

Duncan Foley is a Leo Model Professor at The New School for Social Research

The Economy Needs Agent-Based Modelling

By J. Doyne Farmer & Duncan Foley

In today's high-tech age, one naturally assumes that US President Barack Obama's economic team and its international counterparts are using sophisticated quantitative computer models to guide us out of the current economic crisis. They are not. The best models they have are of two types, both with fatal flaws. Type one is econometric: empirical statistical models that are fitted to past data. These successfully forecast a few quarters ahead as long as things stay more or less the same, but fail in the face of great change. Type two goes by the name of 'dynamic stochastic general equilibrium'. These models assume a perfect world, and by their very nature rule out crises of the type we are experiencing now.

As a result, economic policy-makers are basing their decisions on common sense, and on anecdotal analogies to previous crises such as Japan's 'lost decade' or the Great Depression (see Nature 457, 957; 2009). The leaders of the world are flying the economy by the seat of their pants.

The economy needs agent-based modeling This is hard for most non-economists to believe. Aren't people on Wall Street using fancy mathematical models? Yes, but for a completely different purpose: modelling the potential profit and risk of individual trades. There is no attempt to assemble the pieces and understand the behaviour of the whole economic system.

There is a better way: agent-based models. An agent-based model is a computerized simulation of a number of decision-makers (agents) and institutions, which interact through prescribed rules. The agents can be as diverse as needed — from consumers to policy-makers and Wall Street professionals — and the institutional structure can include everything from banks to the government. Such models do not rely on the assumption that the economy will move towards a predetermined equilibrium state, as other models do. Instead, at any given time, each agent acts according to its current situation, the state of the world around it and the rules governing its behaviour. An individual consumer, for example, might decide whether to save or spend based on the rate of inflation, his or her current optimism about the future, and behavioural rules deduced from psychology experiments. The computer keeps track of the many agent interactions, to see what happens over time. Agent-based simulations can handle a far wider range of nonlinear behaviour than conventional equilibrium models. Policy-makers can thus simulate an artificial economy under different policy scenarios and quantitatively explore their consequences.

Why is this type of modelling not well-developed in economics? Because of historical choices made to address the complexity of the economy and the importance of human reasoning and adaptability. The notion that financial economies are complex systems can be traced at least as far back as Adam Smith in the late 1700s. More recently John Maynard Keynes and his followers attempted to describe and quantify this complexity based on historical patterns. Keynesian economics enjoyed a heyday in the decades after the Second World War, but was forced out of the mainstream after failing a crucial test during the mid-seventies. The Keynesian predictions suggested that inflation could pull society out of a recession; that, as rising prices had historically stimulated supply, producers would respond to the rising prices seen under inflation by increasing production and hiring more workers. But when US policy-makers increased the money supply in an attempt to stimulate employment, it didn't work — they ended up with both high inflation and high unemployment, a miserable state called 'stagflation'. Robert Lucas

and others argued in 1976 that Keynesian models had failed because they neglected the power of human learning and adaptation. Firms and workers learned that inflation is just inflation, and is not the same as a real rise in prices relative to wages.

The cure for macroeconomic theory, however, may have been worse than the disease. During the last quarter of the twentieth century, 'rational expectations' emerged as the dominant paradigm in economics. This approach assumes that humans have perfect access to information and adapt instantly and rationally to new situations, maximizing their long-run personal advantage. Of course real people often act on the basis of overconfidence, fear and peer pressure — topics that behavioural economics is now addressing.

But there is a still larger problem. Even if rational expectations are a reasonable model of human behaviour, the mathematical machinery is cumbersome and requires drastic simplifications to get tractable results. The equilibrium models that were developed, such as those used by the US Federal Reserve, by necessity stripped away most of the structure of a real economy. There are no banks or derivatives, much less sub-prime mortgages or credit default swaps — these introduce too much nonlinearity and complexity for equilibrium methods to handle. When it comes to setting policy, the predictions of these models aren't even wrong, they are simply non-existent (see *Nature* 455, 1181; 2008).

Agent-based models potentially present a way to model the financial economy as a complex system, as Keynes attempted to do, while taking human adaptation and learning into account, as Lucas advocated. Such models allow for the creation of a kind of virtual universe, in which many players can act in complex — and realistic — ways. In some other areas of science, such as epidemiology or traffic control, agent-based models already help policy-making.

There are some successful agent-based models of small portions of the economy. The models of the financial market built by Blake LeBaron of Brandeis University in Waltham, Massachusetts, for example, provide a plausible explanation for bubbles and crashes, reproducing liquidity crises and crashes that never appear in equilibrium models. Rob Axtell of George Mason University in Fairfax, Virginia, has devised firm dynamics models that simulate how companies grow and decline as workers move between them. These replicate the power-law distribution of company size that one sees in real life: a very few large firms, and a vast number of very small ones with only one or two employees.

Other promising efforts include the credit-sector model of Mauro Gallegati's group at the Marche Polytechnic University in Ancona, Italy, and the monetary model developed by Robert Clower of the University of South Carolina in Columbia and Peter Howitt of Brown University in Providence, Rhode Island. These models are very useful, but their creators would be the first to say that they provide only a tentative first step.

To see in more detail how an agent-based model works, consider the model that one of us (Farmer) has developed with Stefan Thurner of the University of Vienna and John Geanakoplos of Yale University to explore how leverage affects fluctuations in stock prices (published in a Santa Fe Institute working paper). Leverage, the investment of borrowed funds, is measured as the ratio of total assets owned to the wealth of the borrower; if a house is bought with a 20% down-payment the leverage is five. There are four types of agents in this model. 'Noise traders', who trade more or less at random, but are slightly biased toward driving prices towards a fundamental value; hedge funds, which hold a stock when it is under-priced and otherwise hold cash; investors who decide whether to invest in a hedge fund; and a bank that can lend money to the hedge funds, allowing them to buy more stock. Normally, the presence of the hedge funds damps volatility, pushing the stock price towards its fundamental value. But, to contain their risk, the banks cap leverage at a predetermined maximum value. If the price of the stock drops while a fund is fully leveraged, the fund's wealth plummets and its leverage increases; thus the

fund has to sell stock to pay off part of its loan and keep within its leverage limit, selling into a falling market.

This agent-based model shows how the behaviour of the hedge funds amplifies price fluctuations, and in extreme cases causes crashes. The price statistics from this model look very much like reality. It shows that the standard ways banks attempt to reduce their own risk can create more risk for the whole system.

Previous models of leverage based on equilibrium theory showed qualitatively how leverage can lead to crashes, but they gave no quantitative information about how this affects the statistical properties of prices. The agent approach simulates complex and nonlinear behaviour that is so far intractable in equilibrium models. It could be made more realistic by adding more detailed information about the behaviour of real banks and funds, and this could shed light on many important questions. For example, does spreading risk across many financial institutions stabilize the financial system, or does it increase financial fragility? Better data on lending between banks and hedge funds would make it possible to model this accurately. What if the banks themselves borrow money and use leverage too, a process that played a key role in the current crisis? The model could be used to see how these banks might behave in an alternative regulatory environment.

Agent-based models are not a panacea. The major challenge lies in specifying how the agents behave and, in particular, in choosing the rules they use to make decisions. In many cases this is still done by common sense and guesswork, which is only sometimes sufficient to mimic real behaviour. An attempt to model all the details of a realistic problem can rapidly lead to a complicated simulation where it is difficult to determine what causes what. To make agent-based modelling useful we must proceed systematically, avoiding arbitrary assumptions, carefully grounding and testing each piece of the model against reality and introducing additional complexity only when it is needed. Done right, the agent-based method can provide an unprecedented understanding of the emergent properties of interacting parts in complex circumstances where intuition fails.

A thorough attempt to understand the whole economy through agent-based modelling will require integrating models of financial interactions with those of industrial production, real estate, government spending, taxes, business investment, foreign trade and investment, and with consumer behaviour. The resulting simulation could be used to evaluate the effectiveness of different approaches to economic stimulus, such as tax reductions versus public spending.

Such economic models should be able to provide an alternative tool to give insight into how government policies could affect the broad characteristics of economic performance, by quantitatively exploring how the economy is likely to react under different scenarios. In principle it might even be possible to create an agent-based economic model capable of making useful forecasts of the real economy, although this is ambitious.

Creating a carefully crafted agent-based model of the whole economy is, like climate modelling, a huge undertaking. It requires close feedback between simulation, testing, data collection and the development of theory. This demands serious computing power and multi-disciplinary collaboration among economists, computer scientists, psychologists, biologists and physical scientists with experience in large-scale modelling. A few million dollars — much less than 0.001% of the US financial stimulus package against the recession — would allow a serious start on such an effort. Given the enormity of the stakes, such an approach is well worth trying.