Meltdown modelling

Could agent-based computer models prevent another financial crisis? **Mark Buchanan** reports.

It's 2016, and experts at a US government facility have detected a threat to national security. A screen on the wall maps the world's largest financial players — banks, governments and hedge funds — as well as the web of loans, ownership stakes and other legal claims that link them. High-powered computers have been using these enormous volumes of data to run through scenarios that flush out unexpected risks. And this morning they have triggered an alarm.

Flashing orange alerts on the screen show that a cluster of US-based hedge funds has unknowingly taken large ownership positions in similar assets. If one of the funds should have to sell assets to raise cash, the computers warn, its action could drive down the assets' value and force others to start selling their own holdings in a self-amplifying downward spiral. Many of the funds could be bankrupt within 30 minutes, creating a threat to the entire financial system. Armed with this information, financial authorities step in to orchestrate a controlled elimination of the dangerous tangle.

Alas, this story is likely to remain fiction. No government was able to carry out any such 'war room' analyses as the current financial crisis emerged, nor does the capability exist today. Yet a growing number of scientists insist that something like it is needed if society is to avoid similar crises in future.

Financial regulators do not have the tools they need to predict and prevent meltdowns,

says physicist-turned-sociologist Dirk Helbing of the Swiss Federal Institute of Technology Zurich, who has spent the past two decades modelling large-scale human systems such as urban traffic or pedestrian flows. They can do a good job of tracking an economy using the statistical measures of standard econometrics, as long as the influences on the economy are independent of each other, and the past remains a reliable guide to the future. But the recent financial collapse was a 'systemic' meltdown, in which intertwined breakdowns in housing, banking and many other sectors conspired to destabilize the system as a whole. And the past has been anything but a reliable guide of late: witness how US analysts were led astray by decades of data suggesting that housing values would never simultaneously fall across the nation.

Likewise, economists can get reasonably good insights by assuming that human behaviour leads to stable, self-regulating markets, with the prices of stocks, houses and other things never departing too far from equilibrium. But 'stability' is a word few would use to describe the chaotic markets of the past few years, when complex, nonlinear feedbacks fuelled the boom and bust of the dot-com and housing bubbles, and when banks took extreme risks in pursuit of ever higher profits.

In an effort to deal with such messy realities, a few economists — often working with physicists and others outside the economic mainstream — have spent the past decade or so exploring 'agent-based' models that make only minimal assumptions about human behaviour or inherent market stability (see page 685). The idea is to build a virtual market in a computer and populate it with artificially intelligent bits of software — 'agents' — that interact with one another much as people do in a real market. The computer then lets the overall behaviour of the market emerge from the actions of the individual agents, without presupposing the result.

Agent-based models have roots dating back to the 1940s and the first 'cellular automata', which were essentially just simulated grids of on-off switches that interacted with their nearest neighbours. But

they didn't spark much interest beyond the physical-science community until the 1990s, when advances in computer power began to make realistic social simulations more feasible. Since then they have found increasing use in problems such as traffic flow and the spread of infectious diseases (see page 687). Indeed, points out Helbing, agent-based models are the social-science analogue of the computational simulations now routinely used elsewhere in science to explore complex nonlinear processes such as the global climate.

"We have had a massive failure of the dominant economic model."
— Eric Weinstein

ILLUSTRATIONS BY JESSE LEFKOWITZ

That is why he is eager to bring social and physical scientists together to develop computational 'wind tunnels' that would allow regulators to test policies before putting them into practice. "The idea is to invest a lot in science," he says, "and thereby save hundreds of times as much by avoiding or mitigating future crises."

Just more theory?

That notion is a tough sell among mainstream economists, many of whom are less than thrilled by offers of outside help. "After any crisis," says Paul Romer of Stanford University, California, a leading researcher in the economics of innovation, "you hear recommendations to recruit scientists from other fields who can purge economics and finance of ideology and failed assumptions. But we should ask if there is any evidence that more theory, developed by people who don't have domain experience, is the key to scientific progress in this area."

Others think some fresh thinking is long overdue. "We have had a massive failure of the dominant economic model," says Eric Weinstein, a physicist working in mathematical finance for the Natron Group, a hedge fund in New York, "and we're trying to find the right people to deal with this failure. At least some of those people are likely to be unfamiliar voices and come from other parts of science."

At least some economists agree. The meltdown has shown that regulatory policies have to cope with far-from-equilibrium situations, says economist Blake LeBaron of Brandeis University in Waltham, Massachusetts. "Even fairly simple agent-based models can be used as thought experiments to see if there is something that hasn't been considered by the policy-makers."

LeBaron has spent the past decade and a half working with colleagues, including a number

of physicists, to develop an agent-based model of the stock market. In this model, several hundred agents attempt to profit by buying and selling stock, basing their decisions on patterns they perceive in past stock movements. Because the agents can learn from and respond to emerging market behaviour, they often shift their strategies, leading other agents to change their behaviour in turn. As a result, prices don't settle down into a stable equilibrium, as standard economic theory predicts. Much as in the real stock market, the prices keep bouncing up and down erratically, driven by an ever-shifting ecology of strategies and behaviours.

Nor is the resemblance just qualitative, says LeBaron. Detailed analyses of the agent-based model show that it reproduces the statistical features of real markets, especially their susceptibility to sudden, large price movements. "Traditional models do not go very far in explaining these features," LeBaron says.

Another often-cited agent-based model got its start in the late 1990s, as the NASDAQ stock exchange in New York was planning to stop listing its stock prices as fractions such as $12\frac{1}{4}$ and instead list them as decimals. The goal was to improve the accuracy of stock prices, but the change would also allow prices to move by smaller increments, which could affect the strategies followed by brokers with unknown consequences for the market as a whole. So before making this risky change, NASDAQ chief Mike Brown hired BiosGroup, a company based in Santa Fe, New Mexico, to develop an agent-based model of the market to test the idea.

"Over ten years on the NASDAQ Board," says Brown, "I grew increasingly disappointed in our approach to studying the consequences of proposed market regulations, and wanted to try something different."

Once the model could reproduce price fluctuations in a mathematically accurate way, NASDAQ used it as a market wind tunnel. The tests revealed that if the stock exchange reduced its price increment too much, traders would be able to exploit strategies that would make them quick profits at the expense of overall market efficiency. Thus, when the exchange went ahead with the changeover in 2001, it was able to take steps to counter this vulnerability.

Agent-based models are also being used elsewhere in the private sector. For example, the consumer-products giant Procter & Gamble of Cincinnati, Ohio, has used agent-based models to optimize the flow of goods through its network of suppliers, warehouses and stores. And Southwest Airlines of Dallas, Texas, has used agent-based models for routing cargo.

Despite such successes, however, financial regulators such as the US Securities and Exchange Commission (SEC) still don't use agent-based models as practical tools. "When the SEC changes trading rules, it typically has either flimsy or modest support from econometric evidence for the action, or else no empirical evidence and the change is driven by ideology," claims computational social scientist Rob Axtell of George Mason University in Fairfax, Virginia. "You have to wonder why Mike Brown is doing this, while the SEC isn't."

Risk of the new

A big part of the answer is that agent-based models remain at the fringe of mainstream economics, and most economists continue to prefer conventional mathematical models. Many of them argue that agent-based models haven't had the same level of testing.

Another problem is that an agent-based model of a market with many diverse players and a rich structure may contain many variable parameters. So even if its output matches reality, it's not always clear if this is because of careful

tuning of those parameters, or because the model succeeds in capturing realistic system dynamics. That leads many economists and social scientists to wonder whether any such model can be trusted. But agent-based enthusiasts counter that conventional economic models

also contain many tunable parameters and are therefore subject to the same criticism.

Familiarity wins out, notes Chester Spatt, former chief economist at the SEC. Regulators feel duty-bound to adhere to generally accepted and well-vetted techniques, he says. "It would be problematic for the rule-making process to use methods whose foundation or applicability were not established."

Still, agent-based techniques are beginning to enter the regulatory process. For example, decision-makers in Illinois and several other US states use computational models of complex electricity markets. They want to avoid a repeat of the disaster in California in 2000, when Enron and other companies, following market deregulation, were able to manipulate energy supplies and prices for enormous profit. Rich computational models have made it possible to test later market designs before putting them in place.

"We've had a lot of success in developing these models," says economist Leigh Tesfatsion of Iowa State University in Ames, who has led the development of an open-source agent-based model known as the AMES Wholesale Power Market Test Bed. "It has worked

"We still implement new economic measures without any prior testing."

— Dirk Helbing

because we've focused on all the details of the real situation and can address questions that policy-makers really care about," she says.

Other models have successfully simulated financial markets. At Yale University, for example, economist John Geanakoplos, working with physicists Doyne Farmer of the Santa Fe Institute and Stefan Thurner of the Medical University of Vienna, has constructed an agent-based model exploring the systemic consequences of massive borrowing by hedge funds to finance their investments. In their simulations, the funds frequently get locked into a self-amplifying spiral of losses (see page 685) — much as real-world hedge funds did after August 2007.

At the University of Genoa in Italy, meanwhile, Silvano Cincotti and his colleagues are creating an agent-based model of the entire European Union economy. Their model includes markets for consumer goods and financial assets, firms that interact with banks to obtain loans, and banks that compete with one another by offering different interest rates. Based on real economic data, the model currently represents some 10 million households, 100,000 firms and about 100 banks, all of which can learn and change their strategies if they find more profitable ways of doing business.

"We hope that these simulations will have an outstanding impact on the economic-policy capabilities of the European Union," says Cincotti, "and help design the best policies on an empirical basis."

This is the kind of ambition that has inspired Helbing. He doesn't pretend to be an economic modeller himself: since the early 1990s his own work has focused on simulations of human behaviour in relatively small groups — how traffic ebbs and flows through a road network, for example, or how crowds crush towards a door in a panic situation — as well as on experiments to test his predictions with real data. But that work has given Helbing a keen appreciation for the way complex collective phenomena can emerge from even the simplest individual interactions. If pedestrians can organize themselves into smoothly flowing streams just by trying to walk

through a crowded shopping centre — as he has shown they do — just imagine how much richer the emergent phenomena must be in a group the size of a national economy.

Crisis logic

That observation acquired fresh force for Helbing after last year's global financial meltdown made it clear that a regulatory system based on conventional economic theory had failed.

"It's remarkable," he says, "that while any new technical device or medical drug has extensive testing for efficiency, reliability and safety before it ever hits the market, we still implement new economic measures without any prior testing."

To get around this impasse, he says, researchers need to reimagine the social and economic sciences on a larger scale. "I imagine experts

from different fields meeting in one place for extended periods of time," he says, "so that their complementary knowledge could 'collide', creating new ideas, much as particle colliders create new kinds of particles." Ultimately, such an effort would bring together social scientists, economists, physicists, ecologists, computer scientists and engineers in a network of large centres for socioeconomic data mining and crisis forecasting, as well as in supercomputer centres for social simulation and wind-tunnel-like testing of policy.

That is a large ambition, Helbing admits — especially as he has only recently got tentative approval for a one-year grant from the European Commission to develop the idea. But now, in the aftermath of the meltdown, may be the time to start.

Axtell endorses that view. "Left to their own devices," he says, "academic macroeconomists will take a generation to make this transition. But if policy-makers demand better models, it can be accomplished much more quickly."

"The revolution has to begin here," agrees Weinstein, who helped organize a meeting in May at the Perimeter Institute for Theoretical Physics in Waterloo, Canada, that assembled the kind of interdisciplinary mix of experts that Helbing envisions. "And I think ideas from physics and other parts of science really have a chance to catalyse something remarkable." ■

Mark Buchanan is a science writer based in Cambridge, UK. After writing this story, he was involved in reviewing grant proposals on the topic of agent-based modelling.

See Editorial, page 667, and Opinion, pages 685 and 687.

"Experts' complementary knowledge could 'collide', creating new knowledge."
— Dirk Helbing