

INTERNATIONAL MONETARY FUND

Sandcastles and Financial Systems: A Sandpile Metaphor

Francesco Luna and Luisa Zanforlin

WP/25/41

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate.

The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

**2025
FEB**



WORKING PAPER

IMF Working Paper
Institute for Capacity Development

**Sandcastles and Financial Systems:
A Sandpile Metaphor**
Prepared by **Francesco Luna^{*}** and **Luisa Zanforlin[†]**

Authorized for distribution by Selim Elekdag
January 2025

IMF Working Papers describe research in progress by the author(s) and are published to elicit comments and to encourage debate. The views expressed in IMF Working Papers are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

ABSTRACT: Social welfare costs from bank resolution, including contagion and moral hazard, are often thought to be minimized when supervisors can direct the merger of a failing bank with a sound, healthy one. However, social losses may become even larger if the absorbing institutions fail themselves. We ask whether social welfare losses are indeed lower when supervisors intervene rather than not. We use the sand pile/Abelian model as a metaphor to model financial losses which, as sand grains that fall onto a pile, eventually lead to a slide/failure. When capital in the system is insufficient to absorb the failing institution there will be welfare losses. Results suggest that, over the longer-term, social costs are lower when supervisors manage mergers. Additionally, financial networks that have a structure that minimizes social losses also minimize crises frequency. However, the bank employed resolution strategy will determine which financial network structures are associated with the minimum average loss per bankruptcy event.

JEL Classification Numbers:	C63, C90, G21, G28,
Keywords:	Applied Abelian model, Banking, Financial Networks, Financial Economics, Banking Crisis, Bankruptcy, Bank Resolution, Self-organized Criticality, Bank Regulation and Supervision
Author's E-Mail Address:	FLuna@imf.org ; Lzanforlin@imf.org

[†] The authors are very grateful to Selim Elekdag and German Villegas Bauer. The simulation code was written in python and the notebooks are available upon request.

WORKING PAPERS

Sandcastles and Financial Systems: A Sandpile metaphor

Prepared by Francesco Luna and Luisa Zanforlin

Contents

I. Executive Summary	3
II. Introduction.....	4
III. Merging Failing Banks on Both Sides of the Atlantic	8
IV. The Model.....	11
A. Laissez-faire.....	15
B. Supervisors' Directed Mergers (the Managed case).....	17
C. Comparing system performance under the two strategies	18
V. Searching for the Optimal Network Structure	21
A. The relevance of symmetric relationships.....	23
B. Increasing the Network Dimension	24
C. Changing Regulators' Objectives while Managing Failures: the Cynical Supervisors	26
VI. Analysis of Results	27
A. Increasing the number of simulations	30
B. Comparing Different Network Structures	31
C. Avoiding Symmetric Links.....	32
V. Conclusions and further research plans.....	35
Annex Tables	37
Annex 1. Laissez Faire case, OLS: Average loss on Network size, Average cluster coefficient, and Average Shortest Path.....	37
Annex 2. OLS on combined sample.....	38
Annex 3. OLS Combined sample without symmetric links	40
References.....	42
FIGURES	
Figure 1. US bank failures 1975-1990	10
Figure 2. Number of Banks Failing in a Given Year (1975-1990)	11
Figure 3. Generic financial system as directed graph	16
Figure 4. Matrix representation of the financial system network	17
Figure 5. Simulation Outcomes	19
Figure 6. Total social welfare loss.....	20
Figure 7. Average welfare loss. Four-node networks.....	22
Figure 8. Average welfare loss. Four-node networks without symmetric links.....	24
Figure 9. Network landscape, 4 nodes.....	25
Figure 10. Comparison of pre and post optimization variables	29
Figure 11. Total loss distribution	30
Figure 12. With and without symmetric links. Comparison of results.....	33

I. Executive Summary

Among bank resolution strategies, private sector solutions entailing merging the failing bank with a strong healthier institution have often been supervisors' preferred option. This has been the case for many reasons, not least because of absence of disruption to the rest of the financial sector, to depositors, to the payment system and, importantly, to public expense. Such mergers prevent social losses associated with bank liquidation proceedings, prevent fiscal costs, and contagion. In addition, this bank resolution strategy enhances financial stability by minimizing moral hazard. However, critics have pointed out that larger institutions may just delay the problem and would magnify social welfare costs when they themselves eventually fail; see Beck (2004) and (2011), Hoggarth et al (2004).

We use simulations from a model based on the metaphor of a sand pile to investigate whether this strategy indeed minimizes social welfare losses from failing banks. The model generates losses for banks that, as sand grains, accumulate and eventually trigger a sand "slide" or "failing bank". We assume supervisors follow one of two possible bank resolution policies when a bank is failing, either they let it to its shareholders to deal with the problem or they "arrange" a merger with the strongest institutions they can find at that moment. We ask which of the two supervisory strategies minimizes social welfare losses stemming from bank failures. We assume that banks fail when accumulated losses reach a regulatory threshold, as is the case across many jurisdictions. Social welfare losses will arise in those cases when shareholders do not have enough capital to support their bank. When supervisors follow a strategy of allocating losses from failing banks to the strongest institutions in the system, shareholders are still called to support their institution when it fails after having supported other banks' failures. We assume that this occurs because the banks in the network have become systemic and therefore no other bank in the network can absorb it. "Systemic" banks fail when significant losses have already been absorbed by the system through the mergers. The pattern of welfare losses that follows occasional events of a bank failing, uses standard assumptions for the social welfare function. The simulations suggest that when supervisors direct mergers total welfare losses over the longer term are lower than if they do not intervene, even if failures are, on average, more costly when they occur.

We also search for the characteristics that a financial system (network) should have to minimize the average welfare loss per failing institution, whether a merger is sought or not; and for the network characteristics which would minimize the frequency of failing bank events (rather than minimizing welfare losses), when supervisors manage the resolution of a failing bank via a merger. We find that the characteristics of the optimal network depend on the chosen bank resolution policy and on supervisors' preferences between minimizing the frequency of failing bank events and minimizing total welfare losses when they choose to manage failing banks. The analysis of the characteristics of the different financial networks would allow supervisors to assess how inherently prepared (robust) their system is to withstand failures. Thus, these results could inform the supervisors on the relative efficiency or cost of their chosen resolution strategy.

II. Introduction

The inherent fragilities in financial intermediation and banking activities—such as asymmetric information and imperfect contract enforcement—have manifested themselves time and time again, under different semblances, to cause banks to fail. The question of how to manage failing banks has never lost its relevance and reemerges with force, albeit sporadically, at each new instance.¹ The interconnected nature of banks and financial institutions, their role in safeguarding savings and ensuring the smooth functioning of the payment system underscore how the appropriate management of a single institutions' demise is a necessary element to avoid the occurrence of bank runs, large-scale failures by contagion, and the impairment of the credit flow to the economy.

Among the different bank resolution policies designed to address an illiquid and potentially insolvent banking institution, the preferred ones are those that minimize public costs, disruption to the system, and moral hazard. In this respect, mechanisms where supervisors manage the take-over of the failing banks by a stronger healthier institution fare well when compared to alternatives (Beck 2004 and 2011).² Evidence of this can be found in an early study by Hoggarth, Reidhill, and Sinclair (2004) where they evaluate resolution policies adopted in 33 banking crises over the world during 1977–2002 and find that, when faced with individual failing banks, authorities have usually sought a private-sector resolution policy where arranged takeovers avoided losses to be passed onto taxpayers. While analyzing bank resolution strategies during the Global Financial Crisis, White et al (2014) specifically evaluate the different strategies in terms of feasibility, system disruptions and fiscal costs. Indeed, they find that the private sector solutions where failing banks are taken over by healthier institutions as those with the lower overall costs.

These resolution options go under the more general name of mergers or partial mergers. In its most simple form, the process entails seeking for a “white knight”, a stronger healthier, financial institution that would take over the failing one. In a more general application of the concept, the merger may be designed to be whole or just partial. In cases where it is partial, shareholders of the failing bank are made sound to the extent there is remaining value in the parts of the institution which are not sold, while the healthier institution that takes over part of the failing bank assets and deposits is compensated by the franchise value of the absorbed assets. Such procedures are commonly referred to in the US by the name of purchase and assumption (P&A) or outright sale of business. In the broad review of bank resolution strategies that followed the GFC, the practice has been often praised albeit considered as of limited benefit in cases of a systemic crisis (White 2014). Nevertheless, it is explicitly referred to in the Financial Stability's Board Principles for Systemic Bank Resolution (FSP 2014) as one of the strategies to resolve systemic banks.

Only very recently in the US, we have witnessed a well-known pattern unfold again: a bank is considered unsafe, depositors rush to get their money out, and authorities have to step in to assuage fears and to resolve an illiquid and insolvent bank. The causes may be different and the regulation may be adjusted to try to prevent similar occurrences going forward, but a very common immediate reaction, as in the case of Silicon Valley Bank and Credit Suisse, is to distribute remaining good assets and deposit-like liabilities of the failed entity across larger, more robust and healthy institutions in the system.

¹ FSB (2014) “Key Principles for Systemic Bank Restructuring”.

² White et al. (2014)

From a theoretical point of view, supervisors' practice of supporting the merger of a failing bank with a stronger healthier one appears to offer a rapid, feasible, solution but has been seen to have two, potentially opposite, effects on financial stability. On the one hand, it may boost financial stability by minimizing moral hazard because it provides bankers with incentives to be solvent so as to profit from their competitors' failure. On the other hand, the creation of larger institutions may spoil financial stability as it could potentially magnify the problems in the long term by creating "Systemically Important Financial Institutions" (Perotti and Suarez 2002) and V. Acharya, T. Yorulmazer (2007). (A. Pais, P.A. Stork (2011)). Such institutions either increase risk-taking behaviour and, thereby, the cost of failures or reduce overall liquidity in the system and increase the risk of connected failures (Acharya and Yorulmazer (2008)). However, the effect of increased bank concentration on financial stability in itself is debated (K. Schaeck, M. Cihak, S. Wolfe (2009)) as larger banks are thought to help with economies of scale in the provision of financial services and to contribute in reducing excessive risk-taking behavior from excessive competition. ³ In seminal work, Acharya and Yorulmazer (2008) show that even accounting for some downside effects on system-wide liquidity following the mergers, supporting the merger of failing institutions with stronger healthier ones is a resolution option preferable to others, to the extent that it is worthwhile for the public sector to support it with its own funds.

The intent of this paper is to analyze whether bank resolution strategies where supervisors direct mergers entail lower welfare losses than strategies where failing banks are left to their shareholders to resolve. It also questions whether the structure of a financial sector, in terms of its network characteristics, influences the welfare losses stemming from the chosen resolution strategy. In this respect, it follows a small literature in financial economics that models financial systems in terms of banks as nodes connected in a network through their balance sheet relationships. In particular, Allen and Gale (2000) use a network model with equal nodes to analyze how random liquidity shocks can propagate and spread through financial systems and analyze which characteristics of the network increase the possibility of contagion. The financial system as a network is used to identify those institutions that are "too big to fail" in Alentorn et al. (2007). Nier et al (2007) analyze the transmission of financial shocks through interbank exposures depending on the characteristics of the network, the size of the nodes and the size of the exposures. Elliott, et al (2014) develop a general model that simulates financial contagion and cascades of failures among organizations linked through a network of financial interdependencies (cross-holdings) before analyzing how the characteristics of the cross-holdings influence results. Acemoglu et al (2015) study the network's role as a shock propagation and amplification mechanism, highlighting the implications of the network's structure on the extent of financial contagion and systemic risk.

Differently from the above-mentioned papers, we use the Abelian "sand pile" model (also known as chip-firing game) to simulate the effects of the welfare losses from the chosen bank resolution mechanisms in the financial system. In the late 90s as member of the Center for Computable Economics at UCLA, one of the authors was invited to participate in a research workshop at the Santa Fe' Institute. In that occasion, they were exposed for the first time to the sand pile model by Scott Page. In his presentation, professor Page talked about a few experiments that his group had run on the sand pile model with the following twist: the sand grains, rather than falling randomly on the surface of the pile, were artificially addressed to the spot where the probability of a slide was the smallest at that time. According to his account, the simulations showed that the slides were rarer, but more pronounced when they eventually occurred.

We use the Abelian model to construct a metaphor: the financial sector is represented by spots on a sand pile (nodes) which represent the banks. The behavior of economic agents generates losses for the banks. These losses can be either random, generated by episodes of borrower defaults, or "directed" because supervisors are merging institutions that fail with the strongest banks in the network. Within the network itself, any bank

³ Allen and Gale (2004), Beck and DeJonghe (2013) Noy (2017)

fails after accumulated losses reach a threshold. This will be equivalent to the sand pile collapsing thus generating a “slide”. When a “slide” occurs, the shareholders will have to absorb the losses from the failure. If there are not enough shareholders to absorb the full amount of their accumulated losses, then the system will suffer a welfare loss (for example a loss of savings deposits) which we will call a “crisis”.

Our model is built on a directed graph, a network of nodes tied by directed links. The number of nodes in the system is constant. Each node represents a bank and the links represent shareholder relations. We assume that only financial institutions can own other financial institutions in a strict “fit and proper” requirement. Shareholders are randomly assigned to banks and are represented by the number of links of each node. In the network, all banks of the system are initially endowed with a random amount of capital at the beginning of each simulation. They all will be exceeding the minimum capitalization ratio. If a bank accumulates financial losses in excess of 4 it will be failing. In other words, after the losses surpass the threshold, they will lead to a “slide”—an example of a self-organized criticality.

For pure computational reasons, a simplifying assumption had to be made on how losses from failing entities are distributed to their shareholders, which does not change the overall behavior of the system. We assume that supervisors will only allow shareholders to take on assets and liabilities of their failing owned institutions equivalent to one unit of their capital. This will occur independently of how many shareholders there are (i.e. how much each shareholder has contributed to the initial capital). We assume this could embody a supervisor strategy of limiting systemic risk through network and that banks will not be aware of this until the moment a related institution fails. This also will require that, where there are more than 4 shareholders, the supervisor will have to allocate randomly the losses from the failing bank. In this way, both losses stemming from borrower defaults and losses stemming from failing controlled institutions are also random in this system. However, when shareholders are less than 4, there will be social welfare losses in the system because there will be losses that cannot be absorbed by the shareholders.

Each time an owned bank fails, a shareholder will be allocated $\frac{1}{4}$ of the failing bank balance sheet. This implies it will have to absorb assets, liabilities in excess of the remaining capital of the bank that failed and thus it will have to put up 1 unit of its own capital to absorb them. This will be registered as a reduction in capital by 1 unit. Each time a shareholder absorbs a loss from an owned bank, it grows in terms of assets and liabilities. When there are less than four shareholders, the part of the bank balance sheet that could not be allocated will have to be liquidated and thus will generate a social welfare loss. This will occur for example, because deposits are lost.

In a first scenario, we consider the financial losses accruing to the banks to be credit losses which, as sand grains, randomly fall on the surface of a sand pile and accumulate in banks, nodes, because borrowers are hit by random shocks and become unable to repay their loans. When the amount of accumulated credit losses surpasses the minimum capital threshold (4), the bank will become “failing or likely to fail” and thus a resolution event will be triggered. In our metaphor, the threshold will be established by the supervisors, and will be equivalent to a sort of minimum capitalization ratio. In the first scenario, when resolution events are triggered, supervisors do not intervene and leave shareholders absorb the failing institution, where there are not enough shareholders, there will be social welfare losses as described above.

In a second scenario, we replicate the dynamic of a resolution mechanism where supervisors direct the strongest institutions in the system to absorb the failing banks. When, a bank reaches the minimum capital threshold, the supervisor distributes the losses among the most robust nodes in the system (the ones with the

highest level of capital at that moment). This is equivalent to the process of directing the sand grains to the “safest” current position on the sand pile surface. However, as discussed above, each bank requested to become a “white knight” will be able to absorb only one unit of loss, and therefore the resolution strategy is equivalent to conducting P&A operations, each one of them equivalent to $\frac{1}{4}$ of the failing bank’s balance sheet. The white knight will thus receive assets and liabilities from the failed bank corresponding to one unit of its capital. In this respect, this strategy spreads the losses among a wider numbers of banks implicitly creating larger and larger institutions. We ask whether this strategy of directing losses/sand grains is able to reduce the welfare losses when the slides/failures of the “white knights” eventually occur: Would the resulting “sandcastle” be more resilient to these shocks, less prone to avalanches?

In the second scenario, welfare losses will eventually erupt when the shareholders of the institutions that had been identified as the strongest, our white knights, fail themselves, and there is no bank in the network that is able to absorb their losses. Thus, the losses generated by their failure will have to be borne by depositors/taxpayers. In short, the supervisors’ strategy of allocating the losses to the strongest institutions allows for losses to accumulate (an increase in pressure) before “crises” are triggered, and a domino effect is unleashed in a “systemic failure” in terms of welfare losses.

To our knowledge the Abelian model has not been used before in the context of financial contagion analysis. In part, this exercise builds on that approach to address policy questions related to procedures and desirable structural features of the financial system. However, we believe this is the first time a self-organizing criticality model is employed to analyze the impact of policy options to address bank failures.

Our results show that when the supervisors direct mergers, events where the white knight (or their shareholders) eventually fail—i.e. when there is not enough capital in the system that can take over their losses—occur indeed less frequently but each event is on average more onerous in terms of social welfare because the accumulated financial losses are large, as would be expected. However, most interestingly, we find that in the medium/long term the total social welfare loss is smaller if the authorities support mergers rather than leaving shareholders to deal with the failing banks themselves. We also find that the bank resolution strategy also determines which financial network structure is optimal in the sense that it minimizes average losses per failing institution.

The rest paper is organized as follows: Section I recounts a series of historical events that put into question the efficacy of the strategy of “merging” ailing banks with stronger ones to show that in several cases such measures delayed the inevitable and possibly lead to much worse results. Section II presents in detail the network model and describes in what way it is useful to represent two bank resolution strategies that authorities may choose to pursue; then compares the outcomes of the two bank resolution strategies across different networks. Section III presents a methodology to seek for the best performing network assuming that regulators can have different optimization strategy; when the networks have higher dimension and when supervisors have different optimization strategy. Section VI discusses the results of the model when the number of simulations is increased, when the characteristics of the network are different and in the case symmetric link among institutions would be disallowed; Section V briefly discusses the results and concludes with policy implications and possible extensions of the work.

III. Merging Failing Banks on Both Sides of the Atlantic

One of the historic notable examples of the repeated implementation of merger strategies toward failing banks is Austria during the interwar period.⁴ Historians report that, at the beginning of World War I (WWI), in Austria there was a network of banks of approximately 300 small, specialized banks and ten universal banks. Among the smaller banks were the Sparkassen (savings institutions) and entities with small operations. Profitability issues in the corporate productive sector at the end of the war rapidly led to non-performing assets and distress in the banking sector. Two of the universal banks were acquired by foreign competitors. Many of the smaller entities that ended up falling into distress were either merged into competitors or were liquidated. To be noted in particular, one of the universal banks, the Bode the Boden-Credit-Anstalt absorbed a multitude of small- and medium-sized banks. Those universal banks which suffered from high levels of NPLs were merged and reorganized as needed and by 1927 only four remained.⁵ In 1927, three of the remaining universal banks experienced an increasing amount of non-performing loans. In 1930, the Union Bank UB and the Verkehrsbank (VB) were merged into the B the Boden-Credit-Anstalt (BCA). In 1931, the BCA itself had to be absorbed by the largest universal bank remaining, the CreditAnstalt (CA), (which ended up with 90 million schillings of capital, but 140 million schillings of accumulated non performing loans). Unfortunately, but not surprisingly, at the end of 1931 the CA also collapsed. The total losses of the bank were estimated at 1,070 million schillings, 7 times the original amount of assets absorbed. Many historians agree that the absorption of the other three universal banks directly contributed to the CA's demise. Finally, the large scale exceptional liquidity support awarded to CA by OENB out of its own gold reserves, is sometimes thought to have contributed to the end of the gold standard in Europe.⁶

On the other side of the Atlantic, the political debate in the US Congress during the pre-WWI period had favored developing the so-called unit banking model. In this setup, banks are restricted in size and geographic breadth so as to become local in nature, the intent of the legislators was to avoid contagion and systemic losses from individual failures. The oversea experiences in Austria gave additional ammunition to politicians against the universal banking model. The McFadden Act of 1927 (Banking Act, 1927), as amended by the Banking Act of 1933, prohibited banks from establishing branches across states. The states soon followed suit by establishing that branches should be limited to the central office county or to counties adjacent to it or by restricting branching altogether. States in the latter category are referred to by historians as Unit Banking states. The landscape of US banking after the Great Depression remained fairly unchanged with reduced merger and very limited acquisition activities, in part due to a 1963 US Supreme Court decision which held that a merger was illegal if it resulted in a firm controlling an undue share of the relevant market (373 U.S. 321 and "Clayton Act" 1914).

However, competitive and technological pressures were building toward a de-regulation of the banking system. At the end of the seventies, regulators had to confront with the inefficiencies of the unit-banking model, what is now called the "Savings and Loans (S&Ls) crisis". The need to address promptly episodes of banking distress

⁴ Marcus (2011) Macher, F. (2018)

⁵ - the Verkehrsbank (VB), the Unionbank (UB), the Boden-Credit-Anstalt (BCA), and the CreditAnstalt -

⁶ The money that had to be printed by the Austrian government to face domestic depositors' withdrawals in CA and the drop in foreign reserves generated by the pull back of foreign creditors, eventually caused Austria to abandon the gold standard and the rest of creditor countries, beginning with the UK followed suit.


pushed legislators towards relaxing restrictions on banking activities and, eventually, the passing of the Garn-St. Germain Act of 1982 (Depository Institutions Act, 1982), which permitted limited interstate expansion by allowing failing institutions to be acquired by out-of-state institutions.⁷ By the end of the 1980s the risks posed by geographic lending concentrations were well understood, so attempts were made to eliminate the remaining legal impediments to full interstate banking. In fact, several state actions had enabled many banking firms to use bank holding company affiliations to circumvent geographic restrictions. In 1983, the rising amount of non-performing loans in the banking sector lead to estimate that it would cost roughly \$25 billion to pay off the insured depositors of failed institutions, however the thrifts' insurance fund, known as the FSLIC, had reserves of only \$6 billion and the authorities declined to close down distressed institutions. As a result, the total assets of the S&L industry continued to grow rapidly until 1985, when the whole industry came under distress. According to historical records between 1980 and 1994 more than 1,600 banks insured by the FDIC were closed or received FDIC financial assistance. ⁸From 1986 to 1995, the number of federally insured savings and loans in the United States declined from 3,234 to 1,645. In 1989 Congress instituted the Resolution Trust Corporation (RTC). This agency was created to deal with failures and had the task of dealing with distressed assets and customer deposits of the S&Ls. Estimations put the cost to taxpayers through the RTC to be as high as \$124 billion.

Figure 1. presents the number and average size of failed institutions in the US in the period 1970 to 1990. On the right vertical axis the average size of bank failures in a given year which were resolved through purchase and assumption transactions, in other words, by re-distributing assets and deposits across the system. Specifically, P&A are a process whereby the failed banks' assets are purchased by other banks and liabilities are assumed. On the left vertical axis, we included the total number of bank failures in that year:

⁷ More specifically, during the bank distress years of 1920-1939, 15 states relaxed their branching restrictions, while in the four decades that followed (1939-1979) only four states relaxed branching limits. When bank distress returned in the 1980s, 15 states relaxed their branching rules (Mengle 1990, Calomiris 2000, pp. 63-7).

⁸ Federal Reserve History

Figure 1. US bank failures 1975-1990




Source: FDIC Dataset on US Bank Failures and Assistance and authors calculations.

The chart displays the average size of banks, in terms of total assets, that failed in any given year and were resolved through Purchase and Assumptions (P&A) and includes the estimated loss per bank (as presented in FDIC statistics).

In our terminology, this would represent the increasing losses that are generated as the assets of the failed banks are distributed across the system. We chose this specific and limited period because of the significant overhang of bad assets in the system at the time which would make it closer to the closed model of sandcastles we are examining later as a theoretical model.

Figure 2. presents the relationship between bank size and number of failures in a given year of distress. It can be noted that the larger the number of banks that failed in a year, the larger their average size (Figure 2). It is important to note the data here below do not include a dynamic pattern though. Together with the previous graph, FDIC data for the distress period appears to suggest that both the average size of banks and the average amount of losses per bank tend to increase in the sample of banks resolved through P&A as the situation of distress persists:

Figure 2. Number of Banks Failing in a Given Year (1975-1990)



Source: FDIC Dataset on US Bank Failures and Assistance and authors calculations.

The chart displays the average size of banks, in terms of total assets, that failed in any given year and were resolved through Purchase and Assumptions (P&A) and the total number of failed banks.

In our terminology this represents the outcome the authorities achieve through “distributing” the assets of the failed banks to the remaining banks in the system and thus achieving larger and larger banks that in turn, when they fail, generate higher and higher losses.

IV. The Model

The sand pile model is a simplification of the original Abelian⁹ setup which generates a self-criticality and could be referred to as an example of catastrophe theory. In the model, sand grains, fall randomly on the surface of a sand pile, thus a “dune”, and as they do so they potentially trigger “slides” i.e. they cause the sandpile to collapse. While traditionally the sand pile is modeled to grow into a dune by the grains of sand falling randomly, the sand grains can also be modelled to be directed to the location in the pile that is less likely to collapse.¹⁰ In these cases, the pile becomes a “sand castle”. We ask whether, the castle obtained would be able to endure more shocks than a randomly obtained dune. In addition, we ask whether, if the castle were eventually to collapse, will its slide be broader than that of a dune. Would the slide be wider than the sum of all the slides that would have occurred if the pile had not been tampered with? Although the idea was proposed in various

⁹ The name, as suggested by Kumaraswamy Velupillai, could be linked to the Norwegian mathematician Niels Henrik Abel probably for his work on group theory. So far, however, we must confess, we have not been able to confirm this hypothesis.

¹⁰ We give the source of this “inspiration” in the next section. However, we have learned that such a twist has been studied and classified in the combinatorics literature as “chip-firing games” on a graph. See for example Spencer (1986) and Bjorner et al (1991)

contexts (see Holroyd et al 2008), the model is often referred to as the Bak–Tang–Wiesenfeld model (Bak, et al. 1987) and proposes a dynamic framework exhibiting self-organized criticality. For the purpose of this exercise, rather than operating in a two-dimensional grid, we set up a network or, more precisely, a directed graph.

In our model, borrowers are hit by random shocks, they cannot repay their loans, and thus generate credit losses that accumulate in banks reducing their capital. When capital breaches the minimum regulatory level the losses generate a “slide” or a failing bank event. In the first scenario supervisors do not intervene (we dub this “laissez faire”) and leave shareholders deal with their institution, thereby the system mimics the dynamics of a “dune” in our metaphor.

In a second scenario, when a bank’s capital falls below the regulatory minimum, supervisors assign losses from failing banks to the strongest elements in the network: the “white knights”. The allocation of losses to healthier ones with capital in excess of regulatory limit makes the system able to withstand the losses without generating a crisis. In our metaphor, we are moving away from the “dune” to generate a “sandcastle”.

It is important to note that in both scenarios the system is hit by the same number of random shocks. We ask which of the two strategies, bank resolution policies in our case, is most advantageous for the system in terms of minimizing social welfare losses.

As mentioned above, the banking system represented in the model is composed of banks endowed with random levels of capital above the minimum required. Institutions also differ in the number of shareholders. The shareholders of each bank are represented by the number of connections each entity has to the rest of the network. It is assumed only banks can own other banks, sort of a stringent version of a “fit and proper” requirement. The number of directed links in and out of each node (bank) is randomly determined when the network is created. The number of shareholders determines the amount of support each bank has before it generates a welfare loss for the system. Banks that have an “insufficient” number of shareholders when they fail, will generate a welfare loss for the system of the magnitude of the capital shortfall that its current shareholders cannot absorb.

Also, the higher the number of “shareholders” the larger the financial conglomerate each banking institution belongs to. In this respect, we assume that each financial network can be defined in terms of: (i.) size or number of total number of links/relationships across financial institutions that belong to the network; (ii) average clustering coefficient¹¹, or number of shareholding relationships among institutions that are connected within a “financial group”; and iii) the average minimum path of the network, which is the average number of shareholding relationship that link any two nodes in the network. This latter measure resembles the notion of “completeness” in the literature on financial networks introduced by Allen and Gale (2009) whereby “complete”¹² networks are those where each group is connected to every other group in the system. As mentioned above, the networks we will be analyzing may in some cases be incomplete.

While randomly generated networks can have reciprocal connections, in most regulatory frameworks, cross-ownerships among deposit-taking financial institutions are either not allowed or have to be deducted from

¹¹ A clustering coefficient is a measure of how nodes in a graph/network tend to cluster together. For a precise definition of this measure we refer the interested reader to https://en.wikipedia.org/wiki/Clustering_coefficient

¹² In graph theory these are called “strongly connected” graphs

regulatory capital. These types of regulations acknowledge the issue that cross-holdings actually implies less capital for the aggregated financial group¹³. To make the network simulation closer to a banking system with restrictions on cross-ownerships, a second set of simulations will be run on a network without reciprocal connections and the implications for the results will be discussed subsequently.

The network is closed, so that it represents banks in a specific region. However, the network is not fully connected as to present the semblance of large and interconnected banking groups. In this respect, it may happen that the network is not “complete”¹⁴ as discussed in Allen and Gale (2000) whereby it may not be possible to find a connection between any two financial institutions in the network. This is consistent with the idea of large financial groups that operate separately in the financial sector. Also, differently from the setup in Allen and Gale (2000) in our case links are representing ownership relationships and not interbank deposits/credits.

More specifically, we assume that at the beginning of time banks are randomly endowed¹⁵ with a min of 8 and a max of 12 capital units. We assume that prudential regulations in the model require banks to hold a minimum 8 units of capital (comparable to the 8 percent of minimum capital requirement as a share of total assets) and thus each bank will be able to withstand up to 4 units of losses before failing (depending on its initial endowment). When a bank goes below the 8 percent level of capital, it is categorized as “failing” and a resolution strategy is followed.

We also assume that the banking network is subject to effective prudential oversight by the network’s banking regulator. This implies that any expected losses the banks may have at any moment in time have to be fully covered by adequate provisions. Thus, the only losses that can occur to the banks in the network are unexpected, thus random, and will have to be covered by capital¹⁶.

In our simulations, borrowers from the banks in the network are assumed to default randomly and unexpectedly. Thus, the banks in the network that own the defaulted loans will have to recognize the loss against their capital as they have not had a warning to make a provision. For simplicity, each loan that becomes non-performing (NPL) in the model has a value of one unit. The occurrence of a new, unexpected, NPL will generate a loss of capital of one unit for the bank. This implies that at each new NPL, the total amount of capital of the banking sector is reduced by the total amount of the new random NPLs generated. After the occurrence of a credit loss, the number of shareholders of each bank remains unchanged.

In the Abelian model, credit losses are the equivalent of the grains of sand falling onto the pile. All grains (hence NPLs) are assumed to be equally sized. Following the sand pile metaphor, once the number of grains piled up in the same place has reached a threshold, a “slide” occurs and the grains move to the downstream nodes in the directed graph. If a bank goes below the 8 minimum capital level is declared as *failing or likely to fail* and a resolution event is triggered.

In a first set of simulations, supervisors do not intervene and leave shareholders to absorb the losses. Shareholders however will be allowed to absorb only one unit of losses for each bank they own which fails. When they take over the failing bank, they will each receive its assets and liabilities equivalent to the ¼ of the assets and ¼ of the liabilities of the failing institution and they will have to put up 1 unit of their own capital to

¹³ FDIC “...Reciprocal Cross Holdings: Under the agencies’ capital rule, a banking organization must deduct from regulatory capital”

¹⁴ In other words, the graph underlying the network may not be “strongly connected”.

¹⁵ The initial distribution is random to depict some heterogeneity in the starting situation.

¹⁶ As per usual banking regulations, expected losses have to be covered by provisions while unexpected losses will be covered by capital. In this setup, if the bank gets hit by an unexpected NPL its capital will be reduced by 1.

absorb them. If there are more than 4 descending connections, we assume that shareholders will be randomly selected to take over the losses.

By construction of the original network, there may be fewer than four (4) shareholders (remember that the links are random so it may well be that a node does not have 4 descendants/shareholders). It may also be the case that some shareholder bank has already failed by reaching its minimum capital threshold and is tagged as under distress, and isolated from its controlled institutions (the upstream nodes). It is thus important to note that shareholder banks in distress cannot support their affiliated banks: all “incoming links” for such banks are broken for that period.

When there are banks with less than 4 shareholders, it is assumed nobody else in the system wants to acquire the remainder part of the portfolio of the failing institutions and supervisors do not intervene. Therefore, the equivalent proportion of loans has to be recalled (to repay the depositors). This, in turn, will generate a welfare loss in that (credit/production/output) will have to be reduced proportionally to the total amount of the capital shortfall. We call welfare losses only those losses the network does not absorb and, as a result, explicitly fall on the general public (the deposit insurance corporation or eventually taxpayers). Historically, this operation has been carried out by the deposit insurance corporation (DIC) that provided resources to pay out depositors and, in this way, would spread the losses across the system.

When there is a welfare loss in the system as a result of a failing bank, i.e. there are not enough shareholders to absorb all the losses of the failing institution, we call it a “crisis”. In this respect, it is important to note that there can be bank failures that do not generate “crisis”, this will happen if there are at least 4 shareholders and each of them has enough capital to support the institution.

Also, since each shareholder will be requested to absorb $\frac{1}{4}$ of the losses of each failing entity they own, there will be instances when the n th failing entity will cause one of its “shareholders” capital to fall below the minimum and thus fail itself¹⁷ in a domino effect. So, in the more general form, we associate “slides” with failing bank events, where the extent of the total losses to the system from each event will depend on the characteristics of the network and the extent of accumulated losses in the network. Therefore each “slide” or “failing bank event” will be characterized by the amount of failing institutions and the total welfare loss.

In a second set of simulations, supervisors direct the strongest institutions (nodes) in the system to absorb the losses of a failing bank¹⁸. In this setup, the stronger institutions, the “white knights” are defined to be those with a higher amount of remaining capital. If a bank is called to absorb a failing one, it will suffer a “loss”, in the sense it will have to use its own capital to support the acquisition of assets and liabilities from the institution it absorbs.

To recall, we have assumed, that the mergers are never “total” but rather P&As for $\frac{1}{4}$ th of the balance sheet. Thus, every institution called to become a “white knight” will only purchase assets and liabilities which it can support with one unit of capital, and thus will see its own capital reduced by one as a result of each of the P&A operation it is called to participate in. In this way the supervisors’ strategy implicitly generates larger and larger banks to avoid banks’ failures.

We also assume that, to begin with, an institution can only participate in taking over a failing institution if its capital is at or above 9. However, because of the nature of the simulations, eventually the stronger institutions

¹⁷ The status “under distress” will last for one day. Since the objective is to gauge the characteristics of a specific network, the same network will be used on each day in the life of a generation. The structure itself can be modified only across generations, as we will discuss later.

¹⁸ Note that these chosen banks may not be the shareholders of the failing one, hence the losses are distributed more “widely” across the system.

will have a capital of 9 and thus the n th failing institution will generate a social welfare loss because of the accumulated losses already in the system. Eventually, the “white knights” themselves will fail, and no other institution in the system will be strong enough/willing to absorb them and thus losses from the failure will either burden their shareholders or become welfare losses as those described above.

Also in this second scenario, we also expect “slides” i.e. shareholder bank failures to be associated with other banks failing if they also happen to be the shareholders of the original shareholders of the “white knight” bank. In this respect, the loss allocating algorithm is generating “systemic banks” i.e. banks the failure of which is eventually associated with a large number of connected failures.


A. Laissez-faire

As discussed above, the first set of simulations assumes that supervisors do not intervene when banks are failing. In this case, the banks’ shareholders will have to take over the failing institution to the extent possible. They acquire a share of their assets (loans) and liabilities (deposits) and will have to provide the supporting capital on account of their larger balance sheet. We will refer to this scenario as “laissez faire”, i.e. where the regulators do not intervene in the process such that the random structure of the system (i.e. shareholders) take care of the failure. In this first scenario, supervisors take a passive role and only require related parties to have capital above minimum regulator levels ¹⁹.

For these simulations, we have fixed the dimension of the directed graph at 15 nodes and a typical network would look like the one depicted in Figure 3.

¹⁹ Since a failed bank has deposits in excess of assets as a result of the random occurrence of 4 NPL shocks, each of its shareholding connections are called to absorb its losses. It is assumed that supervisors will distribute the assets and the liabilities of the failed bank in such a way that each shareholder bank can acquire at most 1 capital unit of the failed entity and thus put up at most 1 unit of their current capital resources. This occurs because it is assumed that each entity can absorb a loss of 1 unit of its capital and still remain above minimum regulatory requirements--i.e. all institutions are assumed to have a max of 12 capital units to begin with against a minimum of 8 in the capital regulation, so that absorbing 1 unit of the total loss of the failing entity would bring them down by 1 unit closer to the 8 threshold, which is still acceptable to supervisors. Thus, the additional capital necessary to support the failed institution balance sheet is put up by other institutions in the conglomerate so that the original loss is shared across institutions.

Figure 3. Generic financial system as directed graph



It is important to stress that the graph is not fully connected: the number of directed links in and out of each node is randomly determined when the network is created. This implies that power relations are not symmetric, so the shareholders and controlled institutions of each bank (node) are well identified. In this case, there could be cross-ownership and in Figure 3, the symmetric relations (when they exist) are represented by lines with arrow points at both extremes like the one linking node 6 and 11 (on the lower right corner). On the other hand, the link between 10 and 11 is only unidirectional: from 10 to 11 (in our model this is to be interpreted as “bank 11 is a shareholder of bank 10”).

The directed network is first built by generating an asymmetric square matrix (15 x 15) with zeros along the diagonal (to avoid the paradox of a bank transferring a liability to itself). The following matrix lies behind the directed graph depicted in Figure 1. It won't be a surprise that such a representation is very convenient. It is the perfect set up to apply a simple genetic algorithm, that will split and recombine the sequence of zeros and ones. In Figure 4, row n collects the “address” of all shareholders of the n th node. For example, the first row (indexed as zero in python), links node 0 to nodes 3, 4, 7, 8, 9, 12, and 14. It means that banks 3, 4, 7, 8, 9, 12, and 14 are shareholders of bank 0.

Figure 4. Matrix representation of the financial system network

0	0	0	1	1	0	0	1	1	1	0	0	1	0	1
0	0	0	0	0	0	0	1	0	0	0	0	1	1	1
1	1	0	1	0	1	0	1	1	0	0	1	0	0	1
1	0	0	0	1	1	0	1	1	0	1	1	0	1	1
0	1	1	0	0	1	0	1	1	0	1	0	1	1	1
1	0	1	1	1	0	1	0	0	0	0	1	1	1	0
0	0	0	1	1	0	0	0	1	0	0	1	1	0	0
1	0	1	0	0	1	1	0	1	1	0	1	0	1	0
1	1	1	0	1	1	0	1	0	0	1	1	0	1	1
1	1	1	0	1	1	0	0	0	0	1	0	1	0	1
1	1	1	0	1	1	0	1	1	1	0	1	0	0	0
0	1	1	1	1	0	1	0	1	0	0	0	1	1	0
0	0	1	0	1	1	0	1	1	0	0	1	0	0	1
1	0	1	0	1	1	1	1	1	0	1	1	1	0	1
0	1	1	1	0	0	1	1	1	0	0	1	0	0	0

As mentioned above, the threshold that will trigger an event (a slide/failing bank) is set at 4 (this number seems very common among the simulations found on the web; see for example <http://www.natureincode.com/code/various/sandpile.html> and also is designed to be consistent with the banks' regulatory framework we suggested as the requirement is to remain above 8).

The first round of simulations is organized on a set of 300 iterations (they could be interpreted as days) and the performance of the network is calculated at the end of the "year." Initially, the network allows for cross-holdings across institutions. The performance focuses on three dimensions: i) the average time (number of days) between failing bank events, ii) the average number of connected banks failing (slides) in the network each time there is a bank that is failing (we call this a crisis event), and iii) the average welfare loss generated by each failing bank (crisis) event. In our terminology, an event of a bank failing is associated with a "slide" in the model, but such event may not translate into a public welfare loss at every instance of occurrence. For example, if there is enough capital in the system and the shareholders are strong enough, the bank failure is fully absorbed by its shareholders. We call the "slide" event a "crisis" if there is a welfare loss associated with it.

B. Supervisors' Directed Mergers (the Managed case)

We use the model discussed above to analyze the behavior of the system if regulators were to direct losses of a failed bank to the strongest institutions in the network (managed case)²⁰.

Regulators' intervention is reflected by allocating assets and liabilities of a failing bank to the banks (nodes) that are the strongest in the network in terms of remaining capital. This process will then represent closely the process of "looking for an appropriate match". The banks receiving the assets and liabilities will have to put up one unit of their own capital, thus their capital level will be reduced by one. As discussed above, the process of

²⁰ There are, in principle various ways to model this scenario. The simplest would be to direct each new shock to the strongest institution. This would correspond to the "benevolent dictator" case, in the sense that losses would be managed "centrally" from an institution that has full control of the system. Alternatively, the supervisor may intervene only when a specific bank has reached the threshold and goes bankrupt, before involving its official shareholders. This is the option we have chosen, and the results reported in this paper refer to such a case. However, we have also run all the same simulations for the "benevolent dictator" case. The results are qualitatively equivalent, but more "extreme" since the supervisor in that framework, has more latitude and can potentially identify the best allocations. This is the actual experiment performed with the dune by Professor Page and his team: the sand grains are immediately directed to the "safest spot" in the pile.

receiving assets and liabilities from failed institutions equivalent to one unit of capital will imply that these white knights are becoming larger and larger institutions (sort of “systemic banks”)

These simulations are also done for the 300 “days” period as in the previous case. The banks that are called to be “white knights” remain exposed to the same random shocks as the rest of the system and may eventually generate losses to their shareholders if they fail and no other institution in the network is able to support them²¹. The same will be true if their own shareholders fail, no stronger institution would be considered able to absorb them. In this way these simulations replicate, in a simplified way, the welfare effects from a “systemic bank failure”.

Results from the two sets of simulations are then compared to assess whether the process of directing mergers, and thus directing the losses to the strongest institutions, will result in a reduction of the frequency of welfare-loss generating crisis events and to identify possible side effects. In particular, we are most interested in assessing whether these rarer events of crisis, i.e. where shareholders are not able to absorb the losses, are more costly to the system as a whole than those when losses occur randomly and single institutions are left to their shareholders to handle.


C. Comparing system performance under the two strategies

The outcome from the simulations records the average time to a failing bank event (crisis), the average number of failing banks associated with crisis events (like in a domino effect), and the average welfare cost incurred in a crisis.

In Figure 5, these series are collected under the two sets of simulations dynamics on a group of 50 randomly generated networks (on the x axis).

²¹ In this model, there is no “recapitalization option” in the sense that no new capital can be created, but only “consumed” by shocks. This is clearly a simplification that allows us to address more directly the question of “how a network should be structured to face efficiently the shocks”


Figure 5. Simulation Outcomes



Results suggest that when supervisors direct the losses to the strongest institutions, the frequency of crises²² (meaning events that generate welfare losses for the system) is reduced as expected. (because even if each institution fails after 4 losses, a simplifying assumption, only the institutions that each time have more capital are called to be white knights. However, as expected, simulations also show that when institutions eventually fail as a result of the built-up pressure in the system, the average welfare loss at each of the crisis event is larger because it is associated with a higher number of connected institutions also failing (the supervisors' intervention eventually leads to the failing of systemically important financial institutions). Thus, our initial results, support the view that supervisors' interventions that seek to merge failing institutions with stronger, healthier ones lead, on average, to a deeper social welfare loss per outcome when the crisis eventually happens.

However, when the *total* social welfare cost of bank failures to the system is computed over the full period²³ (300 "days") our simulations find that total welfare loss is typically lower in the "managed (shotgun marriage)" case (Figure 6: "tot Loss man" in the graph). This implies that the benefit from the lower frequency of the recorded crises, more than compensates for the higher average welfare cost of each occurrence²⁴.

Figure 6. Total social welfare loss



In this respect, regulators' decision to pursue a resolution strategy by directing mergers, as discussed, appear to yield a superior outcome from a total welfare cost point of view, despite each crisis becoming more costly.

²² Note that the average period of time that leads to a crisis event, sometimes less than two days, may seem too short. This is due to the high number of shocks/losses occurring in the simulations which was not chosen to deliver realistic time to crisis, but rather to study the network characteristics.

²³ The Y axis records the sum of all losses induced by all crises over the period of 300 days. In the "benevolent dictator" case mentioned above, the total losses in the managed case are strictly smaller: the supervisor has control over the whole system and manages to "optimize" the loss allocation. So the contrast between "managed" and "Laissez faire" is even starker than in the case depicted in Figure 6.

²⁴ This is true also when symmetric links are excluded.

V. Searching for the Optimal Network Structure

We now try to ascertain whether certain characteristics of the network structure (financial system) may be generating lower welfare losses upon a failing bank depending on which of the two resolution strategies are followed. The analysis of the network characteristics would allow the supervisory authority to assess how inherently prepared (robust) that very structure is to withstand shocks, in terms of avoiding large welfare losses, given the chosen bank resolution policy. In addition, it would inform the regulatory authority on which structures it could promote or prevent in order to minimize welfare losses from failing banks.

To assess the potential performance of a network, and/or to improve upon the current network structure, it is necessary to identify in the first place the characteristics of the “champions”. These are the networks with structural features that have been consistently exhibiting the “best performance” and that can, hence, safely be employed as benchmarks. In our case these are the networks that minimize total welfare losses, or alternatively, as we will see later, the ones who minimize the frequency of the crisis.

One strategy to identify the best performing networks would be to run the same set of simulations over all possible network structures. Such an analysis is feasible only for rather small networks. For example, we constructed all possible networks of size 4, that is, with 4 nodes. A simple calculation shows that a directed graph with 4 nodes can be constructed in 2^{12} possible ways (4096).²⁵ However, already the equivalent calculation for a system composed of 15 institutions shows that there are 2^{210} possible network structures or 1,645,504,557,321,206,042,154,969,182,557,350,504,982,735,865,633,579,863,348,609,024. This implies that for a realistic financial system, with more than 15 institutions, the optimization problem, i.e. finding the network structure that minimizes total welfare losses, would not be simple. However, with 15 nodes our network may well represent the banking relationships in a small country. For example, in Rwanda there are 16 commercial banks²⁶.


We begin by analyzing the 4 node networks²⁷ and we expose each of them to the same set of 100 shocks. We conduct the simulations first assuming that the shareholders are left to take care of the failing institutions and then assuming that the losses created by each new bankruptcy are directed by the supervisors to the stronger institutions in the network. Figure 7 represents the distribution of results in terms of average welfare loss of crisis events under the laissez-faire and the managed strategy.

²⁵ As mentioned above, links from and to the same node are excluded so there are 12 possible connections. Each of these can be in two directions, hence the total 2^{12}

²⁶ In fact, the financial system in Rwanda is also composed of several microfinance institutions, but *microfinance and rural institution have their own regulator and in general work differently*.

²⁷ For this exercise, since all possible networks are considered, we do not focus on the case where symmetric links are excluded. For a 4 node system, the possible networks without symmetric links are only 729.

Figure 7. Average welfare loss. Four-node networks



The difference in the average welfare loss distribution under the two bank resolution strategies is quite striking, while in keeping with the earlier result that total welfare losses are typically lower in the managed case over a period of time. In particular, when supervisors actively manage the allocation of losses to the stronger institutions, not only there is a higher average welfare loss per occurrence of crisis, as a result of the many losses accumulated in the system, but also a very volatile distribution of average welfare losses across the different networks.

In addition, we find only one configuration that behaves as a “champion” (i.e. minimizes the average welfare loss in a crisis event) when shareholders are left to take over the failing institution (the “laissez-faire” case). It averages a loss of slightly over 3.5 units per crisis (which occurs on average every 2.37 “days” for a total loss of 147.68 over the 100 days). However, we find 64 configurations that lead to the minimum average loss of 7.1 units per crisis event (on average after 4.84 “days” for a total loss of 146.69), for those cases in which the supervisor intervenes managing the resolution process of the failing bank(s) (the “managed” strategy).


It is important to note that, this set of simulations confirms that the “managed” strategy results in smaller losses over the “longer period” and that it may generate a larger number of networks, or “states of the world”, that lead to the best outcome and therefore there is a higher probability that it can be achieved.

A. The relevance of symmetric relationships

As mentioned above, in many countries the regulatory framework for banks either forbids or requires capital deductions for cross-holdings for financial corporations. These regulations are imposed because cross-holdings effectively reduce the total amount of capital available for the financial group (or cluster in these simulations). We sought to account for this by generating a set of networks that have the same number of nodes but disallowing symmetric relationships. We find qualitatively very similar results for the set of networks that exclude symmetric links. The total welfare loss over a period of time remains lower in the case of the “managed” bank resolution strategy²⁸, thus supporting our important finding that the supervisors’ strategy of allocating losses to the healthier banks has beneficial effects in the long term. As in the previous simulations, the average loss per crisis event continues to be lower in the “laissez-faire” simulations than in the managed one despite being the size of the population much reduced.

²⁸ The minimum Total Loss for the managed case is 137 and for the laissez faire case is 141. The average Total Loss is 146.79 for the laissez faire, and 142.70 for the managed case.


Figure 8. Average welfare loss. Four-node networks without symmetric links



B. Increasing the Network Dimension

We now ask whether it would be possible to identify (perhaps converge to) the “champion” network(s) in a higher dimension universe of potential networks that would represent more closely actual banking sectors. We would like to identify the characteristics of the network(s) that would minimize total welfare losses whichever of the two bank resolution strategies is followed. For the case of networks of 4 nodes, Figure 9 suggests that performance in terms of average loss per crisis event is widely different across networks. In particular, we find that each network, represented by one point on the x axis, differs from its immediate neighbors only for one link, but the difference in average welfare loss per crisis event is stark.

Figure 9. Network landscape, 4 nodes



With such a rugged landscape, the search for the network that minimizes the total welfare loss is likely to get stuck in some local optimum for every dimension of the network above 4 nodes, especially in the “managed” case where we see several large plateaus. For this reason, the search for the “champion” will have to employ a procedure that is robust with respect to such very high nonlinearity. We have decided to use a genetic algorithm.

In particular, in our search for networks that successfully minimize the total welfare loss, we have set up the following procedure. First, we create a population of (15-node) networks and expose them to the same set of shocks under both strategies (laissez-faire and managed) for addressing failing banks. At the end of this set of shocks, we identify the network that has had the best overall performance in terms of total welfare loss.

Second, we generate a new population of networks (a new generation) employing a genetic algorithm that exploits only cross-over and is designed to implement a search routine over a highly non-linear landscape²⁹.

Each generation—composed of 26 “individuals³⁰”—is exposed to the same set of shocks and the selection process is repeated 30 times (i.e. 30 generations). At the end of the overall exercise the characteristics of the best network of the first generation are compared with those of the best network of the last generation that, in our set up, has necessarily exhibited an equal or better overall performance. We say “necessarily” because we are using one of the tricks employed in this sort of simulations: our generations maintain a “ratchet”; that is, the best performer of the current generation remains the incumbent champion in the following generations until a better performer emerges. Furthermore, to make sure the results do not depend on the specific set of shocks employed, we also repeat the full set of exercises with a different series of shocks³¹.

The simulations with the same set of shocks and the same initial number of banks in the network (population) are repeated under the two “bank resolution strategies” as above: the case in which related entities take care of the failing institution and the case in which regulators manage the failing banks.

It is important to note that the number of network structures tested in this exercise, a total of 780 (26 x 30), is only a minuscule portion of the whole potential universe of networks for a 15 node system. This is clearly a limitation on the possibility of identifying the absolute optimum, but it should still give an indication of the characteristics of the networks that would be associated with the minimum welfare losses upon the event of a crisis.

C. Changing Regulators’ Objectives while Managing Failures: the Cynical Supervisors

We now assume that rather than minimizing social losses from crises, the regulatory authorities’ primary objective becomes that of preventing any crisis as long as possible, while still assuming that they resolve failing banks by assigning losses to the strongest banks in the network. We consider this to be a reasonable strategy for supervisors that are drawn back (or operate in a world characterized by less than perfect foresight where “total welfare loss” cannot be the argument of the objective function and determine the on-going behavior). Indeed, one can observe that the occurrence of a crisis always suggests supervisory and regulatory failures in the real world and is usually associated with calls for resignation and regulatory overhaul. Thus, we assume that regulators, given the choice, would choose a network that reduces the probability they would have to deal with a failing bank during their tenure, which would also maximize their chances to remain in the job. Such strategy would be in line with a political economy of the banking sector “à la Acemoglu³², i.e. where the

²⁹ As, mentioned above, each network is described by a string of zeros and ones. To obtain a new network via the “cross-over” procedure, we randomly select two points in the string of one network and substitute the middle section with the equivalent segment taken from another string. We could have set up a two-stage approach. The first would have been based on genetic programming to search for an effective combination of routines (i.e. the genetic algorithm we employ or the random exploration, or the “gradient” approach sketched above). The result would have been a routine to apply in the selection process of the “best” network. So, we see genetic programming as a way to select algorithms (using the genetic principle of the survival of the fittest), while a genetic algorithm is only one such an approach to select the solution. We are aware that the distinction is more semantic than substantial. However, it requires a programming competence that we are currently acquiring and could lead to a more efficient search process, hence the application of the metaphor to a larger network.

³⁰ An “individual” in this context is a network structure.

³¹ We are not testing (or comparing) the champion emerged from one set of shocks on (with) another set because the intent here is not the search for some “global” optimum, but rather the identification of the typical characteristics that a good network should have.

³² Acemoglu and Robison (2006) AER “De Facto Political Power and Institutional Persistence”

system is regulated in such a way whereby supervisors would not need to take any action during their tenure leaving the potential negative consequences from their action to the next generation of supervisors. We dub this regime as “supervisory cynicism”. It translates in doing “whatever we can” given the available instruments/resources for bank resolution, but not believing in benefits of any “long-term” planning. In fact, this is the only “managed case” that could lead over time to the selection of “better” structures with only limited “cognitive resources” on the part of the policy maker. This is because the objective function (time to failure) is “local” and observable, whereas the “total welfare function” is observable only at the end of a very long period: it hence requires very long-term planning.

This means that the criterion for selecting the benchmark structures while searching the optimal network has changed. This time, rather than considering the networks that minimize total social welfare losses when the supervisors manage the failing bank, the optimization algorithm is set to identify those network structures that allow for the maximization of the time between crisis events or, equivalently, the *minimization of the number of crises per period* of time.

VI. Analysis of Results

The results from the optimization process described above, are presented for the three cases: 1) where supervisors follow a resolution policy of “laissez-faire”: 2) where they manage failures through directed mergers, and 3) where, while they manage failures, they prefer to have the longest possible time span across episodes (i.e. minimize the frequency of crises) rather than minimizing total welfare losses. Figure 10 captures the starting and end value³³ for several variables: i) average time until a failing bank event (crisis episode), ii) average number of failing banks (“slides”) per crisis episode, iii) average welfare loss per crisis episode, iv) total welfare loss for the period considered. In addition, some structural features of the networks are presented: v) size of the network (total number of links/relationships across financial institutions that belong to the network), vi) average clustering coefficient³⁴, in our metaphor this would indicate the probability that the banks have common shareholders and vii) average minimum path of the network (this is the average of the number of steps to link any two nodes in the network and it is another measure of how interconnected the overall network is), in our metaphor this would indicate how closely connected is the financial network in terms of shareholders holding multiple institutions. An average minimum path equal to one would represent the notion of “completeness” in graph theory (with every node directly connected with every other node in the network) and in the literature on financial networks introduced by Allen and Gale (2009) whereby “complete” networks are those where each group is connected to every other group in the system.

Each of the seven frames of Figure 10 shows 3 pairs of bins. The first bin of each pair is the initial champion’s performance (e.g. average loss in the third frame), while the second bin records the performance of the final champion (after 30 generations created with the GA described above).

³³ The beginning value is the one recorded for the “champion” of the first generation, whereas the end value is recorded for the champion prevailing after 30 generations. In some cases, the value does not change, indicating that the initial selection got stuck on a rather “strong” local optimum or that the global optimum was found right away.

³⁴ A clustering coefficient is a measure of how nodes in a graph/network tend to cluster together. For a precise definition of this measure we refer the interested reader to https://en.wikipedia.org/wiki/Clustering_coefficient

The first pair reports on the performance of the dynamics during “laissez-faire”. The second pair is for the performance when the supervisors allocate losses to strongest institutions, i.e. the “managed” rule is employed. The third pair is devoted to the performance of cynical supervisors, who do not believe minimizing social welfare losses is beneficial to them and/or doable. We assume that regulatory authorities are able to influence the network structure by issuing regulations and thus they can achieve their intentions in terms of their ultimate objective. So, in this third case, the GA algorithm which picks the “best performers” is applied to a different objective function: the minimization of the frequency of crises.


The comparison of the optimization results, between the network structures that minimize welfare losses occurred randomly and the ones generated by the resolution strategy, unsurprisingly evidences a direct relationship between average time to crisis and the magnitude of the average crisis event. Typically, the longer it takes for a crisis to erupt, the larger the welfare loss that that crisis will generate. This result is consistent with the general criticism of resolution policies where supervisors support the takeover of failing institutions by larger healthier ones.

However, as noted above, the optimization results computed over a significant number of “generations” of the network confirms the initial observation of Figure 6, i.e. that total welfare loss is either equal or lower when supervisors allocate losses identifying white knights than when shareholders are left to take care of their institution (fourth frame of Figure 10 (total welfare loss per period))

In addition, when observing the results for the “managed case” across all characteristics of the network that minimize social welfare losses, results suggest that regulators who prefer to minimize total welfare losses (at the end of the period), could improve the performance of their financial system by promoting certain features in the financial sector network through regulatory intervention. In particular, we find that networks that are larger in terms of total number of shareholding relationships, are more closely connected via shareholding relationships (have a lower average minimum path), and where institutions are more clustered, also generate the lowest welfare losses.

This series of results is particularly interesting in the context of understanding the discussion behind the negative externalities brought by the systemic impact of large and interconnected financial institutions: it appears that the chosen bank resolution strategy also influences the structure of the network that would lead to the minimal total welfare loss, and therefore there is no single solution.

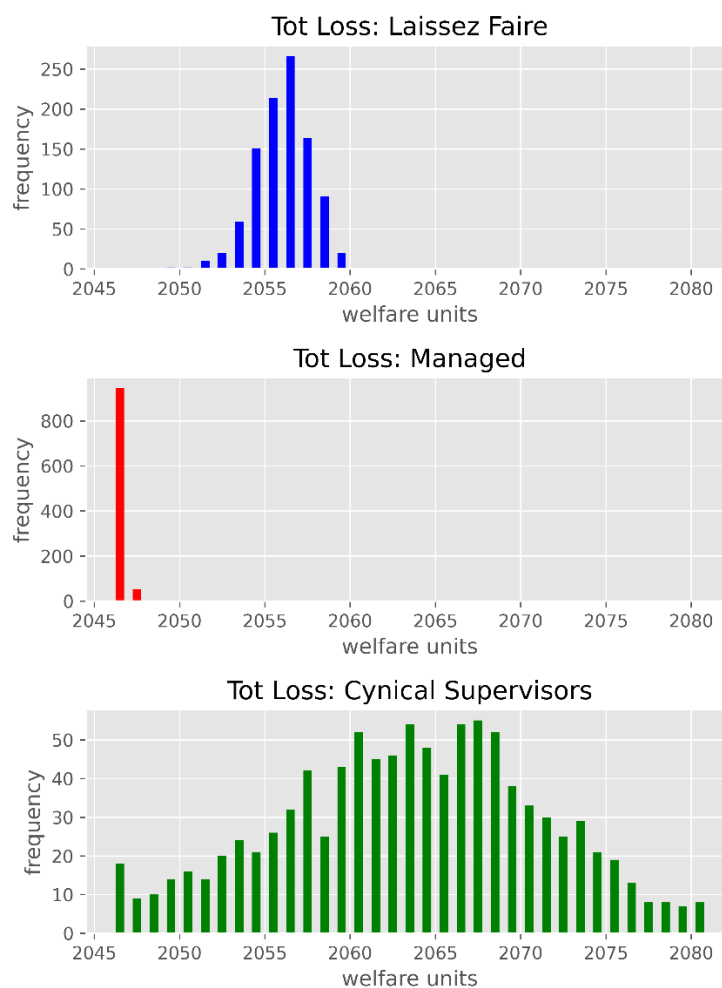
Figure 10. Comparison of pre and post optimization variables



A. Increasing the number of simulations

To further test the hypothesis suggested above of an interaction between regulators' behavior and optimal characteristics of the system, we have run a larger number of simulations. We have performed three sets of 1000 simulations for the three bank resolution policies. Each one of the 1000 simulations starts with a different population and runs over 10 generations. The data of the “champions” at the end of each simulation give 1000 sets of observations for each case. Figure 11 depicts the total loss distribution for each exercise.

Figure 11. Total loss distribution



The results indicate that in the “unmanaged/laissez faire” case—where losses are random and shareholders/related parties take care of the failing institutions—the total welfare loss performance is clearly worse than in the managed case. Notably, the distribution of the welfare loss when supervisors allocate the losses from failing banks is significantly smaller than the “laissez faire” case and almost totally (947 cases out of 1000) concentrated in one outcome (2046). This result appears to confirm the observation that allocation of losses creates some rather flat landscape that could trap the search algorithm in a local optimum. Alternatively, we could conclude that the GA algorithm is so powerful that it manages to find the global optimum in about 95

percent of the cases. Of course, we cannot be sure of this latter hypothesis as the proof would require spanning the full space of 15-node networks, which cannot be done at this time.

Finally, when supervisors are allocating losses to white knights, and the objective function maximizes the time between crises (cynical supervisors), the total welfare loss results are much more volatile. In its best case, (in about 2% of the cases at 2046) they are comparable with the losses obtained for the total-welfare loss minimizing network of the “managed case”. However, the average total loss (at 2063) is larger than the average of the “Laissez faire” case (at 2055.5). Furthermore, it is important to note that only in 4% of the cases, will the total loss under “moral hazard” be smaller than the minimum of the laissez faire regime while in about two thirds of the cases, it will be worse than the “laissez faire” worst case.

Clearly, simulations suggest that “cynical supervisors” would have the easiest time achieving their optimal framework, but the risk involved appears to be quite high. In particular, in the large majority of cases total welfare losses would be in excess of the general “managed case” where regulators are seeking to minimize social welfare losses.

B. Comparing Different Network Structures

We now seek to identify the characteristics that the optimal network should have to guarantee the best results in terms of minimum welfare loss for the event when an institution fails within the network regardless of which resolution strategy is applied. We use OLS regressions on the data generated by the 3000 simulations.

The first exercise has been to regress, for each of the three scenarios, “Total Loss” on: i. the size of the network, ii. the cluster coefficient, and iii. the average minimum path length. However, since “Total Loss” has very low variation, the constant intercept is the only significant coefficient. We do not report those results.

As a second exercise, we decided instead to regress the “Average³⁵ welfare loss” on: i. the size of the network, ii. the cluster coefficient, and iii. average minimum path length. As in the first exercise, the regression was repeated on the simulation results for each of the resolution regimes (laissez faire, managed) and for the case of cynical supervisors. Only in the case of “laissez faire” the coefficients are statistically significant, but the fit is very poor in all cases. (Annex 1).

Finally³⁶, we pooled all the data obtained from the simulations across the three “regimes” and ran instead three separate equations, regressing each time “Average loss” on only one of the independent variables: average size, average cluster coefficient, and average minimum path (we can do this because each of the network characteristics will be independently distributed from the other). The results (in Annex 2) add an interesting twist to the analysis suggesting that the best network structure, in terms of loss minimization “winning structure”, depends on the supervisor’s chosen resolution strategy and therefore is not unique.

In particular, we find that the size of the network in terms of total shareholding relationships plays a positive role in reducing the average loss when supervisors choose the “laissez faire” strategy (which tends to minimize

³⁵ In all scenarios, “laissez faire”, “managed” and “moral hazard”, the correlation between average and total loss is positive, although small, ranging from 0.3, to 0.05.

³⁶ We thank our colleague Carlos de Resende for suggesting this approach, but do not hold him responsible for how we implemented it.

average losses in general). However, we find that when supervisors manage losses and prefer networks structures to minimize the number of crisis events, they will also benefit from networks with a higher number of shareholding relationships

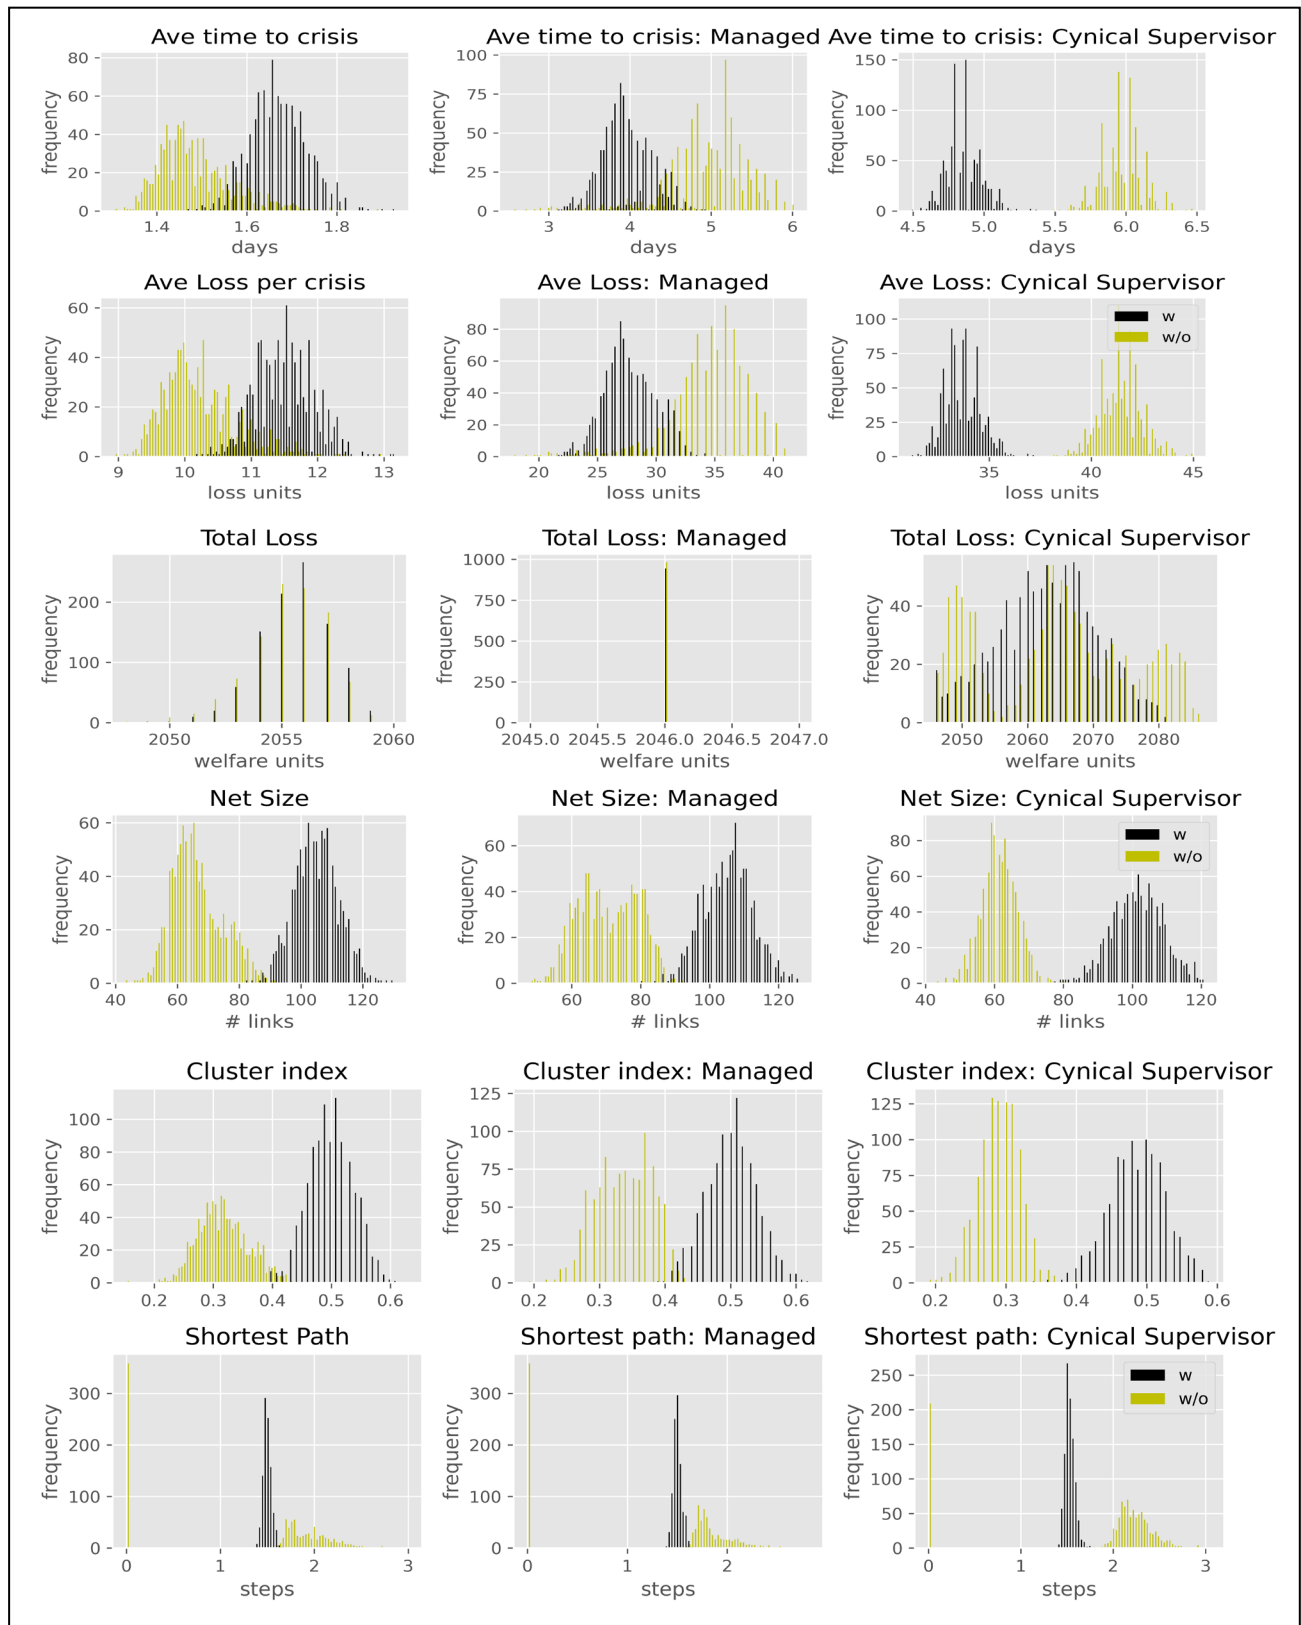
In general, and consistently with the results discussed above, a network that is more connected (as represented by a larger size and a higher cluster coefficient) tend to fare worse when losses are allocated by the supervisors, independently of whether they are minimizing total losses or the time between failure events.

C. Avoiding Symmetric Links

As noted above, many regulatory frameworks for banks either forbid or require capital deductions for cross-holdings for financial corporations. Thus, we generated networks with the same number of nodes, but without any symmetric relationships, such that cross-holdings are completely disallowed for purposes of assessing capital³⁷. Detailed results are reported in Annex 3. Figure 12 presents the distributions of the simulation results to compare the performance of the networks with and without symmetric links (i.e. without or with netting cross-holdings).

³⁷ To be precise, we have eliminated all direct cross-holding. In terms of the directed network this means: if A is connected to B, we made sure that B is not connected to A. However, we did not prevent triangular cross holding like A connected to B, B connected to C, and C connected to A. This aspect should be taken in to consideration in real world calibration which also consider weighted links.

Figure 12. With and without symmetric links. Comparison of results



We find that, when cross-holdings are prohibited, the occurrence of crises is very different depending on the resolution strategy chosen by the supervisors. If the supervisors follow a “laissez-faire” approach crises tend to occur more frequently (i.e. at shorter intervals), than if symmetric links are allowed. This result seems intuitive because there are less shareholders to absorb losses. The average loss is also lower, as one would expect because the system has less time to build “pressure” before the crisis erupts. These findings are also supportive of the intuition behind the regulations, i.e. that by appropriately identifying the extent of supportive shareholders, the system doesn’t overleverage and thus causes institutions to fail more frequently, but also to remain smaller--because we have properly counted the extent of capital in the system.

When supervisors follow a resolution strategy allocating the losses from failing banks to white knights, the absence of symmetric links tends to significantly increase the period between crisis events especially when supervisors are cynical. We assume this depends on the fact that the shareholders of each institution are “stronger” in the sense there is a limited number of institutions they can receive losses from. In effect, it appears to be that the absence of symmetric links associates more closely the number of shareholders with the actual amount of supporting capital and therefore the algorithm picks “better” “white knights”. However, as we found earlier, when crises are less frequent, they are always associated with a higher average welfare loss as the failing institutions tend to be larger, or more numerous in such cases.

In general, the total welfare loss associated with crisis episodes when symmetric relationships are not allowed does not change much with respect to the case in which symmetric links are allowed (as would be expected because the total amount of capital in the system is not different). However, it does appear that the chosen institutions currently with the highest amount of “capital” are somewhat better insulated in the absence of symmetric links and therefore can withstand a longer period of time without crisis. Indeed, the cluster coefficient is lower and the minimum average path is larger, both indicating more sparsely connected networks.

Furthermore, the absence of symmetric links allows supervisors to significantly increase the period between crisis events when supervisors actually choose to do so within their regulatory objectives. As in the cases discussed above, despite the average loss being larger as more institutions are affected in the domino effect, total loss for the system is not necessarily larger, so, once again, disallowing symmetric links is an efficient way of diminishing the incidence of crisis events. Finally, and consistently with the results in the “managed” case, less clustered networks also have a lower frequency of crisis.

Therefore, our results are supportive of the notion that, in the case of an active role played by regulators in managing crises (“managed” and “cynical supervisors” scenarios), prohibiting cross-holdings in financial institutions (or netting cross-holdings from total capital) is beneficial to the system in the sense that such a measure increases the period between crisis event without automatically increasing total welfare losses.

In the process of generating networks without symmetric links, a significant portion of these networks turned out to be “not-strongly connected”, meaning that there is no path that links each node with every other one else; in other words, one or more of the network nodes is isolated or create an isolated subnetwork. Clearly, in these cases it is not possible to calculate an “average minimum path³⁸”. The share of these “not-strongly connected” networks varies between 35 and 50 percent in our simulations. For this reason, the third regression in Annex 3 (average loss regressed on average minimum path) has a total of 1655 observations (rather than

³⁸ In Figure 12 above, these cases are bunched in the “0” bin.

3000 like the other two) as we could use only the strongly connected networks for that regression. This fact led to another interesting observation summarized in Table 1.

TABLE 1 Strongly connected vs not-Strongly connected

	Ave SIZE		Ave CLUSTER		Ave TOT LOSS	
	not-connected	connected	not-connected	connected	not-connected	connected
Laissez Faire	62.74	68.12	0.30	0.30	2055.25	2055.34
Managed	64.54	73.86	0.31	0.31	2046.00	2046.02
Cynical Supervisors	60.22	61.54	0.29	0.29	2063.89	2064.47

The average size and cluster coefficient for Not-Strongly connected networks are significantly smaller than for connected networks. However, these features do not seem to have any impact on the average Total Welfare loss. This would imply that, for optimal networks, the existence of isolated subnetworks does not affect the performance of the overall system, at least from a total welfare perspective.

V. Conclusions and further research plans

In this paper we asked whether supervisors' active management of failing banks reduces social welfare losses from the incidence of failures. We also analyzed whether there are specific characteristics of the financial system structure that may magnify or reduce social welfare losses depending on how supervisors choose to manage failing banks.

We used a sandpile model to replicate the welfare losses stemming from the incidences of failing banks. A first set of simulations assumed losses were randomly generated and the shareholders of any failing bank would have to deal with the entity. A second set of simulations assumed that supervisors allocated the losses to the strongest institutions of the network in terms of capital. We assumed that social welfare losses are generated when shareholders cannot support their own bank. The simulations allow us to compute the social welfare losses and to compare them under both "bank resolution regimes" when using standard assumptions for a social welfare function.

As expected, we found that the social welfare losses are less frequent but larger in those cases where failing banks are actively managed by supervisors who seek to have them taken over by healthier financial institutions. This occurs because pressure accumulates over time in the system and it is released all at once when a white knight itself fails, as it generates a domino effect which deepens the welfare loss to the system. However, when computed over a period of time, we find that total social welfare losses from actively allocating losses to the strongest banks in the system are lower than in the case of "laissez -faire" where shareholders are left to deal with their failing institutions.

Our results validate the theoretical belief that resolving bank through directed mergers in the long run reduces social welfare costs from failures. When we replicated the simulations for a system where cross-ownerships are not allowed, as often required by banking regulations, our results still hold. In this way, our results strengthen the arguments in support of supervisors' active intervention in failing banks.

An additional set of simulations assumed that the objective function of supervisors (not necessarily explicitly stated), when designing regulations that will affect the structure of the banking system, is the minimization of the frequency of crises rather than the minimization of total loss. This strategy was dubbed “cynical supervisors”. We searched for the characteristics of such networks and estimated the extent of social welfare losses. We found that, unsurprisingly, when failure events become less frequent, the average magnitude of each event is larger.

However, in many cases, the total social welfare loss of cynical supervisors is larger than when supervisors’ objective was explicitly to minimize total social welfare losses. Most importantly, we found that those networks which minimize the frequency of failing bank events also minimize social welfare losses independently of which bank resolution regime supervisors follow. In addition, we find that, when supervisors are cynical even if the cumulative welfare losses for the system are sometimes lower than in the case where they do not intervene at all, often they lead to worse results.

Our analysis of the performance of a network, financial system, based on its characteristics (size, clustering, shortest average path), suggests that there does not seem to be one architecture that will always guarantee the minimum social welfare losses. This could be either because we have not yet identified this universally optimal feature or, as suggested in the second round of econometric analysis, because the winning structure depends on the prevailing bank resolution policy followed by supervisors. Features such as size—that leads to a better performance in the *laissez faire* scenario—may be counterproductive in other regimes. Such a result is even more evident when symmetric links are excluded from the admissible network structure.

Our findings suggest several possible extensions. One possible direction would be to search for the best universal structure, by applying the same methodology we have developed, to a network that is structurally similar to a real financial system. This implies building not only a directed graph, but also a weighted one, to capture more precisely the ownership and control relations among financial institutions. However, so far, the restrictions have been imposed quite crudely on the artificially created networks. In fact, only directly symmetric links are excluded. So, if node A is linked to node B, the opposite direction is excluded, but not the case where A is linked to B, B is linked to C which in turns is linked back to A. As part of the extension, it may be interesting to build networks where all shareholding triangulations are prevented³⁹. Given the continuous discussion about the merits of bank capitalization, we could also run simulations with a threshold for “failure” higher than the maintained 8. The goal would be to see whether such higher threshold would have some favorable, non linear results on frequency and, even better, on total welfare loss of crises.

³⁹ Such a construction would lead to a clustering index equal to zero

Annex Tables

Annex 1. Laissez Faire case, OLS: Average loss on Network size, Average cluster coefficient, and Average Shortest Path

Dependent Variable: LOSS_LF

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 08/28/24 Time: 14:43

Sample: 1 1000

Included observations: 1000

LOSS_LF = C(1) + C(2)*SIZE_LF + C(3)*CLUST_LF + C(4)*PATH_LF

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	23.69771	2.625495	9.025997	0.0000
C(2)	-0.022957	0.008981	-2.556009	0.0107
C(3)	-2.472345	0.824850	-2.997329	0.0028
C(4)	-5.659769	1.229812	-4.602140	0.0000
R-squared	0.035325	Mean dependent var		11.49479
Adjusted R-squared	0.032419	S.D. dependent var		0.439353
S.E. of regression	0.432173	Akaike info criterion		1.164008
Sum squared resid	186.0260	Schwarz criterion		1.183639
Log likelihood	-578.0041	Hannan-Quinn criter.		1.171469
F-statistic	12.15726	Durbin-Watson stat		2.095158
Prob(F-statistic)	0.000000			

2] Managed case, OLS: Average loss on Network size, Average cluster coefficient, and Average Shortest Path

Dependent Variable: LOSS_MAN

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 20:27

Sample: 1 1000

Included observations: 1000

LOSS_MAN = C(1) + C(2)*SIZE_MAN + C(3)*CLUST_MAN + C(4)
*PATH_MAN

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	-3.562400	11.28470	-0.315684	0.7523
C(2)	0.025426	0.039429	0.644853	0.5192
C(3)	1.705608	3.944882	0.432360	0.6656
C(4)	18.26081	5.315314	3.435510	0.0006
R-squared	0.067813	Mean dependent var		27.52884
Adjusted R-squared	0.065005	S.D. dependent var		2.165686
S.E. of regression	2.094114	Akaike info criterion		4.320130
Sum squared resid	4367.770	Schwarz criterion		4.339761
Log likelihood	-2156.065	Hannan-Quinn criter.		4.327591
F-statistic	24.15164	Durbin-Watson stat		1.976272
Prob(F-statistic)	0.000000			

3] Cynical supervisors case, OLS: Average loss on Network size, Average cluster coefficient, and Average Shortest Path

Dependent Variable: LOSS_MH

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 21:17

Sample: 1 1000

Included observations: 1000

LOSS_MH = C(1) + C(2)*SIZE_MH + C(3)*CLUST_MH + C(4)
*PATH_MH

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	30.15527	3.972491	7.591022	0.0000
C(2)	-0.004064	0.014649	-0.277464	0.7815
C(3)	1.114338	1.444334	0.771524	0.4406
C(4)	2.239361	1.807025	1.239253	0.2155
R-squared	0.015869	Mean dependent var		33.72440
Adjusted R-squared	0.012904	S.D. dependent var		0.856391
S.E. of regression	0.850848	Akaike info criterion		2.518824
Sum squared resid	721.0458	Schwarz criterion		2.538455
Log likelihood	-1255.412	Hannan-Quinn criter.		2.526286
F-statistic	5.353304	Durbin-Watson stat		2.077381
Prob(F-statistic)	0.001167			

Annex 2. OLS on combined sample

1] Combined Sample, OLS: Average Loss on network size

Dependent Variable: AVELOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 21:55

Sample: 1 3000

Included observations: 3000

AVELOSS = C(4) + C(1)*D1*SIZE+C(2)*D2*SIZE+C(3)*D3*SIZE

	Coefficient	Std. Error	t-Statistic	Prob.
C(4)	27.64821	0.401081	68.93423	0.0000
C(1)	-0.152971	0.003826	-39.98229	0.0000
C(2)	-0.001455	0.003814	-0.381618	0.7028
C(3)	0.059399	0.003954	15.02062	0.0000
R-squared	0.973651	Mean dependent var		24.24934
Adjusted R-squared	0.973625	S.D. dependent var		9.467683
S.E. of regression	1.537591	Akaike info criterion		3.699644
Sum squared resid	7083.105	Schwarz criterion		3.707652
Log likelihood	-5545.466	Hannan-Quinn criter.		3.702524
F-statistic	36903.21	Durbin-Watson stat		1.893216
Prob(F-statistic)	0.000000			

2] Combined Sample, OLS: Average Loss on Clustering coefficient

Dependent Variable: AVELOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 21:59

Sample: 1 3000

Included observations: 3000

AVELOSS = C + C(1)*D1*CLUST+C(2)*D2*CLUST+C(3)*D3*CLUST

	Coefficient	Std. Error	t-Statistic	Prob.
C	27.05958	0.366463	73.83979	0.0000
C(1)	-30.96315	0.735540	-42.09579	0.0000
C(2)	0.864547	0.732593	1.180120	0.2380
C(3)	13.54831	0.752456	18.00545	0.0000
R-squared	0.972738	Mean dependent var		24.24934
Adjusted R-squared	0.972711	S.D. dependent var		9.467683
S.E. of regression	1.564001	Akaike info criterion		3.733704
Sum squared resid	7328.515	Schwarz criterion		3.741713
Log likelihood	-5596.557	Hannan-Quinn criter.		3.736585
F-statistic	35634.00	Durbin-Watson stat		1.884734
Prob(F-statistic)	0.000000			

3] Combined Sample, OLS: Average Loss on Average Minimum Path

Dependent Variable: AVELOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 22:06

Sample: 1 3000

Included observations: 3000

AVELOSS = C + C(1)*D1*PATH+C(2)*D2*PATH+C(3)*D3*PATH

	Coefficient	Std. Error	t-Statistic	Prob.
C	18.10670	0.824059	21.97257	0.0000
C(1)	-4.371290	0.545530	-8.012918	0.0000
C(2)	6.249760	0.546427	11.43750	0.0000
C(3)	10.16341	0.536881	18.93047	0.0000
R-squared	0.979270	Mean dependent var		24.24934
Adjusted R-squared	0.979249	S.D. dependent var		9.467683
S.E. of regression	1.363844	Akaike info criterion		3.459824
Sum squared resid	5572.773	Schwarz criterion		3.467833
Log likelihood	-5185.736	Hannan-Quinn criter.		3.462705
F-statistic	47175.38	Durbin-Watson stat		1.957566
Prob(F-statistic)	0.000000			

Annex 3. OLS Combined sample without symmetric links

1] Combined Sample no symmetric links, OLS: Average Loss on network size

Dependent Variable: LOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 22:20

Sample: 1 3000

Included observations: 3000

LOSS = C(1) + C(2)*D1*SIZE + C(3)*D2*SIZE + C(4)*D3*SIZE

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	25.41284	0.457731	55.51922	0.0000
C(2)	-0.225445	0.006936	-32.50179	0.0000
C(3)	0.122138	0.006513	18.75159	0.0000
C(4)	0.258827	0.007557	34.25095	0.0000
R-squared	0.957433	Mean dependent var		28.59517
Adjusted R-squared	0.957391	S.D. dependent var		13.51284
S.E. of regression	2.789326	Akaike info criterion		4.890810
Sum squared resid	23309.90	Schwarz criterion		4.898818
Log likelihood	-7332.215	Hannan-Quinn criter.		4.893690
F-statistic	22462.51	Durbin-Watson stat		1.428259
Prob(F-statistic)	0.000000			

2] Combined Sample no symmetric links, OLS: Average Loss on Clustering coefficient

Dependent Variable: LOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 22:22

Sample: 1 3000

Included observations: 3000

LOSS = C(4) + C(1)*D1*CLUST + C(2)*D2*CLUST + C(3)*D3*CLUST

	Coefficient	Std. Error	t-Statistic	Prob.
C(4)	26.13648	0.434937	60.09252	0.0000
C(1)	-49.41805	1.383307	-35.72458	0.0000
C(2)	23.41332	1.297799	18.04079	0.0000
C(3)	51.93344	1.514129	34.29922	0.0000
R-squared	0.954809	Mean dependent var		28.59517
Adjusted R-squared	0.954764	S.D. dependent var		13.51284
S.E. of regression	2.874028	Akaike info criterion		4.950638
Sum squared resid	24747.07	Schwarz criterion		4.958647
Log likelihood	-7421.958	Hannan-Quinn criter.		4.953519
F-statistic	21100.02	Durbin-Watson stat		1.469824
Prob(F-statistic)	0.000000			

3] Combined Sample no symmetric links, OLS: Average Loss on Average Minimum Path⁴⁰

Dependent Variable: LOSS

Method: Least Squares (Gauss-Newton / Marquardt steps)

Date: 10/12/24 Time: 22:55

Sample: 1 2075

Included observations: 2075

LOSS = C(1) + C(2)*D1*PATH + C(3)*D2*PATH + C(4)*D3*PATH

	Coefficient	Std. Error	t-Statistic	Prob.
C(1)	22.18174	0.401420	55.25821	0.0000
C(2)	-5.977228	0.205676	-29.06141	0.0000
C(3)	6.834134	0.216716	31.53498	0.0000
C(4)	8.556039	0.180972	47.27819	0.0000
R-squared	0.980187	Mean dependent var		29.79590
Adjusted R-squared	0.980158	S.D. dependent var		13.40085
S.E. of regression	1.887668	Akaike info criterion		4.110487
Sum squared resid	7379.573	Schwarz criterion		4.121355
Log likelihood	-4260.630	Hannan-Quinn criter.		4.114470
F-statistic	34151.51	Durbin-Watson stat		1.216664
Prob(F-statistic)	0.000000			

⁴⁰ As noted above, the sample in this case is smaller as it is run only on those networks that are strongly connected.

References

- Acemoglu, D., Ozdaglar, A. and A. Tahbaz-Salehi (2015) “ Systemic risk and stability in financial networks” *Am. Econ. Rev.*, 105 (2) (2015), pp. 564-608
- Acemoglu, D. and Robinson J. A. (2013) "Economics versus Politics: Pitfalls of Policy Advice." *Journal of Economic Perspectives*, 27 (2): 173-92.
- Acemoglu D. and Robinson J. A. (2006) “ De Facto Political Power and Institutional Persistence” *American Economic Review* May pp:325-330.
- Acemoglu, and D.Robinson. J A. (2000) “Why Did the West Extend the Franchise? Democracy, Inequality, and Growth in Historical Perspective. *Quarterly Journal of Economics* 115(4): 1167– 99Lui.
- Acharya, Viral V. and Tanju Yorulmazer, (2008) Cash-in-the-Market Pricing and Optimal Resolution of Bank Failures, *The Review of Financial Studies*, Volume 21, Issue 6, November 2008, Pages 2705–2742, <https://doi.org/10.1093/rfs/hhm078>
- Acharya, V. and T. Yorulmazer (2007) Too many to fail – an analysis of time-inconsistency in bank closure policies *J. Financ. Intermediat.*, 16, pp. 1-31
- Allen F., and Gale D. (2004) “Competition and financial stability” *J. Money Credit Bank.*, 36 (3), pp. 453-480
- Allen F. and Gale, D. (2000) “Financial Contagion” *Journal of Political Economy*, 108:1, 1-33
- Ahmiud, Miller eds (1998) “Banks Mergers and Acquisitions” Kluwer
- Bak, P.; Tang, C.; Wiesenfeld, K. (1987). "Self-organized criticality: an explanation of $1/f$ noise", *Physical Review Letters*, 59 (4): 381–384.
- Beck T., De Jonghe O., Schepens G.(2013) “Bank competition and stability: Cross-country heterogeneity” *Journal of Financial Intermediation*, 22 (2) , pp. 218-244
- Beck T. (2011) “Bank Resolution: A Conceptual Framework” in *Financial Regulation at the CrossRoads: Implications for supervision, institutional design and trade*. Delimatsis P. and Herger, N. eds. Kluwer.
- Beck T. (2004) “The Incentive Compatible Design of Deposit Insurance and Bank Failure Resolution – Concepts and Country Studies” in *Who pays for Bank Insolvency?* Mays, D. and Luiksilia, A. eds Palgrave McMillan.
- Bjorner, A.; Lovasz, L.; Shor, P.W. (1991). “Chip-firing Games on Graphs”, *European Journal of Combinatorics*.
- Burgee, M. B. (1979). Purchase and assumption transactions under the federal deposit insurance act. *Forum*, 14(5), 1146-1160.
- Calomiris, C.W. (2010) “The political lessons of Depression-era banking reform” *Oxford Review of Economic Policy*, Volume 26, Number 3, 2010, pp. 540–560
- Calomiris, C.W. (2004) “The Political Economy of Bank Entry Restrictions: Theory and Evidence from the U.S. in the 1920s” Manuscript.
- Davidson and Ramirez (2014) Local Banking Panics of the 1920s: Identification and Determinants”, *Journal of Monetary Economics* 66, (2014): 164-177.

- Elliott, M, Golub, B. and M. O. Jackson. (2014). "Financial Networks and Contagion." *American Economic Review*, 104 (10): 3115–53.
- FDIC, History "The Banking Crises of the 1980s and Early 1990s: Summary and Implications"
<https://fdic/history>
- Financial Stability Board (2014) "Market Based Instruments for Systemic Bank Restructuring"
- Hoggarth G., Reidhill J., Sinclair P.. On the Resolution of Banking Crises: Theory and Evidence, 2004 Bank of England Working Paper No. 229.
- Holroyd, A.; Levine, L.; Mészáros, K.; Peres, Y.; Propp, J.; Wilson, B. (2008). "Chip-Firing and Rotor-Routing on Directed Graphs", *In and Out of Equilibrium 2*. 60. pp. 331–364.
- Macher, F. "The Austrian Banking Crisis Of 1931, One Bad Apple Spoils The Whole Bunch, LSE WP 274 February 2018
- Marcus (2011) Fear and Hope in Numbers Financial Data and the writing of history, EUI WP
- Nier E, J. Yang, T Yorulmazer, A. Alentorn (2007) " Network models and financial stability" *Journal of Economic Dynamics and Control*, Volume 31, Issue 6, Pages 2033-2060, ISSN 0165-1889, <https://doi.org/10.1016/j.jedc.2007.01.014>.
- Nouy D. (2017) Too Much of a Good Thing? The Need for Consolidation in the European Banking Sector: Tech. Rep. Speech VIII Financial Forum, Madrid, 27 September (2017).
- A. Pais, P.A. Stork (2011) Bank size and systemic risk *Eur. Financ. Manage.* (2011),
- Perotti, E and J. Suarez (2002) "Last bank standing: what do I gain if you fail? *Eur. Econ. Rev.*, 46, pp. 1599-1622
- Schaeck, K., Cihak, M. and S. Wolfe (2009) Are competitive banking systems more stable? *J. Money Credit Bank.*, 41 (2009), pp. 711-734
- Spencer, J. (1986). "Balancing Vectors in the max norm", *Combinatorica*, 6, pp 55-66.
- Verdier, D. (1997) "Universal Banking and Bank Failures between the Wars" EUI WP97/11
- Weber, Fritz (1995) "From Imperial to Regional Banking: The Austrian Banking System, 1918–1938 in *Banking Currency and Finance in Europe Between the Wars* Charles Feinstein H editor; Pages 337–357
- White, Phoebe and Yorulmazer, Tanju, (2014) "Bank Resolution Concepts, Tradeoffs, and Changes in Practices" (March 31, 2014). *Economic Policy Review*.