

Agent-based Modeling for Complex Financial Systems

James Paulin, Anisoara Calinescu, and Michael Wooldridge

Department of Computer Science, University of Oxford, UK
James.Paulin@cs.ox.ac.uk

The continuing progression of globalization, coupled with the continuing adoption of new computing and network technology, means that today's global financial marketplace is best understood as a complex network of interacting market systems, in which events of world-wide significance unfold on timescales that are barely within the ability of humans to comprehend. Classic economic theories (such as, most famously, General Equilibrium Theory) have proven to be of limited value for understanding and modeling such networks of markets. Many researchers have concluded that the dynamics of networked market systems are better understood as complex adaptive systems [1], in which independent software components interact without centralised control or oversight [2]. A key characteristic of complex adaptive systems is that, even if the individual *micro* behaviors of system components are all perfectly understood, it remains extremely difficult to predict the overall *macro* behaviors that the system might exhibit. One reason for this is that such systems exhibit *emergent properties*, in which a behavior is produced from the overall system which would not be produced by any individual components or subsets of components.

Global financial systems contain thousands of interacting, heterogeneous participants, which operate according to their own private values and incentives. One example of an emergent property in such financial systems is the price series of financial assets that arise from enormous numbers of individual orders to buy and sell placed by participants operating over a wide range of time-scales. Another prominent and extremely worrying type of emergent phenomenon is the *Flash Crash*. A Flash Crash is a sudden, catastrophic drop in asset value, which typically take place over a period of minutes. The first and most famous Flash Crash (and the event that gave its name to the phenomenon) took place over 36 minutes on 6 May 2010, during which time the Dow Jones Industrial Average suffered its largest ever intraday point loss: trillions of dollars were briefly wiped off the world's markets [3]. Since 2010, a number of other Flash Crashes have been reported (see sidebar 1).

Events such as the Flash Crash, and the earlier global financial crisis of 2007-2009, highlight the inadequacy of the established models offered by classical economics. Such models are based on strong assumptions of perfect rationality, perfect information, and market equilibrium, and take no account of the complex web of interconnections that characterise modern financial systems. At the height of the sub-prime mortgage crisis, traditional models proved spectacularly inaccurate in their assessment of the stability of the investment bank Lehman Brothers. The subsequent collapse of Lehman Brothers sent shockwaves through the global economic establishment. Lehman Brothers, apparently financially sound when considered in isolation, had proved vulnerable to distress arising elsewhere in the financial network. When credit dried up in the interbank lending network, Lehman Brothers found itself unable to honour its debt obligations and rapidly collapsed. The experience of the global financial crisis of 2007-2009 vividly illustrates the need for models that more accurately capture and quantify the real

risk of financial systems by treating them as complex adaptive systems, with all the richness and intricacy that implies.

Regulators and researchers are all too aware of the need for new models to further our understanding of the complex reality of financial systems. The stakes are high, as is the difficulty of providing effective monitoring and intervention. Not only is it necessary to provide regulators with tools to assess the health of financial systems in real-time, it is also necessary to provide tools that enable regulators to gain a sound understanding of the consequences of policy interventions and rule changes *in silico*, before they are imposed on markets. In this way, regulators can build additional confidence in their procedures and interventions, ahead of any need to deploy them.

Changes in market structure over recent decades have only added to the challenge facing regulators. The application of computational methods is of course not restricted to those studying financial markets from the outside. Driven by growth in stock exchange data availability (both real-time and historical datasets), recent decades have witnessed the rise to prominence of so-called “algorithmic trading” methods [4]. Here, the business of buying and selling financial securities is delegated to electronic systems of various kinds, united by their key advantage over human traders — speed. The antiquated image of the Wall Street “trading pit” where stock-brokers use hand signals to issue instructions to buy and sell securities on behalf of their clients has been replaced with that of racks of computers in a data centre [5]. The upper extreme of the speed distribution is the domain of *high-frequency traders* (HFTs). HFTs compete with each other on raw speed, utilising advanced communications technologies such as microwave or even laser connections in order to minimise latency between their trading systems and the systems of the stock exchange. This further highlights the need for new modeling paradigms that capture the interconnected ecosystem of interacting market participants of all size and time-scales.

A Simulated Stock Exchange

Classical economic models are predicated on perfect rationality, perfect information, and strong assumptions of market equilibrium. As discussed above, recent crises have demonstrated that these assumptions are simply too strong.

Agent based modeling (ABM) provides an alternative paradigm. In this context, agents are independent software entities that operate according to their own rules and objectives. Through their complex interactions (straightforwardly represented as a network of contractual obligations) they naturally capture the properties of real financial systems — in particular emergence, heterogeneity, and bounded rationality.

The burgeoning field of *agent-based computational finance* [9] has benefited from trends in “Big Data” and computer power. As financial data has become more readily available (though access is still expensive and restrictive), it has been possible to build ABMs that are able to reproduce detailed statistical properties of real financial systems. When properly calibrated and validated (more on this below), these models can be used for “what if?” analyses, allowing regulators to perform *in silico* trials of candidate policy interventions and market rules.

A celebrated example of the successful application of a financial ABM to assess potential rule changes was performed by Darley and Outkin [10]. In 1998 the NASDAQ stock exchange began considering a

reduction in minimum price increment (*tick size*) for orders to buy and sell shares on its trading platform. The change was related to decimalisation, with the minimum price increment reducing from \$1/16th to \$0.01. NASDAQ wanted to understand what impact the proposed change would have on market quality and fairness.

The model Darley and Outkin ultimately produced featured a rich and heterogeneous collection of simulated trading agents. These fell into two distinct groups:

1. Dealers — *market makers* using automated strategies that maintain concurrent offers to buy and sell shares, seeking to profit from the *spread* between the offered buy and sell prices.
2. Investors — possessing a noisy signal based on the true value of the monetary contracts (financial instruments, *e.g.* stocks, loans, and bonds) being simulated.

In the simulation, Investors choose to trade only when prevailing prices appear favourable, according to their private signal. In order to compensate Dealers for their lack of information regarding fair value, Dealers maintain a variably-sized spread and use a variety of heuristic approaches to determine when and where to trade. Agents are stochastically selected to trade according to heterogeneous timescales. The model is adaptive in two senses. Firstly, some classes of Dealer agents are themselves adaptive, utilising Reinforcement Learning to optimise their behavior on-line. Secondly, simulation runs feature evolutionary adaptation. Agents are permitted to go bankrupt, and when this occurs a surviving agent is selected for reproduction with mutation. In this way, the population of agents was observed to evolve, with more successful strategies coming to dominate the population.

The Dealer strategies were calibrated as closely as possible to real trading behavior, by analysing proprietary datasets from NASDAQ. The model was also carefully validated against observed statistical properties of the real exchange data. The results were conclusive — the proposed reduction in tick size, in the presence of so-called “parasitic” Dealer strategies [10], caused the simulated market to suffer a great reduction in its ability to accurately reflect the (exogenously specified) true value of the traded instrument. Furthermore, the presence of parasitic strategies was predicted to give rise to sudden large price movements. These effects were observed in live market data following decimalisation of NASDAQ trading. Darley and Outkin’s work was lauded for the insights it provided during the decimalisation process.

Micro–Macro Modeling

Whilst Darley and Outkin [10] used a *market microstructure* modeling approach, ABMs have also been successfully applied at *macroscopic* levels, where the agents are banks that operate over a timescale measured in days, months, and years. Banks make loans to each other and the *network* of loans facilitates the propagation of financial distress through the system. It has been shown that network properties such as graph density and node degree distribution affect the spread of distress (see Sidebar 2).

However, there are currently very few models that bridge the micro-macro gap. Such a model should fulfil the following criteria: 1) represent the limit order book; 2) account for how market participants

are financed by banks; and 3) compare systemic risk propagation channels both directly, between banks, and indirectly, due to overlapping investment portfolios. Our current research project seeks to provide just such a model. We use an ABM to investigate the propagation of flash crashes between financial institutions connected by a network of shared asset investments. By studying connected networks of banks, funds and assets during a flash crash event, we look at the effect of network topology and agent strategic behavior on the propagation of distress over the extremely short timescales characteristic of flash crashes. The ultimate goal of this project is to provide enhanced computational models of flash crash propagation between assets, institutions and markets, and to ultimately inform new policy insights.

In order to advance this research agenda, we have created a hybrid micro-macro model (see Figure 4). Our model captures key macro-level systemic risk features, by using a network of interbank lending and overlapping portfolios of financial securities held by hedge funds. The novelty in our approach is that, in addition to this network-theoretic framework, we incorporate a realistic micro-level simulation of asset trading on simulated stock markets.

Figure 5 presents visualisations of financial distress propagation through several small representative financial networks in our model. The state of the system is sampled at five-minute intervals, though the microstructure simulation operates on a timescale measured in milliseconds. Although networks of the size illustrated are not calibrated to real-world datasets, the critical importance of network topology is readily apparent. Systemic behavior is also highly-dependent on the trading behavior of the agents in the simulation.

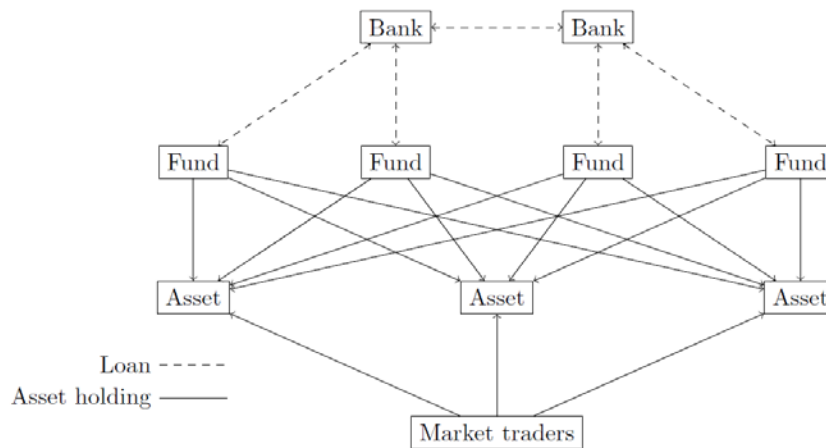


Figure 4: Relationships between the entities in our model. Arrows represent financial relationships, e.g., Loans – between banks and funds, Asset ownership within fund portfolios – between funds and assets or Asset ownership within trader portfolios – between traders and assets.




Figure 5: A temporal visualisation of flash crash propagation through a network of three funds (represented as squares) engaged in algorithmic trading of three assets (circles). As time increases to the right (time in minutes is shown), the impact of a shock to a single asset traverses the network of shared investments, resulting in flash crashes and bankruptcies. An edge from an asset to a fund represents a significant loss due to falling asset price. An edge from a fund to an asset represents distressed selling. a) A sparsely-connected network where fund portfolios do not overlap and distress does not propagate to other funds or assets. b) A densely-connected network which facilitates the ready propagation of financial distress. c) A highly-crowded network where all funds invest in a single asset. Such a network is critically sensitive to price fluctuations of the shared asset. The size of fund nodes represents the total value of Assets Under Management (AUM). The size of asset nodes represents the total invested in the asset, or Market Capitalisation (MCAP).

Challenges and Opportunities

A common criticism of ABMs is that they often have a great many parameters and so are at risk of overfitting empirical data, which happens when the model becomes too complex, and test errors are large – *i.e.*, the model fails to generalise on new data, although training errors are small. The trade-off between the complexity of the model, the training set size, and the generalisation error on new data has been recognised and widely researched.

This problem can be mitigated by reducing model complexity, for example by reducing the number of captured features or by modifying the feature selection algorithm, or by providing larger training data samples. Machine Learning regularisation methods such as dropout (dropping out hidden and visible units in neural networks), and principal component analysis (projecting data into optimally-selected lower-dimensional subspaces) provide robust approaches to controlling model complexity. Data availability is therefore of paramount importance if ABMs are to provide actionable intelligence for financial regulators. Bringing models to data involves two key processes: calibration and validation.

Calibration is the practice of fixing model input parameters using real data (for example the interbank datasets mentioned in the sidebar 2). As calibration can remove many free parameters from models, it is a powerful response to the over-fitting problem. However, the “curse of dimensionality” remains very material for ABMs. Machine Learning methods that can reduce model complexity are starting to find application within agent-based models. For example, *surrogate* models can be used to construct cheap-to-run proxies of full ABMs, with smaller numbers of parameters in need of estimation. Such proxies can be deployed to explore the input parameter space more efficiently than if evaluating the ABM directly. Bayesian methods can further optimise parameter space exploration by suggesting the most promising regions, based on information-theoretic criteria [12].

Validation is the practice of evaluating the model’s ability to reproduce the phenomena of interest. Due to their stochastic nature, it is usually necessary to perform repeated samples of the ABM to build up distributional outputs, an approach known as Monte-Carlo simulation. In financial agent-based models, validation centres around reproduction of a set of so-called stylized facts, for example the fat tails (excess kurtosis) of the price return distribution, mentioned above. Model power is expressed with reference to the model’s ability to jointly produce as many stylized facts as possible.


This need for repeated sampling over high-dimensional parameter spaces makes ABMs computationally expensive to work with. It is not always straightforward to parallelise ABMs and their branching logic is not always well-suited to implementation on GPUs. Machine Learning surrogate methods may offer further advantages here.

Another criticism often levelled at ABMs is that their results do not generalise well. Each model stands in isolation and it is difficult to compare the results of different studies. In a sense, the great power of ABMs, derived from their inherent expressiveness, is also something of a liability and the young science is yet to settle on standards by which to communicate model details or by which to compare and contrast results.

Adding to the challenge facing the adoption of ABMs by regulators, financial data of sufficient granularity to fully calibrate models is often highly proprietary and/or expensive to access.

Regulators recognise the need to support academic involvement and value the contributions from the community as they seek to ensure that the next financial meltdown never happens. Since the 2007-2009 crisis, advances in computational capability have brought huge changes both within markets in the form of algorithmic trading, and also in the academic study of financial systems via the application of network theory and agent-based modeling. Studying global financial markets as complex adaptive systems is a vital and timely endeavour, and one to which the field of Computer Science is uniquely equipped to contribute.

Side Bar 1: Flash Crashes



Sources: Bank of England calculations; Bloomberg.

Figure 1: A comparison of selected flash crash events (from [6]).

During a flash crash, price movements share a characteristic dip and rebound profile, playing out over a time-scale of minutes (see Figure 1). Such rebounds are historically unprecedented and reflect a temporary and profound break in the market's ability to price assets. The disorder and extreme price movement (*volatility*) observed during flash crashes is in part due to traders stepping away from the market, which results in sparsely populated *limit order books* [7] that can expose so-called *stub quotes* (orders placed far from prevailing prices that are not intended to be filled) to execution. In the original Flash Crash of 6th May 2010, the extreme volatility was triggered by a significant batch of sell orders in the SP500 e-mini futures contract — one of the most actively traded financial instruments in the world. Algorithmic traders known as *market makers* reacted to the early part of the price fall by stepping away from the market and either placing less competitive orders or refraining from placing orders entirely. This caused price movements to deepen. HFTs acting as market makers absorbed the early sell orders until their inventory limits were hit. At this point their algorithms were calibrated to

attempt to reduce risk by selling, and HFTs bought from and sold to each other in a frenzy of “hot potato” interactions occurring faster than the blink of an eye [8]. The resulting deluge of market activity data caused significant feed dissemination problems, further undermining investor confidence. Ten minutes later, reflecting a more human-scale decision process, fundamental value investors returned to the market and the added buying pressure returned prices to relative normality. At the lowest point, over one trillion US dollars had been lost from company values.

Although the details differ between flash crashes, the pattern of dip or spike accelerated by algorithmic trading and associated liquidity removal by market participants leading to sparse order books seems to be a characteristic set of features. Flash crashes aren’t restricted to a single asset class. To date, major flash crash events have been observed in futures, equity, foreign exchange, fixed income, commodity and even cryptocurrency markets. A flash crash that spreads between asset classes, instruments or markets has been dubbed a “splash crash” and is a particularly worrisome prospect for those engaged in day to day trading on the global financial markets.

Side Bar 2: Financial Networks




Figure 2: Core-Periphery interbank network structure from Bardoscia et al. [11]. In core-periphery networks node degrees are strongly heterogeneous, with a small number of highly connected nodes supporting a larger set of poorly-connected neighbours. The degree distribution of such networks can often be approximated as a power law, and so core-periphery networks are also known as *scale-free*. Scale-free networks are ubiquitous in natural and social systems, and the global interbank network is no exception.




Figure 3: a) The bipartite graph $G = (X \cup Y, E)$ where edges E represent investments made by institutions (bottom nodes, X) in securities (top nodes, Y). b) a security receives an exogenous shock and falls in price. c) one of the institutions commences distressed selling in response to losses realised from the falling security price. d) price impact due to distressed selling causes the price of further assets held by the distressed institution to fall. e) the second institution, exposed to losses in the single jointly-held security, commences distressed selling in response. f) contagion spreads to the final security, held solely by the second institution.

Banks regularly lend to each other, and such interbank loans have been the subject of intensive academic interest following their connection to the 2008 collapse of Lehman Brothers and the subsequent wider crisis at that time. A typical analysis represents banks as nodes of a network with directed edges between nodes representing loans from one bank to another (see Figure 2). Edges are weighted according to the size of the loan. It is often convenient to scale the loan sizes according to the underlying capital of the lender bank and, in this case, edges represent the *leverage* of banks (leverage is loan-to-value ratio, just like with a residential mortgage). Investments made by banks in sets of financial assets, or portfolios, can also be represented as a network. Here, the network forms a bipartite graph where nodes are partitioned into two disjoint sets representing banks and assets, with edges representing an investment made by a bank in a given asset (see Figure 3).

Network-theoretic tools can be used to analyse the adjacency matrices. For example, the largest eigenvalue of the leverage matrix tells us about the stability of the system to perturbations around a fixed point. Similarly, the node degree distribution and node centrality measures provide metrics that allow researchers to study the effects of network structure on financial distress propagation. Network metrics have successfully identified institutions that were financially vulnerable during the 2007-2009 crisis, for example Lehman Brothers.

It is often the case that the failure (default) of a bank causes financial distress to propagate to creditor banks, and so on, in a cascade of defaults. It is also possible to simulate the hoarding of loans by banks which can again cause financial distress to spread, this time in the opposite direction, to debtor banks. Finally, banks themselves will often hold investments in third-party assets such as stocks or bonds. If banks are forced to sell assets to remain solvent, the price impact of their distressed selling forces down the price of shared assets, triggering losses at banks with overlapping investments. The resulting “fire sale” represents another contagion channel for financial distress.

It is, in principle, possible to calibrate models of the banking system using regulatory data, however banks often treat this data as highly proprietary and release only aggregate lending figures. It is then necessary for researchers to use network reconstruction methods (such as maximum entropy reconstruction) to predict the true interbank network. Public filings are also quarterly or annual, creating further difficulties for researchers who wish to calibrate their models in a timely and robust manner.

References

- [1] J.H. Holland, “Studying Complex Adaptive Systems,” *Jrl Syst Sci & Complex* , vol. 19, no. 1, 2006, pp. 1-8.
- [2] I. Sommerville, D. Cliff, R. Calinescu, J. Keen, T. Kelly, M. Kwiatkowska, J. McDermid, and R. Paige, “Large-scale complex IT systems,” *Commun. ACM*, vol. 55, no. 7, 2012, pp. 71–77.
- [3] D. Easley, M.M.L de Prado, , and M. O'Hara, “The microstructure of the "Flash Crash": Flow toxicity, liquidity crashes, and the probability of informed trading,” *Journal of Portfolio Management*, vol. 37, no. 2, 2011, pp. 118–128.
- [4] N. Johnson, G. Zhao, E. Hunsader, H. Qi, N. Johnson, J. Meng, and B. Tivnan, “Abrupt rise of new machine ecology beyond human response time,” *Scientific Reports*, vol. 3, no. 2627, 2013, pp. 1–7.
- [5] M. Buchanan, “Trading at the speed of light,” *Nature*, vol. 518, no. 7538, 2015, pp.161–163.
- [6] R. Jackson, L. Crowley-Reidy, and A Schrimpf, “The sterling ‘flash event’ of 7 October 2016,” 2017, <https://www.bis.org/publ/mktp09.pdf>. [accessed 2018-02-27]
- [7] M.D. Gould, M.A. Porter, S. Williams, M. McDonald, D.J. Fenn, and S.D. Howison, “Limit order books,” *Quantitative Finance*, vol. 13, no. 11, 2013, pp. 1709–1742.
- [8] A. Kirilenko, A.S. Kyle, M. Samadi, and T. Tuzun, “The Flash Crash: High-Frequency Trading in an Electronic Market,” *The Journal of Finance*, vol. 72, 2017, pp. 967–998.
- [9] L. Tesfatsion “Agent-based computational economics: Growing economies from the bottom up”, *Artificial life*, vol. 8, no. 1, 2002, pp. 55–82.
- [10] V. Darley, and A.V. Outkin, “A NASDAQ market simulation: insights on a major market from the science of complex adaptive systems”, *Complex Systems and Interdisciplinary Science*, vol. 1, World Scientific, 2007.
- [11] M. Bardoscia, S. Battiston, F. Caccioli, and G. Caldarelli, “Pathways towards instability in financial networks,” *Nature Communications*, Vol. 8, no. 14416, 2017, pp. 1–7.
- [12] G. Fagiolo, M. Guerini, F. Lamperti, A. Moneta, and A. Roventini, “Validation of Agent-Based Models in Economics and Finance”. LEM Papers Series 2017/23, Laboratory of Economics and Management (LEM), Sant’Anna School of Advanced Studies, Pisa, Italy, 2017, www.isigrowth.eu/wp-content/uploads/2017/10/working_paper_2017_25.pdf.

Author Biographical Sketches

James Paulin is a doctoral student in the Department of Computer Science at the University of Oxford. His research interests include agent-based modeling, financial system complexity, network theory and artificial intelligence. He holds a Master’s degree (distinction) in Evolutionary and Adaptive Systems from the University of Sussex. His research is supported by Oxford Asset Management LLP. Contact

address: Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD. james.paulin@cs.ox.ac.uk

Anisoara Calinescu is a Senior Lecturer in the Department of Computer Science at the University of Oxford. She holds a DPhil from the Department of Engineering Science of the University of Oxford. Her research interests include complex systems and complexity metrics; agent-based modeling; financial and supply systems. Contact address: Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD. ani.calinescu@cs.ox.ac.uk.

Michael Wooldridge is Head of Department and Professor of Computer Science in the Department of Computer Science at the University of Oxford, and a Senior Research Fellow at Hertford College. His main research interests are in the use of formal techniques for reasoning about multi-agent systems. He is particularly interested in the computational aspects of rational action in systems composed of multiple self-interested computational systems. Contact address: Department of Computer Science, University of Oxford, Wolfson Building, Parks Road, Oxford OX1 3QD. michael.wooldridge@cs.ox.ac.uk.

Image Permissions

Flash Crashes: Bank for International Settlements 2017. All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated. <https://www.bis.org/publ/mktc09.htm>

ABM entity relations: author-created image, no permission necessary

Core-Periphery Network: This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. <https://www.nature.com/articles/ncomms14416>

Bipartite graph: author-created image, no permission necessary

Flash crash propagation visualisation: author-created image, no permission necessary