Ostrich Economics

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 essay 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 essay, I focus on three aspects of financial markets where the conventional thinking in economics proved to be seriously misguided. They are the potential for volatility, the ability to diversify risk, and the lack of smooth adjustment to change. 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, is 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.

The most serious problems for financial institutions arose in the often bizarre world of derivatives. 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 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.

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. The bigger the gap between the two, the smaller will be the price of the derivative itself. 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. So with a big gap, I am not willing to pay very much for that right.

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.

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.

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 was likely to move, the assumption was made in models that the movements in the price followed a pattern known as the Gaussian or 'normal' distribution. 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.

By assuming that the pattern of these changes followed 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-Change 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.

A further set of papers from the statistical physicists established an additional way in which the standard value-at-risk approach was inadequate. 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.

There is no stage villain in the final 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 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.

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.

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.