Have Economists Gone Mad?

Public lecture to be given at the 'Limits to rationality in financial markets' economic analysis workshop, Institute of Advanced Studies, University of Strathclyde, 22 June 2009

Introduction

One of Keynes's most well known phrases refers to the power of ideas. In his 1936 magnum opus, *The General Theory,* he wrote "practical men, who believe themselves to be quite exempt from any intellectual influences, are usually the slave of some defunct economist. Madmen in authority who hear voices in the air are distilling their frenzy from some academic scribbler of years back."

A great deal has been written about the role of 'practical men' in the current financial crisis, whether bankers, regulators or politicians. This lecture focuses on the role of ideas, and specifically of ideas in economic theory.

But in contrast to Keynes' view of the role of ideas in the crisis of the 1930s, our current problems are grounded not in ideas which were advanced by academics 'years back'. They have arisen from ideas which play a prominent role in contemporary academic economics. Far from being 'defunct', these ideas became more and more important in the decade or so leading up to the crash in 2008.

In this talk, I will focus mainly on financial markets, and specifically on why the conventional ways of thinking about them are profoundly misguided. The overall theme of the Institute's workshop is of course financial markets, so it is appropriate to focus on them. However, they are obviously of major importance in their own right.

I will go on toward the end to link the conventional analysis in financial markets into the way mainstream academic economics has been thinking about what in the jargon is called the 'real' economy, meaning in everyday English things like output and employment.

There are three aspects of financial markets where the conventional thinking in economics proved to be seriously misguided. These are:

- the potential for volatility

- the lack of smooth adjustment to change
- the ability to diversify risk

The first of these, involving the arcane world of derivatives, and their even more complex relatives such as collateralised debt obligations and credit derivative swaps, has to be dealt with at considerably greater length than the other two. In part this reflects its relative importance, but also in part the need for description and explanation rather than mere assertion of at least one of these three points. So I will have to ask you to bear with me while I try to explain what derivatives are all about.

Derivatives and volatility in financial markets

What are derivatives?

The most serious problems for financial institutions arose in this often bizarre world. Organisations which made a wrong call on what would happen to a share price, a currency, an interest rate, suddenly found that they had lost not just all of their original stake, but an amount very much larger. How could this be?

The answer goes back about 40 years, to a trio of American academics leading successful but blameless careers. These three discovered ways of applying concepts from statistical physics to financial markets. Fisher Black, described by one of his close friends as 'the strangest man I ever met', soon left academia to make millions at Goldman Sachs before his tragically early death. Robert C Merton and Myron Scholes received the Nobel Prize in economics in 1997 for their findings.

Their highly esoteric discoveries had immense practical significance, enabling the creation of the industry of financial derivatives, worth over $500 trillion according to the Bank of International Settlements. The basic idea of derivatives - so called because their value is derived from, or related to, that of an underlying asset - is very simple. Suppose an investor holds some Vodafone shares. He or she may worry that the price will fall. Someone else may think it will rise. A contract can be struck between them to trade the shares at a specified price at a date in the future. Its price will vary depending in part upon the price of Vodafone shares at any point between now and then.

The crucial feature of derivatives is that their price fluctuates much more than that of the underlying share to which they are linked. The rewards of getting it right can be much bigger, but so too can the losses. Vodafone on the London Stock Exchange has been trading around 125p a share. If I feel optimistic about the company, I can buy the shares now. If at the end of the next month they are 250p, I will have doubled my money. But I could instead buy the right to buy them at, say, 200p at the end of next month for virtually nothing, 1p say, since it is so unlikely that such a big increase will happen in such a short time. If I am proved right and the price really is 250p, I have the right to buy shares at 200p and can then sell them immediately for 250p. My 1p has turned into 50p, far, far more than doubling my money. But if I am wrong and the price stays below 200p, I will lose everything I put in.

The converse is that the person I did the original deal with, who took my money confident in the fact that the price would simply not rise so much, now has to sell me Vodafone shares at 200p when their market price is 250p. He or she made 1p on each trade in the first place, but now stands to lose 50p a trade in return for the 1p I originally paid to him or her.

And this is a description of the simplest possible trade in derivatives. The trades are readily made much more complicated, creating even greater possibilities for magnifying both gains and losses.

In short, derivatives both satisfy and create an appetite for risk. They enable much riskier bets – sorry, considered investment judgements – to be made than if you can just trade in the underlying shares themselves.

As it happens, Merton and Scholes got their comeuppance when they totally misjudged some risks and their financial company, Long Term Capital Management (LTCM), collapsed in 1998 with a loss of nearly $5 billion dollars, being bailed out by the Federal Reserve. As Robin Dunbar put it in his excellent 1999 book about LTCM, *Inventing Money*: 'The last seven days of LTCM's independent existence have a strange feel of their own. Thirty years of finance theory has proven itself useless. Billion dollar track records and Nobel Prizes are now meaningless'. So today's problems are not exactly without precedent.

But the parallels between LTCM and the current crisis run much deeper. The basic way in which risk is assessed in these derivative models is flawed. It was known to be flawed at the time of

LTCM. Yet, despite this vivid example being in front of them, the world's leading financial organisations continued to rely upon this scientifically flawed model.

How do derivatives work?

The subtleties of the Merton-Scholes model of derivatives pricing are to some extent lost when the maths is translated into English. But there are three main factors which determine the price at which I can buy or sell a derivative. And they all involve an assessment of how much risk is involved.

The first is the difference between the price at which the underlying share is being sold now, and the price at which it can be bought or sold in the future under the terms of the particular derivative contract which is on offer today. In the Vodafone example above, there was a large gap. The current price of the share was 125p, and I wanted to be able to buy it in the future at 200p. Now, the bigger this gap, the smaller the chances of the price reaching the level which would make it profitable for me to exercise my right to buy. With a big gap, I am not willing to pay very much for that right. So, the bigger the gap between the two, the smaller will be the price of the derivative itself.

By the same logic, the further into the future is the date at which I can exercise my right to buy, the more the seller of the derivative will want me to pay. Again, in the hypothetical example above, I bought the right to buy the share at 200p at the end of next month, when it currently trades at 125p. There is little chance of this happening in such a short time. But the further we go into the future, the chances increase that at some point Vodafone really will be not just 200p, but more than that. The seller of the derivative therefore wants more money from me now to offset the risk that he or she will actually have to sell me the shares at a price at which the seller makes a loss. In other words the further into the future the contract is the more opportunities there are for the price to reach the level I would want to buy at.

These two factors essentially involve a judgement about how likely it is that the actual share price will ever reach the price at which I can exercise my right to buy it.

But it is the third factor which truly captures the concept of how likely it is that we can expect to see any given price in the future. It is this which is the real key to how derivatives are priced in practice. And it is the assumptions which were made on this which created financial chaos across the world.

Imagine that the Vodafone share price over the past couple of years, say, had been very stable. Purely hypothetically, suppose it had just moved in the range 110p to 140p. Then it seems reasonable to suppose that the chances of the price moving well outside this range, certainly in the immediate future, are pretty small. Now imagine a different world in which the price had fluctuated wildly, between 10p and 300p say. In this world, the chances of the price suddenly zooming up from its current 125p to 200p seem distinctly higher.

The word which is often used in the jargon to describe how much a share price fluctuates over time is its 'volatility'. To solve the Black-Merton-Scholes equations, an assumption has to be made about what this volatility is. The price of a derivative will obviously be influenced by this in a big way. At a current price of 125p of a share which has had low volatility, I am quite happy to sell cheaply the right for someone to buy it from me at a price way in excess of this. But I will want quite a lot more if the underlying share price is very volatile.

Derivatives and the 'fat tail' problem

The above description is a walk-through of a concept in maths known as a partial differential equation. These equations are often fiendishly difficult to solve. Even their very simplest application in the derivatives market, such as the example above, involves hard mathematical concepts.

A major factor contributing to the crisis was that most members of boards of financial institutions were quite unable to understand the mathematics underlying these complex financial instruments. And as a result they were unable to ask the right sort of questions about what was really going on once the lid on the black box started to be lifted.

Lifting the black box reveals a frightening secret. Namely, that the assumption which was widely used about the volatility of share prices was wrong. But more than that. It has been known to be wrong for many years.

When evaluating the range over which a share price is likely to move, a very important assumption is made in models about not just the range over which a share price might move, but the pattern within that range. In the hypothetical example above, Vodafone is currently trading at 125p. Obviously, there is a strong chance that the next trade might be at either 124 or 126p, a difference of just 1p from the current level. There is less chance that it will trade at, say, 145 or 105p, a 20p difference, and even less that at the next trade the price will be 25p or 225p, for example. So in terms of the overall pattern of the changes in the price, we will see lots of small changes, and very few large changes.

By assuming that this pattern changes follows something which is called either the Gaussian – after the great mathematician Gauss who invented it over 200 years ago – or the 'normal' distribution, very large differences between the price now and the price when the share is next traded are effectively ruled out. The normal distribution is given its name for the very reason that there are lots of examples of this distribution in reality, it seems 'normal'. An everyday example is the heights of individuals. Suppose the average height of adult men is 5 feet 9 inches. More men will have this particular height than will have any other. But there will be similar, though slightly smaller, numbers who measure 5 feet 8 or 5 feet 10 inches. The number of men we observe who have any particular height will get less and less the further we move away from the average height.

The normal pattern we observe in heights of men has two features. First, it is symmetric around the average. This means if we count the number of men who are, say, 5 foot i.e. 9 inches less than the average, this will be very similar to the number who are 6 feet 6 inches i.e. 9 inches taller than the average. Second, and crucially in this context, very large differences from the average are never observed in practice. So we never, ever, see an adult male who is just one inch high, and we never, ever, see someone who is 20 feet tall.

In financial markets we do. The assumption that price changes in financial markets – shares, bonds, interest rates, currencies, commodities, - follow this 'normal' pattern appears to be not an unreasonable one to make. For the most part, they do. When we examine the evidence and look at actual price changes, they seem to follow this well-known pattern. But there is a subtle and profound difference. The chances of seeing the share price equivalent of a one inch or 20 foot man are very, very low. But they are not in practice zero. They really do happen.

In the jargon, this sort of pattern is known as ‘fat tails’. The further we move from the average, the more we get into the ‘tails’, in other words the parts of the distribution where the number of times we see such values is very low. We have the bulk of the price changes we observe in the ‘body’ of the distribution, as it were, and just a few examples in the tails, which are only thinly populated. With the normal distribution, this fades away quite quickly, so the ‘tail’ disappears in practice once we move a reasonable distance away from the average. If the tail is ‘fat’, it does not mean we see lots of examples of really big changes. But we do see more than we would if the pattern of changes really did follow this ‘normal’ distribution.

This may seem esoteric. Yet it is at the very heart of the financial crisis. You are on the board of a major financial institution. Your traders have carried out a series of trades not just in simple derivatives, but much more complicated variants on the basic theme. The ‘volatility’ of the price of a financial asset is, as we have seen, crucial to pricing any sort of derivative involving this asset. The traders and their managers have set up an elaborate system for telling you at any point in time what is the risk involved in these derivative trades. Such systems really did exist, sanctified under the name ‘value at risk’. The board could be told at any time how much of the value of the business was at risk.

But the assumption being made was that essentially the underlying volatilities followed the ‘normal’ pattern. Everything seems just fine, and the money rolls in. Until one day, a 20 foot tall man appears. An underlying price changes by an amount which is effectively ruled out by the assumption of normality. Your value-at-risk system is wholly worthless, as indeed your entire company might very well be given the scale of the losses the 20 foot man has caused.

The assumption of ‘normality’ in price changes was used by Long Term Capital Management. This led to a loss of a mere – a mere! - $5 billion. The losses in the current crisis dwarf this figure.

The phenomenon of ‘fat tails’ in price changes has been known since 1900, when Louis Bachelier presented his doctoral thesis in Paris. Admittedly, his work languished in obscurity for decades, but in the final quarter of the 20th century, evidence for the fat tail phenomenon began to pour in. The initial discoveries were by another French mathematician, based in America for much of his life, Benoit Mandelbrot. During the 1990s, the stream became a flood as some of the world's most distinguished statistical physicists began to take an interest in financial markets. Gene Stanley at Boston and editor of the world class journal *Physica A*, Rosario Mantegna at Palermo, Jean-Philippe Bouchaud in Paris, Yi-Cheng Zhang at Fribourg, these plus a host of their fellow scholars and graduate students examined the data on price changes in financial markets. And they found fat tails literally everywhere. Far from being unusual, the exception, fat tails were the norm. Large numbers of top quality academic papers became available on the Internet, each demonstrating the existence of fat tails in some particular aspect of financial markets.

Despite this overwhelming scientific evidence, fat tails were largely ignored in the financial markets. The result was that the potential for volatility, and in particular the potential for large changes in the prices of financial assets, was systematically underestimated. This literature was also ignored in mainstream economics.

Derivatives and discontinuous prices

There is no stage villain in the next point to be considered, no economics Nobel Prize winner whose work has ultimately proved to be misleading. Instead, the finger is pointed at conventional economic theory as a whole.

Standard economic theory relies heavily, indeed almost exclusively, on the mathematical tool of the differential calculus. The concept of partial differential equations discussed above is a particularly difficult example of the genre. At its most basic, calculus is – or, at least, it used to be - widely taught in high schools. Essentially, it is all about how one variable changes when another one does, how fast, and in what direction, up or down. It is a very powerful scientific tool which has many extremely useful applications.

It can even be very helpful in economics. The problem is when economists forget the assumptions on which it is based. To be relevant in any particular context, calculus requires a world in which changes are typically small and smooth. All the main theorems of academic economics essentially rely on the world being like this.

The Black-Scholes equation is but a particular example of this point. Anyone who takes the trouble to follow how the equation for valuing a derivative is obtained – Wikipedia is a good general reference – may see a phrase in English which seems either innocuous or incomprehensible. The equation may be followed in the text by something like: 'This partial differential equation holds whenever the price of the derivative is twice differentiable with respect to the price of the stock and once with respect to time'. Economic theory is replete with phrases such as this.

Such apparently bland qualifications disguise some very strong assumptions. For a relationship – like that between the price of the derivative and the price of the underlying stock - , for this to be 'differentiable' essentially means that across the entire range of values these prices can take, they both vary not just by small amounts at any point in time, but that these changes are smooth. That is to say, there are no 'gaps' in the levels of the prices which are observed. What does this mean in practice?

All scientific theories are approximations to reality. Their usefulness depends upon how accurate these approximations are. For much of the time, the world does appear smooth and continuous. So when, for example, Vodafone shares currently trade at 125p, there are potential buyers or sellers at 124p, 123p, 122p and so on. In other words, for each small, smooth change in the price, differing numbers of people will want to buy and sell.

But in a financial crisis, there may be *no* buyers at 124p. There may even be none at 100p or even 80. Prices do not move either smoothly or in small steps. They make massive jumps. There is no market whatsoever at any intermediate price. So, for example, a strategy based on the idea that a certain percentage of my holding in Vodafone will be sold if the price falls to 124p, and another percentage at 123p and so on, is deeply flawed. In a crisis, the price goes from 125p to 80p, say, in a single leap. And if I try to offload my shares at 80p, I might suddenly find that potential buyers are panicked by this and now will only buy at an even lower price. A strategy with its analytic basis in the tools of calculus will readily get me into very serious difficulties in a crisis.

This, in essence, is the whole point. Economists are not stupid. They do not base their theories on assumptions which are self-evidently wrong. For much of the time, the world really is similar to the sort of world described by economic theory. Things do move smoothly in small steps, we just don't see people either one inch or 20 feet tall.

The theories are nevertheless profoundly wrong. The discrepancies between theory and reality may be subtle, but they exist. The discrepancies are of no mere academic interest. They have led many of the world's leading financial institutions to base their behaviour on assumptions which are at their most devastatingly wrong at the very time they need to be right. When a crisis occurs, the assumptions on which you have based your risk strategy had better be good ones, had better be reasonable approximations to how the real world works. They were not. And we now have the financial wreckage to prove it.

Miscalculating diversification of risk

The final point I want to address on financial markets is that risk across portfolios, across holdings by an institution of assets of different types in different countries, was far less diversified than conventional approaches suggested.

This was established by a further set of papers from statistical physicists. The mathematics here, involving a concept known as random matrix theory, is at least as hard as partial differential equations, and is certainly even harder to explain in English. But the results obtained from applying it to financial markets have very important implications. In the interests of space, these are merely stated and described briefly rather than an attempt being made to explain the underlying logic.

In the illustrative example given above, I own some Vodafone shares. However well or badly I measure it in practice, I know that owning these shares carries a risk. The price may fall. So it seems prudent to spread my risk by buying shares in other companies, British Petroleum say. I may diversify my holdings and buy shares trading on markets outside Britain, or even buy financial assets other than shares, such as government bonds or commodities such as oil or gold.

By doing this, if Vodafone shares fall in value, the other assets in my portfolio might not do. Indeed, if I make shrewd enough choices, some of these assets might tend to rise in value when Vodafone goes down, and vice-versa.

Just as with the concept of the volatility of a particular financial asset, there is a way of measuring the extent to which my portfolio of assets is diversified. In other words, of measuring the extent to which the prices of the individual assets in it tend to move either in the same or in a different direction. The more diversity I can get, the more insulated against risk I am.

The American economist Harry Markowitz, who received the Nobel Prize in 1990, was the brilliant pioneer of the approach widely used to measure this diversity by financial institutions. But, just as with Merton and Scholes, the work of the statistical physicists has shown that this received wisdom, this standard way of doing things, can be very misleading in practice.

Essentially, what random matrix theory analysis shows is that the standard measures of diversification usually over-estimate, sometimes to a very considerable degree, the extent to which portfolios are genuinely diversified. The prices of financial assets have a greater tendency to move together than the standard analysis suggests. To put it crudely, when one market collapses, they all do. Of course, the analysis is considerably more sophisticated than that. But the standard approach to risk made insufficient provision for the possibility that all markets would go down together.

So in financial markets, conventional ideas underestimated the potential for volatility, effectively ignored the fact that in a crisis there may not be a price at which an asset can be sold, and over-estimated the extent to which portfolios had been diversified and the overall risk across a holding of assets thereby reduced.

The 'real' economy and the problem of equilibrium

As everyone knows, the past six months or so have seen tremendous turmoil in the world economy. The most obvious example of this is in financial markets. But overall output is also now falling sharply, in countries like Japan and Germany, faster than at any time since the 1930s.

It seems obvious that these economies are *not* in a state of equilibrium, where everything flows smoothly. We are a long, long way away from any sort of equilibrium. And in the build up to the crash, we were out of sync in the other directions, with huge asset price bubbles, and unsustainable build ups of debt.

'General' equilibrium: the background

What, then, of the real economy, of output, of spending, of employment and unemployment? All scientific theories, to emphasise the point yet again, are approximations to reality. This is true even of, say, quantum mechanics, a theory which has survived the most rigorous empirical testing. But even this is not absolutely identical to reality.

In essence, theories are only as good as the assumptions on which they are based. Since the 1970s mainstream economics has in fact made substantial progress in making its assumptions correspond more closely with the real world. The caricature of the all-seeing, all-knowing Rational Economic Person still exists in the mass market student textbooks. Nowadays, in the leading academic journal articles, he or she might be flawed. In particular, the amount of information available to the decision maker might be far from complete, far from what is required to operate in a totally rational manner.

But economics retains, deep at its heart, the belief that the economy is an equilibrium system. In other words, the way in which decisions are taken – by individuals, firms, governments – and the feedbacks which are created as a result of these decisions together mean that the economy has a natural tendency to be in a state of equilibrium

The price mechanism plays an absolutely fundamental role in ensuring this is the case. If imbalances arise, prices will change. And these changes give signals to people to alter their decisions. In this way, equilibrium will be restored.

As an example, suppose an economy receives an adverse shock of some description, and as a result some workers find themselves unemployed. The supply of labour is now greater than the demand. In an ideal world, provided they are willing to reduce their wages to an appropriate level, employers will once again find it attractive to offer them jobs, and equilibrium will be restored. The wage from the point of view of the workers is the price at which he or she can be bought, or employed to put it less brutally, by a prospective employer. Reducing the price leads to an increase in demand.

One very serious problem for this view of the world until the 1970s was that, demonstrably, there simply weren't enough prices in the real world to guarantee that this process of equilibrium might work. This may seem a strange statement to make. Indeed, it *is* strange, high economic theory is definitely not for the faint hearted.

The difficulty arose from the fact that there were very few markets in existence, and hence very few prices, which dealt with the future. For an entire century, ever since the 1870s, many of the best minds had struggled with the problem of how to incorporate the future into their core theoretical model, the so-called general equilibrium model.

'Equilibrium' for reasons already discussed, and 'general' because the theory purports to describe how equilibrium – supply equal to demand – can arise not just in one or two individual markets, but generally across the economy as a whole, in *all* markets. The technical details need not concern us, but the intellectual and mathematical challenge of this problem was immense. No fewer than 7 out of the first 11 winners of the Nobel Prize in economics received it for their work on general equilibrium. Eventually, by the early 1970s, the problem had finally been solved completely.

I say 'completely' but it is important to realise that the theory in this guise related to an economy with a fixed amount of resources, whether of land, labour, energy or capital. The theory essentially told us about the optimal, the most efficient, allocation of a given set of resources. But it was about a static and not a dynamic, growing economy.

I heard a practical example of the lack of prices about the future given last year at a conference in Warsaw. Axel Leijonhufvud wrote an absolutely brilliant book 40 years ago, *Keynesian Economics and the Economics of Keynes*. As a young lecturer he had spent some time in Buenos Aires at a time when Argentina was undergoing one of its fairly frequent bouts of hyperinflation, with prices rising at a rate of many thousands per cent a year. Axel really liked a corner shop near his flat, where he was well known to the owner. One morning, he went round. It was closed. On the door was a notice 'Closed for lack of prices'. Axel was intrigued, and went to see the owner. It was very simple, he said. He could sell his products today, but prices were rising so fast he didn't know how much it would cost him to re-stock his shop tomorrow. The business had become too uncertain, so he had closed it down. There simply wasn't a price in existence today at which he could strike a contract to re-stock any particular product tomorrow.

Black, Merton and Scholes appeared to have solved this problem. This is why they received the Nobel Prize – Black died before he could be nominated - , not because they had enabled the derivatives industry to be created. They appeared to have solved the problem about how to set prices today for transactions which might take place in the future in a world in which there was uncertainty. In principle, they had filled a very important gap in the availability of prices.

Equilibrium in modern economics

Mainstream economics has moved on from the general equilibrium problem which dominated high, mainstream theory for a hundred years. The major project of the past 30-odd years has been to try to use equilibrium theory to explain the dynamic fluctuations which have been observed in the developed, market-oriented economies ever since the Industrial Revolution.

There are two aspects of this. First, the slow but steady growth in output over time, averaging around 2 per cent a year per head of population. It is this long-run growth which distinguishes capitalism from all other forms of social and economic organisation in human history. This is a major topic in its own right, but I am merely mentioning it here.

The second is the persistent short-term fluctuations in output around this underlying slow growth. From time to time, these fluctuations are severe and output actually falls for a period of time, before growth is resumed. We are living through one of these periods right now – the recession. How modern economics attempts to explain these fluctuations is the final theme of this lecture.

A theory which is based upon equilibrium appears to have an inherent problem when confronted with data such as that in the chart below. This shows annual percentage changes in total output in the UK from 1900 to 2008. It is entirely typical of the Western economies.

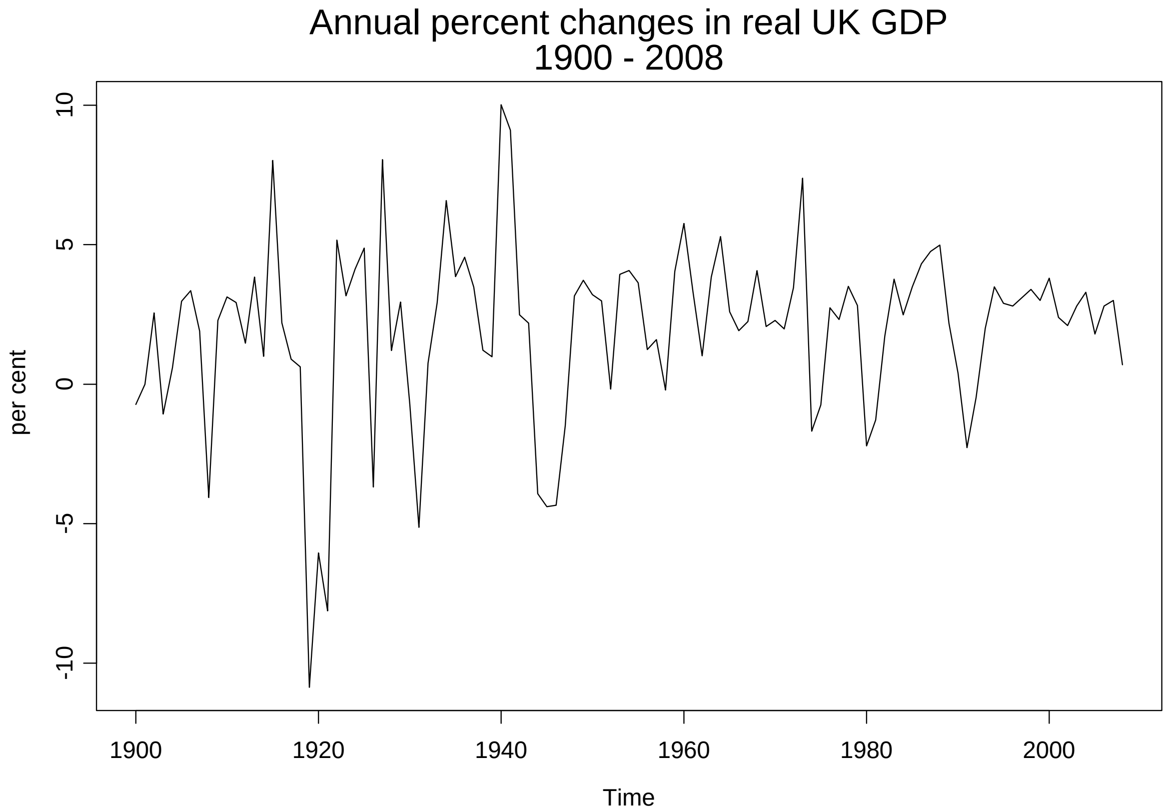

Attempting to understand why these fluctuations take place is a very difficult problem. If it were understood in the same way as, say, building bridges were understood, we would have a pretty complete grasp as to why we, along with all the other developed economies, are in our current predicament. But we do not.

The first serious modern attempts to explain these fluctuations in the context of 'rational' decision makers and equilibrium was something called real business cycle theory, which emerged in the 1970s and 1980s. The 'business cycle', for information, is the jargon used by economists to describe the persistent short-run fluctuations in output of the sort shown in the figure above. For any natural scientists present, I should emphasise that the phrase is purely a descriptive one, economists are aware that the data in the chart do not obviously follow what a scientist would think of as cyclical behaviour.

If decision makers, whether individuals, firms or governments, are rational and gather all available information before making the best possible decision, a key question is: why do they never learn the causes of these fluctuations? If they did so, they would be able to anticipate them and hence be able to smooth them out, perhaps even cause them to disappear completely.

Essentially, real business cycle theory gets round the problem by postulating that cycles are caused by random shocks to the economy. By definition, random events can never be anticipated, even by the cleverest decision maker. Of course, by itself this is not a theory at all, merely a restatement of the problem in a different form. Any movements over time in any series could be 'explained' by invoking a series of random shocks to what would otherwise be a smooth series.

There is far more to the theory than this, and it is here that we need to take a deep breath. Consider for the moment an individual deciding on how to allocate his or her time between work and leisure. At any particular point in time, there may very well be constraints on a person's ability to switch between the two. But if your current employer, for example, insists you work a certain number of hours a week and you want to work less, eventually you can find a firm which will allow you to meet your desire for more leisure. So the assumption that an individual can choose between work and leisure is perfectly reasonable.

Much less reasonable is the assumption which is made that the economy can be analysed by assuming a single so-called 'representative' agent. All differences between individuals are assumed away. At one level, this is not as unreasonable as it might seem. All theories involve simplifications. Individuals do differ in their tastes in terms of their preference between work and leisure, so at one level it is a handy simplification to pretend these do not exist, and that we just have one 'representative' individual. However, as we shall see below, this assumption has profound implications in the current crisis.

Real business cycle theory usually identifies the random shocks with shocks to technology. If technological advance is particularly rapid, the underlying long-term growth rate of the economy is increased. And the opposite applies is we are going through a period when discovery and invention is proceeding only slowly.

The theory postulates that an individual selects the optimal allocation of time between work and leisure given the underlying growth rate of the economy. If it is particularly fast, people will choose to work more. This is because they can earn more in an hour now than they could do if the underlying growth rate was especially slow. Conversely, following a negative technology shock – though no-one has specified satisfactorily what one of these looks like – individuals will choose to work less. Indeed, some may choose not to work at all during such periods. To the untrained observer, it may seem that these people are unemployed, for they are not in work. But to the perceptive real business cycle devotee, they are simply making an optimal choice between work and leisure.

Choosing not to work is an efficient, indeed the optimal, response to a negative technology shock according to real business cycle theory. So if a government, for example, foolishly attempted to create jobs and force apparently unemployed people to work, they would not be doing people a service at all. It may seem that they are better off working and having a higher income, but in the eyes of the real business cycle theorist, they really prefer to be unemployed.

There are further subtleties around real business cycle theory, and unlike most economic theories its practitioners have attempted to confront the theory with empirical evidence.

But the theory never really took complete hold, even within the academic mainstream with its devotion to the idea of optimal decision making and to equilibrium outcomes. A major problem was the experience of the 1930s during the Great Depression. The unemployment rate in Germany rose to 17.2 per cent, one in six of all workers. In Canada 19.3 per cent, one in five, and in America 22.3 per cent, nearly one in four[1]. The idea that as many as one in every four workers were voluntarily choosing not to work, despite the conditions of serious poverty which they experienced, was just too hard for some mainstream economists to swallow. All economists would accept the possibility that some people might prefer to live on benefits than to work, but that one in four might do so was a step too far. Nevertheless, the theory had strong support, and two of its early developers, Finn Kydland and Edward Prescott, received the Nobel Prize for their work.

The more modern development of real business cycle theory goes under the name 'dynamic stochastic general equilibrium' theory, or DSGE for short. Like real business cycle theory, the approach posits that the economy can be understood by reference to a single so-called

[1] The estimates of unemployment rates vary slightly across different sources, these are from A Maddison, *Phases of Capitalist Development*, OUP, 1982

'representative agent', who takes decisions to maximises his/her utility over an infinite time horizon. The concept of random, exogenous shocks is also retained. In short, the approach is based upon the core model of individual behaviour in economic theory, that of the rational agent who forms rational expectations about the future.

This RARE view – rational agent, rational expectations – of how individuals behave retains a very powerful grip on mainstream economists. Almost the whole literature of another social science, psychology, suggests that in general this assumption is a poor approximation to how people really do behave. But the RARE assumption allows economists to manipulate difficult mathematics and thereby sustain the illusion that they really are just as good at science as physicists after all. Certainly, RARE satisfies the demand from clever young men – and the authors of DSGE papers are almost invariably men – to demonstrate their cleverness without the need to have any original ideas.

Without going into the technical details, DSGE adds a lot more detail to real business cycle theory in terms of the institutional constraints under which the 'representative' agent operates – for example, the rules of fiscal or monetary policy.

Although they are difficult to construct, these models have rapidly become very influential in academic economics. For example, the American Economic Association has just launched a new journal. Its title? '*Macroeconomics*'. So this is exactly the place we should be looking if we want to know what academic economists are thinking about what is going on, how we got here.

Economists are notorious in the public mind for disagreeing with each other. But it turns out that, far from being riven by dissent, the academic profession believes it has reached a broad consensus. Indeed, the first issue carries an article by one of the world's leading academic macroeconomists, Michael Woodford, entitled 'Convergence in Macroeconomics: Elements of the New Synthesis'.

The first and most important part of the new synthesis is that 'it is now widely agreed that macroeconomic analysis should employ models with coherent intertemporal general equilibrium foundations'. What does this mean in English? It means a) people and firms act in a rational, coherent manner b) they assess coolly and rationally the future consequences of decisions they take now c) the key driving force underlying the economy is a tendency for it to revert to equilibrium.

Nor is the spread of DSGE models confined to academia. These sorts of models – so-called 'dynamic stochastic general equilibrium models' – have become more and more influential in recent years in central banks and treasuries around the world. A distinguished friend of mine, despite being a highly regarded mainstream macroeconomist, went to a seminar in the Bank of England at the very height of the crisis last autumn. Many of their brightest economists sat down calmly discussing the finer mathematical points of general equilibrium models. He felt rather like Banquo's Ghost, except no-one paid the slightest attention to him at all.

There are two key assumptions to mention in these models. First, they require 'complete markets'. This phrase would be pretty meaningless had we not come across a similar concept in the Black-Scholes derivatives model. This, we remember, helped fill a big gap in the types of markets which exist in practice. Even so, the set of markets which does exist is far from complete. In such a situation, all possible transactions in all possible future states of the world can be constructed with existing assets. Transparently, this is not true.

A key factor in the current crisis is how debtors and creditors react to events. These form two very distinct groups, whose behaviour may very well not be the simple mirror image of each other. Yet the DSGE models posit that the behaviour of the economy as a whole can be modelled with reference to the single, representative agent. So it is not possible to have different behaviour of debtors and creditors in these models. Incredibly, given current circumstances, the way the models are constructed means that money, banks, credit and debt do not affect the real economy – output and employment – in the long-run. No wonder the policymakers received such bad advice!

Final remarks

These will be very brief. I have already covered a fairly large amount of rather technical material. I would just like to say that there *is* a strong tradition in economics which emphasises *dis*equilibrium - Schumpeter, Keynes, Hayek, for example. It certainly doesn't mean that everything they wrote was correct, far from it. But these are some of the most famous names in economics, even though their work has never really been absorbed in the mainstream.

The exciting area of complex systems is a natural one in which to analyse the economy. The economy is made up of large number of individual agents – consumers, firms – who interact with each other at the micro level. From these interactions, the macro features of the economy emerge. The key to progress, however, is to use much more realistic models of behaviour at the micro level. As 2002 Nobel Prize winner Daniel Kahneman put it 'rational models are psychologically unrealistic….. the central characteristic of agents is not that they reason poorly, but that they often act intuitively. And the behavior of these agents is not guided by what they are able to compute, but by what they happen to see at a given moment.'[2]

Workshops like the current one organised by Strathclyde's Institute of Advanced Studies play an important part in the task of building a more realistic discipline of economics, one which is firmly based in the empirical evidence on how people really do behave and not on a set of seemingly logical assumptions about how a 'rational' person should behave.

Finally, am I going to answer the question in the title of the talk? Well, I suppose I retain an element of perhaps not rationality but certainly self-preservation. I already got into a massive amount of trouble with a best-selling book I wrote 15 years ago called 'Death of Economics'. So I have just set out the evidence. It is up to you to draw your own conclusion.

Paul Ormerod June 2009 pormerod@volterra.co.uk

[2] D Kahneman, 'Maps of bounded rationality: psychology for behavioral economics', *Am Ec Rev.*, 93, 1449-1475, 2003