21st Century Economics

This article is based on a public lecture given in the University of Durham, October 2007, when I was a Distinguished Fellow at the Institute for Advanced Study during the 2007/08 academic year. The lecture was intended to be accessible to non-economists.

The title of the lecture is '21st century economics'. But I am going to spend a fair amount of time taking about 20th and even 19th century economics. Only by an appreciation of what went before can we see how the future of the subject is fundamentally different. I say 'fundamentally different', but I do not mean by this that we should reject all previously existing economics. Far from it, the discipline has given us powerful insights into how people behave.

All theories are approximations to reality. Some, such as quantum physics, appear to be extremely good approximations. But all theories require assumptions, simplifications of reality to be made. The question then becomes one of asking whether a particular set of assumptions is a sufficiently good approximation to the real world to be of use in helping to understand it.

I want to argue that we now have the tools to relax the assumptions of conventional economics, and in so doing to develop an even better understanding of how the economy behaves. Conventional economics is not wrong. The strict assumptions it imposes about how humans behave limit its usefulness, limit the circumstances in which its simplifications are a good approximation to reality. The economics of the 21st century is both more realistic, and builds on and develops what went before.

The American writer PJ O'Rourke's excellent book *Eat The Rich*, is a very humorous account of his travels to various countries and his opinions on why some of them are rich and some poor[1]. "One thing that economists do know", he wrote, "is that the study of economics is divided into two fields, "microeconomics" and "macroeconomics. Micro is the study of individual behaviour, and macro is the study of how economies behave as a whole." He goes on: "That is, microeconomics concerns things that economists are specifically wrong about, while macroeconomics concerns things economists are wrong about generally."

I write above of the two 'halves' of economics, micro and macro. But it is microeconomics which is by some considerable margin the more important of the two. The very phrase 'macroeconomics' often invokes a sense of unease within the profession. Indeed, in many circles one of the most devastating things which can be said about a paper on macro is that it 'lacks micro foundations'.

[1] It is also a very interesting book. I would certainly put it high on the list of recommendations for students of economics. It will make them think.

Economics is in essence a theory of individual behaviour. A theory of how individuals gather and process information and how they make decisions. To be more precise, it is a theory of how individual decision making *units* behave. So, yes, much of the theory is concerned with people, how individuals decide how much to save and how much to spend, whether or not to work and, if so, how long to work for, and so on.

But the theory is also applied to firms, which are seen as individual decision making units. Of course, economists are perfectly aware that firms are made up of people, and there are complicated interplays within the company before a decision is made. As a simplification, all this is subsumed, and each firm is portrayed as a single decision making entity. 'Agent' is the jargon phrase for decision making units, whether these are individuals, firms or, in some instances, governments and central banks.

All theories are simplifications of reality, and these particular simplifications are actually useful ones to make. Economists are deliberately abstracting from the internal decision making process within firms or governments, and are concerned with decisions taken which impact on the world outside. As long as we do not lose sight of the fact that some pretty drastic simplifications are being made, this is perfectly acceptable.

Modern economics was founded by Adam Smith in the late 18th century, with his magnum opus, the *Wealth of Nations*. Just like PJ O'Rourke, he was interested in how nations became rich. The difficulty of this question is illustrated by the fact that even now, over 200 years later, we are still far short of having a satisfactory answer to it.

Smith's work is outstandingly original in many ways. But perhaps the greatest of his insights was to formalise not just a theory of individual behaviour, but one which is presumed to be universal, valid across both time and space. Like Darwin's *On the Origin of Species*, the *Wealth of Nations* is packed full of practical examples to support Smith's abstract concepts.

Adam Smith's general theory of behaviour was that, to use modern jargon, agents react to incentives. The homely example of the market stall is one which is frequently used to illustrate the idea. Imagine such a stall selling, say, bananas. The customers in the market will each have his or her own tastes and preferences regarding bananas. Some will relish them, whilst others are not that keen. So the consumers are heterogeneous in their opinions about bananas. As the price rises, fewer bananas will be bought. But consumers will not react identically to a given change in price. Some will be deterred from buying by even a small change in price, but quite a few may continue to buy exactly the same number even if the price goes up a bit. There will be a price at which, however, even the most ardent devotee reduces the number he or she is willing to buy. Consumers are reacting to incentives.

The insight that agents react to incentives is a very powerful one, and it is the one feature which distinguishes economics from the other social sciences. Many sociologists accept it, for example, though by no means all, and there are those in this discipline who deny the very concept of individual agency. But all economists accept the proposition

The banana stall example seems so obvious as to hardly warrant the rather pretentious description that 'agents are reacting to incentives'. But we can see it in many, often quite unexpected contexts. Steven Levitt is an American economist who has done applied work of great distinction. Few have read his academic articles, but millions have read his best seller *Freakonomics*. The book is essentially devoted to illustrating the power of the proposition that agents react to incentives, often with unexpected consequences. So, many people will have read of Levitt's result that more relaxed abortion legislation is an important cause of the sharp falls in crime in America over the past fifteen to twenty years. Most crime is committed by poor, unskilled young men, and children raised by single mothers have a higher probability of growing up into the category. By making abortion easier, the supply of such men into the crime 'market' has been reduced, and as a result crime has fallen.

But Levitt points out just as clearly that the policy of sending convicted criminals to jail for longer has also been in part responsible for the fall in crime. The same effect can be observed in the UK. Crime has fallen over the past decade, yet the jails are packed to bursting. This leads to anguished debates in the quality press as to how this can be. If crime has fallen, why are jails full?

The answer is precisely that it is partly because jails are full that crime has fallen. The Conservative Home Secretary Michael Howard introduced a policy in the 1990s of increasing both the likelihood of a convicted criminal being sent to jail and the length of sentence. The Labour government continued the policy. And the result is that crime has fallen. It is by no means the only reason why, but potential criminals have been reacting to incentives. The 'price' of crime has risen: if you are caught and convicted, there is a greater chance of being sent to prison, and of being sent there for longer than before. So the supply of crime – the willingness to commit it - has been reduced. By no means all criminals have reduced their supply as the price has risen – agents differ in their reactions to a given change in price – but some have.

Levitt's book is a testament to the power of the insight that agents react to incentives. But he relies on this very simple rule. He does not elaborate with assumptions as to how much information agents gather before they take a decision, nor does he specify in detail the exact behavioural rule which agents use when they are processing information. Conventional economic theory does both of these things.

To understand why, it is necessary to step back in time, to the second half of the 19th century in fact. It was then that the first important steps were taken to place the economic theory of individual behaviour onto a formal, mathematical basis. And it was then that the concept of the 'rational' economic agent was created.

The theory reflects the classical, equilibrium physics of that era. Equilibrium physics had, of course, experienced enormous success in explaining many aspects of the physical world. Other sciences, and would-be sciences such as economics, stood in awe of the achievements of physics.

The main mathematical tool then available was the differential calculus, the technique at the heart of classical, equilibrium physics. This branch of maths is particularly helpful when the variables it is used to analyse can be assumed to change in a smooth, continuous manner. And one of its principal uses is to find the maximum (or minimum) value of any particular mathematical formula.

It seemed logical to assume that agents – people, firms – should behave in a way which maximised their own interests. Why should anyone not choose to do so? And, with this assumption, the tools were at hand to enable the maxima to be found.

But how to measure these interests? For firms, the natural assumption to make was that they maximised profits. Each firm would examine the various ways in which it could combine the inputs – machines, labour – into its process of production in order to produce different combinations of goods to sell to consumers. The firm was assumed to choose the combinations of inputs and outputs which, given the prices of these which existed, would maximise its profits.

For consumers, it was slightly trickier to work out what it was they were maximising. The abstract concept of 'utility' was invented. 'Utility' is a difficult word to translate into everyday English, perhaps a combination of 'pleasure', 'satisfaction', 'usefulness' might do the trick. The basic idea was that an individual derives a certain amount of 'utility' from the consumption of a given amount of any particular product.

Each consumer was assumed to have a fixed set of tastes and preferences across all the different kinds of goods which could be made. These preferences differ across individuals, as in our previous example of bananas. Personally, I do not like Marmite, my wife does. Crucially, each person has a limited amount of money to spend, an amount which again varies across individuals. Given the prices of goods, each consumer then selects, given his or her income and his or her preferences, that combination of goods which maximises utility.

At the heart of this model is Adam Smith's assumption that agents react to incentives. Given a particular set of prices, each firm produces the combination of goods which will maximise its profits, and each consumer buys the combination which will maximise utility. If prices change, both companies and individuals will make different choices.

So the model contains a fundamental insight into how the world really does work. And although the model is rather abstract, the assumptions on which it is based do not seem at first sight to be totally unreasonable. People differ in their tastes, and they differ in their incomes. Both these assumptions are clearly true. And, given these, it seems logical that a person should buy the combination of goods which in some way seems best for the individual.

To repeat once more, all theories are approximations to reality. It is a question of how good the approximation is. Physics had discovered that making the assumption that a

system was in equilibrium, in a steady state, was a very good one in many different contexts.

In economics, the concept of equilibrium translated perfectly naturally into the context of markets. In the market stall example above, at any point in time there may be a price of bananas at which the demand and supply are equal. There are no imbalances in the two quantities, the system is in equilibrium.

The real intellectual advance in formalising economic theory in the 1870s was to generalise the problem from the example of a single market to that of *all* markets in an economy at any point in time, the so-called theory of general equilibrium. It seemed obvious that in any given market, a price could always be found which would bring the market into equilibrium, in which supply equalled demand[2]. So the question was posed: can a set of prices be found which will clear *all* markets, so that every single one is in equilibrium and we have a general equilibrium?

The question had a subtle twist. It was not a matter of trying to show that an actual set of prices could be found in any particular economy. This would be an enormous practical undertaking. But even if it could be demonstrated for, say, Britain, this would give no guarantee that a similar thing could be done for France.

The question therefore became: under what conditions can a set of prices be guaranteed to exist which will ensure that the economy is in equilibrium, that supply equals demand in all markets?

Mathematicians will recognise immediately that this is requiring an existence proof, and such proofs can be formidably hard. A more homely mathematical example may in fact help to illustrate the problem. In secondary (high) school, pupils are confronted with the problem of solving quadratic equations, of finding values of a variable which satisfy a formula in which both the variable itself and the variable squared appear. As it happens, a general formula has been proved to exist which ensures that a solution can be found for any example of such an equation. We might usefully think of the problem of general equilibrium as being one of finding the formula or, rather harder, of finding the conditions under which the formula can be found.

Proving the existence of general equilibrium turned out to be an immensely difficult task, one which was not fully worked out until the late 1960s. The importance of general equilibrium theory over the 1870-1970 period is illustrated by the fact that when the Nobel prize in economics was instituted in 1969, four out of the first 11 recipients were awarded it for their work on general equilibrium.

---

[2] As it happens, this is a harder question than the early pioneers realised, though a discussion of this issue would be a diversion from the main theme. For example, there is the problem of dealing with the so-called hog cycle, discussed, for example, in the best selling textbook of the 1960s, Richard Lipsey's *Introduction to Positive Economics*. More generally, even if a price can be proved to exist which clears the market, there is the question as to how it can be brought into existence.

General equilibrium has continued to be the cornerstone not only of a great deal of economic theory, but of the policy advice which economists give. Central banks, for example, in the opening decade of the 21st century have relied increasingly on so-called dynamic stochastic general equilibrium models of the economy. Incredibly, given current circumstances, money and credit/debt play no direct role in such models, and the technical properties of the model ensure that financial markets have no impact on the economy in the long-run[3]. But to discuss this in detail would take us far from our immediate purpose.

The policy implications of general equilibrium go much wider. If all markets are in equilibrium so that supply equals demand in every single one, there are no unused resources anywhere in the economy. So we have an efficient allocation of a given set of resources. If we observe actual markets in which demand and supply are not in balance, an implication is that the efficient workings of the market, the price mechanism which adjusts to ensure equilibrium, are being obstructed in some way or other. The task of policy is therefore to remove these obstructions.

The labour market, where people supply their labour and firms demand it, is a very important market in practice. Yet, supply often seems to exceed demand. In Britain between the two world wars, for example, the unemployment rate averaged 9.4 per cent, peaking at 15.3 per cent in 1932. In the United States, the average was 11.1 per cent, with a high of 24.7 per cent in 1933[4]. *Prima facie*, the labour market was not in equilibrium, with supply vastly exceeding demand.

From the perspective of the general equilibrium mindset, the problem was on the supply side. So, for example, institutions such as trade unions were preventing wage rates (the price of labour to the employer) from falling sufficiently to increase demand. And the unemployed were effectively choosing leisure, to live off the generous (!) benefits of the time, instead of choosing to work.

But the circumstances in which the assumptions of the model are a good approximation to reality are much more limited than the above description might imply. Many of the disputes in economic theory can be thought of as attempts to replace one or more of the key assumptions of general equilibrium with ones which are believed to be more realistic. More precisely, ones which are more realistic in a wider set of circumstances – circumstances in which general equilibrium, at best, sees through the glass darkly.

In constructing the model of general equilibrium, a whole range of assumptions had to be made about how agents behave. Modern economists recognise that most of the time the complete set of these assumptions does not correspond to reality, with the exact nature of

[3] See, for example, Wolfgang Münchau, *Financial Times*, 6 July 2008, http://www.ft.com/cms/s/0/8362b1d0-4b59-11dd-a490-000077b07658.html?nclick_check=1

[4] estimates of unemployment rates in the inter-war period do vary slightly between sources. This data is taken from A Maddison, *Dynamic Forces in Capitalist Development*, Oxford University Press, 1991

the empirical violations varying from context to context. But the culture of the profession appears to involve the belief that relatively minor violations of the assumptions still leave the basic properties of the model, and its policy implications, broadly correct.

Unfortunately, economics itself has shown that this is not the case at all. Even a single violation of the assumptions about behaviour in the general equilibrium model can lead to outcomes which are far removed from those of the model. These are not esoteric findings, buried away in dissident journals, but are published in the leading journals of the discipline. Some important results relating to the behaviour of firms were, for example, surveyed in the top *Journal of Economic Literature* fifteen years ago[5].

There is certainly not the time to go into all the aspects of this. It is very important to note, however, that the past 25 years or so has seen massive growth in two new areas of economics, behavioural and experimental. There is no clear-cut division between the two, indeed they are both essentially concerned with how agents really do behave in practice. General equilibrium assumes from first principles that it is logical for agents to behave in a particular way – they have access to all information and take the best possible decision for themselves on the basis of this.

Behavioural and experimental economics are, in contrast, empirically based. A measure of their importance is the fact that the 2002 Nobel prizes were awarded to the leading practitioners. Their prize lectures are both accessible to the non-economist and contain a vast range of interesting material[6].

The fact that behavioural and experimental economics are empirical means there is often scope for different interpretations of their findings. However, it is fair to say that there is now a substantial body of evidence to suggest that in many situations, the way in which agents behave differs from the assumptions made about them in standard theory.

This is a fundamental way in which the economics of the 21st century will differ from what went before. The discipline will no longer rely on examining the logical implications of purely abstract assumptions about how rational agents ought to behave, but will be much more empirical. The shift has already started, and the leading journals now carry far more purely empirical articles than used to be the case.

The real challenge is to develop theoretical models which are much more firmly grounded empirically in terms of the behavioural rules of agents. I return to this point at the end of the lecture. But, first, a look at two other ways in which the mind-set of general equilibrium needs to be changed.

---

[5] J Silvestre, 'The market power foundation of macro-economic policy', *J Ec Lit*, 31, 105-141, 1993

[6] VL Smith, 'Constructivist and ecological rationality in economics', *American Economic Review*, 93, 465-508, 2003; D Kahneman, 'Maps of bounded rationality: psychology for behavioral economics', *Am Ec Rev.*, 93, 1449-1475, 2003

Economics has already become much more realistic concerning the amount of information which rational agents are assumed to have when they take decisions. I have to say immediately that there is a certain amount of mental gymnastics involved in top flight mainstream economics, whose practitioners are both fully aware of these developments in certain contexts, but appear to set them to one side again when they try to apply general equilibrium models for policy purposes.

In the full-blown proofs of existence of general equilibrium, the proof that a set of prices which clears all markets can always be found, this amount of information is, almost literally, mind-blowing. The interested reader could follow this up in the footnotes below[7].

In practice, what is meant by a rational agent is an agent who gathers all available information which is relevant to a decision, and then on the basis of this information takes the decision which maximises the relevant concept, be it utility or profits. The emphasis here is on the word 'all', all available information.

But it is easy to think of examples where it may be very hard, or even impossible, to gather all available information. Buyers in used car markets, for example, face a problem, especially when buying from an individual rather than a large firm which specialises in used car sales. The seller knows a lot more about the true quality of the car than does the buyer. A large dealership has its reputation to protect, so a used car on offer here might be more likely to be of reasonable quality. But in a one-off transaction between individuals, this is not the case at all.

The American economist George Akerlof examined this problem nearly 40 years ago. In so doing, he developed a formal model in which the agents had different amounts of information – 'asymmetric' information in the jargon. It was a brilliant paper[8] which richly deserved the Nobel prize he was awarded many years later. The example of used cars was merely to illustrate the model he developed.

Asymmetric information can lead to some quite dramatic outcomes. For example, the quality of a used car will vary, some good for their age and make, some bad. The seller has a very good idea about this, the buyer very little. So the best guess for a buyer is that any particular car is of average quality, and so he or she will be willing to pay this amount. But owners of good used cars will not be able to command the price they expect, and so will not put their cars up for sale. As a result, the average quality of cars on the market falls further, which leads to another round of reluctant sellers. And so on, until eventually, where there is pure asymmetry of information, the market will not exist at all. But even if it does, far from giving a good outcome for buyers and sellers, the market system positively drives out the good in favour of the bad.

---

[7] For example, R Radner, 'Competitive equilibrium under uncertainty', *Econometrica*, 36, 31-58, 1968; D Duffie and H Sonnenschein, 'Arrow and general equilibrium theory', *J Ec Literature*, 27, 565-598,1989

[8] GA Akerlof, 'The market for 'lemons': quality uncertainty and the market mechanism". *Quarterly Journal of Economics*, 84, 488–500, 1970

At a stroke, Akerlof extended the realism of economic theory quite dramatically. Again, all theories are approximations to reality, the question is one of how good they are. By permitting different agents to use different amounts of information, Akerlof increased the correspondence between the assumptions and the real world in a wide range of circumstances.

There is now a vast literature which uses the assumption of asymmetric information in many different contexts. But this increase in realism comes at a price. The traditional model of rational behaviour lays claim to being a general model of behaviour, valid at all times and in all places. Agents gather all relevant information and use it to make the best possible decision. Once asymmetry between agents is permitted in the theoretical models, any claim to generality is lost. We need, as a minimum, to specify the nature of the asymmetry involved in any particular situation which is analysed. As Akerlof himself put it in his Nobel lecture 'in this new style [of economics], the economic model is customized to describe the salient features of reality that describe the special problem under consideration.[9]'

The second main development has also been taken into account by economics already, but in a very partial and much less satisfactory way. The description of the theory has involved agents making decisions on the basis of the prices which they face. When prices change, they make different decisions and a different equilibrium is arrived at. This is all well and good.

But it conceals a crucial assumption. Namely, that prices react faster than quantities. People and firms observe prices, and then decide quantities, how much to buy of each product and how much to sell. Imagine now that a shock take place to the system – the price of oil doubles, say – and agents face a new set of prices. They examine the new prices, and only then do they decide the new set of quantities to buy and sell. The key point here is the assumption that no-one is allowed to trade *before* the set of prices has been found which once again ensures that equilibrium holds in all markets. Quantities are not allowed to adjust at all before the set of prices is found which will bring about a new equilibrium.

The process by which the new set of prices is discovered which once again ensures that supply and demand are everywhere equal, has always been an unsatisfactory aspect of general equilibrium theory. In the original formulation of the problem, a mythical Auctioneer was posited who initially called out a set of prices at random, and agents then told the Auctioneer not what they had actually bought and sold at these prices, for only by a miracle would the initial guess clear all markets, but what they would buy and sell if these prices did exist. Eventually, by a succession of iterations of this kind between the agents and the Auctioneer, the ideal set of prices would be found.

General equilibrium theory, for all its mathematical sophistication, has not really moved on beyond this concept of the Auctioneer and his method for finding the prices which

[9] GA Akerlof, 'Behavioral macroeconomics and macroeconomic behavior', *Am Ec Rev.*, 92, 411-433, 2002

will clear all markets. The process rapidly becomes very complicated, involving situations in which, for example, a set of prices does exist, but it can never be found.

In the inter-war years, Keynes mounted a formidable challenge to this economic orthodoxy, culminating in his *General Theory of Employment, Interest and Money* in 1936. Like his contemporaries, he observed the dramatic and persistent levels of unemployment. The conventional view, which we noted above, was that the workings of the price mechanism were being obstructed, unions were preventing wages from falling, or the unemployed were being paid too much benefit.

In essence, Keynes retained all the assumptions of the standard model, but with one fundamental difference. He allowed quantities to change before prices.

So imagine, in Keynes' world, an economy which is in equilibrium. For some reason or other, we need not specify how, a shock takes place which means that the price of labour – wages – is now too high for employers to provide jobs for everyone at that price. But instead of waiting to find out what the new price is which will clear the market, employers make decisions at the existing prices. The labour market no longer clears. The price of labour is too high to generate sufficient demand for workers, so there are people left unemployed. But the problem does not end there. The incomes of those made unemployed fall, and the amount they have to spend is less. So their spending decisions are also based upon 'false prices', as they are known in the jargon, prices which do not ensure that supply balances demand. In this way, the shock is magnified through the system

There is a long theoretical debate following the publication of the *General Theory* about whether mechanisms exist which will eventually bring about an adjustment of prices so that all markets will once again clear, despite the fact that agents are allowed to trade at false prices and adjust quantities before prices.

Keynes was impatient with this debate, which is the source of his famous remark that 'in the long-run we are all dead'. It was not a cynical remark at all. Even if a mechanism could be postulated which in the long-run would conjure up the set of prices which would restore equilibrium, in the meantime we would still see high rates of unemployment. Keynes wanted action immediately, rather than relying upon some esoteric mechanism which might in the long-run restore full employment, because in the long run we are all dead.

The real point of this debate is that it shifts the emphasis from the properties of different equilibria of the system to the path between them. Even if a theoretical model can be built – and enormous intellectual effort has gone into this - which shows that when quantities adjust before prices a full employment equilibrium can eventually be restored, the concept of time, of how long the process takes, becomes crucial.

General equilibrium is in many ways a timeless concept. We use it to compare two situations, both in equilibrium, one of them before and the other after a posited change in

the economic environment. But if a system typically spends fifty years[10], say, far from equilibrium in the transition between two stable states, it is this aspect which becomes interesting and not the position in which we might eventually end up. Even when the equilibrium framework is retained in the theoretical model, it is essential to know both the shape of the path between two equilibria and the time the system spends on it.

My final point is that the real challenge facing economics is to build successful models whose micro-foundations have much firmer empirical foundations, drawn from behavioural and experimental economics.

By 'successful' I mean models which give a better account of empirical phenomena than does conventional economic theory. In particular, models which consist of plausible rules of behaviour of the individual agents which explain emergent macro-phenomena of the system being examined. In other words, the models do not attempt to explain the macro aspects directly. They have rules for agent behaviour, from which it will not usually be obvious what the macro-properties are. These latter emerge from the interactions of the individual agents.

A key aspect of this development is that maximisation is discarded completely. Instead of searching for the best possible outcome, agents use simple rules of thumb which give fairly good outcomes most of the time.

This idea is far from new. It has been around for 50 years, developed in a series of papers by the outstanding polymath Herbert Simon[11]. His description of this type of behaviour is 'satisficing', a term which can lead to confusion. It means agents are looking for, indeed in general will only be capable of finding, a satisfactory outcome rather than the best possible one.

So in addition to operating with imperfect information, there are constraints on the abilities of agents to compute the optimal outcome[12]. This framework extends very considerably the range of situations in which the assumptions underlying a model are reasonable approximations to reality.

Of course, this comes at a price, and a high one at that. Standard economic theory purports to have a model of behaviour which is universal across all agents , all places and

[10] This is by no means an idle speculation. As long ago as 1969, Tony Atkinson, former Warden of Nuffield College Oxford, showed that the transition between equilibrium paths in the standard model of economic growth was of the order of 100 years. AB Atkinson, 'The timescale of economic models: how long is the long run?', *Review of Economic Studies*, 36, 137-152, 1969

[11] See, for example, HA Simon, 'A behavioral model of rational choice', *Quarterly Journal of Economics*, LXIX, 99-118, 1955

[12] Any chess player reading this who doubts this statement should reflect on the number of positions in the game where there is an unequivocal best move, compared to the total number of positions which can actually exist. The proportion is, as an approximation, zero. Of course, it is not exactly zero, but in most positions in chess the 'optimal' move does not exist. Rather, we are unable even with the assistance of modern computers to compute it. Former world champion Gary Kasparov's books on the games of previous world champions, *My Great Predecessors*, are very revealing in this respect.

all times. In any particular situation, it is simply a question of applying this general model. Under this new approach, each model is custom built, and requires careful consideration of the appropriate rule. Models based on these principles often look, by the standards of economic theory, very simple, not to say naïve. But as Daniel Kahneman put it in his Nobel lecture 'psychological theories cannot match the elegance and precision of formal normative models of choice [i.e. rational choice, author's note], but this is just another way of saying that rational models are psychologically unrealistic'.

An important feature of these models is that we are usually reliant on computer simulations to understand their properties. They may appear naïve, but even with a mere handful of parameters, they can give rise to great complexity. In standard economics, analytical solutions are a *sine qua non* of theoretical models. In this more realistic world, it may be possible to discover them, but in general we have to rely on simulation. This does not mean at all that they are in some way inferior. Applied maths in many disciplines relies on computer simulation and not analytical solutions, think, for example, of partial differential equations.

Despite the apparent simplicity of the behavioural rules, there are features of these models which make analytical solutions hard if not impossible to obtain. For example, agents are heterogeneous. In other words, in any given situation, different agents will react differently. This, of course, is a feature of standard economic theory: agents have different tastes and preferences so they will react different to any given stimulus.

The crucial difference in the new approach, however, is that the tastes and preferences of an agent are not fixed, but can evolve over time. So the same agent can react differently to any given change in circumstances at different points in time. This evolution of preferences is a key feature of many real world situations. In financial markets, for example, traders are heavily influenced by the actions of others. A trader might believe, for example, that given the economic fundamentals, the dollar is going to strengthen against the Euro. So logically he or she should buy dollars. However, the decision will almost certainly be influenced by what other people are doing, by the overall sentiment of the market. In identical economic circumstances, the strategy of a trader will differ depending upon what other traders are doing. Tastes and preferences are not fixed[13].

There are many other examples in which it is unrealistic to assume that agents' preferences are fixed. For example, the diffusion of cultural norms and opinions by definition involves changes in preferences[14]. Recent work has shown decisively that obesity, smoking and binge drinking are strongly driven by the spread or containment of social norms across groups of individuals[15]. Another example is cultural services,

[13] Alan Kirman has a brilliant, and very early in the context of the genre, model of this. A Kirman, 'The behaviour of the foreign exchange market', *Bank of England Quarterly Bulletin*, August 1995, 286-293

[14] Two examples from history are P Ormerod and A Roach, 'The medieval inquisition and the suppression of heresy', *Physica A*, 339, 645-652, 2004 and P Ormerod and A Roach, 'Emergent scale-free networks in history: burning and the rise of English Protestantism', *Cultural Science*, 1, http://cultural-science.org/journal/index.php/culturalscience, 2008

[15] NA Christakis and JH Fowler, 'The spread of obesity in a large social network over 32 years', *New England Journal of Medicine*, 357, 370-379, 2007; NA Christakis and JH Fowler, 'The collective dynamics

broadly defined to include popular films, books, music, are an increasingly important part of modern economies. In these markets, almost by definition consumer preferences are not fixed[16]. Until a particular piece of music, film or book actually appears, no-one, critics and fans alike, know whether they are going to like it or not, even if in general they approve of the author or musician or whatever[17]. Yet another example is the diffusion of innovations[18]. A further one involves the explanation of why many social, economic and technological systems routinely display great stability in the presence of continual small failures and shocks that are at least as large as the shocks that ultimately generate a cascade across the system on a near-global scale[19].

The examples given above are just selections from a diverse and rapidly growing literature which is meeting the demand raised by Daniel Kahneman in his 2002 Nobel lecture 'Incorporating common sense psychology of the intuitive agent into economic models is [the] challenge'. Much of this work is being done outside of economics, and certainly outside mainstream economics. But the models are showing that this approach can be both more realistic in its behavioural assumptions and scientifically superior to standard economics in its ability to understand the world.

The challenge of reconstructing economic theory is a demanding one. But it is one which also makes it an exciting time to be an economist. It is the economics of the 21st century.

---

of smoking in a large social network', *New England Journal of Medicine*, 358, 2249-2258, 2008; P Ormerod and G Wiltshire, 'Binge drinking in the UK: a social network phenomenon', European Social Simulation Association conference, Brescia, Italy, September 2008

[16] J Potts, S Cunningham, J Hartley and P Ormerod, 'Social network markets: a new definition of the creative industries', *Journal of Cultural Economics*, 2008, http://dx.doi.org/10.1007/s10824-008-9066-y

[17] DJ Watts and PS Dodds, 'Influentials, networks and public opinion formation', *Journal of Consumer Research*, 34, 441-458, 2007

[18] WB Arthur, 'Competing technologies, increasing returns and lock-in by historical events', *Economic Journal*, 99, 116-131, 1989 is a classic reference, and an excellent survey is C Antonelli, *Localised Technological Change: Towards the Economics of Complexity*, Routledge, 2008

[19] DJ Watts, 'A simple model of global cascades on random networks', *Proc. National Academy of Science*, 99, 5776-5771, 2002; P Ormerod and R Colbaugh, 'Cascades of failure and extinction in evolving complex systems', *Journal of Artificial Societies and Social Simulation*, 9(4), 2006