Seminar Notes On ‘The Physics of Institutions’.

Abstract: Philip Ball traces the development of statistical physics, first proposed by James Clerk Maxwell and Ludwig Boltzman and shows how its principles can be used to understand human systems. He shows how rules of interaction between agents can give rise to such phenomena as phase change and self organised criticality and looks at the use of such models for understanding traffic states and the evolution of business organisations, as well as other social science issues such as the effect of social forces on marriage. Paul Ormerod looks at models used to tackle racial segregation, financial markets and crime studies and suggests how powerful insights into the aggregate properties of human organisations can be gained using quite simple agent characterisation and rules of interaction.

(These notes constitute an edited version of the presentations and discussions. Images referred to in the text are featured in the handouts, but full slide presentations can be found at the end. Because of sound recording difficulties, only the main presenters are referred to by name).

Director of Complexity Research : Eve Mitleton- Kelly
London School of Economics
Houghton Street
London WC2A 2A

Presenters : Philip Ball,
Paul Ormerod,

Compiled For The L.S.E.
by Geoffrey J.C. Higgs 26/3/03

Philip Ball

Suppose you want answers to the following sorts of questions:
'What set of policies would guarantee a party electoral victory?'
'How do companies band together to form alliances and conglomerates?'
'What set of international policies will encourage democracy and discourage conflict?'
'What will the stock market do tomorrow?'
'How will congestion charging affect London's traffic in three years time?'
'How might harsher sentencing affect crime statistics?'
'What is the likely lifetime of a new small business?'
'What are the chances of you and I sharing a mutual friend?'

The ability to predict some of these would be useful, some immensely beneficial. Some would be so valuable that those who possessed the answer would want to keep them secret. All of them are desirable to certain parties, but which might be possible and which are just idle fantasies? In other words which if any aspects of the evolution of society are susceptible to probabilistic estimation and which might be too dependent on the vicissitudes of human behaviour to be accessible to any degree of prediction. Theories about the best way for society to operate reach back as far as Plato's 'Republic', but the notion of approaching such questions using the methods of science, that is, developing a social science worthy of the name dates back to the beginnings of the Enlightenment in the early 17th Century.

That was an age of mechanism, when people like Galileo and Descartes and Newton were starting to propose that nature could be understood like a machine in which forces acting between the component parts give rise to precise mathematical laws that allowed, in some sense, the future to be predicted. This was how Galileo understood the laws governing the motion of objects and it led Newton to formulate his laws of gravitation that allowed scientists not just to empirically predict but to appreciate the underlying basis for the regular motions of the planets. As more and more of nature began to reveal itself governed by physical laws philosophers started to wonder whether such laws applied in the human sphere was well. They began to think of the human organism as a well oiled machine, with gears and levers and pumps controlling it. And if human individuals made up society then maybe there was a physics of society.

It was Thomas Hobbs, who in the 1630's and 1640's used Galileo's physics of motion to derive the conclusion that the best way to govern a nation was by absolute despotism., a notion that was deservedly soon forgotten. But over the last two decades physicists and other scientists have regained an interest in trying to apply the principles of science to social phenomena. Unlike Hobb's they ask not 'How should we govern?' or 'How should we construct our institutions', but rather 'If we set things up according to this or that particular set of rules, can we predict what the likely outcome will be?' Science is not being used to tell us the right or wrong way to do things, but to try to understand which choices lead to which consequences. It is naive to think that we can set up the exact conditions or policies needed to achieve a particular objective though sometimes public policy fails simply because it neglects certain aspects of human psychology. There are also many situations in nature where even if one could in principle account for every relevant factor, the outcome of a particular set of rules or conditions might be quite different from what one expects. And this can be the case in a physical situation such as traffic control where building a new road leads to greater road congestion rather than less. What I want to do is to see whether science can give us some tools that permit a better prediction of the consequences of social decision-making in various contexts.

Now there are certainly people in the audience who know far more than I ever shall about social and economic science and I felt that perhaps the most useful thing that I could do was, as a physicist, to outline some of the central issues and concepts in contemporary physics that seem to have some applicability in these fields. In other words to give you a flavour of what there is out there that might be of some value. In seeking to extend physical science to social science, we might think we are asking 'Are there laws of society in the same way that there are laws of gravitation or electromagnetism', and certainly that's how some of the early pioneers in the field saw it.

The French philosopher August Compte believed that laws like this could be uncovered and he coined the term 'Physics Sociale', and in the book that he wrote in the 1830's: entitled 'A System of Positive philosophy', he argued that this would complete the description of the world that Galileo and Newton had begun. Several other thinkers in the 18th and 19th Century, including Emmanuel Kant, Thomas Buckle and Leo Tolstoy, wondered whether there is some inevitability in the way history advances, such that an understanding of the forces driving it could lead to a more or less certain prediction of its future course.

One of the key observations that led to such positivistic thoughts was the regularity in the statistics of social phenomena. Scientists and philosophers became interested in social statistics in the 17th Century when the London businessman John Graunt began to collect yearly mortality figures for the City of London. Graunt argued that mortality statistics could provide a solid empirical basis for formulating political policy and this view was also shared by the famous astronomer Edmund Halley who was one of Newton's few close friends.

What people began to realise was that there was a certain predictability in theses social statistics. It wasn't just that, on average more, or less the same number of people died each year, or even that a constancy also applied to subdivisions of society such as age or profession etc. It was also that deviations from the averages were interesting and by the early 19th Century mathematicians like Pierre-Simon Laplace, discovered that a whole variety of statistical data could be fitted onto a single mathematical curve. This has become known as the Gaussian or 'bell curve' and describes a probability distribution. It can be seen both as a summary of empirical data and having a predictive function.

Image 4 - The Gaussian distribution

In this diagram h might be the heights of adults in London. We collect the numbers of people at different heights and draw the graph. What it tells us is the probability that any randomly selected individual will have a particular height. The most probable height is the highest point (which is also the average) and the graph falls off fairly sharply for either extreme. To people's surprise they found that Gaussian curves described not only the statistics of births and deaths, over which individuals have little if any control, but also volitional acts. such as crimes or marriages and to some this seemed an affront to the idea of free will. How could it be that supposedly free choices were governed by this mathematical law?

If there is a physics of society it will be essentially a statistical one because mathematical regularities only appear when we look at populations or large data sets. Conversely this means is that in general specific predictions must be probabilistic; we cannot say what will happen to any of the individual components or agents of a system only what the probabilities of the various possible outcomes are. In 1862 John Stuart Mill recognised this statistical aspect of scientific sociology when he said 'very events which in their own nature appear most capricious and uncertain which in any individual case no attainable degree of knowledge would enable us to see, occur when considerable numbers are taken into account with a degree of regularity approaching mathematical'. And this is where the connection with physics comes in.

However Newton had speculated that the trajectories of celestial bodies could be understood and predicted on the basis of the forces of gravity acting between them and anticipated that the same was true of matter at the other extreme, the scale of atoms. They could be understood by the laws of motion determined by the inter-atomic forces, even if no one knew what those forces were, and well in the 19th Century physicists thought of the atomic world as a kind of billiards game in which the atoms were like smooth hard balls that travelled through space and collided with one another according to Newton's laws of motion. The only difficulty was that they couldn't hope to get close enough to see and to measure all the motions and even if they could, atoms are so numerous that it would be impossible to keep track of all the trajectories.

Ironically it was actually the statistical regularities seen in the social sciences that encouraged the physicist James Clerk Maxwell to propose that even if we can't use Newtonian mechanics to formulate a complete description of the atomic scale behaviour of matter, we can anticipate that mathematical laws will arise out of the average interdependent motions of all these particles. He began to think about the probability distribution or atomic motions and he assumed that these would also be described by Gaussian type curves. And this led Maxwell and Ludwig Boltzmann to formulate the science now known as statistical mechanics or now more generally known as statistical physics, in which the bulk scale behaviour of matter, such as the known mathematical relationships between the pressure, temperature, and volume of a gas can be understood to emerge from the microscopic motions, the inscrutable particle motions of the atoms involved. This branch of science is now used to understand just about all the properties of everyday matter from liquids to polymers to superconductors and physicists started to ask whether we might we see in society some of the same phenomena we find in collections of interactive particles. If we can substitute atoms and molecules by people or cars or market traders or businesses can we use statistical physics to understand some of the phenomena that arise in the real world?

Still there's an obvious objection to seeing people as Newtonian automatons and it was expressed by the economist Robert Heilbronner :

"There is an unbridgeable gap between the behavior of subatomic particles and those human beings who constitute the objects of study of social science. Aside from pure reflexes human behaviour cannot be understood without the concept of volition, the unpredictable capacity to change our minds up to the very last moment. By way of contrast the elements of nature behave as they do for reasons of which we know only one thing, that the particles of physics do not choose to behave as they do".

But I think this risks overestimating both the power and the scope of free will. In many social situations it's unrealistic or even meaningless to assume that we can do whatever we want and we often have a very tightly constrained range of choices. In principle if we are driving a car we can steer it anywhere and at whatever speed the vehicle allows, but of course we don't. We tend to drive along a line in the road on the left hand side at a speed appropriate to the circumstances and going from our point of departure to our destination. When we vote we chose one candidate or another generally from a short list of alternatives and we have a similarly limited range of behaviour if we do something like shopping. Our actions which are nominally completely free are constrained by a wide variety of factors, social norms, conventions, economic necessities and so on. So we are far more predictable than we like to believe.

The key factor and one I think that often economics and social scientists have tended to overlook in their models in the past but is intrinsic to statistical physics is interaction. We're affected by one another. People don't drive at 80 miles an hour down Oxford Street because there are others in the way and we normally aim to avoid collisions. We might say that there appears to be a kind of repulsive force between the vehicles that keeps them apart though of course there is no real force that we can measure. Yet if we were to make a model of our behaviour the metaphor would hold. Our choices are influenced by all manner of things and particularly by what our peers do. If everyone on the stock market floor is selling it takes either a very astute or a very slow-witted trader to buck the trend and start buying. This kind of herd-like behaviour is of course well known in economics. Even in elections we might imagine that with secret ballots we're all making our own personal choice, but there turns out to be a very clear signature in the statistics of our collective behaviour; the fact that people are strongly influenced by what others do.

This interactive behaviour shows up in the probability distributions of the statistics. Independent, apparently random events may show up as a Gaussian curve but deviations from it are generally a sign that the agents in the system are not behaving independently, but are feeling the influence of mutual interaction.This is the kind of simple diagnostic tool that is learned from statistical physics and there's an important corollary to this. In the social sciences there's a strong tradition of creating psychological models of phenomena, of trying to understand social behaviour on the basis of individual psychology and social biologists like E. O. Wilson have argued that social science could be made more scientific if theses models were more firmly rooted in the evolutionary biological origins of individual behaviour. Now there's a case for saying that, but it makes the unwarranted assumption that social behaviour is a straightforward extrapolation of individual behaviour and it seems that this is often not the case at all. The behaviour of a human group; how it organises itself into institutions, for example, can't be deduced or predicted from the predilections of an individual. And it's very clear in statistical physics that even in inanimate systems that once the constituents start to interact, completely new collective modes of behaviour can arise. We can study the behaviour of a single water molecule as closely as we like but we would never be able to predict solely from that, that water is a substance that freezes at 0oC and boils at 100o C. We can get that only by looking at the water molecules collectively.

Why water condenses from a vapour to a liquid was explained in the 1870's by Johannes Diderik Van der Waals who showed that it followed from the existence of both attractive and repulsive forces between the particles in a gas. Maxwell and Boltzmann who just treated atoms as hard balls that collided but otherwise did not affect each other could not explain this, but Van der Waals found that when you included the forces the theory predicted that there would exist a liquid state as well as a gaseous state. When you heat a substance its particles, atoms or molecules jig around more frantically and this can overcome the attractive forces that tend to hold the particles together and so we can understand why a substance can change from a solid to a liquid to a gas. But it was less clear why these changes should happen suddenly As ice is warmed to melting point it doesn't get progressively softer and jelly-like. Instead it stays hard until it melts abruptly at zero degrees. And the same with evaporation; water is either a liquid or gas and not something in between. These sudden changes are called phase transitions, in this case between the solid liquid and gaseous states of matter. Van der Waal's theory showed how transitions happen and why they are sudden. In social science and in politics there is a tendency to think that effects happen in proportion to their cause. Small changes have small effects and in statistical physics that clearly isn't always so. Consider density. There is a huge change in density as a vapour is cooled past its boiling point as can be seen from the diagram.

Image 7 - Cooling a vapour past its boiling point

Physicists call this behaviour 'non-linear' since there is no simple straight line relationship between cause and effect and phase transitions like this occur in all areas of modern physics. A magnet offers a second example. If you heat a piece of magnetic iron to 770o C it loses it's magnetism. Below this temperature it's magnetic, above it's non-magnetic so there is an abrupt change from one state to another though it's not a jump in the degree to which it is magnetised. The magnetisation falls smoothly though rapidly to zero at the Curie temperature as shown in Image 8. This is called a 'critical phase transition and the point at which it happens the 'critical point'. Liquids and gases also have 'critical points' at which there ceases to be any difference between the gas state and the liquid state. :

Image 8 - Critical transition

What happens in the case of gas to liquid is that you get curves as in Image 8 but if you increase the pressure the jump gets smaller and smaller until it disappears. Critical points are an important class of phase transitions.

Magnets can be considered to have critical points if they are envisaged as an array of atoms in which each is like a tiny bar magnet or needle which can point in one direction or another; either up or down we could say. This is called the 'Ising Model' of magnetism shown diagrammatically as Image 9 :

Image 9 -The Ising model

If there are more atomic magnets pointing in one direction than the other then they all add up to give an overall magnetisation for the material. In a substance like iron each of theses atomic needles feels the magnetic field of its neighbours and these interactions tend to make all the arrows line up in the same direction. At low temperatures that is what they do and the material is magnetic. As you heat the magnet up, heat randomises the direction of some of these needles and if they become completely randomised then on average they all cancel out and there is no net magnetism. These are the two states but in fact there are two different ways in which this can happen; either the magnets could all point up or they could all point down. So there are two equivalent but different magnetic states and the critical point is when islands or areas of the same magnetic state reach a certain dominance of the whole system. This happens because when an atomic magnet has flipped it exerts a force on its neighbours that tries to make them flip as well. So there's a collective behaviour dependant on the interactions. As we approach the critical point more needles get flipped out of the uniformly aligned state and small patches of the opposite magnetisation start to arise. At first the patches are quite small and do not do much to affect the overall magnetisation but as we get closer and closer to the critical point these regions grow bigger and bigger, but not equally big. What we find is that there are flipped regions of all sizes, from a single atom to whole patches that start to approach the size of the entire system. At the critical point we cannot tell which are the ones in the original state and which are the ones that have been flipped because there are equal amounts of both though they are not evenly distributed'.

Image 10 - The critical point

Physicists often plot the distribution of probability against size on logarithmic scales and we can see that this gives a straight line at the critical point showing that the distribution follows a 'power law'. Distributions following a power law give greater weight to the existence of big fluctuations than a Gaussian distribution.

Image 11 - Critical fluctuations

Questioner 1 : What is the normal distribution?

Philip : That's the Gaussian curve.

Questioner 1: No I mean what would be the distribution approaching the critical point?

Philip : It would be more spread out. An economists often say 'It would be fat tailed'.

The important thing about a critical state of this kind is that it is pretty precarious and the existence of these fluctuations means the system is constantly teetering on the brink of uncertainty between these two choices. If we cool a magnet to just below its critical point then the needles will tend to become aligned in one direction or the other though there's no telling which way it will go. It could go one way or the other on the Ising curve. It just depends on whether one population of patches grows big enough to dominate the system.

Physicists for a long time regarded critical states as unstable, but in the 1980's they found that some systems seemed to adopt critical states that are robust and the canonical example was Per Bak's pile of sand onto which new grains are slowly being poured. Every so often the new grains trigger an avalanche and the avalanches happen at all scales so there's no way of telling when they will occur. If you look at the statistics you find the probability distribution is according to a power law on a log/log plot. So it's a critical state in that sense but it keeps returning to that critical state. For every avalanche the addition of new grains returns the sand pile back to the brink of a landslide. So instead of forever trying to escape from the critical state as a magnet does the sand pile is constantly seeking to return to it. That's why this sort of behaviour is called 'self organised criticality' and it was important because it suggested that there are systems in the world of nature that could be stable in a critical state and remain in it.

Image 12 - Self organised criticality

Per Bak was convinced that economic markets also work in a state of organised criticality and he thought this because the market is constantly experiencing fluctuations that seem to be 'scale free' in that they occur whatever size of sample is taken. Economists have tended to treat these fluctuations as random because they look that way, but it has been known at least since the 1960's that they are not and don't have a Gaussian probability distribution. In a real market index we get more big fluctuations than in Gaussian behaviour. This is significant because it's often the big fluctuations that economists are interested in because they constitute the booms and slumps and crashes. In a Gaussian distribution market crashes would be so rare that they would practically never occur and we know that in the real world they often do. If you try to make market forecasts based the wrong kind of statistical distribution you can be led badly astray. In fact these fluctuation curves are complex, they have fat tails but the shape is also complex and doesn't seem to be following straight power law behaviour.

lf 'organised criticality' is a concept that does provide a general framework for understanding scale free fluctuations and power law distributions are mathematical insights into how they can arise then perhaps we can begin to understand complex systems of interacting components. Such phenomena arise in different physical systems that seem to share nothing in common with each another when described at the level of individual particles or components. The critical points of some magnets can be described mathematically in precisely the same way as the critical point of a liquid gas system and these phase transitions therefore have universal characteristics. Thus self organised criticality has been proposed in systems ranging from mass extinctions of bological species to the formation of solar flares and the statistics of earth quakes. The occurrence of such phenomena doesn't depend on such specifics as exactly what kinds of forces exist between the constituent elements and how big or how small these are. It's not idle speculation to say that you might see similar phenomena in social behaviour I want to look at some instants where theses concepts might be relevant.

Traffic flow I have already mentioned. It turns out that traffic flow does seem to show behaviour that can be regarded as phase transition between different states.

Image 14 -Traffic states

When observations are made on traffic flow we find that the relationship between the rate of flow and the traffic density undergoes sharp changes from free flow where every vehicle can essentially do what it likes, to congested flow where it's all moving but moving more or less at the same rate, to jam where it is barely moving or stationary. These changes between the different states are quite abrupt and we can think of them loosely analogously to the gas, liquid and solid states of matter. From models like this we can understand how such phenomena as 'phantom jams' or those which form without any visible cause occur. And we can predict phenomena that is by no means intuitively obvious, such as the stop and go oscillations between patches of alternating moving traffic and jams. We can understand how a single perturbation of the flow can give rise to this kind of behaviour and we can use these models to test out the effects of various driving regulations or road designs. We might , for example, impose speed limits on certain stretches of road to ease the flow or position exits or entrances onto a motorway in different ways to reduce the chances of crashes. Social physics can be used as a test bed for exploring the consequences of restructuring our rules in one way or another though we have to decide for ourselves which of various outcome is the one that we desire.

Many of the choices that we make are 'binary' in that we only have to chose one of two alternatives and physics-based models have been used to explore how decision making is influenced by peer and neighbour pressure. In business and industry we might have to choose between a PC and a Mac and we might wonder under what conditions a minority product like the Mac can persist indefinitely or under what conditions a market leader will eventually command the entire industry. When two technical standards exist manufacturers might be faced with the decision of backing one or the other as was the case in the early days of video technology with VHS and Betamax systems. And often the choice embeds itself into the culture. The 'qwerty' keyboard configuration, for example, has persisted long after any logical reasons for it.

This setting of technical standards sometimes motivates the formation of alliances. Companies decide that by joining together they are more likely to end up on the winning side. Typically this ends in the creation of just two rival camps which is in a sense the ideal option because it means that every company can join a big camp whilst still staying in a separate camp from its worst rival. The evolution of technical standards for computer operating systems was a classical example . People who used the Unix system which was developed by Bell labs were free to make modifications to it and by the early 1980's there were about 250 different versions of Unix in use and all of them were incompatible with one another, giving rise to an urgent need to standardise. In 1987 Sun Micro-systems and AT&T agreed that they would use the so-called Unix Systems 5 and they formed an alliance which was formalised as Unix International Incorporated. This forced several of their rivals to aggregate into an opposing alliance called the Open Software Foundation which intended to use a different Unix system. The cosequence was that all the other computer companies had to make a choice to go for US5 or OSF?.

Now the interesting question is 'Was there a way for a company to predict what the others might do and so make the best choice for themselves?' Political scientists at the University of Michigan have developed a physics based theory which they call 'landscape theory' and the players in this game are companies that behave rather like gas particles in that they are on the point of condensing into two or more droplets. There's an attraction between them because they want to form an alliance so that they're on the winning side but there is also a force of repulsion between them as rivals. In the landscape model Robert Axelrod and team at Michigan found a rough way to estimate the forces between particles. This was a kind of tailor-made version of James Clerk-Maxwell's gas, except that each particle is unique and its interaction with every other particle is uniquely defined. The principle that governs the final configuration that these particles will adopt is the same one as that in statistical physics and satisfies the most stable arrangement. To find this equilibrium state Axelrod and colleagues defined a total energy for the group, calculated by adding up all the forces of attraction and repulsion between the firms in various coalitions.This defined a kind of landscape of energies for different arrangements of the particles and what was looked for was the lowest energy or biggest dip in the landscape. It's what 'games' theorists call the Nash equilibrium where no particle changes its position or goes to another camp because the lowest energy configuration has been achieved. Finding this equilibrium state, if the number of agents is small, can be achieved by exhaustively calculating the energies of all the different possible aggregations and finding which one is lowest. Even though there was no unique way of assigning relative strength of the repulsions and attractions to the companies concerned there were crude ways of estimating this and the Michigan people found that in the end it didn't matter very much exactly how you calculated it. They predicted that there would be two alliances formed and that this was the most stable arrangement.

Image 15 - Alliance formation

With the exception of where IBM was placed, it corresponded with the two alliances that were formed in reality. The probability of getting this by chance is about 1 in 15 or 16 but the interesting thing is that this prediction was made not on the basis of long term forecasts and cross benefit analyses but simply in the myopic way that each company was looking at every other company and asking: 'How do I feel about them?' 'Do I want to be with them or against them?

Questioner 2 : Just looking for attractive and repulsive forces?

Philip : Yes, you need to find some way of formulating that. How you quantify it is contentious, but it can be done approximately.

I now want to briefly talk about another application and this is to do with how firms grow; that is to say, what controls the size and size distribution of firms in a market. This is something that's been very hard to incorporate into standard economic theories particularly when thinking about markets that are very heterogeneous. We can deal with monopolies and with systems in which there is perfect competition to some extent using game theory and we can deal with oligopolies, but we can't easily deal with a very heterogeneous distribution of firms and predict how they will grow or shrink and what size distribution we are likely to get. This sort of question was looked at empirically by people at Boston University in the 1990's and they simply looked at the data on firm growth rates for US manufacturing companies that were trading between 1975 and 1991. The study encompassed about 8000 firms . What they found was that the distribution of growth rates followed a power law. It's actually a sort of double power law because firms shrink as well as grow. Robert Axtle at the Brookings Institute in Washington did a survey of 20 million US firms and looked at the size distribution, simply how big they were and again he found you get a power law behaviour over a wide range of different sizes.

Axtel has formulated a model of how firms arise that helps us understand where both of these forms of power law come from. Again it's a model that has a lot of agents interacting in one way or another. There is nothing in the model that compels firms to arise but the thing that makes that likely is that each agent is essentially a 'utility maximiser' in that it tries to find the optimum balance of money and leisure and the relative preference for these two things varies throughout the population of agents. There's a mathematical equation that relates an agents efforts to the productivity of the group to which it belongs; how much return it gets for its efforts, and a condition is built in that makes it generally favourable for agents to join together. In a sense it's a kind of increasing return of scale. The interesting thing is that it's not guaranteed in the model, whereas it's often taken for granted in standard economic models of firms growth. Here it's not guaranteed that if you get together with a lot of others that you're going to get an increasing return of scale but simply likely that it will happen. It turns out that this model has no stable Nash equilibrium in that it does not settle into an unchanging state.

Image 18 - Firm turnover

There's constant change in the number of firms in the system over time and we can also see how the size of the largest firm changes over time in terms of rapid expansion and collapse. Most of the firms which arise in this model are as ephemeral as they are in real life and what comes out of it is precisely the kind of statistical distribution in terms of size and growth rate that we see in the real data. If we start to ask questions like 'Why do firms fail?', the model shows that there's a typical trajectory that firms tend to follow. At first they grow more or less exponentially until reaching a peak at which they collapse suddenly and catastrophically to a small size that eventually peters away to nothing:

Image 21 - Typical firm history

This collapse is a consequence of a firm's own success. When it grows big enough it may become a haven for those agents that are free loaders and don't do very much but reap the benefits of what everyone else is doing. When this happens all the other workers in the firm may suddenly wake up to and think 'I can do better elsewhere'.
Not that the agents literally think this; there's no psychology built in at that depth but essentially that describes the process. Agents will leave a firm to seek greener pastures and it's very telling that just before the firm reaches this peak the average productivity per worker plummets.

So from a simple model like this we can learn some revealing things about firms. First of all they are not maximisers. They don't maximise either profit or overall utility as theories of firms often suggest. Individual agents may be trying to do this but it doesn't induce that behaviour in the group as a whole. The firms that do best are not those that aim to make the most profit and the ones that last longest are the ones that are able to attract and retain productive workers.

Questioner 3: It seems you pick examples that can be described in terms of quantity properties, that are amenable to mathematical treatment. But it seems to me that these are terms that apply at random and there is no underlying correcting principle behind it. The only thing that I could discover is that they are all amenable to statistical treatment. And for this there is no need for any kind of reference to physical science because this kind of practice has been going on without it. And the danger is that it misses the essence of human activity and that of animals which is the production of physical and mental states. These are what drive society. And that is what demands the aggregation of things in other words the activity of systems. And it is not the application of professional science to society and social economics but the development of system science which is capable of handling qualitative properties which are predominant in human activity situations and far more important than quantitative properties. I think this is the line of action that is more promising for understanding what drives society and in particular human activity in situations for survival.

Philip : OK, there are various things that follow from that. First of all the use of physics as opposed to just general statistical modelling. You are right in the sense that models like this often lose connection with classical physics in that they are computer models, but I think the important thing is that often what we learn from the physics is that there are generic kinds of behaviour not just specific to the model. I wanted to mention some work that Paul Ormerod has done on a model of marriage: how prevalent it is in society and the factors that determine that prevalence. This is something that doesn't obviously seem to lend itself to quantification but you can formulate a model that looks at the various factors that will influence this socially: things like cultural attitudes and economic incentives. What comes out of this is that you see exactly the same phenomena that you see in the statistical physics of liquids and gases. The same kind of phase transitions and so on. I think it would be surprising if you didn't see these sorts of behaviours in social situations.

Questioner 3 : But what's the conclusion that you draw?

Philip : You can conclude, for example, that if marriage is somehow influenced by economic incentives then if we make them stronger we'll increase the proportion of people that are married in society. And if governments ask 'How can we encourage marriage?' the answer is to put in some economic incentives to help that, but what this model also says is that there isn't necessarily a direct proportionality between economic incentive and more marriages.

Image 22 - Social forces and marriage

So for the same economic incentives you can find either a low proportion or a high proportion of people in a married state depending on where you started from. That's not an obvious conclusion from implementing a policy. If a model like this has any relevance at all then that's one of the things it can tell you; why you might find completely different outcomes for a particular set of policies or rules depending on what the society was like when you began. There are a lot of other things when you start to see the analogy with physics that you can deduce about how that change occurs, where it occurs and where the jumps might be on these two branches. Physics can help and you will see phenomena in these social science models that do have a direct analogy.

But as to whether quantification is necessarily important or whether we need to think about qualitative factors I'm not sure that I need to make that distinction. Because what some of these physical models are saying is that there may well be certain states that a system of interacting agents adopt that are robust and we can model them without resorting to an endless palate of conditions. We can start to understand what those stable states are and how changes between them happen. In a way that seems to be a sort of qualitative statement. It would be unwise to think you could say exactly what conditions will create those changes because the models are very crude approximations of reality. But what is important is that you then avoid thinking you can create whatever state you want to exist and find the conditions under which it will exist. If you can formulate some kind of model that will tell you what states are stable then you can avoid trying to do the impossible.

Questioner 4 : One thing that is occurring to me in this interesting and important exchange is that in some ways it's common sense to understand there are complex systems with different factors interacting in human social systems. The problem is that public policy appears as if that is not understood and policies, especially political policies, seem to be governed by simplistic and naive assumptions about how social systems work which is mechanical compared with complex systems, but given that we are in a scientific age, to have this kind of explanation, may be sufficient to tip the social system of decision and public policy making over into an understanding that it is more complex.

Philip : Well that brings me very nicely to what I wanted to say at the end, that in the mid 19th Century when this idea of statistical social science started that was precisely the point that some people were making, and you are right, it hasn't been heard. William Newmart who was I think the president of the Statistical Society of London said 'the rain and the sun have long passed under the administration of magicians and fortune tellers. Religion has mostly reduced its pontiffs and priests into simple ministers with very circumscribed functions and now men are gradually finding out that all attempts at making or administering laws which do not rest upon an accurate view of the social circumstances of the case are neither more nor less than imposture in one of its most gigantic and perilous forms'. And I agree, I think, you have to know what is possible before you start to think how to get there.

Questioner 5 : What is the force that you apply for being married or not married? Like temperature and the magnet.

Philip : Well Paul can tell you this in much more detail than I, but this model was looking at two factors and I'm sure it wasn't being said that these are the only two factors, but it's interesting to think about isolating those particular factors and seeing what effects they have. One was the social acceptability of marriage; whether it is unfashionable or conversely a requirement of cohabitation. If you accept that there are social factors like that, that make marriage more or less likely, then you can say there is this effect and let's see how the proportion of married people varies depending on the strength of that. And the other one was economic incentives. Of course we're actually looking at three states in this system in that you could be married, separated or single and of course once you are no longer single then you don't go back to it, but you can go between being married and divorced as many times as you like.

Paul : Yes, I don't know much about the physics of magnets but a lot of these models are showing that the bigger the proportion of the population in a particular state, other things being equal, the more probable it is that any individual agent will convert to that state. So at the critical point the system will move towards either one or the other.

Philip : Yes, so the crucial point is that it's looking at interactions. How strongly are people affected by those sorts of pressures. And I think the conclusion is that they don't just do that in a linear way and you sometimes get abrupt jumps and I think that's the value of this. Not that you're going to make a specific accurate prediction, but that you can see that there are types of behaviour that arise that aren't intuitively obvious from the conditions that you've put into the model.

Questioner 6 : Two things, one an observation. I'm fascinated by your alliance formation. I was a managing director of AT&T in Europe from 1984 to 1987 and my observation was there were huge battles within the organisation about what to do. Notwithstanding all the other firms, there were huge political battles and AT&T forced the issue by trying to dominate the market and control everything, lost it and then had to form an alliance. So you can look at the external level but there were things going on within the company as well so it was more complex than your alliance model suggests. And similar things were going on inside Sun Microsystems as well.

Philip : Yes, I suppose in that respect it raises the question that if the model more or less gets the right outcome while neglecting all of that, then can we take a view at a higher level that is going to have some predictive value?

Questioner 6 : But it would be interesting to see an internal model of AT&T because there were several possible outcomes within the organisation.

The question I want to ask is that seeing the same behaviours in social systems and physical systems the thing in common is interacting agents. On the one hand you have physics and on the other people. Isn't that a bit worrying for free will?

Philip : I don't think so. First of all it's very important not to confuse the model with reality in the sense that you might say, for example, that there isn't really a repulsive between people that we measure that prevents us bumping into each other. In a sense what matters is not the reality but the effective result of that. We behave as if there is a repulsive force.

Questioner 4 : But there is. In cultural behaviour you can't stand close to one another.

Philip : Exactly, but you can't measure that as you would a physical force. I suppose with the question of free will. Firstly if it's statistical you're not making pronouncements about what any individual person may do and secondly I think we over-estimate our free will. Voting, demonstrates that. We all think we're making up our own minds but looking at the statistics we're not. That's not surprising since we are affected by each other and perhaps free will can be over-emphasised.

Questioner 5 : Except that as humans we can actually change the rule of interaction.

Paul : Yes and clearly there are attractive and repulsive forces that are not constant.

Questioner 6 : I think this raises a very interesting question: 'At one level in conventional economics, we think in terms of the cognitive ability of agents. At one extreme we have the classical model of economics in which agents are able to gather full information about any particular issue and then process it in a way that will optimise to infinity. That's one paradigm. The paradigm in statistical physics is that agents act purely at random and have no conscious ability to shape the system. I think an interesting question is: 'to what extent can we explain social phenomena better by a near or purely random behaviour as if agents have no cognitive ability rather than the full information economic one. In my view there's a phase transition from situations in which agents have cognitive ability, where it's easy to see what the optimal thing to do is and the rest where it's very very hard and for most decisions in business do we really know in advance what the impact of our strategy is going to be no matter how carefully we research it? We don't and we can model it to some extent as if the change in our strategy was random. It's obviously not because we're acting with conscious intent, but because of the uncertainty of outcome we can model it as if it was random.

Philip : The other thing you can do with these models is to build in an adaptive capability. It's possible to formulate them in a way that certain kinds of behaviour become more successful than others. You can allow for that and you can allow for learning both at the individual and group level.

Questioner 6 : Maybe but I wonder how much agents can actually learn in complex systems once you're outside situations where it's trivially obvious.

Questioner 4 : Well yes, and the question going through my mind when you brought the physics in was 'is adaptability the basic difference?'. Gases and liquids don't have adaptability but living systems do?

Questioner 6 : Well if you think about the evolution (of life) then agents don't act with intent in that respect so can you model it as if (agents actions were random).

Philip : I guess I'd simply say that sometimes it is probably important and sometimes it isn't. I mean we probably don't have a great deal of adaptation going on in terms of how we move around space for example. You can talk about how people walk around that evoke these forces and you can assume that they will stay much the same. Some situations in driving are the same. There's the potential for learning and there's a degree in which children learn not to bump into people, but in general if you want to model, say how pedestrians use a public space, it's not clear that you need to take account of that.

Questioner 7 : The statistical approach is very useful to describe such phenomena but when we come to discuss the role of agents, activities involving 'will' and 'volition' and the rest of it. (You have to go) beyond the mathematics of interacting parts to interacting systems theory which is capable of hopefully answering the role of qualities associated with people.

Questioner 8 : Some of the examples you have chosen have very simple pathology (?) in that the state the system depends on a few things. So the logic leads to simple outcomes but with full complexity such as we see in other social situations it gets very difficult.

Philip : Yes I think that's true and you need to be very selective about where this sort of modelling is going to be useful and where it isn't and not simplify the situation. In any situation where behaviour is volitional you have some predictability that doesn't follow or be immediately obvious from the rules you put in. And that can be surprising and I can see situations where that's useful but I agree there's lots of situations where you have a random scattering and I think that's going to be true of any attempt to use physics in social systems. There's going to be a scattering of situations like these where you might be able to use them usefully.

Paul : Oh I think you're being too defensive there. It doesn't work all the time, but if you take orthodox economics, the set of assumptions required there, is huge in terms of the cognitive ability of agents. We know from other disciplines that except, in very simple situations they simply don't hold, and yet economists have actually been given useful and quite powerful insights into a number of problems using this really quite inadequate model. There is a completely different paradigm about the cognitive ability of agents which can give powerful insights into a very wide range. It's a question of where between the two extremes you might end up with an optimum model. It might be that it's more towards the random one than the full cognitive ability one, so I think you're being too defensive about the range of models which give insights. Economics gives insights into a very wide range of disconnected problems, starting from a few simple principles.

Philip : No I agree. I think there's a temptation to think that you can't get anywhere in modelling unless you have a considerable degree of psychological complexity and I think that's the fallacy of some social science. A lot of the gross behaviour of social systems doesn't depend to any great extent on having a detailed psychological model of the agents. Much more broad brush factors are what determine the overall behaviour.

Paul Ormerod

Complex system models have given valuable insights in a number of areas including those shown below:

Image 1 - Practical examples of applications of complexity

We have done work for a number of organisations including: The British Home Office, the Greater London Authority, The US Department of Defence, The Institute of Complex Additive Systems in New Mexico which is very defence oriented and we've just started a small project with the National Centre for Genome Research in Santa Fé. So there is a very wide range of applications using the Complexity approach. We could hardly have a more disparate polarisation perhaps than that between the GLA and the US Department of Defence and yet it works for both of them.

Financial markets we know are not predictable as far as asset prices are concerned, but where complexity theory has been very powerful, and if there's time I'll talk about this later, is in terms of the volatility which causes a particular problem for economic theory. There is a very distinguished American economist called Kenneth Arrow who formalised the theory of free markets and got a Nobel prize for it and he described the level of volatility in financial markets as an empirical refutation of free market economic theory. Complexity models using quite simple assumptions which show the non predictability of asset prices and are able to generate the large degree of volatility we observe. I may be able to talk a little about crime and what interests me here is that there are such large variations in crime rates at a very fine geographical level. I mean between areas which have very similar socio-economic characteristics. There is a lot of statistical data on crime rates not just in counties in the United States of which there are about 3500, but in individual police precincts. At a very fine geographical level we can find neighbourhoods which are very similar but in which the crime rates are very substantially different. So I want to talk about geographic segregation as a first example. I also want to talk about the ups and downs of the business cycle. Why are economic forecasts so poor? Why is it apparently almost unpredictable? But I also want to talk about it in the same way that Philip was talking, in terms of the distribution of economic recessions. Its not quite a scale free relationship, it's more subtle, but it's as if it's a scale free behaviour and it's got some interesting properties. Again complexity models with simple rules about how firms behave can generate these properties at the overall level.

Philip also mentioned a lot of work by Rob Axtel on company size and growth and I'm interested in this, but also in the extinction of firms. Why should firms become extinct? Lots do and in the United States the death rate of firms on an average is more than 10% in any single year. For the millions firms that are economically active in the United States the scaling relationship that Philip was talking about applies to the extinctions. Mathematically it has an almost identical form to that which has been discovered in the biological record, though obviously a different time scale, for the extinction of biological species. So there may be a general theory of extinction for the way agents interact in this kind of system. It does apply with different time scales. I've got a data base of the top 100 firms and their capitalisation in 1912; their extinction follows the same pattern over the 20th Century.

I want to talk about technologies and ask why sometimes inferior technologies succeed; a major problem if you think that agents possess full information. Why should they adopt an inferior technology or more accurately why should it persist? It ought to have been common knowledge, for example, that Betamax was better than BHS in the early days of video recording. So there are all sorts of areas where this approach seems to give powerful results.

In modelling complex systems like in any science we know that the physical world is fantastically complex yet how can it be, for example, that e = mc2 ? If we think about it it's fantastically simple, but it seems to work. I take the view that although the world is complex and it might be very hard to find the rules out, we want to start from as simple a model as possible and only make it more complicated if we have to. In social science the simplification relates to the rules of behaviour of individuals in the model. In the model of firms extinctions it's as if firms don't know the impact of their strategy changes. It's as if their strategies evolve around them and yet it seems to account for some key stylised facts about evolution and probability of survival with respect to the age of a firm.

We should choose rules of behaviour that can be justified independently of the model. We could discuss doing that, but the key thing is that we have to validate the model not by looking at any particular history or trying to replicate it because by their very nature theses models are probabilistic and there could be many, many alternative histories. The way of validating these models is to ask 'What are the key underlying properties of this particular set of data that we want to replicate by our model', because we know that there's a great deal of contingency in economic and social systems. So Rob Axtel's example of firm size wasn't an attempt to replicate any individual firm but a general model of firm growth which produces the distribution of firm sizes that we actually see. So this raises important questions about methodology in social sciences and I'm very interested in what's happening at the top level in economics which, to be blunt, means America.

There's been a lot of relaxation in economic theory in the last thirty years to take into account the fact that agents operate with imperfect information. In the 1950's and 60's there were some important papers which brought the whole free market theory to an end and forced people to make the models more realistic. So it's accepted that there isn't perfect information and people can say lots of interesting things which can affect the outcome, but the key difference, and this is where the physics approach is interesting, is that agents interact with each other. Even in Axelrod's models of agents with incomplete information which was a big advance, they still had fixed tastes and preferences. An agent was trying to maximise given its own fixed preferences. But the reality is, that because agents interact, those preferences may themselves change. Thus we need a different methodology, and the economists favourite tool of the calculus which tries to optimise and maximise doesn't get us very far, once we understand that the function we're trying to maximise is itself subject to unpredictable change.

This is the key contribution that statistical physics methodology brings. But Axelrod is a really innovative guy and he's got an article out in one of the world's top economic journals looking at outcomes of a schools system in the United States, in particular, concerning the persistence of massive differences in the performance of different racial groups and he starts off by saying that economic theory can't really tell us why. He therefore goes straight to sociology and group behaviour and how people actually form their views by peer influence. And then sets up a neat theoretical model which shows that price economics can't really tell us very much and we need to take account of how tastes and preferences are shaped by social interaction. So this is the way I think that economics is actually going and the physicists are just coming at it from a different perspective.

Properties of the system emerge from individual interactions, characteristically though not necessarily such that the system lacks short-term predictability though there are some underlying regularities. Identifying them may be quite difficult and deciding what we want the models to replicate may be challenging. An analogy with physics might be that a physicist has an hypothesis but some of the experiments to prove it require masterpieces of thought. However it is the regularities that we want to discover. Some of the particular problems of socio-economic modelling are that data series are short, are almost always 'noisy' and agents can vary considerably over time.

I don't know whether it's true or not but Max Planck is alleged to have said to Keynes in the 1930's that the reason he didn't do economics was that the maths was too hard. Perhaps he had an insight into what ought to be done. What I want to do here is to look at how models are set up and how we might think of validating them.   Thomas Shelling, economist and political scientist, looked at American cities and perceived a high level of residential segregation on racial lines and the question was whether or not this was a factor in racial tension. Similar situations occurred in the UK when people discovered that Asians and Whites lived in different areas and Shelling asked whether this meant that people are strongly racially prejudiced. So the property he wanted his model to replicate was a high level of residential segregation between different types of agent. He wanted to test whether, if he gave agents preferences he could generate a model in which people were not really very prejudiced, but would generate the same outcome at the aggregate level. So how did he go about it?

Image 12 - Shelling model (1)

Like all these models it was very simplified and abstract. He started off with a large number of agents (N) and assumed that there were equal numbers of two types. In this model the agent characteristics are fixed (red or blue) and placed on a grid so that each have the same number of neighbours, in fact eight. So initially the agents are placed at random with a small percentage of empty squares to which people can move. Then he asks the question of how people choose to move and how do we define what a neighbourhood is? There are lots of different ways in which this could be defined, but an obvious one would be that an agent looks to see how many others of the same kind live in close proximity. It was also typical of these models that an agent was chosen to move at random to the nearest empty square. There were all sorts of simplifications made in the models, but the aim was to see if it would generate the same behaviour and if it didn't then it could be made more complicated. The model progresses in a series of steps and at each step an agent is chosen at random to move. If you have 2000 agents it is usually a different agent but there's a small probability that it could be the same one. The model is run for as long as desired. In general the agent decides to move if more than a specified percentage of all agents in that neighbourhood, in this instance all the eight squares round it, are of a different kind. The agent then proceeds to the next step and is called at random to decide whether or not to move. Examples of the solutions that I've got here are actually saying that an agent feels comfortable even if it's neighbours are split four and four because it is then in a five to four majority. It will only move if it's in a minority. So it's not strongly prejudiced. These are the simple rules: agents are initially scattered at random with some empty squares to move to and there's a simple rule for deciding which is by definition determined by a low level of prejudice.

Image 15 - Initial configuration of agents

Image 16 - Configuration after only two moves per agent

The solutions show that on average after just two moves per agent we get dramatic segregation. Now each individual solution will be different but it will have the same qualitative characteristic. If we run the model hundreds of times we get the same sort of qualitative property. In the same way that if we're thinking about the business cycles, for example, each individual history is different but the qualitative characteristics are the same. The solutions don't mean that people aren't strongly prejudiced, only that the observed outcome could arise if agents only have a very mild preference in favour of people of their own type. And there are recently some results which extend this in saying, though it's scarcely credible, that you can get segregation even if people are willing to be in a minority in their own neighbourhood. But we can see how the model gives us an insight: here's an important question and using simple rules we can say well we can account for it even if people are only mildly prejudiced. So it's a way of gaining and insight into policy implications for example. It might be the case that if we have other information that in general most people aren't very prejudiced should we be worrying about racial segregation? A mild preference might be very natural.

Eve : I'm not surprised at the blocks of preferences but I am surprised at the blocks of empty squares. They have moved from random positions to be clumped together. Is that something you would expect?

Paul : Well, I'm using this model to say 'This is as simple as you can get'. And you're right we don't observe clumping of open spaces, but the key feature of observed segregation is replicated.

Questioner : Do you find at the end that everybody is in a stable position? Or do you always have some agents who want to keep moving?

Paul : I can't remember. We just programmed the basic Shelley model that gives these results. I think bits do keep on moving but the general pattern is shown.

Questioner : If you model with a higher level of prejudice do you get different blocks of colour?

Paul : Um, well paradoxically in this model, if you have higher levels of prejudice, it's actually much harder to settle the model. Suppose you say 'I'm only happy if everybody is the same as me' then you have to run it a hell of a lot longer to get a pattern. So maybe you could say we will only observe pattern like this if we only have relatively mild preferences. I mean segregation is not as easily delineated as this. Reality is more complicated but this is an insight.

Questioner : I wondered if white squares are on the edge of regions because people don't want to live there.

Philip : Yes, essentially there's a high energy at the interface so it's a way of avoiding interfaces between the two colours.

Questioner 7 : These are the results of operating this model, but what is the actual programming? How do you turn these statement into mathematical form?

Paul : Well it's just a page or two in 'c language'.

Questioner 4: In India people are very prejudiced because the people will say 'I will live only with my own kind', so in villages we have very distinct boundaries around certain groupings so how does that relate to what you said about the more extreme the prejudice the more unstable it is?

Paul : Well if you're trying to model in Indian villages perhaps you have to give the agents some extra rules of behaviour. I mean it's a good question because you set these models up to address a particular question. These models don't claim generality. What we're saying is that we're setting up fairly myopic rules of behaviour for individual agents which in this particular context produce the underlying properties of the data.
Conventional economic theory assumes a general mode of behaviour; agents maximise utility and in certain circumstances that 'as if' will suffice, but here we're abandoning that and saying each model is context specific. We've just done a model on the Tiannaman Square incident thinking about external pressures and phase changes built on some simple rules which an expert in the area knows about concerning the interactions of the Communist Party, the Government state machinery and so on. Quite simple rules and its a way of thinking about the situation though it's hard to validate.

Questioner 3 : Well it's worrying for me because we can't see the constituents, we can't see the assumptions or the relationships and we can't question how things have been quantified.

Paul : Well I've written the rules up in English and you just convert them to computer language. Without the computer we couldn't do this sort of thing because we do rely on lots and lots of calculations and replications. But the rules are transparent. We're simply saying these are the initial conditions and these are the rules that agents follow. Anybody can look at it and say 'Yes I can see what's been done'. the hard bit is thinking of the rules.

I mentioned a model which gave high volatility but what we're looking for in each case are simple but realistic rules of agent behaviour for explaining a financial market which will generate high levels of volatility. I've mentioned some of the factors already but this model is due to Alan Kirman.

Image 9 - Typical solution of Kirman model

The model has the following properties: short term non predictability and it can generate high levels of volatility and we can illustrate this in a qualitative sense. In the same way that we could formalise the degree of segregation mathematically in the case of the Shelling outcome, we can formalise the degrees of volatility in these models. As before there are N agents and the model evolves in a series of steps. The difference is that in the Shelling model an agent was either red or blue and didn't change colour but in this one it can change its attributes. In the program we described it as zero or one, but in this case an individual could be a zero one minute and a one the next and there are rules for describing how an agent changes behaviour.

Agents are traders on financial markets and the rule is well grounded because in general we can describe traders as operating in one or two molds: either as a 'fundamentalist' or a 'chartist'. A fundamentalist tries to form a view on the underlying profitability of holding that particular asset. If it's the exchange rate he or she is trying to form a view on the underlying features of an economy which would take in interest rate, inflation rate, stability of its government and so on. In other words taking into account economic fundamentals in trying to form that view. A chartist, on the other hand, looks for patterns in the past recent behaviour and extrapolates. So if a price is going up he or she will say 'I think its going to continue to go up', or vice versa. You can be more sophisticated than this, but essentially you're relying solely on the history of the asset price.

This is a fair characterisation of how people behave and they switch between these states. So again at each step an agent is drawn at random and decides whether or not to change and the model evolves in that way as in the Shelling model. In this case there are two rules for changing. First of all there's a fixed probability (e) of change which can be specified in the model. But, and this is key, it changes with an additional fixed probability (b), the proportion of the total number of agents which are in the other state at that time. So there are more people that are acting say, as chartists, and you're a fundamentalist who's saying 'I think shares are under valued'. But as shares keep going down at some point you lose your nerve and believe the market is going to crash so you've switched to a chartist.

Image 10 - Relative amounts of time for different percentages of Chartist traders high propensity to switch behaviour

Questioner 8 : Is it (the probability of change?) fixed for each agent? It doesn't change over time?

Paul : Yes, it's very simple but you can see how we can start making the model more complicated . We could put a distribution on this for agents or we could draw at random for each agent. If we were trying to generate the results that would satisfy the observed phenomena that's one of the things we might have to do. The Shelling model is unusual in that the agent properties don't change, but a general feature of complexity models addressing socio-economic systems is that agents will in some way have their behaviour altered to accommodate what the real people do.

Eve: Would it be possible to actually have emergency behaviour? In other words people don't just switch between chartists and fundamentalists but decide that a third type of behaviour is more appropriate.

Paul : No, you can have as many states of the world as you want and you can switch between them but this is a minimal set of assumptions that will generate the required degree of volatility of financial markets. So if this model will do it why make it more complicated?

The diagram is actually just showing the percentage of agents who are chartists and how people switch around for a particular epsilon and a particular beta and we can generate different levels of volatility depending on how we set those parameters. If you look at it and say 'Yes, it looks like a plot of share prices' and in terms of long run regularities if there's a high propensity to switch behaviour then paradoxically we find that. This is showing the relative amounts of time the system will explore at different mixes. So here nobody's a chartist and here they're all chartists so if there's a high degree of switching the system most of the time will give a roughly equal split between the two types of trader. If there's a low propensity to switch behaviour it actually looks completely the opposite way round. When may seem paradoxical but the reason is that if there's a low propensity to switch behaviour it may take a long time for the system to drift to an extreme but once it gets there it stays for a very long time. If 98% of people are in one mode rather than another it's very hard for it to move back. The current model is not a perfect explanation but it gives qualitatively the short term non predictability and the sufficient degree of volatility.

Questioner 9: Is there any dependency on N, the number of agents. I mean is there a critical mass where you don't get these patterns?

Paul: Not really unless you get a very small N. I think it's something like ten or twelve. But there's an important features of realism that the model doesn't have. It assumes that a trader knows what everybody else is doing. which might seem reasonable in a financial market because there's such a huge amount of information, but it turns out that, although there are many many traders, most individuals are on a particular social network where they usually monitor a small number of sources. So you need to model a particular topology of connections. And some of the interesting stuff that is now being done uses different mathamatical formulae for particular kinds of social network that people operate on and it can have quite different properties in terms of how things disseminate across a network. The way in which viruses spread and the typical period for which they persist depends a great deal on the kind of mathematical structure that models the way the agents are connected. Everybody here is connected to everybody else in that they can observe the overall outcome, whereas normally they might just observe a few. I mean it turns out that computer viruses persist on the web for much longer than standard epidemiology would predict, and that's because the theory assumes that each agent has an equal probability of meeting every other one. But it turns out that the web doesn't have that property; some sites are more important than others and there seems to be 'near scaling' behaviour in terms of the structure of distribution of connections. Models with that characteristic also generate the property that viruses will persist for much longer than standard epidemiology theory suggests.

What I've tried to do here is to raise some general issues about what the aims of these models are, how we try to be as simple as possible and how we validate them.

I mean they are simple. If you saw this model written down it just a little page of maths. It's not like an economics book where you might find ten pages of really hard differential equations.

Questioner : 3 Is it right to say what you have done is to produce the computational algorithm? Because the problem with algorithms is that the interactions and the individual objects are lost. And the other thing is that when you put in these rules mathematically, everything else which is essential in trying to model the real world is lost.

Paul : Well I don't think we're going to agree because we have a different view of how we should do modelling.

Questioner : No, I'm not saying you shouldn't. All I'm saying is that having established that we've got an algorithm there's the danger I have described.

Paul : Well we're giving individuals very simple rules of behaviour which you can write down.

Questioner 3 : Like in a chess program?

Paul : Yes, but presumably the rules in a chess program are much harder. We're trying to explain emergent phenomena at the aggregate level of the system as a whole which we observe, from simple but plausible rules of individual behaviour in this particular context.

Questioner 3 : But the individuals are not in the algorithm.

Philip : Yes they are. The algorithm is applied to each individual. If you have a car going down the road and you say 'If there's another car within a certain distance in front, them slow down and if not speed up' and you do that for each car then the algorithm is being applied to the individual elements.

Questioner 10 : I really want to come back to the question you posed at the beginning which was : 'Given a set of rules can we predict the outcome?' Thinking of the rules is really hard. What I do in life is design institutions as they relate to policy making in two ways. One in a very general way in, for example, a new companies act and I want to ask 'If there's a new law, i.e. a new rule, then what would be the outcome?' that's one issue. And the other is very specific in that if I'm trying to redesign an organisation in terms of its rules then these are not simple like this and I'm wondering if there's any advice you can give me about how we go about institution building.

Paul : Well I think these are hard questions in social science and what you're trying to do is even harder. I mean working out how to improve an organisation is pretty difficult. But the whole thing about science which includes economics is to try to get something which is simple but has still got a world of explanatory power. So you can then understand what's going on. You have to make fantastic simplifications about the world and if they give us a reasonable account of what's going on then it gives us an insight.

Questioner 10 : I think what I'm saying is that in a practical sense, in this work you cant simplify in this sense. The rules you actually make are complicated.

Eve : Well there may be another way of doing it than trying to imagine rules.

Paul : One thing that we did with the Chinese situation was say, 'Here's the Communist party and here's the state machinery and here's the peasantry'. And coming into the system is some pressure from the outside world and social unrest emerges from it. So we just home in on those connections. I mean , and this is just off the top of my head, but we might ask: 'If the Communist party strengthens does this strengthen the economy and will that reduce social unrest?' So we model by simple connections just using positive or negative feedback. In other words just thinking qualitatively in order to simplify the particular problem.

In the Crime model (see below), we knew that most crime was committed by young unskilled men. So we made that our population and could then say we have: people who do not commit crime, people who are susceptible and may commit the odd crime, a small number of people who habitually commit crimes and people who are in prison. Then we have data to say that 70 % of offenders will re-offend, but for some crimes like burglary people do it for a bit and then give it up without any obvious deterrent. In general we can qualitatively conceptualise the categories we're interested in and set up some simple directional flows of influence between them which might all be at the same strength and either positive or negative. You can get some surprisingly complex dynamics from that. We got phase transitions in the China model with those kinds of connections .

Obviously the more information you've got the more you can start quantifying. If you're feeling ambitious you can set up a system of differential equations. So you can progress with different levels of mathematical sophistication depending on how confident you are about your ability to build in the information you have. But sometimes, especially with management problems, it's very hard to validate the model . Management thinks it knows it's market and thinks it knows how it operates. So you get a conceptual map down and then you run it to see what the implications are and see whether that squares with what management believes. If not then the mental map has to be refined and when that's done you can then start to think about 'What if I design a system where I remove that link or what if I put some other connections in here' and so on. But again the goal is to try to simplify it as much as possible. If you have twenty categories you probably won't understand what's going on. So it might be more realistic at one level but useless for finding out what the key links are. As it happens in the case of the crime model a very simple framework enables us to say that once you get inside the criminal justice system you can change the parameters and see what the impact of different prison sentences are and so on.

That's second order in the control of crime and the real key thing is the social influence.

Image 19 - Crime model (1)

Image 20 - Crime model (2)

Image 22 - Crime model (3)

Questioner 4: I just wondered whether you had any examples of organisations that you have worked with that have used this type of modelling with something they have decided to do.

Paul : Well I'm not a management consultant, but I have suggested the approach to a number of companies over the years, first as a way of capturing individual knowledge within the firm by writing it down and exploring any inconsistencies.

Questioner 4 : Well have I understood correctly that out of doing the crime model were you saying that the most important relationship was between the susceptible population and the ones that became criminals? What would you have told us about the usefulness of this model?

Paul : We used data from over a fifty year period and found that peer group pressure is the key and we could put parameters in the model which represent deterrents. However there's an awful lot of empirical work done on trying to quantify deterrents and people do get wildly different answers. So the impact of deterrents is not known. But what we can say here, is that if there is a connection between deterrent and the average prison sentence we can look at what level that starts to become important. Again the model is based on epidemiology theory, and what causes an epidemic to dry up once it has taken hold is a lack of susceptible agents.

Philip : Can I just add to that and say that one thing that comes out of this model is that it shows that changing the severity of the criminal justice system doesn't have much effect on crime rate whilst at the same time it shows it can have a dramatic effect. If you look at the graph it depends where you are on that curve. So it discourages a simplistic application of a particular study to the general problem.

Questioner 12: What this shows me is that these models are useful for the physical world, like production lines and factories, but for social interaction I see this technique as trying to illustrate a situation that has happened in the past. These models will not predict what will happen next. I can learn from it about what might happen and still not make a decision.

Paul : What you really get out of this is a probability distribution of what might happen. So you have to decide in the Kirman model whether people have a high or low propensity to switch behaviour. Then a prediction of the model is that if you have a high propensity to switch then you'll observe an outcome like this. If not you'll have a completely different outcome.

Questioner 12 : What the model shows me is that the traders interaction might have an impact on the market, but if you want to take a decision in real time there is a boundary here in trying to apply this in the business world. I'm not saying it is wrong but it is looking for patterns; some rules which will explain some things.

Paul : This is an important point if you try to run a business using this sort of approach. It might help you by revealing strategies with only a low probability of success but it only gives probabilistic outcomes. So for example we did a model on congestion charges and found there is a probability that congestion will get worse over a three or five year time line. Now maybe our model is wrong but the rule we used seemed reasonable. So we will have to look at that more closely.

The way things are spread by word of mouth is question of epidemiology. You might say 'What's the formula for a successful film?' You might think 'Well you get some famous actors and have a big budget which costs you a lot and you put it on simultaneously here and in the United States and people here come out of the cinema and say 'That's complete rubbish'. And you get twenty million dollars in the first week and then nothing and you've lost forty million dollars. So there isn't a sure-fire formula and you cant generalise to other situations. And management might not like it because they're used to the world of control. It's like the Soviet Union and the Five Year plan. It gives the Illusion of control

Eve: But models are just one tool among many and when you come to 'designing' organisations there are other ways which arise from complexity theory such as creating enabling infrastructures where the modelling is part of it, but it's not the only tool that you use.

Questioner 12 : You have both shown these graphs that show that you can have a certain value of the variables which give a number of different states. So is there any thing that can be done to help people figure out where you are because then you could possibly use more trivial cause and effect things within a tiny frame

Paul : Well, let me put in my plea of mitigation. All these models are relatively new. The Shelling model was thirty years ago and then nothing happened for ages. Kirman was 1995 and almost everything that has been done in this area has been done since 1995. Most of the models here have been produced since 1998, so it's new and there are lots of things we don't have the answer to. It's empirically hard to determine where you might be in terms of starting conditions. These studies were done by perfectly respectable people on different data sets and where people get different results it might sometimes just be bad methodology, but there are lots around which are perfectly reasonable. Maybe there are different phases you can be in. In a different country or a different time you'll be on a different part of the curve. And maybe as we get more experienced, people will get better ideas about how to identify where you might be, but at the moment it's tricky.

Eve : And there are experiments like the project that Thames Valley Police are doing on what is called 'restorative justice', which would go almost contrary to what a model like this would show or predict, where they bring together face to face the victim with the offender, not in order to blame, but to acknowledge and become aware of the consequences of their actions. Not only on the victim but also on the family of the victim as well as the family of the offender.

Paul : Yes, but that is a good example of how we would use the model. The model would give us insight into what might influence what but then as a separate judgement you might say 'Well we've got a lad who's robbed a pensioner and what we don't want is to make him a hard core criminal, so what's the best way of stopping that?' The model doesn't tell you what you should do in each case. What you find is a certain relationship and you may want to change that but the model is not telling you how to change it.

Questioner 4 : Well for me it's plugging the 'S' (susceptible) back into the 'N' (not susceptible). I mean for me it's not so much the predictive ability or the quantitative, but it's just such a useful way of analysing complex social situations so much more clearly.

Paul : What you might want to have is hierarchies of models. Suppose you say 'Well, this is a reasonable model because we want to focus on this thing' we might then require a model that will look at the evolution of this aspect, but what you wouldn't do is try to do it all at once. But I must say that designing an organisation is even harder. I mean in terms of world history how hard has it been to generate capitalism? We've had some tens of thousands of years and we've finally hit upon a system that works.

Questioner 4 : Well sort of.