Utopia Theory: Towards a Physics of Society

Philip Ball
p.ball@btinternet.com

> ‘The centre of politics has shifted…. The neoliberal thinking that has dominated the industrial world for nearly 30 years has led to a financial crisis, which in turn caused the global downturn…. Clearly, there can be no turning back to the failed and discredited politics of old. Instead, we need to use this time of emergency to aim for a different future and to get there by different means…. this is not a crisis of capitalism, but a crisis of a society and democracy that have failed to regulate the market.
>
> Neal Lawson & John Harris, New Statesman 9 March 2009

> “Many of the problems our economy faces are the result of the use of misguided models. Unfortunately, too many [economic policy-makers] took the overly simplistic models of courses in the principles of economics (which typically assume perfect information) and assumed they could use them as a basis for economic policy… We need a new balance between market and government.”
>
> Joseph Stiglitz, New Statesman 16 October 2008

Comments like these have been heard many times since last summer, and in essence what they are really saying is that we have entered a period in which we are no longer so certain that we know *how to govern a society* based on capitalism. It seems trivial to say that the mere existence of a democratic system doesn’t fully specify the answer to that question, but nevertheless I think it is fair to say that many people in democratic nations have been profoundly shocked by the realization of the structural weaknesses that a democratic economy can have, in which banks can vanish (and your money with it) and entire nations can go bankrupt almost overnight.

I don’t claim to have any answers to the immense and difficult questions that arise from the current crisis. But I am struck by two things. First, it seems that the present world situation is rekindling interest in old questions about governance that many had complacently considered to be settled, particularly in relation to the interplay between laws and institutions and the freedoms of individuals in an economic context. Second, it seems there are at least some people who, like Joe Stiglitz, feel that a part of the solution lies in having a better fundamental understanding – a better model – of how society and its structures operate.

In other words, perhaps we need to start looking again at whether there is a genuine science of society, and at the role such a thing might have in helping us to erect the framework on which society hangs.

In essence, this entails asking questions such as: Which (if any) aspects of the evolution of society can be regarded as inevitable? Which are susceptible to accurate probabilistic estimation? And which are too dependent on the vicissitudes of human behaviour to be accessible to any degree of prediction?

These and other such issues are explored in several recent general and technical books and reviews, among which these are some:

- N. Johnson, *Two's Company, Three Is Complexity* (OneWorld, 2007)
- M. Buchanan, *The Social Atom* (Bloomsbury, 2007); *Ubiquity* (Weidenfeld & Nicolson, 2000)
- P. Ball, *Critical Mass* (Heinemann, 2004)
- C. Castellano & S. Fortunato, *Rev. Mod. Phys.* **81** (Jan-Mar 2009)
- J. H. Miller & S. E. Page, *Complex Adaptive Systems* (Princeton University Press, 2007)
- J. M. Epstein, *Generative Social Science* (Princeton University Press, 2006)

Much of the impetus in this field in the past two decades or so comes from work conducted by physicists, but it has also come from quite different areas of the natural and social sciences: from ecologists, epidemiologists, computer programmers, mathematicians, and from economists, anthropologists, urban planners and sociologists. This, as you might imagine, is an interesting blend that doesn't always mix well. Some of the tensions and difficulties stem from cross-disciplinary differences in language and concepts, but I suspect that the most problematic clashes of culture have stemmed instead from differences in aims and in ethos. One sometimes finds that the questions and objectives of one group are meaningless, irrelevant, or even unwelcome to another.

This is to some degree what one might anticipate when the interests and preoccupations of disparate fields converge. But I want to suggest later that this strand of complexity science is not simply a convergence of distinct streams of enquiry, but a rejoining of ways that parted long before anyone began talking about complex systems, indeed before the emergence of science in its modern form.

To pre-empt some of that discussion, I'll say now that the notion of a society governed by laws akin to and as inevitable as the laws of physics became a dominant theme in the political philosophies of the nineteenth century. And it wasn't just a hypothesis – scientists thought they had discovered *evidence* of such natural laws of society. When they looked at the statistics of births and deaths from year to year, they found that not only were the numbers roughly constant but their variations always fitted onto the same kind of distribution. It looked like this: [**Slide 4**]

We know what this is, of course: a bell curve or Gaussian, the probability distribution of a series of random, independent events. Nineteenth-century scientists were astonished that they found gaussian curves not just for statistics of births and deaths, but also for crimes and suicides. You can't generally do much to alter your birth or death, but crimes and suicides are acts of volition, and so it seemed incredible to them that these acts of apparent free will nevertheless seemed to obey a mathematical law by fitting onto a gaussian curve. Today we find nothing very remarkable about this – the factors behind crimes and suicides are complex, but also largely independent from one case to the next, and so in effect these things happen at random. But to nineteenth-century scientists, they seemed to point to the existence of natural laws of society, which were encoded in the gaussian curve (or error curve as it was commonly known after Gauss's work).

These sorts of demographic statistics in fact provided one of the key drivers for the flourishing of statistics as a branch of mathematics, beginning perhaps with Abraham de Moivre's work on the normal distribution in the 1730s and continuing with its use in social statistics by Fourier, Laplace and Quetelet. But now we know that all manner of social statistics don't follow a gaussian law at all. They show a different sort of statistical behaviour. They look, for example, like this:

[**Slide 5**]

This is the number of companies of different sizes in the USA in 1997. It's a power law, which differs from a Gaussian distribution with the same mode by having a far higher probability of extremal values. What this tells us is that firm sizes aren't simply random. This in turn implies that they are not just growing independently of one another. I'll come back shortly to what that might mean. First, let me show you some other examples of power-law statistics in social phenomena.

[**Slide 6**]

This is the number of incoming connections to web pages on the World Wide Web.

[**Slide 7**]

This is the probability distribution of towns around London – how many towns there are of different sizes. [1981 data]

[**Slide 8**]

This is the probability distribution of wealth in the UK in 1996. You see the same shape for other countries too, but with different slopes depending on how wealthy their economy is. It implies that most of the money is in the hands of just a tiny proportion of the population. It's been estimated for the USA that 1 percent of the population has 40 percent of the wealth. These wealth distributions become power laws towards the rich end of the scale, something that was first identified in 1897 by an Italian sociologist named Vilfredo Pareto. This power-law portion of the distribution is called the Pareto law.

[**Slide 9**]

This is an interesting one: it shows the relationship between the size of human conflicts, in terms of human fatalities, and the frequency of conflicts of that size, for conflicts ranging from small skirmishes to world wars. It was first assembled from data collected by the British physicist Lewis Fry Richardson between the 1920s and the 1950s. Richardson was a Quaker, and he hoped that these data would help elucidate the causes of war and thus promote world peace. This sort of analysis has recently been extended to individual wars (such as those in Iraq and Colombia), allowing comparisons to be made of their qualitative nature and the way they are changing over time (**Slide 10**).

[**Slide 11**]

This shows the distribution of times taken by Einstein and Darwin to respond to their correspondence. One finds similar statistics for the time taken to complete other 'queued' tasks.

[**Slide 12**]

This is the distribution of votes gathered by over 10,000 candidates for state deputies in Brazil's 1998 general election.

And so I could go on. In the 1940s, the American social scientist George Kingsley Zipf [G. K. Zipf, *Human Behavior and the Principle of Least Effort* (Addison-Wesley, Reading, MA, 1949)] first pointed out that many social statistics show these power-law distributions. But he had no real explanation for them. He proposed that this behaviour was somehow related to what he called the 'principle of least effort': that it reflected people's tendency to look for ways of doing things that would, in their subjective estimation, cost them as little effort as possible. But he had no way of proving that a universal principle like this was really what lay behind the power laws.

Physics has, however, now provided explanations for power-law behaviour. One of the most popular theoretical frameworks for understanding power-law distributions is self-organized criticality, typically explained with reference to a pile of sand that is building up as new grains are poured onto the peak [**Slide 13**]. Every so often, the growth of the pile will be interrupted by a landslide of grains running down the slopes. These avalanches can be of any size ranging from just a few tumbling grains to the collapse of more or less the entire slope: The frequency distribution of the sizes of avalanches looks like this:
[**Slide 14**]
– another power law (at least in simple theoretical models). So the sand pile is constantly undergoing seizures, or fluctuations, of various sizes, with a power-law probability distribution. Fluctuations with those statistics are characteristic of systems in a critical state, like the density fluctuations at the critical point of a liquid and gas. Traditionally, critical states were regarded as very unstable, and liable to collapse. But the critical state of the sand pile keeps reforming after every avalanche. That is why it is said to be self-organized. The ubiquity of power laws in social statistics has led to proposals that social structures, such as international relations or economic markets, exist in self-organized critical states. The truth is that it is very hard to find good evidence for this, but it does raise the intuitively appealing notion that society displays this mixture of robustness and instability, making it constantly subject to disturbances of all sizes.

Per Bak was convinced that economic markets work in a state of self-organized criticality. And indeed markets do seem to constantly undergo scale-free fluctuations. If you look at the medium-term behaviour of an economic index, it is very erratic. Economists have generally treated these fluctuations as random, because they look that way: they are assumed to be just a kind of noise in the system. But it has been known since the work of Mandelbrot in the 1960s that these fluctuations aren't random at all. That's to say, they have a non-gaussian, so-called fat-tailed probability distribution, which looks like this (**Slide 15**). One of the important features of this distribution is that there are more big fluctuations than would be expected from pure randomness. This is very significant, since it is often those relatively rare big fluctuations that economists are interested in: the booms, slumps and crashes. In a gaussian distribution, market crashes are so rare as to be negligible. In reality, of course, they are far from that, as we are now painfully aware. So if you try to make market forecasts based on the wrong statistical distribution, you can go badly astray. And yet even today, some economic theories, like the celebrated Black-Scholes model for options pricing, insist on using gaussian statistics. This is mathematically convenient, but can, in that particular case, give disastrously misleading indications of risk.

Power-law behaviour is a clue that the fluctuations arise from collective behaviour within the system: in other words, that the dynamics stem from strong interactions between their component parts.

What does this imply for social power laws? Well, let's go back to voting statistics. If they are non-gaussian, that implies that the votes *aren't* independent.

And in fact, that's just has been argued. Bernardes *et al.* have devised a model [**Slide 16**] of the electorate in which each voter is represented as a magnetic spin that may point in one of several directions, with each direction representing a different local candidate. In this magnetic model, the voters are assumed to exert an influence on their neighbours to try to make them align their political allegiance in the same way. If a voter is surrounded by many others that all have the same alignment, there will be a correspondingly strong force persuading that voter to vote the same way as all the others. The model is determined by the rules governing realignments, which I've shown here – basically a kind of majority rule. This is in fact the so-called Sznajd ('Schneid') model proposed in 2000 [**Slide 17**], which has become one of the most popular for studying opinion dynamics. Bernardes et al. found that if the web of social connections between voters has the same kind of structure as a real social network, their model predicted that the final voting statistics would follow a power law.

Now, opinion dynamics is one of the oldest forms of modern social physics: physics-like voter models go back to the 1970s with the work of Wolfgang Weidlich, and Serge Galam introduced the parallel with lattice models of magnetism, in particular the Ising model, in the 1980s. The Sznajd model has been studied in considerable detail, but one can in principle postulate all kinds of local opinion-forming rules and topologies in these lattice models. I should say that this work on the Brazilian elections has been criticized both from the point of view of the empirical statistics and the model predictions- it seems, for example, that non-trivial behaviour, different from a complete consensus on a single candidate, is only transient. But I think that the key *sociological* point is clear regardless of the details. We learn, in case we needed telling, that Margaret Thatcher was wrong – there *is* such as thing as society. Society is not just the sum of independent individuals; rather, when we put all those individuals together, they interact with one another and start to show some form of collective behaviour, and the result is something that looks quite different from what we'd predict from a collection of isolated individuals. In a system like this, there is a big potential for knock-on effects which can sweep, like a sand-pile avalanche, through the system. This means that small influences can have inordinately large effects. That might be cause for concern in, say, an electoral system where different parties have markedly different campaigning budgets or levels of press support.

Particle people

Now, already I have taken the audacious, perhaps even insulting, step of creating a kind of social physics by pretending that people are just like atoms responding to forces.

That sits very uneasily with what many sociologists have long assumed. When we look at society in all its variation and bewildering activity, we are tempted to imagine that this richness and complexity of behaviour must result from the richness and complexity of the human mind. This has often been the view of the social scientist: to assume that the only way you can explain human behaviour is with models that attribute a wealth of psychological complexity to individuals. The economist Robert Heilbroner puts it like this [**Slide 18**]:

> there is an unbridgeable gap between the 'behaviour' of [subatomic particles] and those of the human beings who constitute the objects of study of social science... aside from pure physical 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: the particles of physics do not 'choose' to behave as they do.

It's a valid concern, but it risks overestimating both the power and the scope of free will. In many social situations, it is unrealistic or even meaningless to assume that we can do whatever we want. We often have only a very tightly constrained range of choices. If we are driving a car, we can in principle steer it anywhere at any speed within the vehicle's capability, but of course we don't: left to our own devices, we will all tend to drive in a line along the road, in Britain on the lefthand side, at a speed roughly appropriate to the context, between our departure point and our destination. When we vote, we choose one candidate or we choose another candidate, generally from a very short list of alternatives. Our actions, nominally completely free, are constrained by a wide variety of factors: social norms and conventions, economic necessities, restricted range of choice. We are far more predictable than we like to believe.

Even so, we might imagine that, from the palette of available options, we select freely. But it becomes rapidly clear once we look more closely that we do not. The key factor – and this is what social and economic scientists have tended to overlook in their models, while it is intrinsic to statistical physics – is interaction. We are affected by one another. People don't drive down the local high street at 80 mph because there are others in the way, and we normally aim to avoid collisions.

There is a strong tradition in social sciences of creating psychological models of phenomena – that is, trying to understand social behaviour on the basis of individual psychology. Sociobiologists like E. O. Wilson have argued that social science could therefore be made more scientific if these models are more firmly rooted in the evolutionary biological origins of individual behaviour. This is probably true, but it makes the often unwarranted assumption that social behaviour is a straightforward extrapolation of individual behaviour. It seems that this is often not the case at all: that the behaviour of a group – how it organizes itself into institutions, for example – cannot be deduced or predicted from the predilections of an individual. It is very clear in the statistical physics of complex systems that once individual agents – in physics they could be atoms or electrons, say – once they start to interact, completely new collective modes of behaviour can arise. We can study a single water molecule as closely as we like, but we would never deduce from that that it can adopt three different states of matter, or that it turns from gas to liquid at 100 ^{o}C. We can get that kind of understanding only by looking at water molecules collectively.

Sociologists have long acknowledged that we can be influenced by others, of course, but it's been very hard to get a quantitative handle on the effects of that. But in 2006 Duncan Watts at Columbia and his coworkers devised an ingenious experiment [**Slide 19**]. They enlisted over 14,000 volunteers who could listen online to songs recorded by 48 unknown rock bands, and download ones they liked. Social influence on these choices was studied by presenting the information to the participants in different ways. Some were simply shown a list or grid of the songs in random order. The number of subsequent downloads was then assumed to be a measure of the 'objective quality' of the songs – not in any critical or musicological sense, but in terms of how this was perceived by the study group. Other subsets of the group were provided with information about the song's popularity, in terms of the number of times it had been already downloaded. This information was supplied in two different ways: either as a bare number accompanying the randomly gridded songs, or as a ranking that determined the order in which the songs were listed vertically. The latter makes the 'social information' much more transparent, and thus increases the strength of the social influence.

This influence was found to have a strong effect on the choices made. In particular, the stronger the social interaction, the more 'inequality' there was in the outcome: popular songs were more popular, and unpopular ones less so. In other words, the choices reinforced one another when they were known to other members of the group. Moreover, in the case of social interaction, 'quality' (as measured by the group that made decisions independently) became a less reliable indicator of a song's 'success' as this interaction got stronger. While 'bad' songs never did particularly well, and 'good' ones rarely did poorly, all things seemed possible in between: social interaction made it more likely that mediocre songs could become runaway successes.

Pedestrian dynamics

It is now becoming clear that all kinds of seemingly complex behavioural patterns can be reproduced in models that have very little psychology in them at all, where the resulting dynamics are determined primarily by the nature or simply the fact of interactions between the agents.

One of the simplest examples concerns the question of how people move around space. Dirk Helbing, now at ETH, and his coworkers have devised a model of how people move around space that represents those people as atom-like particles that interact through forces of attraction and repulsion. If you want to see the force of repulsion at work, take a look at a crowded beach [**Slide 20**]. People space themselves a more or less even distance apart, which of course is precisely what particles do when they repel one another. People aim to avoid invading one another's personal space by coming too close – unless of course they are in a group, a family or friends, in which case it's as if an attractive force binds them together.

Helbing has carried out simulations of the motions of these particle people on a computer, in various different situations. For instance, moving down a corridor: some people are programmed to try to move in one direction, and others in the opposite direction. Each person in this model will, if they can, reach a preferred walking speed; but they will slow down if necessary to avoid colliding with someone else. With no more ingredients than this, the particle people show some interesting behaviour. Here's a snapshot of the simulations [**Slide 21**], and you'll see that what has happened is that they've become organized into counterflowing streams: blue is going in one direction, red in the other. There's nothing in the rules of the model that specifies that these streams must form – they just appear, as indeed they do in real life. It makes sense for the crowd to organize itself this way, because it makes the chances of collisions smaller. But these people particles don't *know* that; they don't really have any intelligence at all, they are like robots. Yet they show seemingly intelligent, anticipatory behaviour.

If, however, the people particles are programmed to want to exit through the door much more quickly, something different happens. Then their desire to move fast can overwhelm the repulsive force that keeps them apart, and they start to come into contact. If they are slightly rough, sandpaper-like particles, so that there is a force of friction when they touch one another, they can become jammed in the doorway, preventing people from getting out [**Slide 22**]. Helbing and his colleagues think that this is what happens when a crowd panics. The jamming can mean that, even though the people are trying to move faster, the room actually empties more slowly. Effects like this can undermine architects' calculations for how many emergency exits a building needs. So computer simulations like this might help to improve safety in the design of public buildings. Michael Batty and his coworkers at University College here in London have used simulations like this to help design crowd-control and safety measures for the Notting Hill carnival [**Slide 23**] – here's a simulation of part of the carnival route [**Slide 24**]. Models like this are now being used by architects and urban planners: a London-based company called space Syntax used them to design the pedestrianization of Trafalgar Square, for instance, and they have been used to understand visitor behaviour at the Tate Britain art gallery.

Pedestrian simulations like this can also be used to understand how people wear down trails over open grassy spaces. Here's a trail system that has been trodden down at the University of Stuttgart [**Slide 25**]. No one planned this trail – it's the product of many feet, with each walker simply going about their business of walking from one university block to another. The striking thing about these trails is that they don't represent the most direct routes between any two of the common points of entry or exit. They arise from a compromise between directness and peoples' tendency to walk where they can see others have walked. Helbing and colleagues have modelled the evolution of trail systems like this, and find that they evolve from direct, linear pathways initially [**Slide 26, 27**] to the kind of gently curved trails we tend to find in reality [**Slide 28**]. Here the entry and exit points are at the four corners of the square. Park planners tend to think linearly, or at least in terms of precise geometric lines, and they sometimes find that the paths they've planned are subverted by trails that people wear down spontaneously. How much better it would be if they could predict where people will *want* to walk, and build their paths accordingly. The same is true for town planners, who have all too often planned public spaces according to how they think people ought to move around them, rather than how people would naturally like to do so.

Traffic

It is not a big step to go from pedestrians to drivers. Indeed, road traffic is in some ways even simpler to model, because it is even more constrained – we simply drive in single or multiple file down the road in one direction or the other [**Slide 29**]. Using much the same assumptions as for pedestrians – that drivers try to reach a certain preferred speed if they can, but will slow down to avoid collisions – it's possible to mimic many of the kinds of traffic flow that we experience on the roads. If traffic is fairly dense, a single small perturbation such as one driver braking rather suddenly can trigger the formation of a jam or even a series of jams, creating stop-and-go traffic in which drivers have to pass through several waves of jams. Here's an example [**Slide 30**]. Each of these lines traces out the motion of a single vehicle, and a constant slope indicates a constant speed. A single, transient glitch here causes these jams, in which drivers have to slow right down. For drivers over here [right], these jams seem to have no apparent cause – and I'm sure we are all familiar with that.

These traffic simulations show another common feature of social physics: abrupt switches between global modes of behaviour, which occur in a way precisely analogous, indeed sometimes formally equivalent to, equilibrium phase transitions [**Slide 31**].

Traffic simulations, guided by input from a few monitoring points on specific roads, are now being used to predict traffic flow in cities of the Rhine-Westphalia region in Germany and in Dallas, Texas. They can also help planners predict how to make accidents less likely and how best to synchronize traffic signals. Ultimately they might even become incorporated into automated driver-control systems – autopilots, if you will – in vehicles. Simulations have shown that such systems could potentially dissolve traffic jams that would otherwise form under purely human guidance, because of the more intelligent driving and braking that autopilots can exhibit: in effect, this smoothes away the jam-nucleating fluctuations.

Segregation

The general idea that individual behaviour might not be a good guide to the behaviour of a group – or conversely, that the way a group acts might tell us rather little about how an individual think – was demonstrated well before the current emergence of social physics, in famous and in some ways controversial work in the 1970s by the US economist Thomas Schelling. Schelling wanted to understand the occurrence of segregation in a society: a common and sometimes problematic demographic phenomenon, at that time exemplified by the ghettoization of black neighbourhoods following ‘white flight’ from American city centres.

A phenomenon like white flight could easily be read as an expression of substantial racial prejudice and intolerance. But Schelling showed that a lattice model in which two types (‘colours’ – grey and black below) of agent move about on a grid of sites offers a quite different perspective [**Slide 32**]. In this model, each agent will move to an empty (white) site if more than some critical proportion – say, a third – of its neighbours are of a different ‘colour’. With these rules, an initially well-mixed population quickly evolves to one that is highly segregated.

The segregation is motivated by a preference of agents to be among a majority of their ‘own type’, but the degree of intolerance is not especially high – they are happy to accept up to one third of their neighbours being a different ‘colour’. (Indeed, segregation can develop even if the threshold is even higher than 50 percent) To look at it another way, one would be wrong to infer from the highly segregated society that is produced that the agents in that society must be extremely prejudiced. The segregation is a collective result of interactions, and would not be anticipated simply by examining the ‘rules’ governing the choices made by individuals. Quite how one interprets and acts on these findings in a social context is a complicated question – but one implication is that, once you set up the possibility for exercising preferences that bring together like with like, for example in school or hospital selections, you are very likely to create the conditions for rapid and pronounced segregation. Politicians seem not to have yet assimilated this message.

Alliance formation

The formation of alliances is common in national and international politics, business and industry. It may arise, for example, in the setting of technical standards. The case of the QWERTY keyboard reminds us not only how long these issues of technical standardization have been around but also how the outcomes can be hard to predict and, once locked in place, even harder to change, even if they are sub-optimal – a phenomenon economists call ‘lock-in’.

This sort of competition in business motivates the formation of alliances: companies may figure that by joining together, they are more likely to end up on the winning side. Typically this ends up creating just two rival camps. Two alliances supply the ideal option of allowing each company to be part of a big group while still actively opposing its main rival.

Robert Axelrod, a political scientist at the University of Michigan, and his coworkers have developed a physics-based theory for studying this kind of situation, which they call landscape theory. The players in this game – the companies – are like gas particles on the point of condensing into two or more droplets. They are drawn to one another by a kind of attraction, yet are also kept apart by repulsions [**Slide 33**]. Out of this push and pull emerge configurations in which the particle-like agents aggregate into alliances. The principle that governs their final configuration is the same: what is the most stable way to arrange them? In other words, what is the equilibrium state?

Each possible configuration has an associated energy, and this defines an ‘energy landscape’. In the lowest-energy, equilibrium configuration, no firm can bring about any further stabilization by switching from one camp to another. This is what game theorists call the Nash equilibrium.

The challenge, then, is to find this equilibrium state. If the number of agents is small, the search can be done exhaustively by calculating the ‘energies’ of all possible aggregates and picking out that with the lowest.

Axelrod has applied this model to the standardization of Unix computer operating systems in the 1980s, where he was able to use the properties of US computer companies at that time to predict retrospectively how they would divide into two camps [**Slide 33**]. And he has applied the landscape model to the formation of alliances between seventeen nations in the approach to the Second World War, again finding a pretty close match to the way the countries divided into Axis and Allied powers.

Firm formation

How do firms grow, and what limits this growth? The economist Robert Gibrat proposed in 1931 that a firm grows at a random rate, amplified by the existing size of the firm. This is Gibrat’s now-celebrated Law of Proportionate Growth, an it leads to a log-normal size distribution. But as we saw earlier, the true distribution looks more like a power law.

Robert Axtell of the Brookings Institute in Washington DC, has offered an explanation of these distributions in terms of an agent-based model in which firms arise by the aggregation of many interacting agents each following their own agendas for balancing profit and leisure.

This model has no stable Nash equilibria. It can never settle down into an unchanging state. There is constant flux as firms boom and go bust [**Slide 34**]. But the probability distribution of firm sizes is a power law. And the model suggests some interesting things about what makes a firm successful. The firms that do best are not those that aim to make the most profit. Rather, longevity in a company stems from being able to attract and retain productive workers. A firm fails not when its profit margins are eroded but when it is infiltrated by slackers.

Game theory

In many of these models, agents simply try to find the best option – to maximize their utility, in economic terms – in the face of what all the other agents are doing. But one of the key components of game theory is that agents try to *anticipate*, rather than just respond to, the choices of others, and act accordingly. That’s so in the case of the Prisoner’s Dilemma, and it is also the case in the game-theoretic model that has probably been most studied by physicists and complexity theorists, namely the minority game

[**Slide 35**]. Briefly, this is a model in which agents use diverse algorithms to anticipate which of two choices the others will make, and then to choose the opposite: the aim is always to be in the minority. One might see this, for example, as an analogue of market traders' wish to be a seller in a buyer's market, or vice versa.

Econophysics

This is just one aspect of the now rather immense field of econophysics [**Slide 35**]. I'm going to say no more about this, which is a topic in itself, except to point out that it is perhaps going to be the prime testing ground for social physics models in the years to come – not least because traditional economics has now revealed itself as so desperately in need of fresh ideas.

Who needs social physics?

I wasn't to return finally to the broader questions that I raised at the outset: Is there a physics of society, and if so, how should we use it?

Theories and ideas about the best way for a society to operate reach back at least as far as Plato's *Republic*. But the notion of approaching such questions using the methods of science – that is, of developing a social science worthy of the name – dates to the early beginnings of the Enlightenment. That was, of course, an age of mechanism, when Galileo, Descartes and Newton were starting to propose that nature can be understood like a machine in which forces acting between the component parts give rise to precise mathematical laws that allow future behaviour to be predicted. As more and more of nature began to reveal itself as governed by physical laws, philosophers started to wonder if such regularities applied to the human sphere too. The first attempt to develop that idea was made by Thomas Hobbes in the 1630s and 40s, when Hobbes used Galileo's physics of motion to derive the unappealing conclusion that absolute despotism was the best way to govern a nation [**Slide 36**].

We have to remember that Hobbes developed his ideas in the approach to the English Civil War, and his most famous book, *Leviathan*, was published while Hobbes, thought to be a Royalist sympathizer, was still in exile from Cromwell's England in Paris in 1651. His argument that absolute monarchy was the best form of government has to be seen in the context of the desperate desire that many felt to avoid such civil conflict at almost any cost.

And although the Scottish philosopher David Hume argued that Hobbes' politics promoted tyranny, Hobbes did something important. He suggested that, in making our laws and institutions, and in studying our behaviour, there is room for rationality and logic. It is appropriate to see Hobbes as a part of the Enlightenment tradition that ultimately helped to liberate society from antiquated and arbitrary ideas about how it should be run, and which gave rise to the liberal political philosophies of John Locke and John Stuart Mill.

The notion of a form of natural social law was developed by thinkers ranging from Immanuel Kant to Pierre-Simon Laplace, and found particularly explicit statement in the work of the French philosopher Auguste Comte (1798-1857), who coined the term *physique sociale*, social physics, in the early nineteenth century [**Slide 37**].

Others at this time, such as Henry Thomas Buckle in England and Leo Tolstoy in Russia, wondered whether there is some degree of inevitablity 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. Something of that aspiration can still be perceived among Marxist historians today.

And to the early social physicists, if we can call them that, the appropriate analogy for a physics of society was Newtonian mechanics, in which mathematical laws governed the motion of every individual body in the universe, so that all these motions and trajectories could be calculated. When in the nineteenth century physicists began to think of the atomic world as a kind of billiards game in which atoms were like smooth, hard balls that moved through space until they collided with each other and bounced off according to Newton's laws of motion, it was the statistical regularities seen in *social* sciences that encouraged James Clerk Maxwell to propose that, even if we can't use Newtonian mechanics to formulate a complete description of atomic-scale behaviour, we can anticipate that mathematical laws will arise out of the average, interdependent motions of all these invisible particles.

I think it is important to remember that these early thinkers in social physics were very, even primarily, interested in the social, political and moral implications of their ideas. Today, in contrast, social physics is often regarded as a kind of playground for exploring new phenomenology of complex systems. That's a good motivation in itself, but it would be a shame if the field doesn't display the courage of its convictions that it may have something to tell us about real-world behaviour.

For example, social physics might be seen as a way of putting choice *back* into social modelling, which has in the past tended to draw on traditional economic models and thus on notions such as rational utility maximization which threaten far more to portray people as automata. There is not always a single, rational 'best choice' in a given situation; or even if there is, people do not always take it. Physics-based models explicitly allow for a multiplicity of choices, which are typically assumed to be made not on rigidly deterministic grounds but on a probabilistic basis. All they really assert is that those choices are rarely independent: we are influenced by what others choose. Once that happens, the outcome can be non-intuitive.

I would argue that the primary value of this topic is often to challenge entrenched preconceptions about how human society works. Policy makers are all too prone to linear thinking: they assume that if we understand how an individual tends to think or behave, we can understand what a population will do. It is surely time to move beyond this position and to acknowledge that the interactive nature of society makes it a truly complex and nonlinear system. Physics-based modelling tells us not only that interactions may change the picture entirely, relative to what lone individuals will do. It also shows that this injection of complexity does not necessarily make the problem impossible, because (as is the case with road traffic) there are likely to be robust modes of collective behaviour that are relatively insensitive to the fine details and idiosyncrasies of individual actions and responses.

Most of all, I think it is important to recognize that this approach to understanding human social behaviour is not about, and must not be about, using science to tell us what to do – how to plan and govern and make laws and institutions. It is a tool to help us understand the likely consequences of particular choices. No science can be a substitute for the wisdom and compassion needed to create better societies. We should be bold in the questions we ask, but humble in the answers we offer.