What do you get when you throw too much financial information at the irrational human brain? A surprising insight into global economics, says **Mark Buchanan**

It's the economy, stupid

End-of-year celebrations at the Chicago Mercantile Exchange. But how much rational reasoning supported the financial traders' decisions?

YOU are, on the whole, quite irrational. Often you act on instinct, rather than cold, logical reasoning. Putting it simply, you're human. And your humanity is proving to be a problem.

To economists, that is. Orthodox economics has long operated under the notion that people make economic decisions in a perfectly rational and calculated way. One of the results of this is that, in orthodox economic theory, markets are inherently unpredictable. That's because stocks, for example, should always take values based on the company's realistic prospects for making profits and paying share dividends in the future. Stock prices should only go up or down when those prospects genuinely change – through technological breakthroughs, the emergence of competitors and so on. And because these can't be predicted, market prediction should be impossible.

A tidy picture – except that detailed studies have demonstrated systematic shortcomings in this theory. "Everyone now agrees that stock prices are at least partly predictable," says economist Richard Thaler of the University of Chicago. "Real financial markets do not resemble the ones we would imagine if we only read finance textbooks."

Now a new breed of financial theorists is claiming this is because of the vagaries of human decision-making. While rationality might allow theorists to produce elegant equations with all the apparent rigour of mathematical physics, it is fundamentally inaccurate. Economics becomes a lot clearer, they say, when considered in a new light: as an exploration of human behaviour in all its realistic complexity, rational or otherwise. Working within this vastly richer domain, they have already explained market characteristics that have puzzled traditional theorists for decades, and are beginning to map out a radical and long-overdue revision of economic theory.

Ten years ago, in a somewhat frivolous thought experiment, economist Brian Arthur of the Santa Fe Institute in New Mexico demonstrated the clear limitations of human rationality. He imagined a popular bar in a small village. On any evening, if only a few people decide to go to the bar, they will find plenty of seats and a nice atmosphere. But if many people decide to go, they will find nowhere to sit and too much noise. Rationally, then, an individual should go out on evenings when most others do not, and stay home when most are out. The trouble is, there is simply no way to make this decision without advance knowledge of what others will do. The same quandary faces people deciding when to go to their workplace canteen or companies trying to differentiate their products in the marketplace.

"Rationality demands a great deal of human behaviour – much more than it can usually deliver," Arthur says. What people really do in such a case, he argues, is discard rationality and adopt a more adaptive form of thinking. One person, seeing the bar empty a few Thursdays in a row, may decide to go only on Thursdays. Another, perceiving a tendency for attendance to alternate from night to night, will make other plans on that basis. Without any perfect, rational basis for their decisions, individuals instead look for patterns, form beliefs and act on them, and then modify those beliefs as they continue to "learn" from experience (*New Scientist*, 24 April 1999, p 42).

A study published earlier this year certainly bears that out. Physicist Yi-Cheng Zhang and other researchers from the University of Fribourg in Switzerland set up an internet-based game in which human traders play a virtual market populated by virtual traders. The human players looked at recent market movements and had to decide whether the next change would be up or down. When the ▶

arket changed in a relatively simple way, the ayers were quite good at spotting the likely ext movement. But when the market ynamics became complex, the players fectively abandoned their attempts to work ut rationally what would happen next (*hysica A*, vol 331, p 651). Many of them just ept repeating the same prediction. "It seems," ays team member Joseph Wakeling, "that the apacity of humans to act logically is limited, nd in more complex situations they try other ethods to make decisions."

And no matter how strange these irrational ecisions might be, they can affect financial narkets. If someone decides they will buy tocks whenever sunspot activity is high, for xample, others may exploit their behaviour, nd so gradually the market might become inked to sunspots. In comparison with iron-lad rationality, this perspective offers a far nore flexible picture of human behaviour.

A new generation of market modellers is picking up on this. Physicist Neil Johnson, director of the University of Oxford Centre for Computational Finance in the UK, runs simulations of financial markets populated by virtual agents that continually make decisions to buy or sell some stock or foreign currency. As the effectiveness of one trader's actions depends on what all the others do, rationality is an unreliable guide. Instead, these agents act like Arthur's bar-goers – they try to identify patterns in past price movements and use them to predict the future, rejecting ideas that do not work, and learning all the while.

The agents use a relatively simple procedure to make their predictions. A string of letters depicting whether the market went up (U) or down (D) each "day" offers a crude record of recent market activity. A six-day record might be UUDUDD, for example. Agents then apply various "strategies" to make their predictions based on this record. A strategy is just a mathematical recipe that takes any market record and from it predicts what will happen next – that the market will either go up or down. Given UUDUDD, one strategy might predict U; another might predict D.

What goes up...

The number of possible strategies is enormous. For strings of N letters, for example, the number of possible records is 2^N. A specific strategy assigns either U or D to every one of these 2^N possibilities. Given these two choices, and the 2^N possible records, the number of possible strategies is then $2^{(2^N)}$, a number that grows extremely rapidly as N gets bigger – for $N = 6$, it is already 1.8×10^{19}. N corresponds to the "brain size" of the agents, in that it reflects how far they look into the past when searching for patterns. Those with bigger brains try to detect subtler, longer-term patterns. Even in this simple scheme, the number of different predictions of the future based on the recent past is immense.

In the model, each agent starts out with a handful of strategies, usually chosen at random. As days go by, they see which of these tend to make good predictions, and which make bad. On each day they use whichever strategy is currently most successful to decide whether to buy or sell. But the agents only trade on a given day if they feel "confident" enough, based on how frequently their predictions have been correct in the recent past. Some, mired in self-doubt, take the day off.

The model generates a market with winners and losers, exciting rallies and crashes, and an emotional mood of its own. In qualitative terms, the market looks realistic. But this is only the beginning of its success: the model also predicts market behaviour with a mathematical precision that no traditional economic theory can match.

For much of the 20th century, economists believed that fluctuations in the prices of stocks or other financial instruments followed the "bell curve" of statistics, with most clustered around the average and increasingly few out on the tails of the curve. Among other things, this would imply that extreme fluctuations, such as a one-day price jump of 20 per cent, would be exceedingly rare. Yet over the past few decades, economists have learned otherwise: market fluctuations do not follow the bell curve, but another pattern known as a power law, which indicates that extreme jumps in value are far more common than traditional thinking would predict.

Although these patterns have been discovered in real markets of all kinds, nothing in traditional economic theory can explain them. In contrast, a whole slew of virtual markets based on pattern-searching and learning from experience can mimic the behaviour of real markets quite easily. "I've been getting very accurate agreement with financial facts," notes economist Blake LeBaron of Brandeis University in Waltham, Massachusetts, another pioneer in this area.

Agent-based models also reveal how the buying patterns of one investor can easily influence those of another, resulting in mass panics, waves of euphoria, and other collective movements. "These models are very powerful in explaining the statistical facts of financial markets," says economist Frank Westerhoff of the University of Osnabrück in Germany. But such models can go well beyond reproducing just the statistical properties of real markets – they can predict the actual prices.

It's all down to the way the human researchers

"If someone decides to buy stocks when sunspot activity is high, others may exploit their behaviour. Gradually the market might become linked to sunspots"

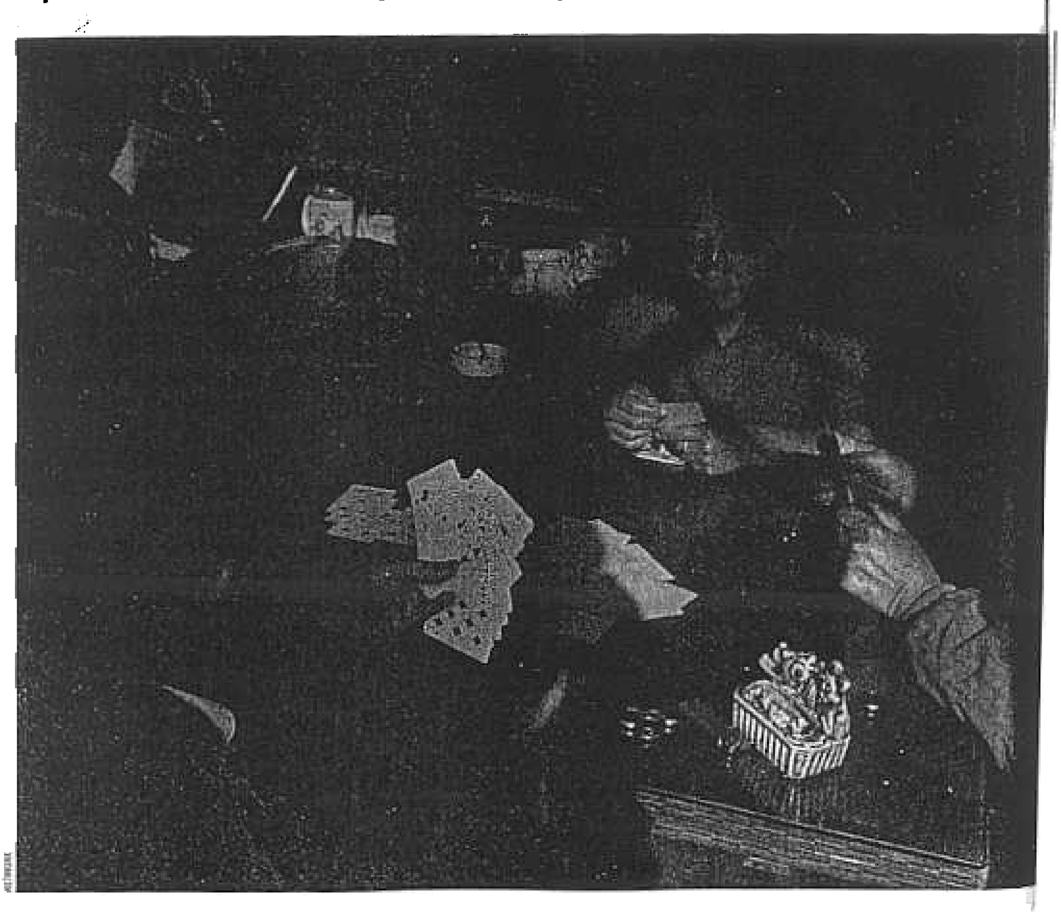

"Once you have a model that accurately mimics market traders' actions, you can start to model how certain policies might help to stabilise the global economy"

assign the virtual traders their strategies. Even if N is only 5 or 6, there are a vast number of possible choices, and this gives the researchers a useful tool. By playing with the choice of these initial strategies, they can "tune" their virtual market until it reproduces the price movements seen recently in some real market.

The implications of this are enormous. For a start, running the tuned model forward into the future might then show where the real market is going, be it the New York Stock Exchange or the foreign currency market for the Euro. In a number of tests, one team (who prefer not to be named because they are involved in launching a company to exploit their work) has shown the model can consistently beat the 50 per cent mark – therefore doing better than chance – in predicting the ups and downs of foreign exchange markets.

More exciting yet is the hope that this method will detect really big movements in advance – changes of, say, a few per cent or more. Identifying these moments beforehand is the dream of any speculator. The technique does seem to predict many large movements in advance, but unfortunately it is too early to rule out the possibility that this success is simply due to luck: events that are rare by definition don't succumb quickly to statistical analysis. In the New York Stock Exchange, for example, a one-day change of five per cent takes place only about once a year. Within a few years, however, it may become clear that these models can tune into a market's behaviour and accurately predict significant movements.

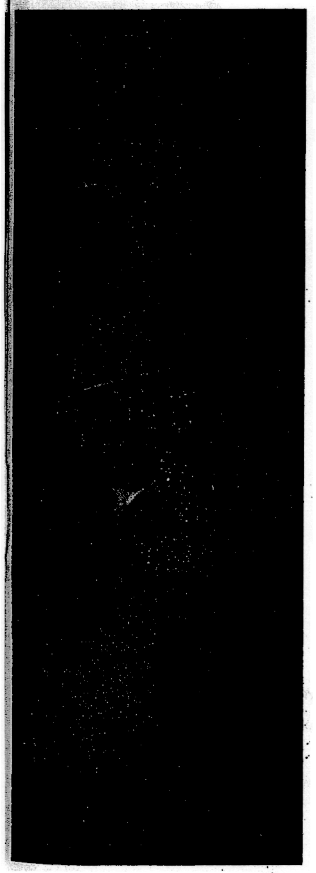

Being forced to make decisions based on hidden or incomplete information provokes people into all kinds of irrational strategies

Self-fulfilling prophecy

Of course, if that does happen, the implications may be a little bizarre. The method is sure to be adopted by many large investors, which would change the nature and philosophy of the markets themselves, and perhaps undermine the technique's ability to make predictions. It's also possible that those predictions might be self-fulfilling. If the model predicted an imminent 5 per cent jump in the New York Stock Exchange, investors armed with this knowledge would naturally flock to the market, buying stocks to make a profit and thereby driving prices up – perhaps by five per cent.

But price predictions may be just the small fry. If fine-tuning can make the model reproduce the market record of ups and down, it means the model is somehow getting under the skin of the market itself. Hidden within the virtual agents' strategies – which are, of course, available and open to analysis – could lie the theories and expectations that are driving the real market, the raw material for generating a realistic economic theory sometime in the future.

There are other, more immediate pay-offs, too. Once you have a model that accurately mimics market traders' actions, you can start to model confidently how fundamental changes in law might affect the markets. You can see, for instance, how certain policies might help to stabilise the global economy.

Most economists agree that excessive speculative investment in foreign currencies plays a major role in creating financial crises, such as the crash that humbled the economies of Thailand, Indonesia, Malaysia and Singapore in the late 1990s, following three decades of incredible sustained growth. By putting pressure on exchange rates, powerful investors can make a killing when a currency is devalued, causing untold misery in the affected nation. Twenty-five years ago, economist James Tobin of Yale University suggested that governments could counter this risk by placing a small tax – less than one per cent – on every transaction involving currency conversion to deter overspeculation. As well as reducing the frenzy of speculation, a "Tobin tax" would raise money for governments. But would it work?

Existing theory is simply inadequate to give any definitive answers. Here the problem has been not so much narrow assumptions about human behaviour, but the sheer complexity of international markets involving millions of investors. For this reason, economists have long argued over the merits of the Tobin tax without reaching a consensus.

But Westerhoff has found a way to bring real-world markets into the laboratory. He has developed a realistic virtual market where he can impose taxes and probe their consequences directly, and he says it is answering the question of the Tobin tax.

His model involves not one but two distinct markets, which he can tax independently, but with traders having knowledge of both markets. Within each, agents soon learn to alter their behaviour in the face of a Tobin-style tax, as speculative trading becomes less lucrative. Westerhoff has found that the tax tends to decrease the volatility of the market in which it is introduced, while increasing the volatility of the second market. This, he argues, may be a benefit. If a few important world markets were to institute a Tobin-type tax, other markets would probably follow suit to avoid becoming more volatile. Importantly, the model shows that if the tax is applied to both markets, both become less volatile. "I used to be unsure of the benefits of a Tobin tax," says Westerhoff, "but now I believe it would work."

It's an impressive result, and it may be just the beginning of what is possible. The range of questions that can be addressed with virtual markets seems almost boundless. And we may soon be able to create the ultimate simulation, according to Robert Axtell of the Brookings Institution, a public policy research organisation in Washington DC. "Creating entire mini-economies in silicon will be possible," he says.

The challenge is to create virtual agents who are driven by exactly the same forces that drive real traders. Gone are the days when we modelled markets with traders who simply knew a good price when they saw one. We now want agents who set up firms, procreate, engage in politics and write constitutions. And, Axtell adds, we'll also need agents who bribe other agents for votes. Maybe then we'll see what economics is really made of. ●

Mark Buchanan's book, *Small World: Uncovering nature's hidden networks*, is published by Orion