There is a huge amount going on outside economics. Social and economic questions are being analysed in a quantitative way by computer scientists, anthropologists, sociologists, mathematicians, physicists, biologists.

Not everything they do makes sense, but there is an awful lot which economists could learn from these disciplines. My first theme today is on what economics can learn from physics. A separate discipline has emerged over the past 10 years, that of econophysics. It has its own community, conferences, websites. Very few economists engage with it.

One key thing which economics can learn from physics is that empirical discoveries are important and motivate the development of theory. In physics as in economics, theoretical advances gain kudos. But serious empirical discoveries also attract high status, which isn't really the case in much of economics.

Well, to some extent what I've just said is only partially true. Vernon Smith and Daniel Kahneman got the Nobel Prize in 2002 for their work in experimental economics, how people actually behave. Kahneman concludes in his Nobel lecture that 'humans reason poorly and act intuitively'. But we don't actually see this in much of economic theory, rational agent models of the macro economy, were all the rage in academia, central banks and regulators right up to the credit crunch in autumn 2008.

Physicists made two very important empirical discoveries about financial markets in the late 1990s, and those making the discoveries got a lot of respect from their fellow econophysicists for them. Essentially, these discoveries completely undermined the conventional theory of risk and returns on portfolios, the so-called mean-variance analysis for which Harry Markowitz got the Nobel Prize in 1990.

Changes in financial asset prices do not follow the normal, or Gaussian, distribution, but are fat tailed. In other words, the normal is a good approximation over most of the distribution, but large changes, whilst rare, are much more likely than under the standard assumption that the changes follow a normal distribution over the whole range.

The second related to the covariance amongst financial assets, whether individual assets or classes of assets. Empirical co-variance matrices are poorly determined, dominated by noise rather than true information, and far from being constant, the individual correlations fluctuate all the time.

Theory still isn't able to account for many of these findings in a satisfactory way, they are very hard problems, but at least we know what has to be done.

The second general point I want to make is drawn from a wide range of disciplines. This relates to what mathematicians refer to as graph theory, but everyone else calls networks. Essentially, this is about how ideas/behaviour/choices spread across networks of individuals, firms, or whatever agents we are focussing on.

The crucial implication of this work is that the tastes and preferences of agents are not, as standard economic theory assumes, fixed, but can be altered directly by observing the behaviour of others. There are examples, mainly in mature fast-moving consumer goods markets, where the assumption of fixed preferences is reasonable. Most people either like or hate Marmite, and their preferences don't really alter on this. But in general, this assumption is simply not true.

Now, economists have done work in this area, especially in financial markets. They often refer to the phenomenon as 'herding'. But there is a huge literature outside economics on both the theoretical and empirical properties of networks, of structures across which the tastes and preferences of agents can be altered.

The people involved in this work come from a wide range of disciplines, they don't carry the economists' baggage of having to assume fixed tastes and preferences in order to get analytical solutions to maximising equations. Instead, they investigate the world as it actually is, not as it is deemed to be.

The particular type of network which is relevant to any given situation will vary. And the theoretical results show that the properties of different types of networks can vary dramatically. For example, an influential popular book was Malcolm Gladwell's *Tipping Points*. This popularised the concept of something called a scale-free network. These are networks in which there are a few key players, agents with large numbers of connections, and in contrast most agents have a small number of connections. The structure of the Internet is approximated by a scale free network, a few sites get lots of hits, most get relatively few.

The highly connected agents – the 'influentials' – exercise a decisive weight in whether or not a particular kind of behaviour spreads across the network or not. But there are also many subtle features.

Major marketing departments are spending huge amounts trying to discover the 'influentials'. But they might not exist. Many networks amongst individuals are 'small world', which is essentially overlapping sets of 'friends of friends'.

The point here is that a very large amount of theoretical and empirical work has been done on social and economic questions using the network approach. Tastes and preferences are not in general fixed. Economists need to know about and use this literature.

An important class of decision rule in such contexts was described as long ago as 1973 by Thomas Schelling, Nobel Prize winner in 2005. He called it 'binary decisions with externalities'. It is binary in the sense that you either buy a product or you don't, you adopt a particular mode of behaviour or you don't. It has externalities in the sense that the behaviour of any single agent in itself can alter the behaviour of others. It is a simple but very powerful concept which has been usefully applied in areas such as fashions, riots, crime, competing technologies, the spread of innovations, the emergence of social norms to name but a few.

Duncan Watts, a mathematical sociologist (yes, they do exist!) at Columbia, and now head of consumer research at Yahoo, had a brilliant paper on this in 2002. Agents are initially in state zero of the world – they haven't adopted an idea or bought a product. A few are selected at random to choose state one of the world. Agents are allocated at random a level of 'persuadability': what proportion of people you are connected to in this context – people you take notice of – need to be in state one for you to switch to it from state zero, and vice versa.

It is simple, but a good approximation to many practical situations. A key question Watts examined was: how do cascades spread across this network, how far do people switch to state one of the world.

He programmed the model and obtained large numbers of solutions. In each one, the initial conditions were identical: a small number of agents is selected at random to be in state one of the world. Yet the outcomes are massively different. Most of the time, the cascade is small, few change their behaviour. But occasionally, a cascade on a global scale takes place.

This is a profound piece of work. In networked systems, the common sense view of the connections between the scale of an event and its consequences no longer holds. Small events, the initial choice by a few agents of state one, can have dramatic consequences. Further, systems which are exposed to continuous shocks for long periods of time may suddenly and seemingly inexplicably exhibit a large cascade.

A clear implication is that empirical analyses of any situation in which networks are important which fails to take account of phenomena such as this are inherently flawed. Changes may be wrongly ascribed to external drivers when they arise as inherent features of the network structure. Of course more complicated models can be built, in which they are both external drivers (e.g. price) and network effects. But the principle still holds.

The final point follows from the technology which Watts used, and which is a distinguishing feature of such models. The personal computer frees us from the constraints of requiring analytical results in order to understand the implications of a hypothesis. Thirty years ago, this was a reasonable, indeed almost the only, way to proceed. But now, agent based models can be readily programmed and their implications explored using simulation techniques. There is an explosion of interest in this methodology in other disciplines, and economics risks getting left behind. Analytical results are nice to have, and some of the smart people can occasionally obtain them in these contexts as well. But we don't rely on the abacus or the slide rule any more. So we should use simulation rather than be constrained by relying on models which have analytical solutions.

So, in summary:

- Accord much more respect to purely empirical discoveries
- Incorporate the vast literature on networks, and make fixed tastes and preferences the special and not the general case
- Use simulation techniques rather than be constrained by only using models with analytical solutions
-

Where do we go from here?