
Growth and collapse: An agent-based banking model of endogenous leverage cycles and financial contagion

Robert De Caux^{1*}, Frank McGroarty² and Markus Brede¹

Electronics and Computer Science¹ and Southampton Business School²,
University of Southampton, Southampton SO17 1BJ

E-mail: robertdecaux@hotmail.com

E-mail: F.J.McGroarty@soton.ac.uk

E-mail: Markus.Brede@soton.ac.uk

*Corresponding author

Abstract: We create an agent-based banking model that allows the simulation of leverage cycles and financial contagion. Banks within our model adapt their investment strategies in an evolutionary manner according to the success of their competitors, creating an endogenous interbank loan network and a dynamic asset market as they try to maximise profit by adjusting their leverage. The system exhibits periods of slow risk growth and fast insolvency cascades, allowing us to assess both the size and frequency of those cascades over a long time-frame.

We demonstrate that banks endogenise systemic risk into their leverage behaviour when the asset market is subject to either low or high levels of volatility, but are less successful for medium volatility when the number of bank insolvencies is maximised. We also show that in a low volatility environment, banks are more susceptible to systemic contagion. While the majority of insolvencies occur through the asset side of the balance sheet due to fire sales causing a rapid depreciation in the asset price, failures through the liability side of the balance sheet tend to be correlated, acting as an amplification mechanism to create far more serious cascades.

By creating realistic cycles of growth and collapse, the model provides a suitable framework for performing further policy tests.

Keywords: financial contagion; agent-based model; bank insolvency; evolutionary game theory; simulation; networks.

1 Introduction

The financial crisis of 2008 revealed that the process of maintaining stability in the banking system through microprudential regulation of individual banks was insufficient and a fallacy of composition (Brunnermeier et al., 2009). A failure of the models and regulators to appreciate the effect of a high level of system interconnectedness led to a dramatic underestimation of the total *systemic risk* (Gai, 2013).

Since the crisis, a wave of literature has concentrated on gaining a better understanding of systemic risk. Using techniques from complex network theory, there has been a strong focus on how interconnectedness and risk sharing can amplify contagion (Gai and Kapadia, 2010; Nier et al., 2008; Ladley, 2013) and which network topologies are most effective for stability (Glasserman and Young, 2015), with more specific questions addressing the difference between targeted and random shocks to the network (Roukny et al., 2013), the effect of leverage on systemic risk (Kuzubas et al., 2014; Greenwood et al., 2012) and which banks within a network are most prone to failure (Battiston et al., 2012). One particularly interesting aspect of investigation has been the interplay between different contagion channels; while losses through interbank market contagion alone are likely to be small (Glasserman and Young, 2014), a combination of liability and asset-side contagion channels working together gives a much higher probability of system-wide insolvency cascades (Caccioli et al., 2015).

However, all these papers only perform analysis on static networks, concentrating purely on the ex-post management of a contagion event, rather than an ex-ante understanding of how such networks evolve in the first place. Recent papers on dynamic network formation (Bluhm et al., 2013; Keiserman, 2014; Aldasoro et al., 2015) do demonstrate how such networks form endogenously through heterogeneous bank incentives, but they still assume that those incentives are fixed, rather than capturing how bank strategies evolve.

In fact, as was seen in the years preceding the recent financial crisis, banks do adjust their strategies over time. A period of prolonged stability created a feedback loop of increased borrowing and increased asset prices (Adrian and Shin, 2010), with banks exhibiting herding behaviour (Cipriani and Guarino, 2008) and failing to endogenise the effect of their heightened leverage on systemic risk (Daníelsson and Zigrand, 2008). Ultimately, the debt burden caused the system to suffer a “Minsky Moment”, with rapid bank deleveraging causing asset prices to tumble (Aymanns and Farmer, 2015).

In order to address this gap in the literature, we create an agent-based model of the banking system that allows the micro-level simulation of both growth cycles and systemic contagion. This builds on the approach of De Caux *et al.* (2017), but whereas their abstracted banks are represented entirely by a single “risk” number, we explicitly model the bank balance sheets, allowing them to adapt their strategies so that leverage and borrowing requirements can grow or shrink over time. Our methodology also allows contagion to occur through both asset and liability channels. An agent-based approach reflects the idea that the banking network behaves as a complex,

adaptive system (Haldane, 2009), with bank interactions creating a complex dynamic of peaks and crashes.

Rather than utilising myopic optimisation to maximise profitability, our banks are social learners, adapting their strategies according to the success of others (Kirman, 2006). Social learning is a particularly good paradigm to use, as it captures the evolutionary pressure that banks are under to increase their leverage (Thurner et al., 2012) and allows the development of herding in bank behaviour (Lux, 1995; Shive, 2010).

Given that an optimal system of financial agents accounts for the trade-off between risk sharing and system-wide collapse (Leitner, 2005), running our simulation over a long time period allows the investigation of the key system parameters that affect systemic risk and their effect on other factors such as bank profitability.

2 Methodology

Our approach develops the multi-agent models of Iori et al.(2006), Georg (2013) and Ladley (2013) by allowing each of a set of N banks to split their portfolio between three different markets in order to maximise a profit function. The three investment options are cash, an interbank market for loans and a risky asset, allowing banks to take leverage to increase their profitability. However our work differs from the aforementioned models in two key ways. Unlike the risky asset in the Georg and Ladley papers (Iori et al. use a risk-free asset), ours is explicitly modelled on a stock market, meaning that its price is determined endogenously and we can investigate the interaction between asset and liability channels for insolvency. In addition, in our model banks continually adapt their strategies using social learning to maximise their profit-based “fitness”, rather than through internal optimisation (Iori et al., Georg) or by using a genetic algorithm to find a steady state for the system (Ladley).

We distinguish between processes at two separate timescales. Behaviour is assessed at an *interday* level, where banks change their strategies depending on their own profitability compared to other banks in the system. Profitability is calculated by simulating market dynamics at the *intraday* level, where banks attempt to proactively optimise their portfolio to meet their strategic requirements.

2.1 Bank setup

Each bank is characterised by its balance sheet and investment strategy, and bounded in its behaviour by several regulatory restrictions. The balance sheet consists of the elements shown in Table 1.

Assets	Liabilities
External assets (a)	External deposits (d)
Interbank deposits ($\sum_{k=1}^N r_k$)	Interbank loans ($\sum_{k=1}^N p_k$)
Cash (c)	Equity (e)

Table 1 **Bank balance sheet.** Breakdown of bank assets and liabilities, with equity treated as a liability for presentational purposes, r_k representing an interbank deposit made with bank k and p_k similarly representing a loan from bank k .

Both cash (c) and external deposits (d) receive the risk-free interest rate, which we set to zero. The external asset (a), which represents a composite market of bonds and stocks, pays out a regular distribution (which we represent as a dividend) that is on average greater than the interest rate receivable on interbank deposits, but is susceptible to stochasticity. Interbank deposits ($\sum_{k=1}^N r_k$) and loans ($\sum_{k=1}^N p_k$) are not defined ex-ante, but evolve in line with the banks' need for cash to implement their respective investment strategies. Therefore a network is formed, with the banks as nodes and interbank relationships as the weighted, directed links between them. Our interbank market represents a conflation of all short-term funding markets, meaning that banks use it to not only manage liquidity fluctuations, but also as a structural source of funds for leveraging and potentially as an investment opportunity depending on returns. All interbank loans and deposits mature after a single day.

The strategy of a bank is defined by the tuple $\{A, C\}$. A represents the multiple of its equity that is to be invested in the external asset, so effectively represents leverage¹. C represents the multiple of e to be held as cash, so effectively represents a liquidity buffer. A bank's strategy also completely determines its funding requirements. If we define D to be $\frac{d}{e}$, then:

$$(A + C - 1 - D)e = L \quad (1)$$

1. Banks do not borrow funds to lend on the interbank market in our model.

where L is the net borrowing or lending requirement for the bank to implement strategy $\{A, C\}$. A positive value for L means that it needs to borrow on the interbank market, while a negative value means that it can be a lender.

The fitness of a bank is determined by its return on equity over a defined time horizon, as described in Appendix A.1. A parameter η controls the relative profit weighting that banks give to a regular income stream in the form of dividends and interest, versus the volatile returns from changes in the asset price.

There are several regulatory criteria that must be satisfied, with these taking priority over the banks' strategic considerations. Following Basel rules (Basel Committee on Banking Supervision, 2015), banks must maintain a Capital Adequacy Ratio above 8%, which necessitates a natural leverage cap at $A = 12.5$ (see Appendix A.2). Unlike other models (Montagna and Kok, 2013), breaching the Basel requirement will not instantly cause a bank to fail, but it will attempt to bring itself back into line with regulation as soon as possible through the intraday portfolio allocation process described below. We also enforce a no shorting restriction on the external asset, and stipulate that a loan to any one counterparty must be for less than 25% of the lender's equity, in line with the Basel framework.

2.2 *Intraday*

During the course of each day, banks follow the stages set out in Figure 1. All activities generate a series of debits and credits on the account of each bank rather than immediate cash settlement, as this makes it more straightforward to resolve creditor payments if a bank defaults.

All loans mature at the start of the day creating payables and receivables, while banks receive a daily dividend from the external asset, which is subject to stochasticity but is on average greater than the return available on overnight loans (see Appendix B.1). Banks also declare and pay a dividend equal to any profit that they make above a multiple μ of their own equity, in order to constrain their equity base.

As with other models in the literature (Georg, 2013; Iori et al., 2006), each bank is susceptible to depositor shocks, which affect its liquidity in the short term and potentially cause more significant structural changes to its funding mix in the long term (see Appendix B.1).

Banks periodically adjust their portfolio allocation to meet their target strategy for the external asset A , using a process described in Appendix B.2 As is conventional in the literature (Hommes and Wagener, 2008), banks are assumed to fix any asset purchases or sales in advance of the market

clearing. In order to reflect the fact that banks would not want (or may not be able during times of stress) to adjust their investments on a daily basis, the probability of portfolio updating for any given day is controlled by the parameter χ .

Unlike models where the number of participants is closed (Hommes and Wagener, 2008; LeBaron, 2012), we treat the banks as part of a wider market for the asset. We therefore introduce a noise trader to provide a level of stochasticity in market making and to control the price power of the banks. As with other similar models (Aymanns and Farmer, 2015), the noise trader represents the population of other market participants, such as hedge funds, whose behaviour is not explicitly modelled. In aggregate, they behave as a weak trend follower, with additional daily volatility provided by less sophisticated investors. The clearing price for the external asset is determined by supply and demand across both bank and noise trader orders (see Appendix B.2).

The insolvency stage occurs before the interbank market and payment settlement, as neither of the latter cause bank equity to change. Insolvency is handled differently in our model to the standard method in the literature. Rather than a bank only becoming insolvent when it has a negative equity, we follow the idea of Elliott (2014) that if a bank's equity hits a failure threshold $\theta > 0$, it discontinuously loses further value immediately and becomes insolvent. This is intended to reflect the fact that a bank's insolvency can be inevitable even while it is technically solvent, for reasons such as a crisis of confidence in the market and subsequent short selling, or through increased margin calls on bank positions.

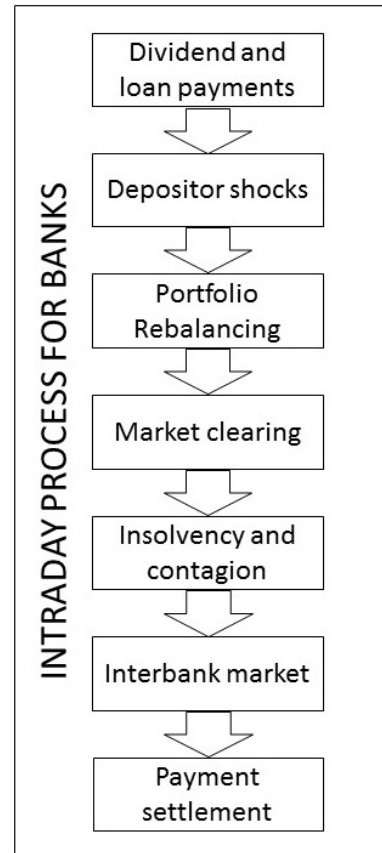


Figure 1. The intraday process. Flow diagram showing the steps that each bank takes during the course of a single time-step.

Once a bank is declared insolvent it enters the liquidation process described in Appendix B.3, whereby it sells all assets, pays off creditors and is removed from the simulation. This subjects other banks to two potential forms of contagion – a devaluation of their asset portfolio and a devaluation of their loan receivables – and so all banks are reassessed to see if their equity has breached the failure threshold. Further rounds of liquidation continue until no new banks become insolvent.

After all insolvencies have been resolved, loan orders are placed in the interbank market. Banks assess their current cash position and calculate their loan order using the process specified in Appendix B.4. Interest rates comprise a base rate driven by market demand, as well as a credit spread based on the riskiness of the borrower. In order to ensure that banks cannot become cash-flow insolvent, we introduce a central bank as a lender of last resort, which offers loans that must be 100% collateralised by the external asset².

To create enough supply for the interbank market, we create a proportion N_c of “core” banks that are structurally long liquidity by giving them more deposits. As these banks do not have additional equity, they will always have a large supply of funds that they cannot invest in assets without breaching their Capital Adequacy Ratio requirement. The remaining banks in the simulation form the “periphery”, replicating observed interbank market topology (Craig and Von Peter, 2014).

Finally, a payment vector determined by the Eisenberg and Noe algorithm (2001) settles the accounts of all banks remaining in the simulation. Such a payment vector will always exist, as there cannot be any newly insolvent banks at this stage of the day.

2.3 Interday

Between each day, banks are given the opportunity to adjust their strategy $\{A, C\}$ according to an evolutionary learning rule, which is dependent not only on their fitness f , but also the fitness of the other banks in the system. This is based on the methodologies of Kirman (1991) and Westerhoff (2009), where banks have a certain probability of switching to whichever strategy has been the more successful over a certain period of time. The process of fitness adjustment is described in Appendix A.2

In order for the system to run for an indeterminate period of time, a new bank is introduced interday for every bank that is liquidated. Each new bank has the same starting equity as the initial banks, with its depositor value

2. For simplicity, the central bank does not take a haircut against the collateral.

determined by whether they are replacing a core or a periphery bank (so that the ratio of the two is maintained). New banks enter with a “safe” leverage ($A \leq 2$) and must purchase any external assets required for them to fulfil their strategy³.

3 Results and discussion

In order to study the dynamics of the model, we choose appropriate parameters to allow long term leverage cycles of slow risk growth followed by “Minsky Moments”, where the asset rapidly collapses in value and potentially causes insolvency cascades.

3.1 Simulation setup

We simulate 500 banks to create a diverse pool of heterogeneous strategies and to allow for the possibility of significant insolvency cascades. All banks begin with the same amount of equity, but differing deposits depending on whether they are core or periphery. Banks are heterogeneous in terms of their strategies, which are assigned randomly at the start of the simulation with A (or leverage) varying between 0 and 12.5.

In terms of parameter calibration (see Table 2), the deposits for periphery banks are slightly more than three times equity, whereas for core banks they are around twenty five times equity to provide plenty of liquidity to the interbank market. There is significant variability in these ratios across different models (Bluhm et al., 2013; Keiserman, 2014), but our ratios are not atypical.

The expected dividend return and volatility are slightly lower than in other papers (Keiserman, 2014; Aldasoro et al., 2015) as we want consistent returns from the asset that are susceptible to being beaten by those of high-risk loans. Banks will prioritise a steady income from dividends over asset speculation with $\eta = 0.7$. The interbank base rate and loan spread parameters are calibrated based on empirical testing, in order for bank funding costs to become unsustainably high with extreme leverage. Market depth is calibrated to allow large scale fire sales to have a significant impact on asset price, but not to the point where all banks will automatically become insolvent.

The strategy evolution parameters, based on the process in Appendix A.2, are calibrated to allow bank strategies to evolve slowly, creating long periods

3. Because A is low, even multiple new banks buying assets will not significantly move the market price.

Category	Notation	Description	Value
Bank	d_p	Depositors (periphery)	100
	d_c	Depositors (core)	800
	e	Equity	30
	μ	Dividend Multiple	2
	ζ	Depositor inertia	0.1
	κ	Depositor shock noise	0.1
	χ	Portfolio update frequency	0.1
	η	Dividend:Capital gain ratio	0.7
	θ	Insolvency threshold	10
Simulation	N	Number of banks	500
	N_c	Proportion of core banks	0.1
	σ	Noise trader volatility	3,000
	ρ	Noise trader reversion	0.01
Asset and loan returns	ξ	Expected dividend return on external asset	0.04
	ω	Dividend volatility on external asset	0.005
	B	Interbank base rate parameter 1	0.01
	Υ	Interbank base rate parameter 2	10^{-4}
	Z	Interbank base rate cap	0.1
	ι	Interbank loan spread	0.01
	λ	Market depth	10^{-6}
Strategy evolution	τ	Lookback	100
	α	Fitness comparison	0.01
	β	Mutation probability	0.01
	ϕ	Mutation size	0.5
	ψ	Inertia	0.1

Table 2 Parameters for simulation

of steady growth. A look-back period of 100 time-steps allows for the smoothing of bank returns over periods of volatility.

The key parameter that we wish to investigate is the noise trader volatility σ , which provides an exogenous shock to the asset price that can drive bank behaviour in terms of portfolio adjustments. A negative shock will cause equity losses to the banks, which depending on its size can cause rapid deleveraging and further significant drops in asset price.

3.2 Risk and asset price evolution

Initially setting $\sigma = 3,000$, we run the simulation for 800,000 time-steps with this relatively low volatility and observe the long term evolution of bank risk and asset price.

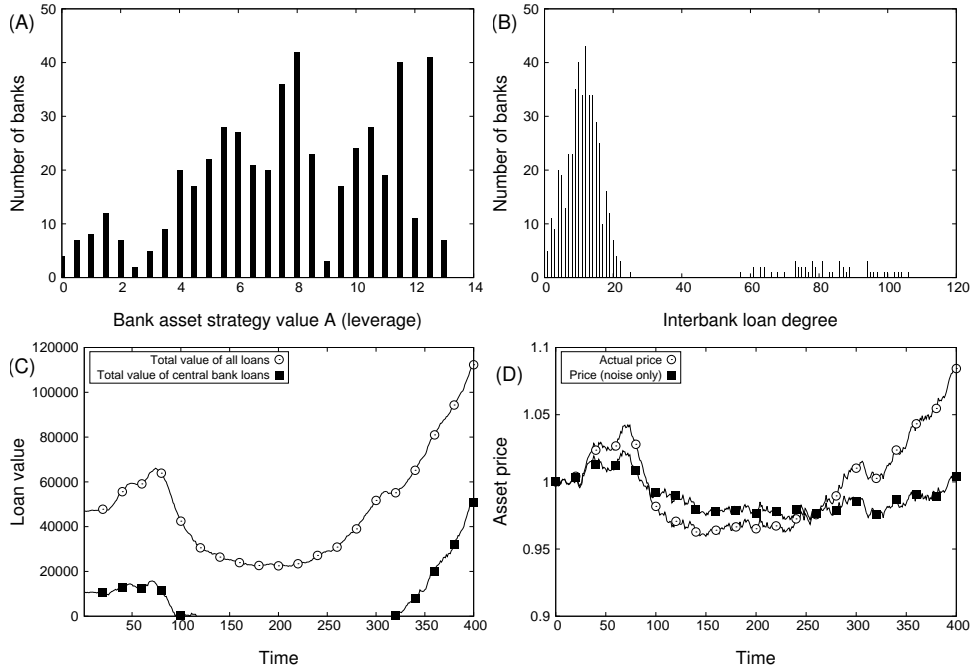


Figure 2. Key simulation statistics. Snapshot at $t = 400$ of a simulation run using the default parameters. (A) shows the range of bank strategies (leverages), while (B) shows the degree distribution for banks in the interbank market. (C) and (D) show the evolution of the interbank market and asset price over the 400 time-steps. (C) shows the split between bank-to-bank lending and central bank borrowing, while (D) shows both the actual asset price and the price as it would be driven purely by volatility from the noise trader.

In order to understand how the system is evolving, we take a snapshot of the system at a point in time, $t = 400$, which can be seen in Figure 2. Banks are employing a range of asset strategies (panel A), with a bias towards high risk-taking although there are still plenty of lower leverage strategies. The topology of the interbank network is core-periphery (panel B), with the average out degree for the core banks being approximately 81 and the average in degree for the periphery banks being approximately 11. Banks are necessarily diversified in their lending due to the counterparty limit of 25% of equity for lenders.

There is a symbiotic relationship between asset price and bank leverage. Banks quickly respond to volatility in the asset price by adjusting their asset

holdings up or down to match their target leverage (see Appendix B.2), but also slowly adjust their leverage *strategy* by copying the most profitable banks in the system, in turn driving the asset price through supply and demand. This can be seen in panel D, where the overall price broadly tracks the noise for 250 time-steps, before becoming driven by an increased demand for assets. Here, banks are copying those competitors with the highest fitness, which due to the rising price of assets and the overall prioritisation of a steady income ($\eta = 0.7$) are those who are taking the highest amount of risk through the highest leverage. The size of the interbank market (panel C) is driven by the systemic demand for assets, with central bank borrowing necessary once the overall loan demand outstrips supply from the core banks.




Figure 3. Evolution of the banking system. The left hand panel shows the average leverage taken by the banks when $\sigma = 2,000$, demonstrating a series of leverage cycles with sudden crashes of varying size. The period between $t = 200,000$ and $300,000$, marked with a box on the left hand panel, is magnified in the right hand panel.

Over the long time period, slow strategy adjustment causes the system to go through periods of leverage and deleverage, or *leverage cycles*, which fluctuate around an intermediate level of risk. Some of the deleveraging is due to self-regulation, whereby banks who increase their leverage too far become unprofitable and copy a less risky strategy. However, the system is also susceptible to sudden crashes, where average leverage drops rapidly due to many simultaneous insolvencies. As can be seen in Figure 3, the size of these insolvency cascades is predominantly small, but system-wide cascades do infrequently occur.

3.3 Insolvencies and contagion

The system is now run across a range of values for the noise trader volatility σ , in order to investigate insolvencies and contagion in more detail. Each insolvency within the simulation can either be direct (during the normal

course of the intraday process) or caused by contagion (either through the fire-sale of assets by other insolvent banks or the non-repayment of interbank loans). Contagion is a difficult concept to define, as wide-scale sales of the asset during regular portfolio adjustment could be viewed as “pre insolvency contagion” that may trigger insolvencies in future time-steps. However, contagion-based insolvencies are defined here to only be those triggered by the actions of other *already* insolvent banks, i.e. those that occur after the first round of insolvencies in a particular time-step. Leverage is simply taken to be the value of A for each bank averaged across both time and the population.




Figure 4. Bank insolvencies and leverage. Average total and contagion-based insolvencies per time-step (left) and average leverage (right). Results for a system of 500 banks run for a period of 200,000 time-steps with varying noise trader volatility. Each σ value is averaged over 50 runs, with error bars representing the value range. A contagion-based insolvency is directly caused by other banks becoming insolvent during the same time-step.

Figure 4 (left panel) shows how the total number of insolvencies grows with noise trader volatility up until around $\sigma = 6,000$, before beginning to drop off. In conjunction with the fact that leverage is monotonically decreasing with volatility (right panel), banks within the system are successfully adjusting their overall leverage to reflect the risk of insolvency for low and high values of σ , but for medium σ they are taking on too much leverage relative to the overall risk of insolvency. Interestingly however, the danger of systemic contagion changes as σ increases, accounting for the majority of insolvencies during low volatility but being reduced to a minor factor when volatility is high and the majority of insolvencies are caused by direct shocks to the asset price instead.

The right panel of Figure 5 shows that the number of cascades peaks around $\sigma = 4,500$, while the average size of cascades is largest when volatility is low (left panel). This is due to banks accumulating more risk through higher leverage when σ is small, leaving them more susceptible to contagion in the rare event of insolvencies.

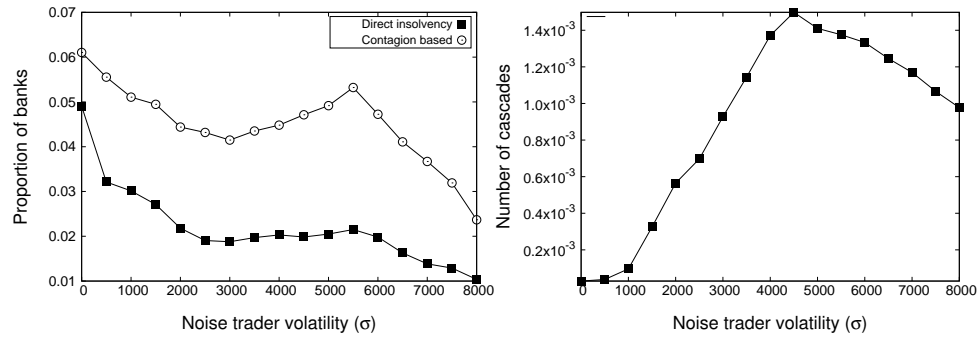


Figure 5. Average cascade size and frequency. For time-steps with at least one insolvency, the left panel shows the proportion of banks that become insolvent directly and the proportion that become insolvent through contagion for different values of σ . The right panel shows the probability of a time-step featuring an insolvency cascade, again for varying σ .




Figure 6. Cascade size distribution. Log-Log plot showing the probability distribution of cascade size for simulations run with $\sigma = 1, 000, 3, 000$ and $7, 000$. Results for each σ have been averaged over 50 runs.

Figure 6 focuses on the distribution of cascade sizes. For the three values of σ shown, cascades occur across all possible length scales. However, although they resemble a power law, the distributions are bimodal with peaks for very small and very large (system-wide) cascades. The low volatility environment has a higher probability of cascades with size 100 or greater, in line with Figure 5.

3.4 Contagion channels




Figure 7. Drivers of contagion and effect of interbank market. The left panel shows the number of contagion-based insolvencies that are attributable to asset side contagion (through liquidation fire sales) and liability side contagion (through non-repayment of interbank loans) for a simulation using the standard parameters. The right panel shows contagion-based insolvencies per time-step for simulations both including an interbank market and with only central bank borrowing.

The left panel of Figure 7 shows that the majority of contagion occurs through the asset side of bank balance sheets by fire sales of the external asset as part of the liquidation process. This is intuitive, as the implicit “asset network” is fully connected, whereas the interbank network is relatively sparse between the core and periphery banks. The results are in line with previous papers that show little contagion occurs through the interbank network alone (Glasserman and Young, 2014; Elsinger et al., 2006). However, as can be seen in the right panel, removing the interbank market completely and only relying on central bank borrowing reduces the total number of contagion-based insolvencies by more than just the original liability-side insolvencies alone⁴. Therefore, it is clear that interdependency between the contagion channels is a major contributor to system-wide cascades, as found by Caccioli (2015).

4. The base rate for loans is adjusted in the central bank borrowing case to account for the loan supply in the system dropping dramatically with no corresponding decrease in demand.

4 Conclusions

We have designed an agent-based model that allows the long term analysis of an adaptive banking system in terms of the trade off between two competing factors; enhanced profitability through increased leverage and systemic risk leading to insolvency cascades. Channels for contagion are created endogenously by the micro-simulation of two key markets – a traded asset and interbank lending. Banks are able to adjust their strategies in an evolutionary manner by copying the strategies of more profitable competitors, thus creating cycles of leverage growth and collapse.

We investigated the importance of the noise trader volatility and found that banks adjust their strategy effectively for high and low volatility markets, but fail to endogenise systemic risk effectively when volatility is less extreme, leading to the highest level of bank insolvencies. However, the system is at greatest risk of large insolvency cascades when the volatility is low, and the failure of a single bank can cause system-wide contagion.

Asset-side contagion (through fire sales as part of the liquidation process) is a much bigger direct factor in insolvency cascades than liability-side contagion, although there is a clear interdependency between the contagion channels, as removing the interbank market dramatically reduces the overall number of insolvencies by more than the direct liability-side contribution alone.

We intend to extend our work by testing the effectiveness of various regulatory policies on the system, including an analysis of how government intervention (via bail-outs or bail-ins) can be used in an optimal manner on insolvent banks to negotiate the trade-off between profitability and systemic risk.

Acknowledgements

This work was supported by an EPSRC Doctoral Training Centre grant [EP/G03690X/1]. No new data was generated in this research.

References

- Adrian, T. and H.S. Shin (2010). “Liquidity and leverage”. *Journal of Financial Intermediation*, 19(3), 418–437.
- Aldasoro, I., D. Delli Gatti, and E. Faia (2015). “Bank Networks: Contagion, Systemic Risk and Prudential Policy”. *SAFE Working Paper*, 87, 1–43.

- Aymanns, C. and J.D. Farmer (2015). “The dynamics of the leverage cycle”. *Journal of Economic Dynamics and Control*, 50, 155–179.
- Basel Committee on Banking Supervision (2015). “U.S. Basel III Final Rule: Standardized Risk Weights Tool”. URL: <http://www.usbasel3.com/tool/>.
URL <http://www.usbasel3.com/tool/>.
- Battiston, S., et al. (2012). “DebtRank: Too Central to Fail? Financial Networks, the FED and Systemic Risk”. *Scientific Reports*, 2, 1–6.
- Bluhm, M., E. Faia, and J.P. Krahen (2013). “Endogenous banks’ networks, cascades and systemic risk”. *SAFE Working Paper Series*, 12, 1–79.
- Brunnermeier, M.K., et al. (2009). “The Fundamental Principles of Financial Regulation”. *Geneva Report on the World Economy*, 11.
- Caccioli, F., et al. (2015). “Overlapping portfolios, contagion, and financial stability”. *Journal of Economic Dynamics and Control*, 51, 50–63.
- Cipriani, M. and A. Guarino (2008). “Herd Behavior and Contagion in Financial Markets”. *The B.E. Journal of Theoretical Economics*, 8(1), 1–54.
- Craig, B. and G. Von Peter (2014). “Interbank tiering and money center banks”. *Journal of Financial Intermediation*, 23(3), 322–347.
- Daniélsson, J. and J-P. Zigrand (2008). “Equilibrium asset pricing with systemic risk”. *Economic Theory*, 35(2), 293–319.
- De Caux, R., F. McGroarty, and M. Brede (2017). “The evolution of risk and bailout strategy in banking systems”. *Physica A: Statistical Mechanics and its Applications*, 468, 109–118.
- Eisenberg, L. and T. Noe (2001). “Systemic risk in financial systems”. *Management Science*, 47, 236–249.
- Elliott, M., B. Golub, and M. Jackson (2014). “Financial networks and contagion”. *American Economic Review*, 104(10), 3115–3153.
- Elsinger, H., A. Lehar, and M. Summer (2006). “Risk Assessment for Banking Systems”. *Management Science*, 52(9), 1301–1314.
- Finger, Karl and Thomas Lux (2017). “Network formation in the interbank money market : An application of the actor-oriented model à-ž”. *Social Networks*, 48, 237–249.

- Gai, P. (2013). *Systemic Risk - The Dynamics of Modern Financial Systems*. Oxford University Press.
- Gai, P. and S. Kapadia (2010). “Contagion in financial networks”. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 383.
- Georg, C-P. (2013). “The effect of the interbank network structure on contagion and common shocks”. *Journal of Banking and Finance*, 37(7), 2216–2228.
- Glasserman, P. and H.P. Young (2014). “How likely is contagion in financial networks?” *Journal of Banking and Finance*, 50(C), 383–399.
- Glasserman, P. and H.P. Young (2015). “Contagion in Financial Networks”. *University of Oxford Department of Economics Discussion Paper*, 764, 1–76.
- Greenwood, R., A. Landier, and D. Thesmar (2012). “Vulnerable banks”. *NBER Working Paper*, 18537, 1–46.
- Haldane, A.G. (2009). “Rethinking the financial network”. *Bank of England speech, delivered at Financial Student Association, Amsterdam, Netherlands*, April, 1–41.
- Hommes, C.H. and F. Wagener (2008). “Complex Evolutionary Systems in Behavioral Finance”. *Tinbergen Institute Discussion Paper*, 54(1).
- Iori, G., S. Jafarey, and F. Padilla (2006). “Systemic risk on the interbank market”. *Journal of Economic Behavior & Organization*, 61(4), 525–542.
- Keiserman, B. (2014). *Contagion and Regulation in Endogenous Bank Networks*. Ph.D. thesis, EESP-FGV.
- Kirman, A. (1991). “Epidemics of Opinion and Speculative Bubbles in Financial Markets”. In “Money and Financial Markets”, edited by M Taylor. Oxford.
- Kirman, A. (2006). “Heterogeneity in Economics”. *Journal of Economic Interaction and Coordination*, 1(1), 89–117.
- Kuzubas, T., B. Saltoglu, and C. Sever (2014). “Systemic Risk and Heterogeneous Leverage in Banking Network: Implications for Banking Regulation”. *Bogazici University Department of Economics Working Paper*, 1, 1–28.

- Ladley, Daniel (2013). “Contagion and risk-sharing on the inter-bank market”. *Journal of Economic Dynamics and Control*, 37(7), 1384–1400.
- LeBaron, B. (2012). “Heterogeneous gain learning and the dynamics of asset prices”. *Journal of Economic Behavior & Organization*, 83(3), 424–445.
- Leitner, Y. (2005). “Financial networks: Contagion, commitment, and private sector bailouts”. *The Journal of Finance*, 60, 2925–2953.
- Lux, T. (1995). “Herd Behaviour, Bubbles and Crashes”. *The Economic Journal*, 105(431), 881–896.
- Meteor Asset Management (2016). “Counterparty Credit Default Swap Rates”. URL: https://www.meteoram.com/uploads/files/040116_1.pdf.
URL https://www.meteoram.com/uploads/files/040116_{_}1.pdf.
- Montagna, M. and C. Kok (2013). “Multi-layered Interbank Model for Assessing Systemic Risk”. *Kiel Working Paper*, 1873.
- Nier, E., et al. (2008). “Network models and financial stability”. *Bank of England Working Paper*, 346.
- Ötoker-Robe, I. and J. Podpiera (2012). “Explaining Credit Default Swaps Pricing for Large Banks”. *The Capco Institute Journal of Financial Transformation*, 34, 62–75.
- Roukny, T., et al. (2013). “Default cascades in complex networks: topology and systemic risk.” *Scientific reports*, 3, 2759.
- Shive, S. (2010). “An Epidemic Model of Investor Behavior.” *Journal of Financial & Quantitative Analysis*, 45(1), 169–198.
- Turner, S., J.D. Farmer, and J. Geanakoplos (2012). “Leverage Causes Fat Tails and Clustered Volatility”. *Quantitative Finance*, 12(5), 19.
- Westerhoff, F. (2009). “A simple agent-based financial market model: direct interactions and comparisons of trading profits”. *BERG Working Paper Series*, 61.

A Bank fitness and strategy

A.1 Fitness

The fitness of a bank at time t is determined by its return on equity over a defined time horizon $[t - \tau, t]$. The profit \mathbb{P} for any time-step comprises two distinct elements. Firstly, *capital gains or losses* from the external asset and interbank loan write-downs (\mathbb{P}^{cap}). Secondly, *dividends and interest* receivable or payable from the external asset and interbank loans (\mathbb{P}^{int}). A function, parametrised by η , controls whether banks focus on a regular income stream or volatile returns from changes in the asset price, giving a fitness function for bank i at time t of:

$$f_{i,t} = \frac{\sum_{j=t-\tau}^t (1 - \eta)\mathbb{P}_{i,j}^{cap} + \eta\mathbb{P}_{i,j}^{int}}{e_{i,t-\tau}} \quad (2)$$

The look-back period τ is a model parameter.

A measure of the health of a bank at any point in time is its Capital Adequacy Ratio \mathcal{C} , which is defined using the formula:

$$\mathcal{C} = \frac{e}{\mathbf{w}^T \mathbf{a}} \quad (3)$$

where \mathbf{w} is a set of Basel III defined risk weights for each asset class and $\mathbf{a}^T = (\sum_k r_k, a, c)$, a vector giving the amount of each asset class held by the bank. The risk weight for each asset class is calibrated to reflect the riskiness of that class (Basel Committee on Banking Supervision, 2015). Within our model, interbank deposits carry a risk weight of 20%, cash carries a risk weight of 0% and the external asset (which we treat as a debt instrument) carries a risk weight of 100%, giving $\mathbf{w}^T = (0.2, 1, 0)$. Therefore we can rewrite \mathcal{C} as:

$$\mathcal{C} = \frac{e}{0.2 \sum_k r_k + a} \quad (4)$$

Given that $a = Ae$, setting $\mathcal{C} = 0.08$ and $\sum_k r_k = 0$ gives a maximum value of $A = 12.5$.

A.2 Strategy updating

During the interday period, each bank b_1 has a chance to update its strategy $\{A, C\}$ through the following process:

- With a certain probability α , bank b_1 with strategy $\{A_1, C_1\}$ will compare its fitness f_1 to a randomly chosen bank b_2 with strategy $\{A_2, C_2\}$ and fitness f_2 .
- If $f_2 > f_1$, b_1 adjusts its strategy to be $\{A_1 + \psi(A_2 - A_1), C_1 + \psi(C_2 - C_1)\}$, where ψ is an inertia factor to prevent large strategy jumps.

The strategy of bank b_1 can also be subject to mutation:

- With a certain probability β , the strategy of bank b_1 is mutated to either $\{A_1 + \epsilon, C_1\}$ or $\{A_1, C_1 + \epsilon\}$, where $\epsilon \sim U(-\phi, \phi)$.
- Any mutations are adjusted to account for $A_1 = 0$, $A_1 = 12.5$ and $C_1 = 0$ being reflecting barriers.

All of $\alpha, \beta, \phi, \psi$ are exogenous parameters to the model.

B Market events

B.1 Dividends and depositors

The dividend rate, with expected value ξ and variance ω , is on average greater than the return available on overnight loans. The rate is applied to the asset units held by a bank, meaning that asset return is independent of asset price.

Depositors broadly track the equity of a bank, with the process characterised by the following equation:

$$d_t = d_{t-1} + \zeta \left[d_0 \frac{e_t}{e_0} - d_{t-1} \right] + \kappa d_{t-1} \epsilon \quad (5)$$

Noise is provided by $\epsilon \sim \mathcal{U}(-0.5, 0.5)$. d_0 and e_0 are initial deposits and equity respectively, while ζ acts as an inertia term for the time lag between deposits and equity. The reason for this approach is that in our model, equity is a key measure of a bank's health, so it is reasonable to assume that safer banks will attract a larger number of depositors, while distressed banks will lose depositors. Equation (5) will also cause a "mini bank run" if a bank suffers a dramatic loss of equity. From a practical point of view, it means that distressed banks with low equity are not able to rely on their relatively high depositor base to continue running a high leverage strategy, but will instead need to keep the same mix of depositor and interbank funding that they relied on at higher equity levels.

B.2 Portfolio adjustment

In order to adjust their portfolio, each bank places an order of u asset units into the stock market, with a positive u indicating an asset purchase and vice versa. If we assume that the most recent asset price from time $t - 1$ is p_{t-1} , then their order will be:

$$u_t = \frac{Ae_t - a_t}{p_{t-1}} \quad (6)$$

The noise trader places an order into the market concurrently with the banks, with the size of the order at time t determined by the following formula:

$$u_t^{noise} = \frac{-1}{\lambda} \left[\rho \ln\left(\frac{p_{t-1}}{\bar{p}}\right) \right] + \epsilon \quad (7)$$

The first term in the equation represents the aggregation of hedge funds or other market players acting as weak trend followers, creating an order that will move the asset price back towards \bar{p} , a moving price average for the asset over a defined lookback period T . The market depth parameter λ determines the sensitivity of asset price to supply and demand (see equation 9) and $\rho < 1$ determines how quickly the asset price will revert to the moving average (Aymanns and Farmer, 2015). The second term, $\epsilon \sim \mathcal{N}(0, \sigma)$, represents volatility caused by noise traders that is a key driver of long term bank strategy. While the value of this noise is in the form of stock units, plugging these units into the asset price formula below allows the effect to be considered in terms of annualised volatility for the asset. As an example, $\sigma = 2,000$ equates to an annualised volatility of around 3.2%, whereas $\sigma = 8,000$ equates to 12.8%.

In order to clear the market, we use a procedure similar to that of LeBaron (2012) and Hommes and Wagener (2008). At time t we find the net demand for the asset U_t by summing across the orders for all N banks and the noise trader:

$$U_t = \sum_i^N u_{i,t} + u_t^{noise} \quad (8)$$

The price of the asset moves according to the following process:

$$p_t = p_{t-1} e^{\lambda U_t} \quad (9)$$

All asset transactions settle at time t and at price p_t . The asset ledger for each bank is updated instantaneously, while debits and credits are created for later cash settlement.

B.3 Bank liquidation procedure

All unpaid dividends declared by the bank are cancelled and all collateralised loans default, with the central bank taking possession of the external asset units used as collateral. The bank then places an order to sell all remaining units of the external asset immediately. Asset sell orders from all insolvent banks will again be settled using the market clearing process described above. A 10% administration cost is levied on the bank's estate, before the remainder is distributed to the creditor hierarchy of depositors, lenders and equity holders in that order. Creditor banks adjust their expected receivables using the Eisenberg and Noe algorithm (2001), which works to correctly settle a chain of insolvencies including potentially circular relationships between banks. Finally, the insolvent bank is removed from the simulation.

B.4 Interbank market

The order L that a bank places in the lending market is determined by adjusting Equation (1) for the fact that by this stage of the day, a bank will not be able to alter its investment in the external asset. Therefore:

$$L = a + (C - 1 - D)e \quad (10)$$

If $L > 0$, the bank is long cash relative to its required position and is offering the excess as a loan. If $L < 0$, the bank requires a loan to make up that shortfall. As with market clearing for the external asset, we sum across all N banks to establish the balance between supply and demand in the market:

$$\mathcal{L} = \sum_{i=1}^N L_i \quad (11)$$

For simplicity, borrowers and lenders are selected randomly to fill each loan, so the model does not incorporate the possibility of relationship banking (Ladley, 2013; Finger and Lux, 2017), which would be an interesting extension. Banks are also subject to the Basel regulation that they cannot lend more than 25% of their equity to a single counterparty. The rate on the loan charged by the lender comprises the following components:

- A base rate driven by market demand of the form $Be^{-\Upsilon\mathcal{L}}$, restricted to the range $[0, Z]$.

- A credit spread γ based on the riskiness of the borrower.
- A profit margin of ι .

B , Υ , Z and ι are all external parameters. In order to calculate γ , we use the Capital Adequacy Ratio \mathcal{C} as the main determinant. In reality there are a number of factors that are used to derive the credit spread (Ötker-Robe and Podpiera, 2012), but \mathcal{C} is a useful proxy as it includes both the riskiness of a bank's investments and the amount of equity that it has as a loss buffer. γ is broadly analogous to the Credit Default Swap ("CDS") rate for the bank, which is the cost in the market of insurance against the bank defaulting. Based on data linking \mathcal{C} and CDS rates for large banks (Meteor Asset Management, 2016), we use the following mapping:

$$\gamma = \frac{1}{(120\mathcal{C})^2} \quad (12)$$

We assume that γ for each bank is known to everyone in the market, allowing them to correctly price a loan according to the perceived risk. We also make the assumption that banks do not wish to give funds to counterparties who are above a certain risk threshold, requiring banks to have $\mathcal{C} \geq 2.5\%$ in order to raise funds through the interbank market.