Mechanisms of contagion in financial networks

Tiffany Yong

May 2024

1 Introduction

We have seen decades of growth in financial integration, both internationally and domestically. Lane and Milesi-Ferretti (2007) compile a dataset on the foreign assets and liabilities of 145 countries from 1970–2004. From 1998–2004 alone, the sum of countries' external assets and liabilities scaled by world GDP increased by nearly 50%. For industrial economies, this growth in financial integration far outstrips the growth in trade integration. Focusing on the US financial market, data from Furfine (2003) and the Federal Reserve H8 report show that in 1997, 18.2% of all US commercial bank liabilities are to other US depository institutions. Duarte and Jones (2017) improve this estimate by including all financial institutions and balance sheet exposures, and find that in 2016, 22.9% of assets and 47.8% of liabilities of bank holding companies come from within the US financial system.

The increasing interconnectedness of financial markets results in associated systemic risk, due to greater interdependence in the financial system. Financial institutions are directly exposed to one another through contractual claims in the interbank lending market, and indirectly exposed through correlated portfolios. These exposures tie the financial health of banks together, and when one bank experiences a shock, they serve as channels to transmit this shock beyond the initially affected bank which results in greater losses in the financial network. This transmission of shocks to other banks is known as financial contagion. In some cases, the costs associated with one bank defaulting causes its counterparties to default as well, causing a cascade of defaults in the financial system.

The empirical evidence supporting the presence of financial contagion is strong. As the 2008 financial crisis spread across global economy, post hoc analysis by Duarte and Jones (2017) find empirical evidence of default spillovers in the United States in 2008–2012, after the financial crisis. Fry et al. (2010) also find evidence of contagion in the real estate market during the United States subprime crisis, and Luchtenberg and Vu (2015) find evidence of contagion from the United States and other developed economies as the financial crisis grew in severity. This contagion effect is not unique to the 2008 financial crisis either. Kenourgios et al. (2011) find evidence of contagion during the 1997 Asian financial crisis, the 1998 Russian financial crisis, the 2000 US dot-com bubble collapse, and the 2002 Brazilian crisis.

Given the existence and potential severity of financial contagion, we are interested in further understanding contagion behaviour. In this paper, I review different theorised mechanisms for contagion, focusing on contagion through interbank lending and correlated portfolios. I discuss the effect of interbank network topology and portfolio correlation structure on likelihood and severity of contagion. I also discuss theoretical models of the formation of interbank networks and correlation structures, backed by empirical evidence from real-world financial markets. Finally, I develop a theoretical model to explain a stylised fact relating to the relative risks of different mechanisms of contagion for different sizes of shocks.

2 Contagion in financial networks

2.1 Early models of contagion in financial systems

The earliest model of contagion and systemic risk in the financial system was pioneered by Meltzer (1967), who attributed contagion in the financial system to information revelation mechanism. In this model, depositors are assumed to have imperfect information about both the solvency of banks and the government's willingness to bail out troubled banks. The failure of one bank without government intervention acts as a negative signal about both the solvency of other banks in the system and the government's willingness to support banks in crisis. If the liquidity is low at other banks and this negative signal is strong enough, this triggers a bank run in other banks.

Following this model, a rich literature studying financial contagion through information was developed. Many empirical studies like those conducted by Aharony and Swary (1983), Gay et al. (1991), and Peavy and Hempel (1988) were performed on major bank runs worldwide. These event studies used the initial date of public disclosure of problems at banks to analyse the effect of information-based contagion. If there is a significant drop in stock prices of banks immediately after the date of disclosure, this is evidence for information-based contagion. However, in a comprehensive review of this empirical literature, Kaufman (1994) notes that in most cases, contagion does not spread beyond the specific product area of the initial failed bank. He concludes that bank contagion through information is purely a rational correction of depositors learning that similar banks are also insolvent, and hence will not bring down solvent banks and the financial system.

At the same time, other empirical studies conducted by Hasan and Dwyer (1994) and Schoenmaker (1996) find evidence for contagion throughout the entire banking system, albeit by using a different methodology. Using probit and autoregressive models, they find strong intertemporal dependence in number of bank failures, even after controlling for macroeconomic influences. Thus, when the experiment design is not limited to capturing information-based contagion, we find evidence of system-wide contagion. This implies the existence of other mechanisms besides information revelation, which can spread contagion throughout the banking industry.

Subsequently, Rochet and Tirole (1996) initiated the theoretical study of the next alternative mechanism of contagion in the banking system — interbank lending. In their three-period model, banks are partitioned into two types: lending and borrowing banks. Lending banks have a relatively large deposit base compared to the size of their investment opportunities, while borrowing banks have a relatively small deposit base compared to the size of their investment opportunities. Hence, in the first period, borrowing banks make risky investments, while lending banks provide liquidity in exchange for claims on the investment payoffs. In the second period, borrowing banks face a liquidity shock. If the borrowing bank becomes insolvent, the bank fails and lending banks lose their claims. If the borrowing bank stays solvent but is illiquid, they dilute the claims of lending banks to raise liquidity required to withstand the liquidity shock. Otherwise, the lending bank receives the full payoff from their claims.

Using this simple three-period model, they find that there is local interdependency, where a bank is less likely to be liquidated if its debtors and creditors are not. They also find a global interdependency, where as the liquidity shock that one bank faces increases, all other banks are more likely to become insolvent, alluding to a contagion effect throughout the entire banking system.

Although this model deals with systemic risks stemming from interbank lending, they are limited to situations where banks are either borrowers or lenders, not both. This modelling assumption allowed the authors to avoid cycles in the interbank lending network, which simplified the clearing system of payments at the expense of realism in the model.

Building on this work, Eisenberg and Noe (2001) create the first model that explicitly deals with the cyclical interdependence that characterises the banking system and creates an algorithm to solve for the vector of payments in this cyclic system. This allows us to model for bidirectional shocks, and aids us in studying the full contagion properties of the financial network.

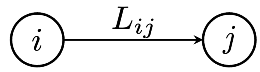

Figure 1: A link in the financial network, where bank i owes L_{ij} to bank j.

2.2 Interbank contagion in financial networks

The Eisenberg and Noe (2001) model begins with a set of rules that guide how firms in default should behave. Specifically, they assume proportionate repayment of creditors, limited liability, and absolute priority of debt. Proportionate repayment implies that when a bank defaults, its assets are split amongst creditors proportionally to their claims on the bank. Limited liability implies that each bank will not be able to pay out more than the total value of their assets. Absolute priority of debt implies that banks will pay off their debts before providing dividends. They also assume that the banking system has positive operating cash flow, implying that assets that are traded between banks are backed by real deposit value from the outside. From these assumptions, they find that regardless of the cyclic nature of the financial system, there exists a unique vector of payments that fulfils all of these assumptions.

Here, financial institutions are represented as n nodes, and financial connections as an $n \times n$ liabilities matrix L, where $L_{ij} \geq 0$ is the amount node i owes to node j as seen in Figure 1 above, and $L_{ii} = 0$ for all i. Banks receive an exogenous operating cash flow $e_i \geq 0$, which is the value of their assets from outside the financial system. Bank i's total assets are hence $\sum_{j=1}^{n} L_{ji} + e_i$, and its liabilities are $-\sum_{j=1}^{n} L_{ij}$, giving it a net worth of $\sum_{j=1}^{n} L_{ji} + e_i - \sum_{j=1}^{n} L_{ij}$.

They solve for the clearing vector using a fictitious default algorithm. They calculate the liability that should be repaid by all firms from interbank loans. If all banks in the system have a non-negative net worth, no bank defaults and the algorithm terminates. Otherwise, we take banks with a negative net worth into our set of "first-order" defaulting banks. We enter the next round assuming only this set of banks default, and calculate liability as such. In the next round, if the algorithm does not terminate, this implies that additional banks have defaulted. We add them to the set of defaulting banks to form a set of "second-order" defaulting banks. Continue in this manner until no additional bank defaults in round k, even when faced with a shortfall of payments from "$k-1$ order" defaulting banks. This allows us to calculate the vector of payments given (L, e), and the set of rules above.

The example in Figure 2 below illustrates the process. In the first round, the vector of payments is represented by the links, where B has A has to pay 10 to B, B has to pay 20 to C and 10 to D, and C has to pay 15 to A and 5 to D. However, here, B has a net worth of -15, and hence defaults. By proportionate repayment, B pays 10 to C and 5 to

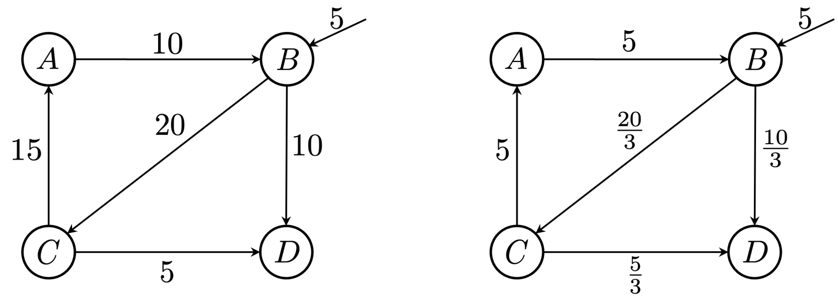

Figure 2: Eisenberg and Noe algorithm example, initial (L, e) on the left and final vector of payments on the right.

D instead. This results in C having a net worth of -10, so it defaults, and the algorithm does not terminate. By proportionate repayment, C pays 7.5 to A and 2.5 to B. But this causes A to default, since it will have a net worth of -2.5. This results in only 7.5 being paid to B, which implies that the net worth of B was actually -17.5, so we must continue running the algorithm until we obtain our unique fixed point vector of payments. The vector of payments in this example is A paying 5 to B, B paying $\frac{20}{3}$ to C and $\frac{10}{3}$ to D, and C paying 5 to A and $\frac{5}{3}$ to D. Here, A, B and C defaulted, while D is the only surviving bank.

This Eisenberg and Noe (2001) model is the working model for most of the literature to come. Other papers have since used links to represent other types of connections, like cross-holding of shares on values of financial institutions as in Elliott et al. (2014), or joint investments as in Erol and Vohra (2022). But fundamentally, the ability to model the sequential default on complex interbank connections arising from one bank-specific shock allowed for new literature pertaining to network topology, and network formation. They also allowed us to come up with another mechanism of contagion, by extending their model of the outside world. We continue by expanding more on these aspects.

2.3 Network topology affecting interbank contagion

The earliest theoretical model studying interbank lending network topology impacting contagion was by Allen and Gale (2000). Their three-period model is similar to that of Rochet and Tirole (1996). Here, banks are partitioned by region, and hold on to a short and a long asset. They can liquidate their short assets at anytime, but liquidating their long asset in the second period is so costly that provides little additional liquidity. Banks must meet demand for liquidity from consumers in their region over all periods to avoid defaulting. There are two types of consumers: early consumers, who will withdraw their deposits in the second period, and late consumers, who withdraw their assets in the third period. The proportion of early and late consumers in each region is random.

This incentivises banks to exchange deposits in the first period as a form of insurance against liquidity shocks. To see this, let the banks in region A and B exchange deposits in the first period. In the second period, suppose region A has a higher than average amount of early consumers, and region B has a lower than average amount of early consumers. Then, the banks in region A can meet their greater demand for liquidity by liquidating their deposits in region B, which also benefits region B as they had an excess supply of the short asset. In the third period, this cross-holding of assets also benefits region B, since they can liquidate their deposits in region A to meet their greater demand for liquidity from late customers. Hence, banks in both regions benefit from the exchange of deposits in the first period. Extending this argument, the more regions that a bank has deposits in, the lower the probability that all regions it is financially linked to will experience a liquidity shock simultaneously. Hence, the more financial links to different regions a bank has, the less susceptible it is to liquidity shocks. This implies that the complete network is the least susceptible to small liquidity shocks.

However, there are also negative effects of insurance through interbank deposits. Suppose that banks in region A have excess demand for liquidity in the second period that cannot be paid back even by liquidating all their short and long assets, and are hence bankrupt. Then, in the second period, banks in region A will be unable to pay back their liabilities, including interbank deposits from banks in other regions. This results in a spillover effect, since banks who hold deposits in region A will see a drop in liquidity supply. If this spillover effect is severe enough, this might cause neighbouring banks to go bankrupt as well. This can set off a chain of bankruptcies, and the shock can spread to the entire connected component that contains region A's banks. This implies that incomplete networks that do not have large connected components are the least susceptible to large liquidity shocks.

We see that when interbank deposits serve as insurance against liquidity shocks, greater interconnectedness as characterised by a larger number of neighbours allows for greater risk-sharing in financial networks. Although banks subject to the initial shock might fail, losses resulting from the shortfall in assets is borne by the entire economy instead of concentrated over a small number of neighbours. However, when the shock is sufficiently large, bigger connected components result in these connections propagating the shocks beyond the initially affected banks. This is known as the "robust-yet-fragile" property, coined by Haldane (2009).

Other authors have since restricted to different classes of networks to identify analytically and using simulations the network structures that are most and least susceptible to contagion. Other than approximately regular networks, a commonly studied class of networks has a "core-periphery structure". This is characterised by highly interconnected large "core" banks, and less-connected small "periphery" banks, which has been empirically observed in many real-world financial networks. This structure has been found in the German interbank market by Craig and von Peter (2014), the Austrian interbank market by Boss et al. (2004), the US interbank market by Soramäki et al. (2007) and Fricke and Lux (2015), the Italian interbank market by Iori et al. (2008), and the Brazilian interbank market by Silva et al. (2016).

Acemoglu et al. (2015) analysed the effect of liquidity shocks on regular network structures of banks with identical assets and liabilities. They define network performance by the social surplus in the economy, equal to the net worth of surviving banks plus liquidation value of defaulted banks. They find that if the shock is below a threshold, the complete network has the best expected and worst-case performance of any regular network, while the ring network has the worst expected and worst-case performance. On the other hand, if the shock is large enough, both the complete and ring network have the worst expected and worst-case performance. Instead, they define a δ-connected network as a regular network where there exists subset of banks S such that for any two banks $i \in S$ and $j \notin S$, the proportion of each bank's total liabilities to the other are at most $\delta \in [0, 1]$. This implies that banks in S are weakly connected to the rest of the network, and they find that for δ small enough, any δ-connected financial network has better expected and worst-case performance than the complete and ring networks. These results are consistent with Allen and Gale (2000), where for small shocks, the complete network performs the best, while for large shocks, networks with small connected components perform the best.

Gai et al. (2011) studied the effect of a randomly distributed small liquidity shock on 2 kinds of random network configurations: the Poisson configuration, where edges are approximately evenly distributed between all banks, and the geometric configuration, which has fat tails resulting in a natural core-periphery structure. When a small shock occurs, the affected bank withdraws interbank assets from its neighbours to stay liquid. Depending on the affected bank's withdrawal patterns, this can result in its neighbours facing a liquidity shock as well, and propagate throughout the system. Their Poisson simulations show that at low and high levels of connectivity, probability of contagion is low, while at intermediate levels of connectivity, probability of contagion is almost 1. At low levels of connectivity, the network is comprised of small connected components, so contagion cannot spread beyond this. At high levels of connectivity, each bank has a high degree, so the risk-sharing effects of financial connections make initial shock transmission unlikely. However, at intermediate levels of connectivity, the network is comprised of large connected components with a low average degree. Here, contagion is not limited by the size of the connected component or risk-sharing effects attenuating the initial shock, so probability of contagion rises sharply. Their geometric simulations show a similar non-monotonic trend of increasing then decreasing probability of contagion as connectivity increases. However, the geometric simulation has a lower probability of contagion at low connectivity than the Poisson simulation. The geometric simulation also requires a higher connectivity threshold before probability of contagion becomes a tail-event. Additionally, targeting the most connected bank for the initial shock has little effect on the Poisson network. However, it results in the probability of contagion in the geometric network being close to 1 for most intermediate levels of connectivity, since its core-periphery structure results in the most connected bank being able to propagate the shock extremely far. Thus, we can conclude that network structures with greater concentration but low connectivity and networks with lower concentration but high connectivity perform well. The first kind of network is akin to a cluster graph comprised of small cliques, while the second kind of graph is akin to a complete network, consistent with Allen and Gale (2000).

Elliott et al. (2014) studied the effects of integration and diversification of random networks on the contagion that results from one bank's asset failing at random. Here, each bank has a proprietary asset, and links are used to represent cross-holding of shares on values of financial institutions. Higher integration is defined as a greater proportion of shares held by other banks, while greater diversification is defined as a greater number of other banks that holds these shares. Both analytically and with simulations, they find that intermediate levels of integration and diversification result in the worst outcome, characterised by the existence of large connected components with a low average degree. They find that greater integration and diversification results in less frequent contagion, since a drop in the bank's own asset value is less likely to trigger a bank failure, characteristic of the complete network. They also find that little integration and diversification results in less contagion, since the cascade is limited to a small connected component. Studying a core-periphery random network, they find a similar pattern to Gai et al. (2011), where targeted shocks to the core result in significant contagion, but shocks to the periphery have little contagion when integration to the core is sufficiently high.

Sui et al. (2020) examined the effect of various network parameters on financial contagion in the core-periphery network. They find that increasing the size of core banks results in a larger phase transition, where the network is more resilient to liquidity shocks in general, but a targeted liquidity shock to a core bank is more likely to result in contagion. Increasing the number of core banks and periphery banks have the same effect. They also find that with a large number of core banks, periphery banks that increase interconnectedness reduce the fragility of the network, and the opposite is true for a small number of core banks.

We can see that there is plentiful theoretical evidence for the robust-yet-fragile property. Overall, we find that for small shocks, greater network interconnectedness in both average degree and size of components results in less contagion. For large shocks, reducing the average size of connected components results in less contagion. The core-periphery structure results in increased resilience of the network as long as the initial shock is not targeted at a core bank. However, Cabrales et al. (2017) finds that depending on the distribution of the size of shocks, different densities and number of components become optimal, so without knowing the distribution of shocks to banks ex-ante, we are unable to conclude which is the network structure that minimises contagion.

Beyond simulation studies, there are little empirical analyses on the effect of interbank lending network structures on contagion. We know that estimations of the network structure can be obtained through techniques like local entropy maximisation, pioneered by Upper and Worms (2004), which are sufficient to conclude a core-periphery structure in the interbank network. However, Craig and von Peter (2014) finds that this aggregated network structure remains relatively stable, even in periods of financial crisis. Exact data on interbank lending is often only available to the central bank and other regulators, so it is impossible to use small shifts in the network structure to infer larger contagion properties. There is also significant heterogeneity between interbank markets of different countries as shown by Allen et al. (2020), so cross-country comparisons are rendered ineffective.

2.4 Interbank lending network formation

As mentioned in Section 2.3, the network structure that most commonly appears in real-world financial markets is the "core periphery" structure. Following this empirical finding, there are multiple theoretical models which endogenously account for the formation of these core-periphery networks. They posit that there are benefits of the existence of a path between two banks, but the formation of links is costly. This results in the stability of a network where a core clique acts as intermediaries to a connected periphery, since this network has a large connected component without the associated costs of having too many edges.

in 't Veld et al. (2020) examined the relationship between trade access by intermediation and the endogenous formation of stable core-periphery structures. They characterise core-periphery networks has having a core bank clique which is fully connected, and a periphery bank independent set which has no pairwise connections. In their two-period model, costly undirected links between banks are formed in the first period, representing the existence of a trading relationship. In the second period, banks are able to trade through these previously-established relationships. Any two banks which are connected, either directly or indirectly through intermediaries, are able to trade. Surplus from trade between two banks is split between the trading banks and their intermediaries, based on the level of competition amongst intermediaries. If there is perfect competition amongst intermediaries, the two trading banks split the full trade surplus. If the intermediaries collude instead, they obtain some positive trade surplus equally divided amongst themselves, and the trading banks have a smaller trade surplus.

In the homogeneous case, they assume that all trades between banks result in the same trade surplus. With this assumption, they find that the core-periphery structure is not always unilaterally stable, meaning that there exists a coalition with incentives to deviate. This incentive to deviate arises in two situations. The first is when one core bank's neighbours are a subset of another core bank's neighbours. Then, the latter core bank does not gain paths to new periphery banks by connecting to the former core bank, and there is incentive to sever this link. The second situation is when there are many periphery banks. For a large enough number of periphery banks, one periphery bank will have an incentive to deviate and enter the core to profit from intermediation benefits. They also simulate a dynamic model which begins with an empty graph and allows banks to deviate to their best response consecutively. They find that for no values of level of competition between intermediaries and cost of maintaining links does the dynamic model converge to a core-periphery network.

In the heterogeneous case, they assume two types of banks, big and small, and model that the trade surplus is split between banks proportionately to their size. With this assumption, given that forming links is sufficiently costly and the trade surplus difference between small and big banks is sufficiently large, the core-periphery structure is unilaterally stable. They modify the dynamic model above, still assuming ex ante homogeneity, but allowing banks to reinvest their trade surplus into growing the size of the bank. They find that this updating of bank sizes endogenously creates of heterogeneity in bank sizes, and results in convergence to the core-periphery network for a large range of link costs and intermediary competition levels. This illustrates a possible model of network formation leading to core periphery networks even without ex ante assumptions of heterogeneity. This mechanism is further supported by the fact that empirical evidence from Section 2.3 shows that core banks tend to be large.

Babus and Hu (2017) study another mechanism, where revealed information of neighbours' actions influences network formation. In their model, banks are evenly divided into borrowing or lending banks at the beginning of each period. Each borrowing bank is paired with a random lending bank, and they can trade with one another. However, there is limited commitment, so banks can renege on liabilities at the end of the period. Furthermore, banks are situated in an information network, and costly links between banks represent bilateral monitoring. Without a link between the two banks, the lending bank is unable to observe the borrowing bank's history of fulfilling contracts. Given these limitations, banks will only utilise self-enforcing contracts, where a path of intermediaries between the lending and borrowing bank can provide credible threats which induce repayment. With a large enough network, trade can only be sustained with intermediaries. However, these intermediaries must be compensated sufficiently to ensure they have no incentive to keep repayments for themselves. In this context, the constrained efficient and stable network structure is the star network, which is an extreme example of a core-periphery network with only 1 core agent. To see this, note that intermediation increases trade and repayment, which increases social welfare. Since each intermediary must be compensated, minimising the number of intermediaries required increases the number of possible trades. However, due to the cost of maintaining links, we also want to minimise number of links. The star network is the acyclic connected graph which minimises the distance between each pair of nodes. When a star network is achieved, no agent has an incentive to deviate given that link costs are small and number of banks is large. Hence, intermediation can result in the formation of a stable core-periphery network without any form of heterogeneity.

Farboodi (2023) builds upon the previous model by accounting for asymmetric relationships in the interbank market, while assuming perfect information. In their three-period model, they account for asymmetric relationships by fixing the set of borrowing and lending banks. Borrowing banks have a probability of being assigned risky investment opportunities in the second period, while lending banks will never be assigned investment opportunities. In the first period, banks commit to interbank contracts, where creditors commit to providing funding to debtors given that the creditor has no investment opportunity, and the debtor has an investment opportunity either directly or through intermediaries. Notice that borrowing banks are most able to obtain loans in the first period, since they might have access to risky investment opportunities in the second period, and can provide a higher expected return to lenders. Furthermore, since having a long chain of intermediaries is costly, lending banks prefer to be connected directly to a borrowing bank which is connected to every other borrowing bank. Thus, when payoffs from intermediation are sufficiently high, borrowing banks are incentivised to intermediate and bear the costs of defaults when risky investments do not pay off. This creates the core-periphery structure, where borrowing banks form the intermediary core and lending banks form the periphery. However, this structure is inefficient, since borrowing banks have an incentive to overconnect to other borrowing banks, resulting in very high default risks which could be otherwise avoided if a lending bank was the intermediary.

Overall, we see that heterogeneity is a key part of the formation of core-periphery structures, whether in bank size, access to information, or access to investments. This results in different payoff structures for banks in the core and banks in the periphery, which allows the core-periphery structure to be the stable equilibrium. Even when banks are assumed homogeneous at the outset, this heterogeneity can arise through the best response dynamics of banks in the system.

3 Outside assets and correlated portfolios

The previous sections deal with shocks to banks in a black box, modelled as exogenous changes to consumer liquidity demand, changes in value of external assets, or realisations of returns to risky investments. These shocks represent changes in banks' operating cash flow as modelled by Eisenberg and Noe (2001), which induces changes in banks' net worth. Glasserman and Young (2016) make this explicit by introducing the concept of an "outside" sector with a representative node. The links between this outside sector node and banks in the network can be seen as banks borrowing to and lending from non-financial institutions, which is the basis of their real deposit value from which they can engage in interbank lending.

We are interested in studying this outside sector for multiple reasons. Firstly, this outside sector is where shocks to banks originate from. Secondly, a large proportion of bank assets are held in this outside sector. Allen et al. (2020) studied the average ratio of interbank loans to total bank assets for banks in the Euro Area and in the US. They find that for all countries, <30% of banks assets are from interbank loans, while the other >70% are external assets, with the US having only 2.4% of its assets coming from interbank loans. Finally, the structure of the outside sector can also result in contagion through other mechanisms beyond interbank lending.

3.1 Contagion through the outside sector

Cifuentes et al. (2005) are the first to introduce a complementary mechanism of contagion, which occurs through assets held in the outside sector. Their model is similar to that of Eisenberg and Noe (2001), where banks engage in interbank lending and have an operating cash flow obtained from outside assets. They extend the model by distinguishing between liquid and illiquid outside assets, where liquid assets have a constant price, while the price of illiquid assets depends on downwards-sloping inverse demand function d^{-1}. Additionally, banks are beholden to a minimum capital ratio, so the ratio of their net worth to the market value of their assets must be above a certain threshold r^*. Below this threshold, they are required to sell their assets to prevent being overleveraged.

When a bank experiences a liquidity shock, the equilibrium solution involves a vector of payments x as discussed before, but also a vector of sales of the illiquid asset s, and price p of the illiquid asset. This solution is such that x is the clearing vector of payments solved in the fictitious default algorithm, s is the minimum amount of illiquid asset sold to satisfy the minimum capital ratio, and p is the general equilibrium price obtained by $p = d^{-1}(\sum_i s_i)$.

In this context, when a bank defaults, it has to sell off all its assets, including their illiquid ones. This increases the demand of illiquid assets, pushing down its price. As the price of these assets are driven down, other banks that hold these assets face a reduction in their net worth and are in violation their minimum capital ratio. If these other banks do not have sufficient liquidity buffers, they will be made to sell their illiquid assets too. But this further pushes down the price of illiquid assets, which induces more banks to sell their illiquid assets. This downward spiral of prices and selling of illiquid assets en-masse is known as a "fire sale". This large devaluation of illiquid assets caused by the initial default of the affected bank can result in other banks defaulting, which is how contagion spreads through correlated portfolios in the outside sector.

3.2 Structure of correlated portfolios

We have studied a new theoretical model which shows that correlated portfolios of outside assets can result in contagion. Knowing this, we are interested in studying what correlation structures arise between banks' portfolios.

Acharya and Yorulmazer (2007) develop a theoretical model where banks have an incentive to correlate portfolios due to bailout policies. Banking regulators have historically intervened in times of financial crisis using the "too-big-to-fail" principle, which argues that some banks are too large and interconnected to be allowed to fail. They extend this to correlated bank failures to create the "too-many-to-fail" problem. This is where regulators find it optimal to bail out banks when the number of bank failures is large, but when the number of bank failures is small, it is optimal for surviving banks to acquire failed banks. In their two-period model, there are two banks and two industries. In the first period, both banks each choose one industry to invest in. In the second period, if a bank's return to investments is high, it can invest in the next period — otherwise it defaults. The "too-many-to-fail" principle applies here, so whenever both banks default, they will be bailed out if the bailout subsidy is less than the cost of liquidating all their assets to non-financial institutions. But if only one bank defaults, the failed bank's assets are sold to the surviving bank, minus an equity share taken by the regulator. In this context, whenever the cost of liquidating banks' assets is high enough, the regulator will bail out both banks when they fail together. Here, if the subsidy received when being bailed out is higher than the gain in value associated with buying out a failed bank, there is incentive for banks to herd their investments. This finding is robust when extended to n banks. Hence, we see that the "too-many-to-fail" principle results in greater correlation in investments. They also extend their model to study how the "too-big-to-fail" principle impacts portfolio correlation structures. They introduce heterogeneity in bank sizes, with one big bank and one small bank. Based on the "too-big-to-fail" principle, if only the big bank defaults, regulators will bail it out given that the cost of liquidating its assets is higher than the cost of the bailout subsidy. On the other hand, if only the small bank defaults, the regulators will sell its assets to the big bank. In this modified model, the small bank has no benefits from being the only survivor, since it is unable to acquire the big bank. Hence, it has an incentive to herd its investments with the large bank. However, the big bank is guaranteed to be bailed out if it fails, so it has no herding incentive. Instead, it prefers to differentiate its portfolio so they can survive to take advantage of the discount of purchasing the smaller bank. Thus, based on the "too-big-to-fail" principle, we see that big banks prefer to differentiate themselves from small banks, while small banks prefer to herd with large banks.

Bräuning and Fillat (2019) studies the portfolios of large banks, and the effects of financial regulation on their degree of correlation. They are specifically interested in how the diversification requirements uniquely applied to large banks in the form of stress tests affects their portfolio correlation. To do this, they analyse how large US banks' portfolios changed after the implementation of stress-testing in the 2010 Dodd-Frank Act. A previous theoretical model developed by Wagner (2010) finds that when banks focus on increasing diversification, their portfolios will be more correlated. Bräuning and Fillat (2019) obtain similar empirical findings, with the degree of similarity between banks' portfolios increasing and the distribution of portfolios narrowing as stress-testing requirements were put in place. They posit that this is due to banks adjusting their portfolios when they perform poorly on the stress test, causing them to diversify and hold a similar portfolio to banks that performed better on the stress test.

Elliott et al. (2021) builds on this work by studying the relationship between the interbank network structure and correlated portfolios. They find empirical evidence from the German banking system: banks which are linked in the interbank network also tend to have correlated portfolios. They then introduce a theoretical model that rationalises this behaviour. In their three-period model, there are n banks and n types of risky investments. In the first period, banks borrow from non-financial institutions, and is obliged to repay them in the last period. In the second period, banks each choose a portfolio of investments and form interbank claims on each others' portfolios. In the final period, if banks are unable to repay their debts from the first period, they default, inducing a default cost. There is limited liability, so when a bank defaults, it only has to pay its market value to its debtors. They model shocks in the system as a drop in the return of investment types with some probability r. Shocks are small with probability p and large with probability $1 - p$, with large shocks being relatively rare. All shocks are larger than the value of a single bank, but with risk sharing through diversification of portfolios and counterparty exposures, banks can withstand small shocks. On the other hand, large shocks are such that at least one bank must default for any given network and portfolio structure. In this model of shocks, the socially efficient network structure is such that banks have no overlap in their portfolios, and the interbank claims network is comprised of d clusters, with greater interconnectedness within clusters and lesser interconnectedness between clusters. This way, if a small shock occurs, risk-sharing prevents banks from defaulting, and if a large shock occurs, it is isolated to one cluster. However, banks have an incentive to deviate from this socially efficient network and correlate their portfolios with their counterparties in the interbank network due to limited liability. When a bank's counterparties default on their debt obligations, the bank has to bear the counterparties' default cost if they survive. By correlating portfolios to fail together, they do not have to bear this default cost because their equity value has already gone to zero. Although this increases the bank's risk of default, it still results in a higher expected utility, since they do not have to internalise negative externalities from the failure of other banks. This phenomena is known as "risk-shifting", and increases the contagion risk in the banking system.

We see both theoretical and empirical support for the idea that banks correlate their portfolios with their counterparties, and that large banks tend to correlate their portfolios. We also have a theoretical prediction that in the presence of systemic risk, small banks prefer to herd with banks of any size to improve their chances of getting bailed out, resulting in more correlated portfolios. To validate this theory, we are interested in empirically observing the behaviour of banks when contagion risk increases. Kabir (2017) analyses shifts in portfolio herding behaviour as a result of the 2008 financial crisis, and finds an increase in intentional herding of commercial banks in periods of higher volatility. Luengnaruemitchai and Wilcox (2004) finds similar herding behaviour in the volatile US banking market during the 1980s. Barron and Valev (2000) specifically differentiates between small and large banks during this same volatile banking period, and finds that the international lending patterns for small banks follows that of large banks. Thus, we also have empirical evidence supporting the idea that small banks prefer to herd when faced with high systemic risk.

3.3 Structure of interbank network and correlated portfolios affecting contagion

We now examine both how the interbank network topology affects this new mechanism of contagion through correlated portfolios, and how the structure of correlated portfolios as discussed in Section 3.2 affects probability and extent of contagion.

In their introduction of contagion through outside assets, Cifuentes et al. (2005) also ran simulations of the effects of a bank default on a network of identical banks, to study the parameters which affect contagion through correlated portfolios. They find that beyond a certain threshold of liquidity buffers held in each bank, no contagion through the outside sector is observed. They also observe that contagion through correlated portfolios is minimised in either the empty network or the complete network, and maximised when each bank has an intermediate number of connections. This is in line with what we have concluded in Section 2.3.

Gai and Kapadia (2010) studied the effect of one affected bank failing at random contagion in a Poisson network, in the presence of asset price contagion. They assume that initial assets are comprised of 80% outside assets and 20% interbank assets, and there is one illiquid outside asset that all banks hold. When a bank defaults, all its outside assets are sold in the market, following the asset price equation $p = e^{-\alpha x}$. They find that the incorporation of this additional mechanism of contagion increases the upper threshold of average degree beyond which probability of contagion is close to zero. It also significantly increases the extent of contagion when it occurs, especially for networks with a low average degree. Even with the incorporation of liquidity risk, the robust-yet-fragile property continues to be observed, with complete and empty networks having the lowest probability of contagion, but complete networks spreading contagion the most.

Shen and Li (2020) examined the effect of a liquidity shock to a random affected bank on contagion behaviour in regular and core-periphery networks, in the presence of correlated portfolios. In their simulations, instead of holding one outside asset, banks hold a portfolio of securities, which they sell when they face bankruptcy to repay its creditors. They vary both degree of interconnectedness and degree of portfolio overlap to study how contagion behaviour evolves with these parameters. In their analysis of regular networks, they find that in the presence of correlated portfolios, the severity of contagion in sparsely connected networks sharply increases, corroborating results from Gai and Kapadia (2010). For all regular financial networks, they find that as degree of portfolio overlap increases, the probability of contagion first increases, then decreases. Initially, as degree of overlap increases, a new channel of contagion through correlated portfolios becomes possible, increasing probability of contagion. However, when degree of overlap is high enough, no bank is too affected by the devaluation of any particular security.

In their analysis of core-periphery networks, they study a shock to the core bank and to a random periphery bank separately. When the core bank is shocked, the probability and extent of contagion increases monotonically as the size of the initial shock increases. When the initial shock is small, as degree of portfolio overlap increases, we see that the probability of contagion first increases, then decreases. When degree of portfolio is overlap is low, although periphery banks have very little capital buffer to withstand a devaluation in securities, periphery banks are unlikely to share common securities with other banks, which limits the probability of contagion. As portfolio overlap increases, this channel of contagion is more significant, and periphery banks are more likely to be affected by contagion and spread the contagion to other periphery banks. When portfolio overlap is high, losses from devaluation of securities is evenly spread across all periphery banks, and hence contagion is again limited. When a random periphery bank is shocked, as degree of portfolio overlap increases, the probability of contagion first increases, then decreases. The probability of contagion never reaches zero, unlike in the complete network case, because of the susceptibility of periphery banks to devaluation of securities.

Overall, the addition of contagion through correlated portfolio increases systemic risk. The relative susceptibility of different interbank network structures to contagion remains consistent to the results we have obtained in Section 2.3. We also find that as degree of portfolio overlap increases, risk of contagion first increases, then decreases. When portfolio overlap is low, risk of contagion is also low, due to the lack of correlation to transmit shocks. When portfolio overlap is high, risk of contagion is also low due to the risk-sharing effects of greater financial integration. However, due to the rarity of financial crises and idiosyncracies between banks, there is limited empirical data on the effect of correlated portfolios on contagion.

4 Relative impact of mechanisms of contagion

We see that the addition of contagion results in default cascades occuring more frequently and with higher severity. Consequently, we are also interested in studying the relative contributions of the two contagion mechanisms to overall systemic risk.

Mikropoulou and Vouldis (2023) explore this by conducting simulations of shocks on the interbank market and studying how the magnitude of the shock impacts the relative risk from different contagion mechanisms. They use the terminology of "direct" contagion for contagion through interbank lending, and "indirect" contagion for contagion through correlated portfolios. They model banks as having interbank exposures to one another, exposures to outside sectors, and liquid instruments as a buffer. In their simulation, they construct their interbank market network and portfolio position using actual bilateral exposures from a supervisory dataset on large euro area banks. They observe that the reconstructed network has heterogeneity in bank size, and exhibits a core-periphery structure. Two different kinds of shocks are studied: a liquidity shock to banks, and a drop in the value of assets in an outside sector. They find that in both cases, as the size of the shock increases, indirect contagion increases to a larger extent than direct contagion. There are observed inflection points where indirect contagion sharply rises beyond a certain threshold for the size of shock. Overall, we find that when shocks are large, the risk from correlated portfolios is greater than from interbank links. We attempt to develop a theoretical model to explain this phenomenon.

4.1 Theoretical model for relative risk of contagion

We have a set of banks, $N = \{1, \ldots, n\}$. There are two types of banks: k big banks and $n-k$ small banks. Interbank lending exposures exhibit a core-periphery structure as described in Section 2.3. The set of large banks C forms the core clique, the set of small banks P is an independent set that forms the periphery, and each periphery bank is connected to at least one core bank.

Banks also hold a portfolio comprised of assets from a set of different sectors of the economy $S = \{1, \ldots, s\}$. Here, sectors represent geographical regions or business sectors. Given the small size of periphery banks, we model their portfolios to be undiversified, and they each hold assets from only one sector. Notice that periphery banks in the same sector have perfectly correlated portfolios, and periphery banks in different sectors have no portfolio overlap. In Section 3.2, we observe that banks prefer to herd with their counterparties, and small periphery banks prefer to herd with large core banks. We assume that banks have no gains from purchasing the assets of defaulted banks, so core banks prefer to herd with their periphery bank counterparties too. We know that large banks herd indirectly with one another through diversifying their portfolio. Hence, we assume that core banks do not deliberately correlate their portfolios, and their portfolio overlap arises from having periphery bank counterparties in the same sector. Overall, the portfolio of periphery banks consists of assets from one sector, and the portfolio of core banks consists of assets from each sector their periphery bank counterparties are in.

Like Gai and Kapadia (2010), we assume that 80% of a bank's initial assets are held in the outside sectors, and the other 20% are interbank assets. Beyond a bank's size and exposure to different sectors, we assume that banks are ex-ante identical. This allows us to

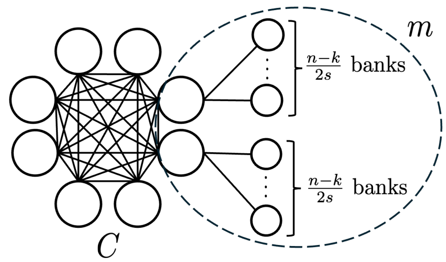

Figure 3: Our model of the network structure.

make various simplifying assumptions. We let the size of interbank deposits between banks of the same type be the same. We can define λ such that

$$\lambda = \frac{\text{size of interbank deposits from periphery to core bank}}{\text{size of interbank deposits between core banks}}.$$

Suppose that the $n-k$ periphery banks be evenly split into s sectors, so each sector has $\frac{n-k}{s}$ periphery banks. We also assume that sector has 2 core banks exposed to it, and the $\frac{n-k}{s}$ periphery banks in the sector are evenly divided to be counterparties of these 2 core banks. A graphical representation of our model is in Figure 3 above — for clarity, we only show the periphery banks in sector m. With these assumptions, we have that $s \geq \frac{k}{2}$, and each core bank is linked to $\frac{n-k}{k}$ periphery banks.

As in Elliott et al. (2021), the units of assets available in each sector is normalised to 1. For each sector, let the proportion of assets that periphery banks hold be ℓ, and the remaining $1-\ell$ is held by core banks. We assume that these assets are evenly distributed between banks of the same type. Hence, each periphery bank in sector m will hold $\ell \cdot \frac{s}{n-k}$ proportion of the assets, and each core bank will hold $\frac{1-\ell}{2}$ proportion of the assets.

We suppose that periphery banks have sufficient liquidity to withstand a drop in value of proportion β, and core banks have sufficient liquidity to withstand a drop in value of proportion τ, beyond which they fail. As in Gai and Kapadia (2010), we assume zero recovery from failed banks. We also use their inverse demand function, so the price of the asset from sector m, denoted as p_m, is given by

$$p_m = e^{-\alpha_m x_m}$$

where $x_m > 0$ is the fraction of assets from sector m that have been sold on the market, and $\alpha_m > 0$ is calibrated such that the asset price p_m falls by proportion β when β of the assets in sector m have been sold. Thus, firms can raise liquidity by selling x_m assets, to obtain

$$L = x_m e^{-\alpha_m x_m}.$$

The shock to the system is modelled as in Mikropoulou and Vouldis (2023), where there is an increase in liquidity demand to each bank in sector m of size $-\frac{\ln(1-\delta)}{\alpha_m} \cdot \frac{s}{n-k}$, with $\delta \geq \frac{5}{4}\beta$. Since periphery banks are undiversified, they can only generate liquidity by selling their assets from sector m. Our magnitude of the initial shock is chosen such that when all $\frac{n-k}{s}$ periphery banks simultaneously sell sufficient assets at original prices to cover the increase in liquidity demand, this results in a drop in price of asset by δ. Since 80% of assets are held in the outside sector, a drop in asset price of at least $\frac{5}{4}\beta$ results in a drop in bank value of at least β. All periphery banks exposed to this sector fail, so they default on their debts to core banks and sell all their assets.

Contagion of this initial shock to one of the core banks i in sector m through interbank lending is of proportion

$$C_{interbank} = \frac{0.2}{t_i} \cdot \frac{(n-k)\lambda}{(n-k)\lambda + k^2},$$

where t_i is the number of sectors the core bank i is exposed to. The second term represents the proportion of i's interbank assets held by periphery banks.

Contagion through correlated portfolios is caused by all the periphery banks selling their assets, which results in a further drop in p_m. Thus, the contagion effect to core bank i is

$$C_{correlated} = \frac{0.8}{t_i} e^{-\alpha_m \ell}.$$

On the other hand, we assume that as in Bräuning and Fillat (2019), core banks have to meet diversification requirements, and prefer to generate liquidity by selling their assets from other sectors. To meet the liquidity demand above, core bank i sells an equal amount of assets in all other sectors, generating a drop in prices ε, where

$$1 - \varepsilon = e^{-\alpha_m\left(-\frac{\ln(1-\delta)}{\alpha_m} \cdot \frac{s}{n-k} \cdot \frac{1}{t_i}\right)} = (1-\delta)^{\frac{s}{n-k} \cdot \frac{1}{t_i}}.$$

Here, we notice that as the size of the initial shock increases, the fall of asset prices in other sectors ε increases. When the initial shock is large enough such that $\varepsilon = \beta$, this causes periphery banks in other sectors exposed to core bank i to fail, which by our calibration of α_m is when

$$\beta = -\frac{\ln(1-\delta)}{\alpha_m} \cdot \frac{s}{n-k} \cdot \frac{1}{t_i},$$

or equivalently, when the initial liquidity shock to each bank in sector i is multiplied by a factor of t_i. This is actually an upper bound — if both core banks in sector m are also the core banks in another sector v, the combined sales of their assets results in a fire sale much sooner.

In order for the core bank to stay liquid as periphery banks in the first sector fails, the core bank's liquidity buffer should be such that

$$\tau \geq \frac{0.2}{t_i} \cdot \frac{(n-k)\lambda}{(n-k)\lambda + k^2} + \frac{0.8}{t_i} e^{-\alpha_m \ell}.$$

The core bank is less susceptible to contagion when the number of sectors it spreads its assets across t_i is high, or the proportion of its claims in periphery banks are low, or the proportion of assets held by periphery banks ℓ is low. In order for the core bank to stay liquid as periphery banks in all sectors fail, the requirement is now

$$\tau \geq 0.2 \cdot \frac{(n-k)\lambda}{(n-k)\lambda + k^2} + 0.8 e^{-\alpha_m \ell}.$$

Our theoretical model provides justification for inflection points where contagion through correlated assets sharply rises, as observed in Mikropoulou and Vouldis (2023). When faced with a liquidity shock in a sector, core banks attempt to dissipate the shock by selling assets in other sectors to raise liquidity. They are subject to diversification requirements, so they sell equal amounts of assets from each sector they are exposed to. If the shock is small, this results in no further bank failure. However, when the shock is large, this results in a fire sale in many sectors all at once, due to many periphery banks defaulting through indirect contagion.

5 Conclusion

We have provided an overview of contagion in financial networks, through mechanisms of information revelation, interbank lending and correlated portfolios. We examined both theoretical models of these mechanisms, and the empirical evidence of their existence. We find empirical evidence supporting the idea that contagion through information is a limited concern.

For contagion through interbank lending, we discussed the clearing mechanism applied to solve defaults in an interconnected system. We review the types of network topologies that minimise contagion. For small shocks, network topologies that maximise average degree minimise direct contagion, while for large shocks, network topologies that minimise size of connected components minimise direct contagion. We then explored empirical evidence on the structure of network topologies, and the theoretical models of network formation justifying observed structures. Interbank networks exhibit a core-periphery network, characterised by a highly interconnected core and and a less-connected periphery. Various theoretical models of core-periphery formation exist, which all use heterogeneities – in bank size, access to information or to investment opportunities – to justify the existence of a core and a periphery.

For contagion through correlated portfolios, we mention fire sales, which is where large increases in supply of illiquid assets drives down its price, reducing the net worth of other banks. We also look at portfolio correlation structures and find that indirect contagion is minimised portfolio correlation is very small and very large. Social welfare is maximised when portfolio correlation is very small. We then explored empirical evidence on the structure portfolio correlation, and theoretical justifications. We find that banks prefer to correlate portfolios with their interbank lending counterparties, and small banks to correlate portfolios with large banks. Large banks are subject to diversification requirements, which results in correlated portfolios.

Finally, we study the relative contributions of contagion from interbank lending and correlated portfolios when faced with a liquidity shock. We discuss a simulation which shows that as the size of the liquidity shock grows, indirect contagion increases to a larger extent than direct contagion, with observed inflection points where contagion rises sharply. Our novel contribution is creating a simple theoretical model that justifies this behaviour using diversification requirements and the core-periphery structure of networks.

Further work here includes removing simplifying assumptions. We are interested in modelling sectors of different sizes, or an uneven number of banks in each sector. The bounds we obtained above in our model of fire sales are also overestimated. For mathematical simplicity, we only took into account the price of the asset when deriving a bound on the net worth of a periphery bank. A more sophisticated model would find an exact solution for net worth as assets are sold, so both assets held and price decreases. Greenwood et al. (2015) has also created a measure of a bank's fragility, "systemicness", which is associated with the quantity of illiquid assets the bank holds, and the leverage of banks which also hold that asset. It would be interesting to enrich our theoretical model by varying the leverage of banks that hold each asset, and seeing its effects on the spread of indirect contagion.

References

Daron Acemoglu, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. Systemic risk and stability in financial networks. *American Economic Review*, 105(2):564–608, February 2015. doi: 10.1257/aer.20130456. URL `https://www.aeaweb.org/articles?id=10.1257/aer.20130456`.

Viral V Acharya and Tanju Yorulmazer. Too many to fail—an analysis of time-inconsistency in bank closure policies. *Journal of financial intermediation*, 16(1):1–31, 2007.

Joseph Aharony and Itzhak Swary. Contagion effects of bank failures: Evidence from capital markets. *The Journal of Business*, 56(3):305–322, 1983. ISSN 00219398, 15375374. URL `http://www.jstor.org/stable/2352800`.

Franklin Allen and Douglas Gale. Financial contagion. *Journal of Political Economy*, 108(1): 1–33, 2000. ISSN 00223808, 1537534X. URL `http://www.jstor.org/stable/10.1086/262109`.

Franklin Allen, Giovanni Covi, Xian Gu, Oskar Kowalewski, and Mattia Montagna. The interbank market puzzle. *Available at SSRN 3540973*, 2020.

Ana Babus and Tai-Wei Hu. Endogenous intermediation in over-the-counter markets. *Journal of Financial Economics*, 125(1):200–215, 2017. URL `https://EconPapers.repec.org/RePEc:eee:jfinec:v:125:y:2017:i:1:p:200-215`.

John M Barron and Neven T Valev. International lending by us banks. *Journal of Money, Credit and Banking*, pages 357–381, 2000.

Michael Boss, Helmut Elsinger, Martin Summer, Stefan Thurner, et al. An empirical analysis of the network structure of the austrian interbank market. *Financial Stability Report*, 7: 77–87, 2004.

Falk Bräuning and José L Fillat. Stress testing effects on portfolio similarities among large us banks. *Federal Reserve Bank of Boston Research Paper Series Current Policy Perspectives Paper*, (19-1), 2019.

Antonio Cabrales, Piero Gottardi, and Fernando Vega-Redondo. Risk sharing and contagion in networks. *The Review of Financial Studies*, 30(9):3086–3127, 2017.

Rodrigo Cifuentes, Hyun Song Shin, and Gianluigi Ferrucci. Liquidity risk and contagion. *Journal of the European Economic Association*, 3(2-3):556–566, 2005. URL `https://EconPapers.repec.org/RePEc:tpr:jeurec:v:3:y:2005:i:2-3:p:556-566`.

Ben Craig and Goetz von Peter. Interbank tiering and money center banks. *Journal of Financial Intermediation*, 23(3):322–347, 2014. URL `https://EconPapers.repec.org/RePEc:eee:jfinin:v:23:y:2014:i:3:p:322-347`.

Fernando Duarte and Collin Jones. Empirical network contagion for us financial institutions. 2017.

Larry Eisenberg and Thomas H. Noe. Systemic risk in financial systems. *Management Science*, 47(2):236–249, 2001. ISSN 00251909, 15265501. URL http://www.jstor.org/stable/2661572.

Matthew Elliott, Benjamin Golub, and Matthew O. Jackson. Financial networks and contagion. *American Economic Review*, 104(10):3115–53, October 2014. doi: 10.1257/aer.104.10.3115. URL https://www.aeaweb.org/articles?id=10.1257/aer.104.10.3115.

Matthew Elliott, Co-Pierre Georg, and Jonathon Hazell. Systemic risk shifting in financial networks. *Journal of Economic Theory*, 191:105157, 2021.

Selman Erol and Rakesh Vohra. Network formation and systemic risk. *European Economic Review*, 148:104213, 2022. ISSN 0014-2921. doi: https://doi.org/10.1016/j.euroecorev.2022.104213. URL https://www.sciencedirect.com/science/article/pii/S0014292122001222.

Maryam Farboodi. Intermediation and voluntary exposure to counterparty risk. *Journal of Political Economy*, 131(12):3267–3309, 2023.

Daniel Fricke and Thomas Lux. Core–periphery structure in the overnight money market: evidence from the e-mid trading platform. *Computational Economics*, 45:359–395, 2015.

Renée Fry, Vance L. Martin, and Chrismin Tang. A new class of tests of contagion with applications. *Journal of Business & Economic Statistics*, 28(3):423–437, 2010. ISSN 07350015. URL http://www.jstor.org/stable/20750850.

Craig H Furfine. Interbank exposures: Quantifying the risk of contagion. *Journal of money, credit and banking*, pages 111–128, 2003.

Prasanna Gai and Sujit Kapadia. Contagion in financial networks. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 466(2120):2401–2423, 2010.

Prasanna Gai, Andrew Haldane, and Sujit Kapadia. Complexity, concentration and contagion. *Journal of Monetary Economics*, 58(5):453–470, 2011. ISSN 0304-3932. doi: https://doi.org/10.1016/j.jmoneco.2011.05.005. URL https://www.sciencedirect.com/science/article/pii/S0304393211000481. Carnegie-Rochester Conference on public policy: Normalizing Central Bank Practice in Light of the credit Turmoi, 12–13 November 2010.

Gerald D Gay, Stephen G Timme, and Kenneth Yung. Bank failure and contagion effects: Evidence from hong kong. *Journal of Financial Research*, 14(2):153–165, 1991.

Paul Glasserman and H. Peyton Young. Contagion in financial networks. *Journal of Economic Literature*, 54(3):779–831, September 2016. doi: 10.1257/jel.20151228. URL `https://www.aeaweb.org/articles?id=10.1257/jel.20151228`.

Robin Greenwood, Augustin Landier, and David Thesmar. Vulnerable banks. *Journal of financial economics*, 115(3):471–485, 2015.

Andrew G Haldane. *Rethinking the financial network*. Springer, 2009.

Iftekhar Hasan and Gerald P Dwyer. Bank runs in the free banking period. *Journal of Money, Credit and Banking*, 26(2):271–288, 1994.

Daan in 't Veld, Marco van der Leij, and Cars Hommes. The formation of a core-periphery structure in heterogeneous financial networks. *Journal of Economic Dynamics and Control*, 119:103972, 2020. ISSN 0165-1889. doi: https://doi.org/10.1016/j.jedc.2020.103972. URL `https://www.sciencedirect.com/science/article/pii/S0165188920301408`.

Giulia Iori, Giulia De Masi, Ovidiu Vasile Precup, Giampaolo Gabbi, and Guido Caldarelli. A network analysis of the italian overnight money market. *Journal of Economic Dynamics and Control*, 32(1):259–278, 2008.

M. Humayun Kabir. Do investors herd during the financial crisis? evidence from us financial industry: Do investors herd during the financial crisis? *International Review of Finance*, 18, 07 2017. doi: 10.1111/irfi.12140.

George G Kaufman. Bank contagion: A review of the theory and evidence. *Journal of Financial Services Research*, 8:123–150, 1994.

Dimitris Kenourgios, Aristeidis Samitas, and Nikos Paltalidis. Financial crises and stock market contagion in a multivariate time-varying asymmetric framework. *Journal of International Financial Markets, Institutions and Money*, 21(1):92–106, 2011.

Philip R Lane and Gian Maria Milesi-Ferretti. The external wealth of nations mark ii: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of international Economics*, 73(2):223–250, 2007.

Kimberly F. Luchtenberg and Quang Viet Vu. The 2008 financial crisis: Stock market contagion and its determinants. *Research in International Business and Finance*, 33: 178–203, 2015. ISSN 0275-5319. doi: https://doi.org/10.1016/j.ribaf.2014.09.007. URL `https://www.sciencedirect.com/science/article/pii/S0275531914000518`.

Pipat Luengnaruemitchai and James A Wilcox. Pro-cyclicality, banks' reporting discretion, and "safety in similarity". 2004.

Allan H Meltzer. Major issues in the regulation of financial institutions. *Journal of Political Economy*, 75(4, Part 2):482–501, 1967.

Christina Mikropoulou and Angelos T Vouldis. Financial contagion within the interbank network. 2023.

John W Peavy and George H Hempel. The penn square bank failure: Effect on commercial bank security returns—a note. *Journal of Banking & Finance*, 12(1):141–150, 1988.

Jean-Charles Rochet and Jean Tirole. Interbank lending and systemic risk. *Journal of Money, credit and Banking*, 28(4):733–762, 1996.

Dirk Schoenmaker. *Contagion risk in banking.* Citeseer, 1996.

Peilong Shen and Zhinan Li. Financial contagion in inter-bank networks with overlapping portfolios. *Journal of Economic Interaction and Coordination*, 15(4):845–865, 2020.

Thiago Christiano Silva, Sergio Rubens Stancato de Souza, and Benjamin Miranda Tabak. Network structure analysis of the brazilian interbank market. *Emerging Markets Review*, 26:130–152, 2016. ISSN 1566-0141. doi: https://doi.org/10.1016/j.ememar.2015.12.004. URL `https://www.sciencedirect.com/science/article/pii/S1566014115000709`.

Kimmo Soramäki, Morten L Bech, Jeffrey Arnold, Robert J Glass, and Walter E Beyeler. The topology of interbank payment flows. *Physica A: Statistical Mechanics and its Applications*, 379(1):317–333, 2007.

Peng Sui, Sailesh Tanna, and Dandan Zhou. Financial contagion in a core-periphery interbank network. *The European Journal of Finance*, 26(7-8):691–710, 2020.

Christian Upper and Andreas Worms. Estimating bilateral exposures in the german interbank market: Is there a danger of contagion? *European Economic Review*, 48(4):827–849, 2004. URL `https://EconPapers.repec.org/RePEc:eee:eecrev:v:48:y:2004:i:4:p:827-849`.

Wolf Wagner. Diversification at financial institutions and systemic crises. *Journal of financial intermediation*, 19(3):373–386, 2010.