failures in the best-case equilibrium.

6.4. The Balanced Complete Bipartite Subgraph (BCBS) problem

Our hardness result is based on the hardness of finding a maximum balanced clique in a bipartite graph. This is known as Balanced Complete Bipartite Subgraph (BCBS) problem.

Definition 1 (BCBS). Given a bipartite graph $G = (V_1, V_2, E)$ with $|V_1| = |V_2| = n$, the *Balanced Complete Bipartite Subgraph* (BCBS) problem is to find the largest integer K such that there exists sets $C_1 \subset V_1$ and $C_2 \subset V_2$ with the properties that $|C_1| = |C_2| = K$, and the induced graph on $C_1 \cup C_2$ is a complete bipartite subgraph of G.

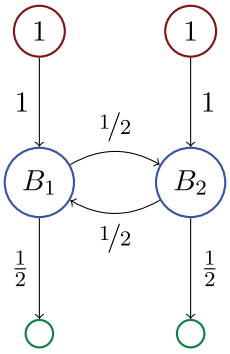

Fig. 4. Banks are in blue, external shareholders are in green, and the asset is shown in red. In this example $\mathbf{D}\vec{p} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$, and $\mathbf{C} = \begin{bmatrix} 0 & \frac{1}{2} \\ \frac{1}{2} & 0 \end{bmatrix}$. If $\underline{v}_i = 2$, and $b_i = 1$, then two equilibria are: $\vec{v} = \begin{bmatrix} 1 \\ 1 \end{bmatrix}$, and $\vec{v} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$.

The BCBS problem is known to be NP-hard (Gary and Johnson, 1979; D.S. Johnson, 1987). This provides strong evidence that there is no scalable algorithm that it that can find the size of the maximum balanced clique in a bipartite graph.

In fact, there is significant evidence that even *approximating* the size of the largest balanced clique is hard. Feige showed that for some $\delta > 0$ it is Random 3-SAT hard to approximate BCBS to within a factor of n^{δ} (Feige, 2002, Theorem 3).

Feige and Kogan showed that if BCBS can be approximated to within a factor of $2^{(\log n)^{\delta}}$ for every $\delta > 0$ then 3-SAT can be solved in time $2^{n^{3/4}+\epsilon}$ for every $\epsilon > 0$ (Feige and Kogan, 2004, Theorem 1.3).

Feige and Kogan go on to conjecture that for some $\delta > 0$, there is no polynomial-time algorithm to approximate BCBS to within a factor of n^{δ} (Feige and Kogan, 2004, Conjecture 1.1).

Our primary hardness result shows that there are networks where calculating the maximum number of bank failures that can result from a small shock is equivalent to solving the BCBS problem. Thus there are financial networks where even approximately estimating the number of bank failures that can result from an arbitrarily small drop in asset prices is a computationally intractable problem.

6.5. The complexity of calculating the maximum number of failures

In this section, we give our main result concerning the computational complexity of estimating the

maximum number of failures that can occur given a small drop in values of the underlying assets.

In particular, this hardness result applies finding the number of failures that occur at equilibrium when the institutional cross-holdings are fixed, and completely known, but there is some small uncertainty in the prices of the underlying assets.

In the situation that the cross-holdings and the assets are fixed, it is straightforward to calculate the market values of the institutions at equilibrium (Elliott, Golub, and Jackson, 2014, Section 3.2.3). Note that the network we construct is acyclic.

Theorem 3. For every bipartite graph G on $2n$ nodes, and every $\epsilon > 0$, there is a financial network with $\Omega(n)$ institutions, and a $d > 0$ such that computing the maximum number of institutions that could fail following a shock of $d\epsilon$ in asset prices is as hard as solving the BCBS problem in G.

Proof. Let $\ell > 0$ be any integer.

Our starting point is the following hardness result for the BCBS problem. Given an $n \times n$ balanced bipartite graph G, it is hard to decide whether the largest balanced bipartite clique size in G is at least $K \times K$ or at most $K/g \times K/g$ for some gap function g. For instance, $g = 2^{(\log n)^{\delta}}$ under the assumption that 3-SAT $\notin DTIME\left(2^{n^{3/4+\epsilon}}\right)$ for some $\epsilon > 0$.

Given an $n \times n$ balanced bipartite graph G, we will construct a financial network with $(2 + \ell)n$ institutions such that if G has a balanced bipartite subgraph of size K, then a drop in asset prices by $K\epsilon$ can cause at least $(2 + \ell)K$ failures. On the other hand, if the largest balanced bipartite subgraph of G is of size $\frac{K}{g}$, a drop in asset prices of $K\epsilon$ can cause at most $K + \frac{K}{g}(\ell + 1)$ failures.

This shows that estimating the maximum number of failures induced by a fixed drop in asset prices is at least as hard as estimating the size of the maximum balanced bipartite clique. Without loss of generality, we assume that every vertex in G has degree at least K.

Let D denote the maximum degree of any vertex in G. Let $0 < \epsilon < 1$ be an arbitrary parameter, and let $0 < r < 1$ denote the minimum amount of self-holdings of the institutions in the network we are constructing. (The reduction will hold for any choices of $0 < \epsilon, r < 1$.)

For each node in the graph G, we will associate a financial institution. Let institutions $1, \ldots, n$ correspond to the left-hand nodes of G, and institutions $n + 1, \ldots, 2n$ correspond to the right-hand nodes of G.

We will also generate n underlying assets, labelled $a'_1, \ldots, a'_n$ and institution i will complete own asset a'_i for $i = 1, \ldots, n$. Institutions $n + 1, \ldots, 2n$ will own none of the underlying assets. All assets will initially be valued at 1. Thus

$$\mathbf{D} = \begin{bmatrix} 1 & & & & & \\ & \ddots & & & & \\ & & 1 & & & \\ & & & 0 & & \\ & & & & \ddots & \\ & & & & & 0 \end{bmatrix}$$

Define $N = \frac{D}{1-r}$ (recall D is the maximum degree of G, and r is an arbitrary parameter that will determine the *integration* of the resulting financial network). We will use $\Gamma(j)$ to denote the neighbors of vertex j in G. Notice that our definition of N ensures that

$$1 - r = \frac{D}{N} \geq \frac{|\Gamma(j)|}{N}$$

for all $j = 1, \ldots, 2n$. This means that if institution j sells an equal $\frac{1}{N}$ stake in itself to all of its neighbors, it will be left with at least an r fraction of self ownership. To operationalize this, we define

$$c_{ij} = \begin{cases} \frac{1}{N} & \text{if } i > j \text{ and } (i, j) \text{ an edge of } G \\ 0 & \text{otherwise} \end{cases}$$

For $1 = 1, \ldots, n$ let $\underline{v}_i = 1 - \frac{|\Gamma(i)|}{N} - \epsilon$, so if institution i's asset drops in value by ϵ then institution i will fail. Let the failure penalty $\beta_i = 1 - \frac{|\Gamma(i)|}{N} - \epsilon$ for $i = 1, \ldots, n$. Thus if asset i drops in price by ϵ, institution i fails and its value immediately drops to 0. For $i = n + 1, \ldots, 2n$ let $\underline{rv}_i = \frac{|\Gamma(i)|-d}{N}$. Notice that if all assets are initially valued at 1 then

$$v_i = \begin{cases} 1 - \frac{|\Gamma(i)|}{N} & \text{if } 1 \leq i \leq n \\ \\ \frac{|\Gamma(i)|}{N} & \text{if } n < i \leq 2n \end{cases}$$

This financial network has the following properties:

1. If $j > n$ and d of bank j's neighbors fail, *i.e.*, d of the assets a_i ($i \in \Gamma(j)$) drop in value by ϵ, then bank j will fail.

2. If $j > n$ and the total drop in value of bank j's neighbors is less than $\frac{d}{N}$ then bank j will not fail.

Now, we examine the properties of this system when the assets are allowed to drop in price by total amount $d\epsilon$. A drop of $d\epsilon$ can always cause d institutions on the left-hand side of the network to fail, simply by dropping the value of each of their assets by ϵ.

What happens if the price drop is not concentrated among exactly d assets? Let t denote the number of assets that drop in value by at least ϵ. If $t < d$, then at least $t\epsilon$ from the "shock budget" of $d\epsilon$ was used to lower the price of these t assets, which leaves a budget of $(d - t)\epsilon$ remaining. Now, consider how much this drop can affect one of a right-hand institutions, j. Even if this drop is concentrated entirely among the left-hand neighbors of j, the drop j feels is at most

$$\frac{1}{N}(t + (d - t)\epsilon) < \frac{1}{N}(t + (d - t)) = \frac{d}{N}$$

Thus institution j cannot fail, since the failure of a right-hand institution requires a drop in value of at least $\frac{d}{N}$. This means that in this case exactly $t < d$ institutions fail. Since we are interested in the *maximum* number of failures that can arise from a drop in asset value of $d\epsilon$, we can, without loss of generality, assume that *exactly* d assets drop in value by ϵ, causing *exactly* d failures among the left-hand institutions.

Now, suppose there is a biclique of size K in G. If $d = K$, then causing d failures among the left-hand members of this biclique will cause $d = K$ failures among the right-hand members.

On the other hand, suppose the largest biclique is of size $\frac{K}{g}$. If the failure of d left-hand institutions causes the failure of P right-hand institutions, then each of the P failed institutions on the right must be connected to each of the d failed institutions on the left. Thus there must be a biclique of size $\min(d, P) = P$.

Thus in the "yes" case (G has a biclique of size K) we can cause at least K right hand failures with a failure budget of $K\epsilon$. In the "no" case (the largest biclique in G is of size $\frac{K}{g}$) the maximum number of failures of right-hand institutions is bounded by $\frac{K}{g}$.

Now, to amplify this discrepancy, we add a chain of ℓ institutions connected to each right hand institution. Thus for every right-hand institution, b_i, it will have a chain of institutions $b_i^{(1)}, \ldots, b_i^{(\ell)}$ where $b_i^{(j)}$ owns a $1 - r$ fraction of $b_i^{(j-1)}$ and has no other holdings. Thus if b_i fails, then $b_i^{(1)}$ through $b_i^{(\ell)}$ fail as well (see Fig. 5).

This new network has $(\ell + 2)n$ banks, and has the following properties. Given a failure budget of $K\epsilon$, if the largest balanced clique in G is a $K \times K$, a drop in value by $K\epsilon$ causes at most $(2 + \ell)K$, but if the largest biclique is of size $\frac{K}{g}$, then a drop in asset values of at most $K\epsilon$ can cause at most $K + (\ell + 1)\frac{K}{g} = (g + \ell + 1)\frac{K}{g}$ failures.

Note that when the gap, $g = n^{\delta}$, choosing $\ell = poly(n)$, we obtain a gap of $((\ell + 2)n)^{\delta'}$ for some $\delta' < \delta$.

Applying a result of Feige and Kogan (2004), we obtain the following Corollary.

Corollary 2. If 3-SAT $\notin DTIME\left(2^{n^{3/4+\epsilon}}\right)$ for some $\epsilon > 0$, then there exists a $\delta > 0$ such that there is no polynomial time algorithm that can calculate the maximum number of failures in a financial network caused by a drop in asset prices of $d\epsilon$ to within a factor of $2^{(\log n)^{\delta'}}$ for some $\delta' > 0$.

7. Conclusion

This work highlights two distinct sources of instability in financial networks, instability arising from fluctuations in cross-holdings and instability arising from fluctuations in asset prices. More specifically, we show that there are networks where small fluctuations in cross-holdings or asset prices can have striking consequences, and these consequences have numerous implications.

Our first result (Corollary 1, Theorems 1, 2) shows that the effect of small changes in cross-holdings is strongly tied to the integration of the network. In highly integrated networks small changes in cross-holdings can have potentially unbounded effects on market valuations, while in networks with low integration, changes in cross-holdings have more tightly bounded effects on the market values of the institutions in the network. These results can be interpreted in many ways. From a regulatory perspective, in a highly integrated network, any regulator must know the entire cross-holdings network to a very high degree of accuracy in order effectively understand the market values of the institutions. From an institution's perspective, small changes in investment by individual institutions can have their effects greatly magnified throughout the network. From a

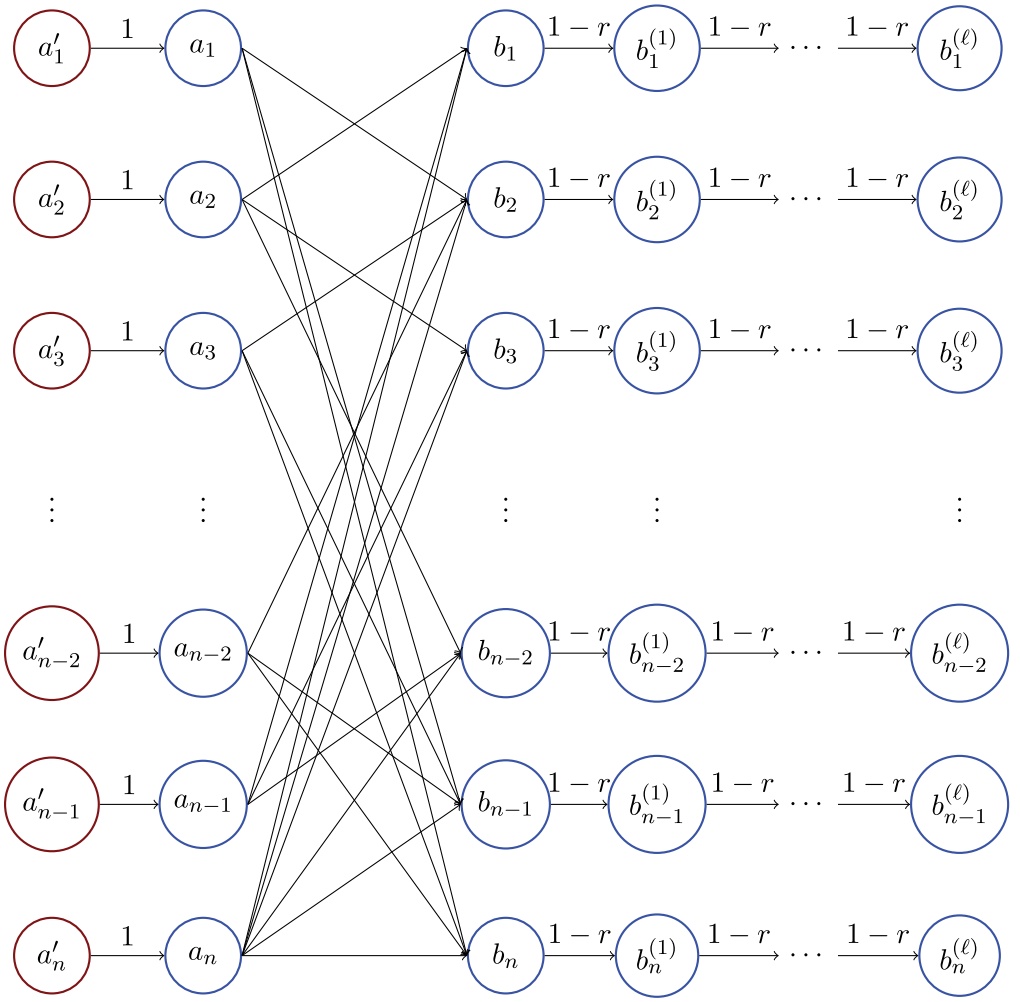

Fig. 5. There are n assets, $\{a'_i\}$, shown in red. Each asset a'_i is fully owned by institution a_i. Institutions $\{a_i\}$, $\{b_i\}$ correspond to a hard instance of balanced bipartite clique. Institutions $\left\{b_i^{(j)}\right\}$ serve to amplify the failures that occur in the first level.

predictivity perspective, if one wishes to forecast market values into the future, any small forecasting uncertainty in the cross-holdings can have enormous effects on the (predicted) market values of the institutions. From a privacy perspective, institutions cannot maintain any privacy in their investment portfolios without compromising the ability of outsiders (e.g. other institutions, outside investors or regulators) to calculate the market value of the institutions. These problems arise only in highly integrated networks (Theorem 1), and they can all be mitigated by imposing a cap on integration. If institutions are required to maintain some fixed percentage of their ownership outside of the network, the sensitivity to changes in cross-holdings can be drastically reduced.

Our second result shows that small changes in the prices of the underlying assets can have unpredictable effects on the number of failures in the system. Specifically, we show that there are networks where it is *computationally intractable*, even with perfect information about the cross-holdings, to estimate the number of failures that can occur after some small drop in asset prices. This result too can be interpreted from different perspectives. This result implies that a regulator (with perfect information about the network cross-holdings) who believes there may be some bounded fluctuation in asset prices cannot be expected to distinguish a network where these fluctuations cause a small number of failures, from a network where fluctuations of the same magnitude can cause a massive number of failures. Institutions

face the same computational challenge. An institution may wonder whether its investment portfolio will protect it from a bounded shock in asset prices, and our results show that even with complete information about the investments of all other agents in the system, it may be infeasible to determine whether a specific portfolio is safe.

Previous works have imposed specific probabilistic models on fluctuations in asset prices, and for a given probabilistic model the number of failures can usually be estimated. But what if the model is incorrect (e.g. the asset prices do not fluctuate independently)? Our results show that there are situations where changes in the distribution of fluctuations (but not their overall magnitude) can have huge and unpredictable effects on the number of failures.

Moving forward, it is an important research question to understand what constraints on the network will allow institutions and regulators to perform these stability analyses.

References

Acemoglu, D., Ozdaglar, A., Tahbaz-Salehi A., 2015. Systemic risk and stability in financial networks. The American Economic Review 105(2), 564–608.

Allen, F., Gale D., 1998. Optimal Financial Crises. In: *The Journal of Finance* 53(4), pp. 1245–1284. DOI: 10.1111/0022-1082.00052

Allen, F., Gale, D., 2000. Financial Contagion. In: *Journal of Political Economy* 108(1), pp. 1+. ISSN: 00223808. DOI: 10.1086/262109

Arora, S., et al. 2010. Computational Complexity and Information Asymmetry in Financial Products. In: *ICS*. Tsinghua University Press, pp. 49–65.

Arora, S., et al. 2011. Computational Complexity and Information Asymmetry in Financial Products. In: Commun. *ACM* 54(5), 101–107. ISSN: 0001-0782. DOI: 10.1145/1941487.1941511

Blume, L., et al. 2011. Which Networks are Least Susceptible to Cascading Failures? In: *FOCS*. IEEE, pp. 393–402. ISBN: 978-0-7695-4571-4. DOI: 10.1109/focs.2011.38

Braverman, M., Pasricha K., 2014. The Computational Hardness of Pricing Compound Options. In: *ITCS '14*. Princeton, New Jersey, USA: ACM, pp. 103–104. ISBN: 978-1-4503-2698-8. DOI: 10.1145/2554797.2554809

Brioschi, F., Buzzacchi L., Colombo M.G., 1989. Risk capital financing and the separation of ownership and control in business groups. In: *Journal of Banking & Finance* 13(4-5), 747–772. ISSN: 03784266. DOI: 10.1016/0378-4266(89)90040-x

Diamond, D., Dybvig, P., 1983. Bank runs, deposit insurance, and liquidity. In: *The Journal of Political Economy* 401–419.

Dwork, C., 2006. Differential Privacy. In: *ICALP '06*. Vol. 4052. Springer. Chap. 1, pp. 1–12. ISBN: 978-3-540-35907-4. DOI: 10.1007/11787006_1

Dwork, C., et al. 2006. Calibrating Noise to Sensitivity in Private Data Analysis Theory of Cryptography. In: *TCC*. Ed. by Shai Halevi and Tal Rabin. Vol. 3876. Springer Berlin/Heidelberg, pp. 265–284. ISBN: 978-3-540-32731-8. DOI: 10.1007/11681878_14

Eisenberg, L., T.H. Noe, 2001. Systemic Risk in Financial Systems. In: *Manage-ment Science* 47(2), pp. 236–249.

Elliott, M., Golub B., M.O. Jackson, 2014. Financial Networks and Contagion. In: *American Economic Review* 104(10), pp. 3115–3153. DOI: 10.1257/aer.104.10.3115

Fedenia, M., Hodder J.E., Triantis A.J., 1994. Cross-holdings: estimation issues, biases, and distortions. In: *Review of Financial Studies* 7(1), pp. 61–96. ISSN: 1465-7368. DOI: 10.1093/rfs/7.1.61

Feige, U., 2002. Relations Between Average Case Complexity and Approximation Complexity. In: *STOC '02*. Montreal, Quebec, Canada: ACM, pp. 534–543. ISBN: 1-58113-495-9. DOI: 10.1145/509907.509985

Feige, U., S. Kogan, 2004. *Hardness of approximation of the Balanced Complete Bi-partite Subgraph problem*. Tech Rep MCS04-04. Weizmann Institute of Science.

Flood, M., et al. 2013. *Cryptography and the Economics of Supervisory Information: Balancing Transparency and Condentiality*. Tech Rep 0011. Office of Financial Research.

French, K.R., J.M. Poterba, 1991. Were Japanese stock prices too high? In: *Journal of Financial Economics* 29(2), pp. 337–363. ISSN: 0304405X. DOI: 10. 1016/0304-405x(91)90006-6

Gai, P., S. Kapadia, 2010. Contagion in _nancial networks. In: *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* 466(2120), pp. 2401–2423. DOI: 10.1098/rspa.2009.0410

Gary, M.R., D.S. Johnson, 1979. *Computers and Intractability: A Guide to the Theory of NP-completeness*. WH Freeman and Company, New York.

Glasserman, P., H.P. Young, 2014. How likely is contagion in financial networks? In: *Journal of Banking & Finance*. ISSN: 03784266. DOI: 10.1016/j.jbankfin.2014.02.006

Gouriéroux, C., Héam J.C., Monfort A., 2012. Bilateral exposures and systemic solvency risk Expositions bilatérales et risque systémique pour la solvabilité. In: *Canadian Journal of Economics/Revue canadienne d'_economique* 45(4), pp. 1273–1309. DOI: 10.1111/j.1540-5982.2012.01750.x

Johnson, C.R., 1982. Inverse M-matrices. In: *Linear Algebra and its Applications* 47, pp. 195–216. ISSN: 00243795. DOI: 10.1016/0024-3795(82)90238-5

Johnson, D.S., 1987. The NP-completeness column: An ongoing guide. In: *Journal of algo-rithms* 8(3), pp. 438–448.

Morris, S., 2000. Contagion. In: *Review of Economic Studies* 67(1), pp. 57–78. DOI: 10.1111/1467-937X.00121

Narayan, A., A. Papadimitriou, A. Haeberlen, 2014. Compute Globally, Act Locally: Protecting Federated Systems from Systemic Threats. In: *Proceedings of the 10*th *Workshop on Hot Topics in System Dependability (HotDep'14)*. Broomfield, CO.

Poole, G., T. Boullion, 1974. A Survey on M-Matrices. In: *SIAM Review* 16(4), pp. 419–427.

Willoughby, R.A., 1977. The inverse M-matrix problem. In: *Linear Algebra and its Applications* 18(1), pp. 75–94. ISSN: 00243795. DOI: 10.1016/0024-3795(77)90081-7 ISSN

Appendix

A. A flow model of the system

It is often instructive to create a stochastic matrix representing the system, which models money flowing through the network.

There are n financial institutions, and we introduce n addition entities, the "shareholders." In this model shareholder i owns a $\hat{C}_{ii}$ fraction of institution i and has no other financial ties to the system. This defines a $2n \times 2n$ matrix, $\mathbf{A}$, defined as

$$\mathbf{A} = \left[\begin{array}{c|c} \mathbf{C} & \mathbf{0} \\ \hline \hat{\mathbf{C}} & \mathbf{I} \end{array}\right]$$

The augmented matrix $\mathbf{A}$ is now column stochastic, *i.e.*, its columns sum to 1. The ijth entry of the matrix $\mathbf{A}$ represents the fraction of "agent" j owned by agent i. There are $2n$ agents because there are n financial institutions (agents $1, \ldots, n$) and n collections of shareholders (agents $n+1, \ldots, 2n$). The lower right corner of $\mathbf{A}$ is the identity matrix because the shareholders are completely self-owned. The columns of $\mathbf{A}$ sum to one because for each agent, j, its entire value is owned by the other financial institutions or the external shareholders.

Initially, the banks are assumed to have some intrinsic value $\mathbf{D}\vec{p}$, which indicates the value of underlying assets owned by each institution. Without loss of generality, we will assume that all assets are completely owned by the institutions in the network (otherwise, we can simply rescale the value of each asset). This assumption is equivalent to saying that $\mathbf{D}$ is column stochastic.

The $2n \times 2n$ matrix $\mathbf{A}$ allows us to view the market valuations of each institution as the steady state of a dynamical process. Money flows into the system from the underlying assets, and at each time step, the value residing in each financial institution is distributed to its stakeholders according to their stake. This process terminates when all the money in the system (coming from the underlying assets) has been distributed to the external shareholders. Algebraically, this process can be viewed as follows: given a vector $\vec{W} \in \mathbb{R}^{2n}$, where W_i denotes the value of the underlying assets owned by i, for $i = 1, \ldots, n$ and $W_i = 0$ for $n < i \leq 2n$, then

$$\lim_{t\to\infty} \mathbf{A}^t \vec{W} = \lim_{t\to\infty} \mathbf{A}^t \left[\begin{array}{c} \vec{w} \\ \hline \vec{0} \end{array}\right] = \left[\begin{array}{c} \vec{0} \\ \hline \vec{v} \end{array}\right] = \left[\begin{array}{c} 0 \\ \vdots \\ 0 \\ \hline v_1 \\ \vdots \\ v_n \end{array}\right] \quad (4)$$

where v_i will denote the market value of institution i, *i.e.*, v_i is the value of institution i owned by its external shareholders. Now, basic linear algebra tells us that

$$\mathbf{A}^t = \left[\begin{array}{c|c} \mathbf{C}^t & \mathbf{0} \\ \hline \hat{\mathbf{C}}\left(\mathbf{C}^{t-1} + \mathbf{C}^{t-2} + \cdots + \mathbf{I}\right) & \mathbf{I} \end{array}\right]$$

Throughout this work, we will use the standard operator norm for a matrix (in the L_1 sense),

$$\|\mathbf{A}\| = \|\mathbf{A}\|_1 = \sup_{\vec{x}} \frac{\|\mathbf{A}\vec{x}\|_1}{\|\vec{x}\|_1}.$$

Because the columns of $\mathbf{A}$ sum to one, we have $\|\mathbf{A}\vec{x}\|_1 = \|\vec{x}\|_1$ for any $\vec{x}$. In particular, this means that

$$\left\| \mathbf{A} \left[\begin{array}{c} \mathbf{D}\vec{p} \\ \hline \vec{0} \end{array}\right] \right\|_1 = \|\vec{p}\|_1 \quad (5)$$

This is the algebraic statement that the total market value of all the institutions is exactly the total value of the underlying assets in the system.

To analyze Equation 4 we recall the standard fact about matrix series

Lemma 2.

If $\mathbf{C}$ is a matrix with $\|\mathbf{C}\| < 1$, then $\mathbf{I} - \mathbf{C}$ is invertible and

$$(\mathbf{I} - \mathbf{C})^{-1} = \sum_{k=0}^{\infty} \mathbf{C}^k$$

Proof. First, note if $\vec{v} \neq \vec{0}$, then $\|\mathbf{C}\vec{v}\| < \|\vec{v}\|$, so

$$\|(\mathbf{I} - \mathbf{C})\vec{v}\| \geq \|\mathbf{I}\vec{v}\| - \|\mathbf{C}\vec{v}\| > 0$$

so $\mathbf{I} - \mathbf{C}$ has a trivial kernel, and hence is invertible.

Then, the result follows just as in the scalar case: Letting $\mathbf{S} = \sum_{k=0}^{N} \mathbf{C}^k$, we have $\mathbf{S}(\mathbf{I} - \mathbf{C}) = \mathbf{I} - \mathbf{C}^{N+1}$. Thus

$$\sum_{k=0}^{N} \mathbf{C}^k = (\mathbf{I} - \mathbf{C})^{-1}(\mathbf{I} - \mathbf{C}^{N+1})$$

Letting $N \to \infty$, and noting that $\mathbf{C}^{N+1} \to 0$ gives the result.

This leads to the closed formula used in (Elliott, Golub, and Jackson, 2014):

$$\vec{v} = \hat{\mathbf{C}}(\mathbf{I} - \mathbf{C})^{-1}\mathbf{D}\vec{p} \tag{6}$$

Where the $n \times 1$ vector $\vec{p}$ represents the values of the underlying assets and the vector $\vec{v}$ represents the market value of the financial institutions (as measured by their external shareholders). Equations 4 and 5 then tell us that $\|\mathbf{D}\vec{p}\|_1 = \|\vec{p}\|_1 = \|\vec{v}\|_1$, which is just the simple statement that the total value of the agents is the same as the total value of the underlying assets, in other words value is never created or destroyed by the network.