



Office of Financial Research
Working Paper #0003
December 21, 2012

Using Agent-Based Models for Analyzing Threats to Financial Stability

Richard Bookstaber¹

¹ Research Principal, OFR, richard.bookstaber@treasury.gov

The Office of Financial Research (OFR) Working Paper Series allows staff and their co-authors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official OFR or Treasury positions or policy. Comments are welcome as are suggestions for improvements, and should be directed to the authors. OFR Working Papers may be quoted without additional permission.

www.treasury.gov/ofr

Using Agent-Based Models for Analyzing Threats to Financial Stability

Richard Bookstaber

Abstract

Existing models of financial instability tend to be based on top-down, partial-equilibrium views of markets and their interactions; they are unable to incorporate the complexity of behavior among heterogeneous firms or the tendency for all types of firms to change their behavior during a crisis. This paper argues that agent-based models (ABMs)—which seek to explain how the behavior of individual firms or “agents” can affect outcomes in complex systems—can make an important contribution to our understanding of potential vulnerabilities and paths through which risks can propagate across the financial system.

Introduction

The financial crisis revealed important weaknesses in the risk models used by financial institutions, supervisors, and financial economists to understand and respond to risks in the financial system. The current generation of models is unable to model financial vulnerabilities, the shocks that might expose these vulnerabilities, and the process by which such shocks might propagate through the financial system. Rectifying these weaknesses is a critical first step in developing the capability for dealing with threats to financial stability. The Office of Financial Research has a responsibility to contribute to the development of models that can be used to identify and analyze shocks that may affect financial stability and to facilitate a rapid response to those shocks.

Much of the criticism since the crisis has focused on Value-at-Risk (VaR) models, which measure the risk of loss of a portfolio over a given time horizon and probability.¹ VaR is the most broadly applied risk management tool, underlying Basel II and most industry risk management systems. Although VaR has demonstrated value in gauging risk during typical day-to-day market moves—which reflect the same distribution as the recent past—it does not perform well in times of market dislocation. Given its reliance on historical statistical relationships, it should not be expected to do so.

Another widely applied approach to risk measurement is stress testing. In the last decade stress testing came into its own for regulatory and supervisory purposes.² Yet even as stress testing became standard fare in the period up to and following the 2008 crisis, its limitations grew increasingly apparent. Stress tests did not anticipate the extreme shocks that occurred during the crisis, failed to shed light on some of the sectors and risk factors that were instrumental in the development of the crisis, and ignored the dynamics between the sectors that ultimately were affected.³ To take one example, the IMF offered a

¹ Indeed, the failure of VaR models was the subject of a Congressional hearing in 2009. See Bookstaber and Taleb (2009).

² Basel II requires institutions to perform stress tests for credit, market and liquidity risk. Following the Asian financial crisis, the IMF included stress tests in its menu of approaches to examine the soundness of banks and the financial sector in the Financial Sector Assessment Program. In 2009, the Federal Reserve's Supervisory Capital Assessment Program turned the supervisory function of stress testing up a notch by applying it as a tool for policy action during a crisis as opposed to evaluating the financial landscape during stable times.

³ See Borio, Drehmann, and Tsatsaronis (2012).

reassuring assessment not long before Iceland’s systemic breakdown: “The banking system’s reported financial indicators are above minimum regulatory requirements and stress tests suggest that the system is resilient.”⁴

A key drawback of both VaR and typical stress testing practices is that they are guided by history. The volatility and correlations for VaR are drawn from an historical sample, and stress test scenarios typically replicate historical events or express extreme “tail events” based on an historical distribution, even though it is well known that the nature of crises is to have unanticipated shocks and unexpected interrelationships where the past offers limited guidance. Similarly, both VaR and stress tests implicitly model the risk of a crisis as just a “bad draw” from the historical distribution. In doing so, these methods manifest a second failing: a crisis comes from the unleashing of a dynamic that is not reflected in the day-to-day variations of pre-crisis times. The effect of a shock on a vulnerability in the financial system—such as excessive leverage, funding fragility, or limited liquidity—creates a radical shift in the markets similar to what is observed in traffic jams or the panic of crowds. Economic relationships change during these times of stress. Thus the extreme event reflects the inappropriateness of the risk model, not an extreme draw from it.⁵

New measures of financial stability have been proposed since the crisis to improve our understanding of financial shocks and vulnerabilities. The OFR, in its first Annual Report, offers a preliminary empirical evaluation of the performance of these measures during four historical crises.⁶ But absent a model that can chart the course of events during financial disruptions, it is difficult to assess the value of these measures, especially given the many changes in the financial landscape that occur over time which lead to new vulnerabilities and paths through which shocks can propagate across the system. And existing models of the financial system are partial equilibrium models that are not built for this task. These models are designed to rebound back from shocks and return to the equilibrium state rather than running off the rails.

To make up for these weaknesses, a model of the financial system needs to be flexible, in order to provide results without extensive historical data either for parameterization or for testing, and in order to describe markets that fall into disarray.

Of course, the demands of a model of financial crises extend beyond measuring risk and highlighting vulnerabilities to evaluating policy actions. If a crisis is underway, the question naturally arises of what actions to take to stem its financial and economic costs. Here too, an equilibrium model governed by applying historical relationships is likely to be inadequate. We need to have a way to specify the world with the critical non-linearities, complexities, and instabilities, and to project forward likely paths for how the world might evolve given various policy actions.

One promising approach is agent-based modeling. This paper surveys agent-based models (ABMs), looking specifically at their application for financial markets and institutions, and even more specifically at their application for understanding potential vulnerabilities and paths for risk in those markets and institutions. This is an overview of the field rather than a comprehensive survey of the literature. The discussion aims to illustrate how ABMs might help fill the gaps that remain in standard

⁴ See IMF (2008), page 5; and Alfaro and Drehmann (2009).

⁵ See OFR (2012), Chapter 3.2 for a discussion of the limitations of current approaches to stress testing.

⁶ See OFR (2012), Chapter 3.1. A more comprehensive survey of financial stability measures is presented in Bisias and others (2012).

risk tools such as VaR and stress testing—in particular, their inability to deal with crisis dynamics—and to see how ABMs can be used as a tool for crisis management and to gauge the effects of policy actions.⁷

The next section introduces ABMs and illustrates their application in social interactions and vehicular and pedestrian traffic flow, areas in which they have led to valuable insights. Section II highlights the differences between traditional economic models and the ABM approach, focusing on the assumptions that can be weakened as one moves from traditional models to the ABM approach. Section III gives some examples of applications of ABMs in finance and economics. ABMs are outside the academic mainstream in economics and finance; as a result, the literature is sparse, with much of what has been done coming out of industry and central banks rather than academia. Section IV discusses some of the practical issues in implementing ABMs. Section V concludes with a discussion of how ABMs might fit into the agenda of the OFR and be used to address general questions of policy analysis.

I. What is Agent-Based Modeling?

Agent-based models use a dynamic system of interacting, autonomous agents to allow macroscopic behavior to emerge from microscopic rules. Depending on the application, the agents might represent biological organisms, social groupings, institutions, or physical entities. In a financial application the agents might include asset management firms that invest in various assets and banks that lend to these agents to finance their investments.

The models specify rules that dictate how agents will act based on various inputs. Each agent individually assesses its situation and makes decisions on the basis of its rules. Once the model has specified the initial conditions and the agents' rules, the "world" is let loose and all subsequent events are driven by agent interactions. These interactions—the attempts by agents to express actions within their world—are determined dynamically by the agents' rules, which encompass the agents' preferences and data-processing methods, and the information the agent receives about the state of nature. After each time step, the new environment and values for the variables resulting from the interaction of the agents is fed into the simulated world for the next-step iteration. The model can be run repeatedly, in effect creating multiple worlds, in order to generate a distribution of results.

The agents are free to act within their computational world, just as their counterparts do in the real world. They plug their knowledge about their local environment into rules that are designed to mimic the behavior of real-world agents. That is, the rules are not necessarily based on optimization and assumptions of rationality, but may instead reflect the heuristics that are typically applied based on observations of real-world bounded rationality.

Examples of Agent-based Models

⁷ This can extend even further to using ABM in an early warning system, but doing so requires adding to the identification of vulnerabilities the anticipation of shocks that might occur, a task that in itself is fraught with uncertainty, and that is made all the more complex because the policy and market responses cannot always be anticipated. For example, in the 2008 crisis central banks took unprecedented actions that would have been difficult to anticipate, as would the effects of those actions.

Birds in flight. The flight of a flock of birds appears to operate as a system, yet is based on the decisions of the individual birds, each reacting to the other birds in the flock.⁸ Building a macro, top-down model of the behavior of the flock will miss the reality of the situation, because at the macro level the movements of the flock are complex, non-linear, and not based on any system-wide program. Yet the task turns out to be remarkably easy if it is viewed, more realistically, as the aggregation of local interactions by the individual birds in the flock. Reynolds (1987) models the movement of a population of artificial “boids” based on three simple rules:

1. Separation: Don't get too close to any object, including other boids.
2. Alignment: Try to match the speed and direction of nearby boids.
3. Cohesion: Head for the perceived center of mass of the boids in your immediate neighborhood.

Reynolds' method is an early example of ABM, modeling the interaction of the birds atomistically rather than modeling the flock at a macro level.

Social Interaction. Schelling's (1971) model of neighborhood segregation is one of the earliest and best known ABMs based on movement in a spatial network. Two groups of agents are randomly distributed on a lattice; individuals move to empty locations if the number of in-group neighbors falls below a certain threshold. The threshold sets the extent to which each individual is segregationist, that is, his or her preference to stay within the group: strongly segregationist individuals relocate in response to small drops in in-group neighbors, while weakly segregationist individuals only move to avoid being in the minority. The model shows that extreme segregation tends to arise even in populations whose members are only weakly segregationist. Later models by Latané (1996) and Axelrod (1997) use rules where each agent mimics his or her neighbors. Starting with a random distribution for the types of agents, a population filled with mimics might be expected to converge on a single profile. However, perhaps surprisingly, the system achieves stable diversity.⁹

Traffic Flow and Crowd Dynamics. ABM has found real-world application in modeling traffic flows. Researchers at Los Alamos National Laboratory have developed an ABM traffic simulator in which agents act based on rules that are drawn from actual patterns of daily activity, such as the time of day and the routes taken to work, shopping, and recreation. For instance, when a trip takes too long, people find other routes, leave at different times, or decide not to drive to an activity at all. Because the ABM tracks individual travelers' locations and routes, it is a useful tool for determining how changes in transportation policy will affect reliability and congestion. The model was first applied in Portland, Oregon, where it requires 120,000 links and 1.5 million travelers and simulates the movement of cars across the transportation network on a second-by-second basis.¹⁰

⁸ An example of the complex flight dynamics, with a flock of many thousands of birds moving as a single wave is that of starling murmuration. Keim (2010, 2011) provides an overview of recent research regarding the empirical and mathematical properties of this flocking behavior and links to examples.

⁹ Stable minority subcultures persist because of the protection of structural gaps created by cultural differences that preclude interaction, thereby insulating agents from homogenizing tendencies.

¹⁰ The interactions of individual vehicles produce realistic traffic dynamics that analysts use to estimate vehicle emissions and to judge the performance of the transportation system. The model forecasts how changes in transportation policy or infrastructure might affect those activities and trips, capturing the important interactions between travel subsystems, such as an individual's activity plans and congestion on the transportation system. Toroczkai and Eubank (2005) describe the Los Alamos National Laboratory model, called TRANSIMS, along with a variant of the model that simulates epidemics, called EPISIMS.

The application of ABM for modeling flow dynamics has also been applied to congestion and crowding among pedestrians. This can help design venues to minimize the risk of panic when a crowd tries to push out the door at the same time.

Epidemiology. Epidemiologists have also found a use for ABMs.¹¹ A simple model of an epidemic in a human population can be constructed with a few simple steps that are reminiscent of Conway's Game of Life. Agents are characterized as being either healthy or ill. Two rules are defined that determine which of these states each agent is in. At each tick of the clock, all agents make a random movement. If two agents collide, and if the variable "healthy" has the same value for both, nothing happens; if "healthy" is false for one agent and true for the other, the healthy one is infected by the sick one.

The random movement rule can be replaced with a more intricate rule that mirrors the path of people throughout the day, similar to a flow model like that of traffic. The rule of becoming ill when agents collide can be made more realistic by having the contagion be stochastic and based on factors such as age and length of time the ill and healthy agent are in the same vicinity.¹²

In all of these examples—birds in flight, social interaction, traffic flows, and the spread of epidemics—the essence of the problem being modeled lies in the complex interactions of individual agents. By employing a model that focuses on that fact, the macro behavior, which is often not intuitive, becomes manifest.

II. Agent-Based Models in Economics and Finance: A Comparison to Traditional Approaches

As these examples show, ABM takes a bottom-up perspective. Aggregate properties and the macro outcome are the result of the dynamics of the individual agents. This section contrasts this perspective with the top-down perspective of traditional neoclassical models, in which the bottom level typically comprises a representative individual and is constrained by strong consistency requirements associated with equilibrium, rationality, and the regularity conditions required for mathematical tractability.¹³

The Standard Neoclassical Model

Dynamic stochastic general equilibrium (DSGE) models build on a neoclassical microeconomic foundation and address the issues of system interaction that the OFR would demand from an ABM. These models therefore provide an example of the application of the standard approach to the same sort of problems that could be addressed through an ABM.

The DSGE framework employs interrelated models of key economic sectors: a representative household, which consumes; a set of firms, which produce goods; and the government, which sets

¹¹ For example, ABMs have been applied to the use of vaccines in controlling pandemic influenza in an intra-community network (Glass and others, 2005) and to model the real-time spread of the H1N1 virus (Bajardi and others, 2009).

¹² Los Alamos National Laboratory researchers have applied their traffic model to this purpose. Diseases are transmitted through the air between two agents who spend time in one another's proximity. By tracking people using the traffic flows they create a contact network of people nodes and location nodes, with a time stamp of each person at each location. There is a likelihood of disease transmission when two people are at the same location at the same time.

¹³ For a treatment of the limitations of traditional equilibrium models and a discussion of alternative models that deal with these limitations, see Farmer and Geanakoplos (2009).

monetary policy. The models determine the supply and demand of goods based on the optimization of a utility function for the households and a production function for the firms. The models are subjected to random exogenous shocks, which give them their stochastic nature. DSGE models thus contain the critical elements of neoclassical models in a multi-period setting—representative and homogeneous agents acting rationally based on the optimization of a mathematically well-behaved objective function, a pricing mechanism that drives the system toward a market-clearing equilibrium, and uncertainty that is exogenously introduced to the system. The models specify functional forms that allow for analytical tractability and that allow the model to conform to the assumptions of rationality.

The simplifying assumptions of a DSGE model, which are typical for traditional, top-down analytical models, facilitate mathematical tractability but at the same time limit their appeal for modeling market crises. For example, in a DSGE model, the risk is introduced through well-specified exogenous shocks that do not change through the actions taken by the agents, whereas in a real crisis, the risk tends to come from the actions of the agents themselves, such as the pulling away of liquidity, the fire sales due to forced liquidation, and the withdrawal of sources of funding. Agents react, adapt in their behavior, and in doing so create endogenous uncertainty.

Furthermore, the assumption of rationality and optimization—and with it the ability of the actors to solve a multi-period decision—is not realistic during a crisis (if it is ever realistic). During crises, historical relationships no longer hold, the course of events depends on feedbacks among agents, and the key determinants of those feedbacks are unknown—who is leveraged, what positions they might have to liquidate, how the banks might alter their funding with shifts in the value of their collateral. The very notions of a representative agent and of an equilibrium model itself are inherently inappropriate for modeling a crisis.

Solow (2010) has this to say about the DSGE models:

They take it for granted that the whole economy can be thought about as if it were a single, consistent person or dynasty carrying out a rationally designed, long-term plan, occasionally disturbed by unexpected shocks, but adapting to them in a rational, consistent way....The protagonists of this idea make a claim to respectability by asserting that it is founded on what we know about microeconomic behavior, but I think that this claim is generally phony. ...

The DSGE school populates its simplified economy – remember that all economics is about simplified economies just as biology is about simplified cells – with exactly one single combination worker-owner-consumer-everything-else who plans ahead carefully and lives forever. One important consequence of this “representative agent” assumption is that there are no conflicts of interest, no incompatible expectations, no deceptions.

...the basic story always treats the whole economy as if it were like a person, trying consciously and rationally to do the best it can on behalf of the representative agent, given its circumstances. This cannot be an adequate description of a national economy, which is pretty conspicuously not pursuing a consistent goal.

Solow’s sentiments could be extended to other top-down models of the aggregate economy.

The Agent-based Model

Whatever the benefits and failings of the traditional, neoclassical approach represented here by the DSGE models, these models are at their weakest when the world is careening out of control. The

structure of ABMs overcomes some of the problems that arise in applying DSGE and other traditional models to financial systems in times of crisis. As is evident from the applications discussed in the last section, ABMs focus on determining the way that even simple and predictable local interactions can generate complex global patterns. This makes ABMs particularly appropriate for studying dynamic processes that lack central coordination, that are “pretty conspicuously not pursuing a consistent goal.” Because of the bottom-up approach at the core of ABMs, the building blocks of ABMs are best described through the fundamental characteristics of the agents. These characteristics include:

Agents are autonomous. The system is not directly modeled as a globally integrated entity. Systemic patterns emerge from the bottom up, coordinated not by centralized authorities or institutions (although these may exist as environmental constraints) but by local interactions among autonomous decision-makers. This process is known as “self-organization” (Kaufman, 1996).

Agents are heterogeneous. An ABM need not use a representative agent. The agents can have different rules and heuristics, endowments, and objectives. This heterogeneity readily allows models to incorporate gaming behavior and informational asymmetries.

Agents have bounded rationality and operate based on behavioral heuristics. Economists have come to recognize that behavior is not necessarily consistent with the assumptions of optimization and rationality that underlie the neoclassical models of the DSGE world, but may instead reflect the heuristics that are typically applied based on the real-world bounded rationality, where the environment of economic agents is too complex for rationality and optimization to be a viable simplifying assumption. Perhaps agents might have some local and partial principles of rationality (e.g., myopic optimization rules), but more generally, agents in ABMs are assumed to behave with bounded rationality and operate using heuristics.¹⁴

Agents are interdependent. Agents have direct interactions. The decisions undertaken today by an agent directly depend on the past choices made by other agents in the population, thus creating strong path dependence.¹⁵

Interdependence may involve iterative processes through which agents influence other agents, who in turn influence others. Interdependence may also be indirect, as when agents’ behavior changes some aspect of the environment, which in turn affects the behavior of other agents, so that the consequences of each agent’s decisions depend in part on the choices of others. Thus the relational assumptions imbedded in the social structure of the agents’ interactions are combined with behavioral assumptions based on the heuristics of the agents.

Agents adapt. As a result of interdependence and bounded rationality, agents adapt to the changing environment based on the rules that govern their behavior.¹⁶ In economics, agent adaptation is the core

¹⁴ There is a broad literature on the use of heuristics; one recent set of work edited by a major proponent is Gigerenzer and others (2011).

¹⁵ See Page (2006). Other references that discuss ways in which an agent’s decisions are influenced by others include: Granovetter (1978); Golub and Jackson, (2010); DeMarzo and others. (2003); and Jackson and Yariv (2007).

¹⁶ The rules that lead to adaptive behavior might themselves be fixed, or might also be dynamic. One common approach used in ABMs for creating dynamic rules is the use of genetic algorithms, which adjust the rule’s parameters; another less frequently used approach is genetic programming, which adjusts the functional form of the rules. Both of these

of Goodhart's law, a variation of the Lucas Critique, which states that once policymakers have identified a policy target, market participants will change their behavior and the target will lose its value. And thus the system evolves, generally in a complex fashion. Aggregate properties are thought to emerge out of repeated interactions among simple entities, rather than from the consistency requirements of rationality and equilibrium imposed by the modeler. As the agents adapt, their interaction can generate a "complex adaptive system" along the lines of Holland (1975, 2006). Because the essence of an ABM is to have agents interact and adapt, it focuses attention on the implications of market changes on model dynamics.

Having adaptable agents is particularly important when dealing with the changes in the marketplace that occur during periods of crisis. For example, a fund will operate on different rules if put in an environment of forced liquidations. The usual rules of fine tuning positions based on expected returns will be replaced by rules that focus instead on the cost of speedy liquidation. Banks that follow rules for maintaining stability in margin and haircut levels in normal times will replace these with rules that treat margin and haircuts levels as critical decision variables. With these sorts of adaptations in rules will also come adaptations in the information the agents acquire and process.

An ABM thus describes a system from the perspective of its constituent units. A priori constraints on agent interactions can be dictated by the realities of the problem being addressed rather than being imposed based on equilibrium conditions, homogeneity assumptions, or mathematical regularity conditions that are required by analytical frameworks. An ABM can fit characteristics of individual behavior that are evident in financial markets such as memory, path dependence, and the non-Markovian behavior and temporal correlations that come from learning and adaptation. An ABM also reflects the sort of information set and behavior rules typical in financial markets because it allows each agent to be individually specified. The dynamics of the market depend on agents approaching the market with different interpretations of market events, objectives, and routes for adaptation. Furthermore, the result of the agents' interaction leads to a system that may exhibit complexity at the global or macro level even when the agents are operating using simple rules, such as in the "boids" model.

These characteristics give an ABM essential features needed for modeling emergent phenomenon during crises:

Emergence is a process that leads to the appearance of structure not directly described by the defining constraints and instantaneous forces that control a system. Over time 'something new' appears at scales not directly specified by the equations of motion. An emergent feature also cannot be explicitly represented in the initial and boundary conditions. In short, a feature emerges when the underlying system puts some effort into its creation. (Crutchfield, 1994).

Emergent phenomena result from the interactions of individual entities. By definition, emergent phenomena cannot be reduced to the system's parts: the whole is more than the sum of its parts because of the interactions among the parts. An emergent phenomenon can have properties that are decoupled from the properties of the parts. For example, a traffic jam, though the result of interactions among individual vehicle drivers, may move in the direction opposite from that of the cars that cause

are designed to imitate an agent's heuristic, trial-and-error approach to learning. Yeh (2007a, 2007b) discusses the difference between adaptive and evolutionary computational approaches to learning.

it, for example, through “rubbernecking.” ABMs offer a useful approach to modeling emergent phenomena because an ABM simulates the behavior of the system’s constituent units (the agents) and their interactions in a non-equilibrium setting, capturing emergence from the bottom up when the simulation is run.

A canonical application of an ABM to emergent phenomenon is simulating the crowd interactions that lead to life-threatening stampedes in such disparate situations as fires or the rush for seats at a concert—recent examples of such stampedes include the panic in Harare, Zimbabwe, in the Roskilde rock festival, and in Mecca during the Hajj. By simulating the crowd’s emergent behavior, designs and procedures can be innovated to reduce injuries. Simulations in an ABM have provided valuable insights into the mechanisms of the uncoordinated responses induced by such crowd panic and have suggested practical ways to reduce the stampedes.

One well-known example looks at escape panic such as might occur during a fire in a movie theatre or a concert hall. In a simulation with one exit, it turns out that more people escape when a pillar is posted a few feet in front of the exit. Intuitively, one might think the column will slow down the outflow of people. However, ABM simulations, backed by real-world experiments, indicate that the column provides structure to the flow, leading to fewer injured people and a significant increase in the flow.¹⁷ This result is an example of a counterintuitive consequence of an emergent phenomena. After all, who would think of putting a column in front of an emergency exit?

There is an analogy between these crowd dynamics and the circuit breakers and other mechanisms that are designed to regulate the flow of liquidation in financial markets. In the face of a liquidity or maturity mismatch, a market shock can lead to a literal rush toward the exits. This leads to a feedback in the market akin to that of crowd dynamics; when sources of liquidity dry up, even those who are unaffected by the mismatch join in the rush to exit the market. This dynamic has little in common with a stable equilibrium environment.

III. Applications of Agent-based Models in Economics and Finance

An ABM can endogenize many characteristics that are taken as exogenous in equilibrium models. To identify market vulnerabilities and to know where to put policy guard rails, we need to know where people tend to veer off the road. Doing so requires reflecting on the set of choices facing the individuals and on the effects of feedback and adaptation on their decisions.

One example of the insight that might be gained through this is the so-called volatility paradox, which is the tendency for volatility to drop when market risk and leverage are rising, luring investors into complacency. A bottom-up approach to this procyclical dynamic would model the effect of individual agents’ actions on volatility and thereby model volatility as an endogenous part of the market. This might uncover clues to the potential for a crisis. For example, it may help answer whether the low volatility resulted from increasingly levered agents who are supplying liquidity based on small price advantages.

¹⁷ See Helbing and others (2000). In a simulation with 200 people in a room and with 45 seconds to escape, 44 people escape and 5 are injured when there is no column, whereas with a column 72 people escape and no one is injured.

An ABM approach might have similar benefits for understanding shifts in the correlation between asset returns. Similar to volatility, correlations tend to be low in calm periods and rise markedly once a crisis hits. Low correlations lead to the same sort of procyclical effects as low volatility because they suggest the ability to reduce risk through diversification, what might be termed the diversification paradox. As with volatility, if we look at correlations as endogenous to the dynamics of the market, they might provide early warning that market participants are finely differentiating one asset from the other and are searching out exotic hinterland markets. If we can treat volatility and correlation as endogenous rather than as exogenous statistics of market risk, we may find that they are informing us about market vulnerability.

Systemic risk—risks to the operation of complex systems—can be an emergent phenomenon that arises from the interactions of individual actors, generating collective behavior at a system-wide level whose properties are not obvious from the decision rules of each of the individual actors. Typically the individual actors believe they are acting prudently; the systemic risks occur because no one properly understands (or perhaps cares) how their behavior will affect everyone else and because there are nonlinear feedbacks that are not properly taken into account by anyone’s models.

Leverage is a prime example of this pathway to systemic risk. Thurner (2011) applies an ABM to this problem in which individual banks lend with smaller and smaller collateral requirements—that is, at increasing leverage—as they feel safer; the growth in credit fuels asset prices. On the way down, as each bank gets more nervous about the rising uncertainty in the world, it stiffens collateral requirements, reducing leverage and pressuring asset prices. Rarely are banks in a position to take into account that they may all be behaving similarly, and that as a result they will actually create the catastrophe they are each trying to avoid. Thus we see that leverage has a dramatic effect on the economy, but also that the economy has a dramatic effect on leverage.

This dynamic played out during the LTCM failure in 1998, which Bookstaber (2000) uses to describe the liquidity crisis cycle. The proximate cause of LTCM’s collapse was the failure of the Russian bond market. While LTCM did not have much exposure in Russia, it did have high exposure in other markets, such as Danish mortgage bonds. Other investors were heavily invested and leveraged in both markets. When the Russian market failed and these firms received margin calls, their next step was to liquidate what they could, which meant, among other things, selling out of their positions in Danish mortgage bonds. Because this created a contagion into a market where LTCM was heavily invested and leveraged, LTCM was caught in an avalanche. If LTCM’s risk managers had been asked to do a stress test for the firm’s direct Russia exposure, the warning light would have likely remained green. It is only by working through the interactions of heterogeneous agents that the risk would be manifest.

The following brief sections describe ways that ABMs are already being used in finance and economics to analyze market functions and systemic risk: funding and leverage from banks to asset managers, housing and mortgage prepayments, payment systems, market microstructure, and macroeconomics.¹⁸

Banks and Asset Managers

¹⁸ Because the objective of this survey is to illustrate work in the field rather than to provide an academic synopsis or synthesis, this paper features only one or two applications in each of these areas. This is not meant to imply an endorsement of any one approach or result, nor to suggest that any excluded works are less valuable.

Given the institutional characteristics of the LTCM crisis and other episodes, one good place to start for modeling crisis behavior is to look at the relationship between those demanding funding for leverage and those who provide that funding. Funding demanders can be viewed as hedge funds and providers can be viewed as banks.

Turner and others (2010) provide one model for the interaction of banks and hedge funds.¹⁹ The model employs four types of agents: Noise traders who trade more or less at random but with a bias of driving prices towards a fundamental value; leveraged value investors in the form of hedge funds that hold a stock when it is underpriced and otherwise hold cash; investors who decide whether to invest in a hedge fund; and a bank that can lend money to allow hedge funds to leverage their investments. In this model the presence of hedge funds usually dampens volatility, pushing stock prices towards their fundamental value. However, if the price of a stock drops precipitously, the hedge funds' wealth falls and leverage increases; thus, highly leveraged funds have to sell stock into a falling market in order to keep within leverage limits.

This ABM simulates interactions across agents to show how the behavior of hedge funds amplifies price fluctuations, in extreme cases causing crashes, and how the standard ways banks attempt to reduce their own risk can create more risk for the system. This model thus shows qualitatively how leverage can lead to crashes and generates price behavior that conforms with what is observed in reality.

Housing

The conventional top-down model for housing is typically an econometric model based on trends in incomes, interest rates, and housing supply and demand. But the actual dynamics that drive the housing market play out at the micro level through the interactions of individual buyers and sellers. These micro level dynamics can be specified in an ABM.

Gilbert and others (2009) employ an ABM to analyze and investigate shocks in the English housing market, looking at the interactions among buyers, realtors, and sellers. They simulate the interaction of individual households who seek to buy and sell properties; house prices emerge from this process. The model simulates a world in which households can move when forced (for example, due to death or job relocation), when their mortgages become too expensive for their income and they need to trade down, or when their incomes rise to the point where they can trade up.

In the ABM, sellers put houses on the market based on valuations provided by realtors, and then reduce the price if they do not sell.²⁰ Buyers look for a house that best fits their income level and make an offer, but become discouraged and leave the area over time if they cannot find a property that matches their income. The interaction among agents occurs because in the British system a buyer can only buy the seller's house either if it is already empty or once the seller has vacated and moved to another house. Offers only go through if this "property chain," a common feature of the British

¹⁹ Brunnermeier and Pedersen (2009) provide a similar dynamic, but without employing an ABM approach.

²⁰ In the English housing market, realtors play a pivotal role in determining the sale price for a house. In Gilbert's model, this is done by looking at the median price of houses in the geographical area of the house being put up for sale (similar to what an appraiser will do in the U.S. market). Thus the ABM includes a spatial dimension. If a house fails to sell during a period, the selling price is reduced and it continues to be on the market for the next period and until it is sold or demolished.

housing market, does not break down. The model works out which chains of offers remain intact and which offers fail as chains collapse.²¹

Mortgage Prepayments

ABMs find a natural niche with mortgage prepayments. Prepayments are determined by the heterogeneity and the often-non-optimal actions of individual homeowners. Some homeowners are slow to react to lower rates, leading to a residual of mortgages that fail to prepay—a behavior that is termed burnout. Also, homeowners differ in the variables governing their prepayment decisions, such as issues of retirement, relocation and divorce. This, coupled with an abundance of homeowner-specific data, has led the industry to use ABMs for modeling mortgage prepayments.²²

Payment Systems and Credit Risk

The main difficulty in representing payment systems for analytical purposes is that the behavior of a system as a whole is more than the sum of the behavior of its parts. For example, how and in what order participants transfer funds to each other is as important for the outcome of the settlement process as the amount, originator, and destination of each transaction.

The Bank of England pioneered the application of ABMs to payment systems by central banks.²³ The Bank of Italy introduced one of the most recent and detailed ABMs of a “plain vanilla” real-time gross settlement system, described in Arciero and others (2009). That model includes a number of distinct types of agents and events: banks, the central bank, payment requests, defaulted operations, interbank loan requests, crisis events, and banks hit by such events that are then unable to perform any operation for a given time.

Disaster is simulated through the introduction of an agent of the “disruptive event” breed. A bank picked at random stops payment requests and interbank market activity. Other banks become aware of this over time, and when they do, they stop payment and loan activity with that bank. They adjust their liquidity expectations, and money market activity slows down and losses accumulate. Whether the system goes back to normal functioning depends on the relative impact of the disruptive event compared with the total amount of liquidity in the system.

Market Microstructure

Many of the most puzzling issues in finance deal with problems of behavioral heterogeneity, and the dynamics that result from that heterogeneity. Empirical features such as trading volume, sudden price changes, kurtosis, and market crashes occur because of the heterogeneity among market participants.

²¹ Khandani and others (2010) analyze how three apparently favorable market conditions acting in concert—lower interest rates, higher home prices, and easier access to mortgage loans—can increase risk in the housing market. They test for the effect that refinancing to extract home equity has on the residential housing market, and gauge the magnitude of the refinancing ratchet effect by creating a numerical simulation of the U.S. mortgage market calibrated to the existing stock of real estate, and using a simple heuristic for the typical homeowner’s equity extraction decision. Although their model omits a few common features of ABMs, such as agent interaction or adaptive behavior, it illustrates the explanatory power of a bottom-up approach.

²² Geanakoplos and others (2012) describe a prepayment model, using actual loan data for millions of households, in which burnout at the aggregate level is a result of agent heterogeneity at the micro level.

²³ Galbiati and Soramaki (2008).

ABMs should help address these features because they are inherently structured to accept heterogeneous agents.

LeBaron (2006) models market behavior under heterogeneity using a single risky asset, which follows a random walk growth process, and a risk-free asset. Agents chose among a set of portfolio strategies that determine the fraction of wealth in the risky asset. Agents must evaluate rules using differing amounts of past information to evaluate rules. In this way the model implements two behavioral features. First, agents clearly exhibit bounded rationality; they cannot determine the entire state space of the economy. Second, they employ heterogeneous strategies with a small sample bias because they don't all choose to use as much data as possible. The strategies of the various agents are used to numerically form a market excess demand function with the price determined by market clearing.

A number of measures of the returns generated from this simple ABM share characteristics found in real market returns. For example, the model returns show strong evidence for leptokurtosis, a slow decay in autocorrelations that is a characteristic of market returns, as well as cross-correlations between trading volume and absolute values of returns with the strong positive contemporaneous correlation present in actual equity time series.

Macroeconomics

One macroeconomic ABM is EURACE.²⁴ It is characterized by a set of interrelated markets and different types of interacting agents, modeled according to a balance-sheet approach. The EURACE model encompasses price-making agents, a heterogeneous production sector and an explicit modeling of the household consumption behavior. It represents an integrated economy consisting of three sectors: the private sector (consumption goods, investment goods, and labor market), the financial sector (credit and financial markets), and the government sector. A next-generation macro-oriented ABM project that is underway is the Complexity Research Initiative for Systemic Instabilities (CRISIS), a multi-year endeavor which seeks to develop a large-scale model of the financial markets for the EU countries.²⁵

The dynamics of monetary aggregates are determined by commercial banks and central bank money supply. The dynamics of credit depend on the supply of credit from the banking system and the demand for credit from firms to finance their production activity. The model shows the emergence of endogenous business cycles which are mainly due to the interplay between the real economic activity and its financing through the credit market.

IV. Implementation of Agent-Based Models

Because an ABM is a simulation comprised of many agents and a complex of rules, it can be a challenge to build. There are many parameters to fit and many dimensions of output to monitor and evaluate.

Software Tools and Design

The computer code for many agent-based models requires a large number of parameters and moving parts. To some extent, that is inevitable for a complex model that seeks to describe the behavior of

²⁴ Deissenberg (2008a, 2008b). Cincotti and others (2010) provides a model along the lines of EURACE.

²⁵ Beinhocker and others (2011).

many interacting, heterogeneous agents. But the code can be made more manageable by adhering to standard object-oriented design, which lays out the interactions and parameters in a more understandable way and makes the models easier to follow. This allows other researchers to apply portions of the model in their own work and to enhance components of the model in the process.

Starting with the development of Swarm at the Santa Fe Institute in the mid-1990s, more than a dozen software toolkits have been developed to facilitate ABM, ranging from simple prototyping to scalable system tools.²⁶

Model Calibration and Validation²⁷

An ABM can generate the distributional dynamics for a modeled economy because it generates a period-by-period sample of the world it has specified. The calibration of parameters takes advantage of this by typically using some variant of generalized method of moments. For ABMs of financial markets, such as LeBaron (2006), comparisons are made between the kurtosis, autocorrelation, and volatility of the model and of actual data. Similarly, Thurner, Farmer, and Geanakoplos (2010) look at the results that their model of banks and hedge funds generates for return distributions. The housing model of Geanakoplos and others (2012) compares the output to actual housing foreclosures and other housing metrics. The macro model of Ashraf and Howitt uses manual search to find parameter values that lead model simulations to closely match the values of economic variables including unemployment rate, inflation volatility, and rate of job loss.²⁸

Model validation is about in- and out-of-sample accuracy, but a second objective in building the model is to keep the number of parameters manageable. As with any modeling approach, ABM faces a trade-off between descriptive accuracy and explanatory power. But as with the risk in over-fitting during calibration, the flexibility of ABMs makes this problem particularly acute.

The aphorism for dealing with this trade-off between descriptive capability and explanatory power that is aptly applied in the ABM sphere is KISS—“Keep It Simple, Stupid!”—by beginning with a simple or an existing, well-tested model and adding layers of detail bit by bit, or by starting with the most descriptive model and simplifying it as much as possible.²⁹ The extent to which the ABM can replicate reality plays the crucial role in deciding when the process should stop.

ABMs and Networks

The relationship between ABMs and networks is intuitive and obvious, and is broadly applied in the physical and biological sciences, especially in the area of statistical mechanics. However, it has not been treated explicitly in the financial arena.³⁰ In an ABM for finance, the agents interact through

²⁶ A list of these is available at <http://www2.econ.iastate.edu/tesfatsi/acecode.htm>.

²⁷ In the parlance of ABM, validation is making sure the right model is being built, verification is making sure it is being built correctly, and calibration is estimating the model parameters to fit the real world. For example, Balci (1997) states that model verification “deals with *building the model right*” while model validation “deals with *building the right model*”.

²⁸ ABMs are also validated by comparing their modeled behavior to experiments with real people. A wide-ranging survey of more recent research along these lines can be found in Duffy (2006).

²⁹ Axtell and others (1996) argue that the two hallmarks of cumulative disciplinary research are critical experiment and subsumption.

³⁰ One well-known paper that helps create a link between ABM and networks is Dodds and Watts (2005). Zhao and others (2010) make a stronger case for using ABMs to generate networks.

buying and selling assets, borrowing and lending to fund their business, and receiving and disseminating information. Each of these dimensions of interaction operates along a network, and at any point in time the network can be described to see, for example, how much each particular agent is borrowing from or lending to each other agent.

But whereas an ABM is dynamic, a network is a static view of the world. That is, at any point in the evolution of an ABM one can extract the network of the moment, but as the ABM progresses, the network may change. For the analysis of market dislocations and crises, this dynamic generation of networks is critical. In particular, for modeling crises and vulnerabilities, one question to ask is how would the network change if one node were to disappear (for example through a bankruptcy) or if a set of edges were to break (for example through a set of banks pulling their funding lines). These sorts of changes would typically propagate to other agents as those who are affected by the shock adapt and alter their interactions based on their decision heuristics. Carrying this process through to look at network dynamics can give an assessment of the weak points in the system, and whether the shock will propagate in a way that is of systemic concern.³¹ Although the dynamic might be viewed through a network representation, the machinery and analysis occur in the agent interactions and the network is a snapshot of that phenomenon at a given point in time.

V. Conclusion: Characteristics for an ABM focused on financial stability

ABMs can be valuable to the analysis of financial stability because the agents can represent the actual entities in the financial system. ABMs can identify entities by name (for example, they can have one agent that is Morgan Stanley, another that is Citadel), they can delve into the actual policies and procedure of agents in the face of various shocks, and they can describe each agent's initial conditions based on the data available to regulators.³² For example, how do banks alter their haircuts in the face of higher volatility in the collateral? How do hedge funds decide what to liquidate in the face of a forced sale? Do they tend to liquidate the most liquid assets?

The characteristics of an ABM directed toward threats to financial stability might include:

Key agents. The key agents for analyzing systemic risk are those that provide funding, those on the other side who are leverage, and those who are liquidity providers. The first of these can be represented by the money market funds and the banks, operating in the repo market. The second set can be presented by hedge funds. The third can be longer term, unleveraged investors, such as asset managers and pension funds.

³¹ For example, there is a growing literature on the robustness of the Internet to attacks, specifically understanding how the structure of servers across the globe transmitting information between themselves is robust to targeted attacks of select servers. One influential paper in this stream is Albert and others (1999).

³² By allowing various degrees of detail, one can look inside the box if necessary. Agents can be composed of more elementary agents in various forms of hierarchical organization. For example, a macroeconomic ABM might include a hierarchy of nested agents starting at the higher level of the national economy (financial, business, household, and government sectors); then other agents operating within as institutional entities within the sectors (commercial banks, insurance companies, and broker-dealers); then agents fulfilling functions within those organizations (the various trading desks and control functions). ABM thus provides flexibility to tailor the breadth and depth of the agents to the application at hand. Of course, integrating this into a single model increases the complexity of the model substantially.

Behavior rules for the agents. Keeping in mind that the agents in the model can represent actual entities in the financial system, the rules for the key agents in the face of a market shock can be determined by looking at their policies and procedures and through interviews with the decision makers in the firms, such as the Chief Risk Officer. Rules, for example, could consider how leverage is to be reduced as a hedge fund approaches its margin limit or risk limit; how a bank plans to react in terms of increasing its haircuts to its clients, reducing its inventory and thus market making posture, and liquidating assets; or how the participants in a short term funding market are expected to react to an increase in counterparty or collateral risk.³³

The accessibility to the behavior rules for the key participants in a crisis is notably different from the difficulty in specifying the behavior of the participants in the earliest application of ABMs to finance, namely traders in a market. The behavior rules for trading often are opaque, and are rarely specified, and even when they are, such as for quantitative strategies, they are highly proprietary. By contrast, many of the behavior rules for banks are set out in their policies and procedures, and more generally can be made available to regulators through their supervisory role.

Policy levers. ABMs should be structured to make the policy levers apparent. These levers include minimum haircuts and margin requirements for investors, capital and liquidity ratios for banks, and loan-to-value ratios for mortgages. The policy levers might also include “circuit breakers” which operate to slow down any liquidity and funding demand in order for these to move at a pace closer to that of the decision process for the key liquidity providers and sources of funding.

Shocks and vulnerabilities. The model should allow for the range of shocks that are typical in causing and propagating a crisis. These include a seizing up of liquidity; a fire sale in the face of forced deleveraging with the subsequent funding and liquidity effects; a sudden funding impairment, which is often brought on by a shock to real or perceived credit worthiness or liquidity; or in the extreme case, the failure of a firm posed as an exogenous event.

Data to set the initial conditions. Supervisory data are available for the various agents which can be used to set their portfolio positions, and hence the degree of their crowding and liquidity in the face of a forced liquidation; leverage; counterparty exposures; and risk, liquidity, and capital levels. For hedge funds, this data will come from Form PF, required by the SEC starting in 2012.

Policy Applications

It is self-evident that policy changes do not occur in a vacuum; that policy analysis must take into account the path of subsequent adjustments. ABMs are readily employable for policy experiments because the parameters in an ABM can be designed to mimic real-world key policy variables such as tax rates, subsidies, interest rates, money supply, and other key behavioral measures affecting individual incentives in growth, innovation, or technologically-related policies. The initial conditions in an ABM can represent initial endowments, and the interaction and behavioral rules employed by economic agents can represent alternative institutional, market, or industry structures, along with the differential clock speeds of regulatory lags and industry response.

³³ For hedge funds, we cannot know their portfolio construction rules, because those depend on their “secret sauce.” But in the face of a fire sale liquidation, it is likely that the objective will be to sell based on what is liquid as opposed to what is considered the best trade.

Policy makers can use ABMs to explore major policy changes that diverge far from current policy settings. An ABM with learning and adapting agents can provide a virtual policy experiment, exploring the importance of behavioral adjustments in a given situation. And the features of ABMs make them particularly well suited for analyzing an economy in extreme situations where standard empirical models are likely to fail.

From the preceding discussion, it might be apparent how ABM can be employed to get insights, if not specific quantitative results, for policy issues such as the following:

- ⤴ Optimal leverage. What are optimal leverage policies? What external conditions do these depend on? Should leverage be static or dynamic? If dynamic, how should this be determined? Are proper bank leverage levels fundamentally different from those of hedge funds and other financial institutions?
- ⤴ Crowding. What constitutes crowding in a market, what statistics can provide an indication of crowding, and how does crowding affect the dynamics of a forced liquidation?
- ⤴ Intervention. During the current crisis the government has made a series of interventions in the financial sector. Which types of interventions are effective and which are not?
- ⤴ Interbank network. How should the relationships among banks be structured to craft a more robust interbank network? What should be the borrowing and lending requirements? How does the network structure affect leverage? What are the effects if a particular node breaks?
- ⤴ Time granularity. At what frequency do banks and regulators need to monitor risk in order for risk control policies or interventions to be effective?
- ⤴ Data for ongoing risk analysis. What data should be collected in order to monitor the financial system, both to simulate the possible systemic risks and to create an “early warning system” of vulnerabilities and possible shocks?

ABMs have already been put to the task of addressing policy issues and related strategic issues. For example ABMs have been employed in industrial policy for the coffee market, the pharmaceutical industry, the evolution of the computer and semiconductor industry;³⁴ in market design, to analyze bidding behavior in market environments and simulate the impact of regulatory changes;³⁵ in fiscal policy, to analyze if a Tobin tax can reduce speculation;³⁶ and in social policy, to explore the impact of anti-crime policies.³⁷

Research Directions

The manifest application of ABMs within the OFR is to create models that help assess threats to U.S. financial stability. This objective is tightly bound to the specifics of our world—the banks and asset managers, along with their balance sheets, and risk management policies and procedures for meeting regulatory requirements; the regulatory agencies, along with their areas of responsibility, and even the propensities of those who are in charge in those agencies. This is a different and more detailed starting

³⁴ For the coffee market, see Midgley and others (1997); for pharmaceuticals, see Malerba and Orsenigo (2002); for semiconductors, see Malerba and others (2008).

³⁵ See Duffy and Unver (2008). Neugart (2008) looks at the impact of labor market policies in an ABM where firms in different sectors require workers with different skills. Darley and Outkin (2007) describe an ABM employed by the NASDAQ to explore the effects of decimalization before decimalization was implemented in 2001.

³⁶ See Mannaro and others (2008).

³⁷ See Wilhite (2008).

point than for much financial research, which seeks to develop theories and models that apply across many possible worlds.

This survey has laid out the ways in which ABMs may be a valuable tool for this task of modeling the agents of our financial system, of modeling emergent phenomenon and disequilibria, and of mapping out the effect of shocks on the financial system, and thus of depicting an ensuing crisis.

Although it may be a promising approach to deal with financial vulnerability and crises, ABMs remain outside the mainstream of economic and financial research. One reason is that they present a clear challenge to many of the methods and machinery built up around the neoclassical school. Another is that because agent-based modeling does not represent the world in a set of clean and elegant equations, it does not lend itself to the standard mode of academic publication. Axelrod (2006) has noted the problems of ABMs in mainstream economics: “ABM can be a hard sell. Since most formal theorists equate models with mathematical models, it is not surprising that some of them are hard to convince about the appropriateness and value of an agent-based simulation.”

Further research could develop ABMs to capture the systemic effects of leverage, funding, liquidity and interconnectedness in the financial system. Research could also develop methods for validation, verification, and calibration, and tools for assessing ABM performance so that the efficacy of alternative models can be compared and so that policy makers can determine how much confidence to put into their results. Because there are a number of groups pursuing ABMs for this purpose, it will be helpful to arrive at a common platform—a common set of software tools—to allow different groups to run one another’s models and share modules.

Bibliography

- Albert, Reka, Hawoong Jeong, and Albert-Laszlo Barabasi. "The Diameter of the World Wide Web." *Nature*, (1999): 130-131.
- Alfaro, Rodrigo, and Mathias Drehmann. "Macro Stress Tests and Crises: What Can We Learn?" *BIS Quarterly Review* (December 2009): 29-41.
- Arciero, Luca, Claudia Biancotti, Leandro D'Aurizio, and Claudio Impenna. "Exploring Agent-Based Methods for the Analysis of Payment Systems: A Crisis Model for StarLogo TNG." *Journal of Artificial Societies and Social Simulation* 12, no. 1 (2009).
- Ashraf, Quamrul, and Peter Howitt. "How Inflation Affects Macroeconomic Performance: An Agent-Based Computational Investigation." Working Paper, Brown University, 2008.
- Axelrod, Robert. (1997). *The Complexity of Cooperation: Agent-Based Models of Competition and Collaboration*. Princeton, NJ: Princeton University Press.
- Axelrod, Robert. "Agent-based modeling as a bridge between disciplines." *Handbook of Computational Economics* 2, no. 2 (2006): 1565-1584. North-Holland.
- Axtell, Robert, Robert Axelrod, Joshua Epstein, and Michael Cohen. "Aligning Simulation Models: A Case Study and Results." *Computational and Mathematical Organization*, (1996): 123-141.
- Bajardi, Paolo, Chiara Poletto, Duygu Balcan, Hao Hu, Bruno Gonçalves, Jose J. Ramasco, Daniela Paolotti, Nicola Perra, Michele Tizzoni, Wouter Van den Broeck, Vittoria Colizza, and Alessandro Vespignani. "Modeling vaccination campaigns and the Fall/Winter 2009 activity of the new A(H1N1) influenza in the Northern Hemisphere." *Emerging Health Threats Journal*, (2009).
- Balci, Osman. "Verification, validation and accreditation of simulation models." *Proceedings of the 1997 Winter Simulation Conference*, (1997): 135-141.
- Beinhocker, Eric, Doyne Farmer, Olaf Bochmann, Alessandro Gobbi, and Jacob Grazzini. Complexity Research Initiative for Systemic Instabilities Project Summary. (2011).
- Bisias, Dimitrios, Mark Flood, Andrew W. Lo, and Stavros Valavanis. "A Survey of Systemic Risk Analytics." Working Paper #0001, Washington, DC: Office of Financial Research, January 15, 2012.
- Bonabeau, Eric. "Agent-based modeling: Methods and Techniques for Simulating Human Systems." *Proceedings of the National Academy of Sciences*, (2002): 7280-7287.
- Bookstaber, Richard. "Understanding and Monitoring the Liquidity Crisis Cycle." *Financial Analysts Journal*, (2000): 17-22.

- Bookstaber, Richard, and Nassim Taleb. "Testimony." *The Risks of Financial Modeling: VaR and the Economic Meltdown, Hearing before the U. S. House of Representatives, Committee on Science and Technology, Subcommittee on Investigations and Oversight*. U.S. House of Representatives, September 10, 2009. 111th Congress.
- Borio, Claudio, Mathias Drehmann, and Kostas Tsatsaronis. "Stress Testing Macro Stress Testing: Does It Live up to Expectations?" BIS Working Papers no. 369, Basel: Bank for International Settlements, January 2012.
- Brunnermeier, Markus K., and Lasse Heje Pedersen. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22, no. 6 (2009): 2201–2238.
- Cincotti, Silvano, Marco Raberto, and Andrea Teglio. "Credit Money and Macroeconomic Instability in the Agent-based Model and Simulator Eurace." *Economics: The Open-Access, Open-Assessment E-Journal* 4, (2010): 2010-26.
- Crutchfield, James. "The Calculi of Emergence: Computation, Dynamics, and Induction." *Physica D*, (1994): 11-54.
- Darley, Vince, and Alexander Outkin. (2007). *Nasdaq Market Simulation: Insights on a Major Market from the Science of Complex Adaptive Systems*. World Scientific.
- Deissenberg, Cristophe, Sander van der Hoog, and Herbert Dawid. "EURACE: A Massively Parallel Agent-based Model of the European Economy." *Applied Mathematics and Computation*, (2008a): 541-552.
- Deissenberg, Christophe, Sander van der Hoog, and Herbert Dawid. "Production and Finance in EURACE." In K. Schredelseker, and F. Hauser, *Complexity and Artificial Markets, Lecture Notes in Economics and Mathematical Systems*, (2008b): 147 – 158. Springer.
- DeMarzo, Peter, Dimitri Vayanos, and Jeffrey Zwiebel. "Persuasion Bias, Social Influence, and Unidimensional Opinions." *Quarterly Journal of Economics*, (2003): 909-968.
- Dodds, Peter, and DJ Watts. "A generalized model of social and biological contagion." *Journal of Theoretical Biology*, (2005): 587-604.
- Duffy, John. "Agent-Based Models and Human-Subject Experiments." In L. Tesfatsion and K. L. Judd (Eds.), *Handbook of Computational Economics* 2, (2006): 950-1005. Elsevier.
- Duffy, John, and M. Utku Unver. "Internet auctions with artificial adaptive agents: A study on market design." *Journal of Economic Behavior and Organization*, (2008): 394-417.
- Farmer, Doyne, and John Geanakoplos. "The Virtues and Vices of Equilibrium and the Future of Financial Economics." Cowles Foundation Paper No. 1274, 2009.

- Galbiati, Marco, and Kimmo Soramäki. “An Agent-based Model of Payment Systems.” Working Paper no. 352, London: Bank of England, August 2008.
- Geanakoplos, John, Robert Axtell, Doyne Farmer, Peter Howitt, Benjamin Conlee, Jonathan Goldstein, Matthew Hendrey, Nathan Palmer, and Chun-Yi Yang. “Getting at Systemic Risk via an Agent-Based Model of the Housing Market.” Papers and Proceedings, ASSA Annual Meeting. American Economic Association, 2012, forthcoming.
- Gilbert, Nigel, John C. Hawkesworth, and Paul A. Swinney. “An Agent-Based Model of the English Housing Market.” Association for the Advancement of Artificial Intelligence, Spring Symposium, Stanford, CA, March 23-25, (2009): 30–35.
- Glass, Robert, Laura Glass, and Walter Beyeler. “Local Mitigation Strategies for Pandemic Influenza.” Sandia National Laboratory Staff Report, SAND Number 2005-7955J, 2005.
- Golub, Benjamin, and Matthew Jackson. “Naïve Learning in Social Networks: Convergence, Influence and Wisdom of Crowds.” *American Economic Journal Microeconomics*, (2010): 112-149.
- Granovetter, Mark. “Threshold models of collective behavior.” *American Journal of Sociology*, (1978): 1420-1443.
- Gigerenzer, Gerd, Ralph Hertwig, and Thorsten Pachur. (Eds.) (2011). *Heuristics: The Foundations of Adaptive Behavior*. New York: Oxford University Press.
- Helbing, Dirk, Illes Farkas, and Tamas Vicsek. “Simulating dynamical features of escape panic.” *Nature*, (2000): 487-490.
- Holland, John. (1975). *Adaptation in Natural and Artificial Systems*. Ann Arbor: University of Michigan Press.
- Holland, John. “Studying Complex Adaptive Systems.” *Journal of Systems Science and Complexity*, (2006): 1-8.
- International Monetary Fund (IMF). *Iceland: Financial System Stability Assessment Update*. IMF Country Report no. 08/368, Washington: IMF, August 19, 2008.
- Jackson, Matthew, and Leeat Yariv. “Diffusion of Behavior and Equilibrium Properties in Network Games.” *American Economic Review*, (2007): 92-98.
- Judd, Kenneth. “Computationally Intensive Analyses in Economics.” In L. Tesfatsion and K. L Judd (Eds.), *The Handbook of Computational Economics II*, (2006): 881- 893. Elsevier.
- Kaufman, Stuart. (1996). *At Home in the Universe: The Search for the Laws of Self-Organization and Complexity*. New York: Oxford University Press, USA.
- Keim, Brandon. “Amazing Startling Flocks Are Flying Avalanches.” *Wired Science*, June 16, 2010. <http://www.wired.com/wiredscience/2010/06/starling-physics> (accessed February 15, 2012).

- Keim, Brandon. "The Startling Science of a Starling Murmuration." *Wired Science*, November 8, 2011. <http://www.wired.com/wiredscience/2011/11/starling-flock> (accessed February 15, 2012).
- Khandani, Amir, Andrew Lo, and Robert Merton. "Systemic Risk and the Refinancing Ratchet Effect." Harvard Business School Working Paper Series, No. 10-023, 2010.
- Latané, Bibb. "Dynamic Social Impact: Robust Predictions from Simple Theory." In R. Hegselmann, U. Mueller, and K. Troitzsch (Eds.), *Modeling and Simulation in the Social Sciences from a Philosophy of Science Point of View*, (1996): 287-310. Netherlands: Kluwer Academic Publishers.
- LeBaron, Blake. "Agent-Based Financial Markets: Matching Stylized Facts with Style." In D. Colander (Ed.), *Post Walrasian Macroeconomics*, (2006): 221-235. New York: Cambridge University Press.
- LeBaron, Blake, W. Brian Arthur, and Richard Palmer. "Time series properties of an artificial stock market." *Journal of Economic Dynamics and Control*, (1999): 1487-1516.
- Malerba, Franco, and Luigi Orsenigo. "Innovation and Market Structure in the Dynamics of the Pharmaceutical Industry and Biotechnology: Towards a History-Friendly Model." *Industrial and Corporate Change*, (2002): 667-703.
- Malerba, Franco, Richard Nelson, Luigi Orsenigo, Sidney Winter. "Public Policies and Changing Boundaries of Firms in a History-Friendly Model of the Co-evolution of the Computer and Semiconductor Industries." *Journal of Economic Behavior and Organization* 67, no. 2 (2008): 335-380.
- Mannaro, Katuscia, Michele Marchesi, and Alessio Setzu. "Using an artificial financial market for assessing the impact of Tobin-like transaction taxes." *Journal of Economic Behavior and Organization*, (2008): 445-462.
- Midgley, David, Robert Marks, and Lee Cooper. "Breeding Competitive Strategies." *Management Science*, (1997): 257- 275.
- Miller, John. "Active Nonlinear Tests (ANTs) of Complex Simulation Models." *Management Science*, (1998): 820-830.
- Neugart, Michael. "Labor Market Policy Evaluation with ACE." *Journal of Economic Behavior and Organization*, (2008): 418-430.
- Office of Financial Research (OFR). *Annual Report*. Washington, DC: OFR, 2012.
- Page, Scott. "Path Dependence." *Quarterly Journal of Political Science*, (2006): 87-115.

- Reynolds, Craig. "Flocks, Herds, and Schools: A Distributed Behavioral Model." *Computer Graphics*, (1987): 25-34.
- Schelling, Thomas. "Dynamic models of segregation." *Journal of Mathematical Sociology*, (1971): 143-186.
- Solow, Robert. Building a Science of Economics for the Real World. *Prepared statement to the House Committee on Science and Technology Subcommittee on Investigations and Oversight*. U.S House of Representatives, July 20, 2010. 111th Congress. Washington, DC: Government Printing Office, 2010.
- Thurner, Stefan. "Systemic Financial Risk: Agent Based Models to Understand the Leverage Cycle on National Scales and Its Consequences." Working Paper IFP/WK/FGS(2011)1, Organisation for Economic Co-operation and Development International Futures Programme, January 14, 2011.
- Thurner, Stefan, J. Doyne Farmer, and John Geanakoplos. "Leverage Causes Fat Tails and Clustered Volatility." Cowles Foundation Discussion Paper, No. 1745, January 2010.
- Toroczkai, Zoltan, and Stephan Eubank. "Agent-based modeling as a decision-making tool." *The Bridge*, (2005): 22-27.
- Wilensky, Uri. "NetLogo Home Page." Netlogo, 1999. <http://ccl.northwestern.edu/netlogo/>. Evanston, IL: Center for Connected Learning and Computer-Based Modeling, Northwestern University.
- Wilhite, Allen. "Crime, Protection and Incarceration." *Journal of Economic Behavior and Organization*, (2008): 481-494.
- Yeh, Chia-Hsuan. "The Effects of Intelligence on Price Discovery and Market Efficiency." *Journal of Economic Behavior and Organization* 68, no. 3-4 (2007a): 613-625.
- Yeh, Chia-Hsuan. "The Role of Intelligence in Time Series Properties." *Computational Economics* 30, no. 2 (2007b): 95-123.
- Zhao, Zhenyuan, J.P. Calderón, Chen Xu, Dan Fenn, Didier Sornette, Riley Crane, Pak Ming Hui, Neil F. Johnson. "Effect of Social Group Dynamics on Contagion." *Physical Review E*, (2010).