

The Physics of Institutions

A symposium at the London School of Economics

26 March 2003

Philip Ball

Consultant Editor, *Nature*

4-6 Crinan St

London N1 9XW

p.ball@nature.com

Abstract

In economics, business, society and politics, institutions and structures typically arise from the interaction and consensus of many individuals. Although these modes of organization are usually subject to planning and legal constraints, there now seems good reason to believe that their properties and behaviours may be dominated by factors that arise unbidden and are very difficult to control. Several centuries of economic theory still do not suffice to avoid market slumps and recessions. Most businesses fail or disappear within a few years. Free markets do not seem to flatten wealth distributions. Harsh penal systems do not systematically and unequivocally reduce crime.

In short, there seem to be ‘laws’ in these social systems that have at least something of the character of natural physical laws, in that they do not yield easily to planned and arbitrary interventions. Over the past several decades, social, economic and political scientists have begun a dialogue with physical and biological scientists to try to discover whether there is truly a ‘physics of society’, and if so, what its laws and principles are. In particular, they have begun to regard complex modes of human activity as collections of many interacting ‘agents’ — somewhat analogous to a fluid of interacting atoms or molecules, but within which there is scope for decision-making, learning and adaptation. These ‘interacting agents models’ often reproduce many aspects of the observed behaviour even on the basis of surprisingly simple assumptions about the ‘rules’ governing an individual’s actions and interactions. This suggests that complex human systems have a tendency to become ‘self-organized’ into patterns and structures that no one has planned or foreseen. An ability to understand these structures is a prerequisite for being able to achieve a degree of control over the outcome, or to discover the boundaries of that potential for control.

This seminar will present and discuss examples of agent-based models and other physics-based descriptions of behaviour in areas such as economics, business, demography, voting, crime and decision-making. It will explain some of the underlying and unifying principles of a ‘physics of society’, and show what it can and cannot provide as a tool for guiding policies.

Are there natural laws of society?

Suppose you could predict:

- what set of policies would guarantee a party electoral victory
- how companies will band together to form alliances and conglomerates
- what set of international policies will encourage democracy and discourage conflict
- what the stock market will do tomorrow
- how congestion charging will affect London traffic three years from now
- how harsher sentencing will affect crime statistics
- the likely lifetime of a new small business
- the chances of you and me sharing a mutual friend

Some of these predictive capabilities would be clearly useful. Some would be immensely beneficial. Some would be so valuable that those who possess them might want to keep them secret. All are, in some degree and to some parties, desirable.

But which of them are possible, and which are just idle fantasies?

In other words, 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 are old questions. 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. Let's just recall the context of that period, in the early seventeenth century: it was an age of mechanism, when the likes of 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. This is how Galileo understood the laws governing the motion of objects, and it led Newton to his laws of gravitation that allowed scientists not just empirically to predict but to appreciate the underlying basis for the regular motions of the planets. 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. They regarded the individual human body as a well-oiled mechanism, an assembly of so many levers and pumps and timekeepers. The next step was to progress from individuals to society as a whole: to seek for a physics of society.

The first attempt to develop this tells us perhaps all we need to know about the idea of using science to prescribe our social institutions. It was made by Thomas Hobbes in the 1630s and 40s, when Hobbes used Galileo's physics of motion to derive the conclusion that absolute despotism was the best way to govern a nation. If that's where a truly scientific social science leads us, you might reasonably say 'you can keep your physics of society.'

But over the past two decades years or so, physicists and other scientists have regained an interest in trying to apply the tools of science to social phenomena. Their approach is quite different from Hobbes's. They are asking 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 outcome will be?' The science is being used not to tell us what is the right or wrong way to do things, but to try to understand which choices lead to which consequences.

This is one of the great failings of some types of social and political science: the naïve assumption that the conditions needed to achieve a particular objective are obvious, or conversely, the consequences of a particular set of policies are obvious. It's not hard to think of examples that seem to defy intuition, such as how, in some circumstances, road building can increase rather than decrease congestion. In some cases, failures of public policy to achieve certain end results might arise because of the simple neglect of certain aspects of human psychology. But there are other instances where, even if one has accounted for every relevant factor, still the outcome of a particular set of rules or conditions might be quite different from the one expected.

In essence, I want to look at what science has to say about some of the reasons why that might be so, and how it can lend tools that permit a better prediction of the consequences of social decision making in many contexts. I shall outline some of the central ideas in contemporary physics that seem to have some applicability in these fields – in other words, try to give a flavour of what is out there in physical science that might be of some value to the business and social science communities. In the second part of the seminar, Paul Ormerod will put some flesh on these bones by outlining some of the physics-inspired modelling that he has been conducting of phenomena in the business and economic spheres.

You could say that, in seeking to extend physical science into social science, one is asking 'are there laws of society?', in the same sense as there are laws of gravitation or laws of electromagnetism. Certainly, that's how the early pioneers in this field saw it. The French philosopher Auguste Comte (1798-1857) believed that such laws could be uncovered when he coined the term *physique sociale*, or social physics. In his *System of Positive Philosophy* (1830-42), he argued that this would complete the scientific description of the world that Galileo and Newton and others had begun. He said:

Now that the human mind has grasped celestial and terrestrial physics, mechanical and chemical, organic physics, both vegetable and animal, there remains one science, to fill up the series of sciences of observation—social physics.

Several thinkers in the eighteenth and nineteenth centuries, including Immanuel Kant, Henry 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 these positivistic thoughts was that there was a kind of regularity in the statistics of social phenomena. Scientists and philosophers became interested in social statistics in the seventeenth century, when the London businessman John Graunt began to collect yearly mortality figures for the city. The famous astronomer Edmund Halley, one of Newton's few close friends, also became captivated by mortality statistics. Graunt argued that statistics like this could provide a solid, empirical basis for formulating political policy.

What people began to realise is that there is a certain kind of predictability in these social statistics. It wasn't just that, on average, more or less the same number of people die each year, or even that this constancy applies also to subdivisions of society according to age or profession. It was also the deviations from averages that began to interest them. By the early nineteenth century, scientists like Pierre-Simon Laplace, the great French mathematician and astronomer, had discovered that a whole variety of social statistical data could be fitted to a single mathematical curve (see figure).

This is the well-known bell curve, known to physicists and mathematicians as the gaussian and to statisticians as the normal distribution. This curve describes a probability distribution. You can see it either as a summary of empirical data or as a predictive function. Suppose, for example, that this curve relates to the heights of adults in London. If you measure everyone's height and plot the numbers versus height, you will get a gaussian curve. But once you've ascertained that this is the statistical distribution of heights, you can use the curve predictively to say what the chances are that any randomly selected individual will have a particular height. The most probable height is the average one, and the probability falls off sharply for either extremes of shortness or tallness. In other words, these early social scientists realised that gathering social statistics has a predictive value.

To their surprise, they found that gaussian curves described not only statistics of births and deaths, over which each individual has rather little control, but also those of volitional acts such as crimes and marriages. To some people, this seemed an affront to the idea of free will, since it suggested that supposedly free choices were governed by a mathematical 'law'.

The key consideration, then, is that if there is a physics of society, it will be essentially a statistical one. Mathematical regularities appear only when we look at populations or at large data sets. Conversely, what this means is that in general specific predictions must be probabilistic: we can't say what will happen to any of the individual components or agents in the system, but only what the probabilities of the various possible outcomes are. In 1862 John Stuart Mill recognized this statistical aspect of a scientific sociology when he wrote that

very events which in their own nature appear most capricious and uncertain, and which in any individual case no attainable degree of knowledge would enable us to foresee, occur, when considerable numbers are taken into the account, with a degree of regularity approaching to mathematical.

This is where the connection with physics comes in. 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 element in the system so that all these motions and trajectories could be calculated. But in the nineteenth century a new type of physics arose. Newton suspected that just as the trajectories of the celestial bodies could be understood and predicted on the basis of the forces of gravity acting between them, so could the behaviour of matter at the other extreme of scale – individual atoms – be described by laws of motion determined by interatomic forces. It was just that no one knew what these forces were.

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. The trouble was that they knew they couldn't hope to be able to see or measure these motions. And even if they could, atoms are so numerous that it would be impossible to keep track of all their trajectories at once. It was actually 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. Maxwell began to think about the probability distributions of atomic motions, which he assumed would also be circumscribed by the gaussian curve. This led Maxwell and Ludwig Boltzmann to formulate the science known as statistical mechanics, in which the bulk-scale behaviour of matter, such as the known mathematical relationships between the pressure, temperature and volume of a gas, emerge from the statistics and the averages of inscrutable particle motions.

This branch of science is now used to understand just about all of the properties of everyday matter, from liquids to polymers to superconductors. It has become known as statistical physics, and it is out of this discipline that the new social physics has emerged. Things have come full circle, as physicists have asked 'might we see in society some of the same phenomena that we find in collections of interacting particles?' If we 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?

How can people be treated as particles?

There is an obvious objection to this idea of a science of society, particularly now that the idea that people are mere Newtonian automata has fallen out of fashion. The economist Robert Heilbroner puts it like this:

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, 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 Oxford Street at 80 mph because there are others in the way, and we normally aim to avoid collisions. You might say in this case that there is effectively a repulsive force between cars that keeps them apart. Of course, there isn't really such a force, there's nothing we can measure – but we behave as if there was, and that is sufficient. When we make choices, we 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 an astute or a slow-witted trader to buck the trend and keep buying: this sort of herd-like behaviour is well known in economics. Even in elections, where we might imagine that our secret ballot is purely a matter of personal choice, there is a clear signature in the statistics of collective behaviour: of people being influenced by others. Such behaviour often shows up in probability distributions. A collection of independent, random events shows up as a gaussian statistical distribution of outcomes. If the statistics show deviations from the gaussian form, that is generally a sign that the agents in the system are not behaving independently but are feeling the influence of mutual interactions. This is a simple diagnostic tool learnt from statistical physics.

There's an important corollary to this. 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 statistical physics 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 °C. We can get that kind of understanding only by looking at water molecules collectively.

Key concepts of statistical physics

This is a good place to start exploring some of the central ideas of statistical physics. Maxwell and Boltzmann were essentially interested in gases, and their statistical mechanics couldn't explain why at 100 degrees water vapour suddenly condenses to a liquid. This was explained by Johann Diderik van der Waals in the 1870s, who showed that it follows from the existence of both attractive and repulsive forces between particles in a gas. Maxwell and Boltzmann treated the particles as hard balls that don't interact at all until they touch one another, whereupon they bounce elastically. Van der Waals looked at what happened when he included in the theory the attractive force that acts between the particles over a longer range, and he found that the theory then predicted the liquid state too.

If you heat up a substance, its particles – its atoms or molecules – jiggle around more frantically. This jiggling overcomes the attractive forces that tend to hold the particles together, and so we can understand why substances change from a dense solid to a fluid liquid to a tenuous gas as they are heated. But it is less clear why these changes should happen suddenly. Ice doesn't get progressively softer and jelly-like as it approaches zero degrees; instead, it stays hard until it melts abruptly to water. The same with evaporation: water is either liquid or gas, but not something in between.

These sudden changes are called phase transitions: transitions between the solid, liquid and gas phases of matter. Van der Waals' theory showed how phase transitions happen and why they are sudden. This abruptness is the key point. In social science and in politics, there is a tendency to assume that effects happen in proportion to their cause. If we make a small change to some system or some law, there will be a correspondingly small effect on the system's behaviour. But in statistical physics, that clearly isn't always so. Let's consider the density of a liquid, for example. Most liquids get steadily denser as you cool them (water is an anomalous exception here). So there seems to be a nice, steady, predictable cause and effect here. But the moment you cool the liquid through its freezing point, it all goes haywire. A tiny drop in temperature, which at slightly higher temperatures caused only a tiny increase in density, now induces a sudden, big jump in density as the liquid turns solid (see figure). The same is true for some other properties of the system, such as viscosity: a tiny change in temperature, and the whole system rigidifies. If a system is close to a phase transition, small changes can have major effects. Physicists say that this behaviour is nonlinear, which means that you don't have this simple straight-line relationship between cause and effect.

Phase transitions are central to all kinds of areas of modern physics. But they come in several varieties. They don't always happen in a big jump, as in the case of freezing, melting and evaporation. To understand another important kind of phase transition, we can turn to a different model system: the magnet. If you heat up an iron magnet to 770 °C, it loses its magnetism. Below this temperature it is magnetic; above, it is non-magnetic.

So the change is abrupt in this sense. But there is no jump. If we look at how the strength of magnetism, called the magnetization, changes with temperature, it looks like this (see figure).

This is called a critical phase transition, and the point at which it happens is called a *critical point*. In fact, gases and liquids have a critical point too: it is the point at which there ceases to be any difference between gas and liquid – the point where the transitions of evaporation and condensation disappear. This point is reached by heating a gas or liquid under pressure. For water, the critical point occurs at 374 °C and 218 atmospheres pressure.

To appreciate the importance of critical transitions, we should take a closer look at what happens when a magnet passes through this transition. A magnet consists of an array of magnetic atoms, each of which is like a little bar magnet with a north and south pole. In some magnets these needles can point in any direction, but in others it must choose between only two directions, which we could write as ‘up’ and ‘down’. If there are more of these atomic magnets pointing in one direction than the other, then they add up to give an overall magnetization for the whole array: the material is magnetic. In a substance like iron, each needle feels the magnetic field of its neighbours, and these interactions tend to make the needles all line up. At low temperatures, they all point in the same direction, and the material is magnetic. As the magnet is heated up, the heat randomizes the direction of the needles. If their orientation becomes completely randomized, so that there are as many pointing up as pointing down, then they cancel each other out and the material becomes non-magnetic.

There are thus two distinct but equivalent magnetic states: one in which all the needles point up, and the other in which they all point down (see figure).

If the direction of a needle is flipped by the jiggling caused by heat, it then exerts a force on its neighbours which tempts them to flip too. So the needles have a degree of collective behaviour, owing to the interactions between them. As we approach the critical point, more needles get flipped out of the uniformly aligned state, and small patches of the opposite magnetization start to arise. At first, these patches are small, and they don’t do much to disrupt the overall magnetization. But as we get closer and closer to the critical point, these regions grow bigger. But they don’t all grow equally big. What we find instead is that there are flipped regions of all sizes, from just a single atom to patches approaching the size of the whole system. Exactly at the critical point, we can’t tell any longer which are ‘pristine’ regions and which are ‘flipped’: there are equal numbers and extents of both. In other words, there are up and down regions of all sizes, and on average they cancel out (see figure). What’s more, these domains are constantly shifting, as new ones appear and existing ones grow or shrink. At the critical point, the microscopic state of the system is very dynamic, constantly changing. There’s no characteristic size scale to these fluctuations: they are said to be scale-free.

If we look at the size distribution of these patches, we find that it isn’t random, even though at a glance the state of the system *looks* random. Randomness, remember, would

be signalled by a gaussian distribution of sizes. (Physicists generally plot this distribution using logarithmic axes.) But the actual distribution is quite different: on this kind of graph, it is a straight line. This kind of mathematical relationship is called a *power law*. That's the key point to remember: on a log-log plot, a power law shows up as a straight line.

The characteristic of this kind of probability distribution is that it gives greater weight to big fluctuations, relative to a gaussian distribution. That's to say, the chances of finding a big fluctuation are much greater than they would be if the fluctuations were simply random.

The robustness of criticality

A key characteristic a critical state like this is that it is very precarious. The existence of these big fluctuations means that the system is teetering on the brink of uncertainty between two choices. If we cool a magnet to just below its critical point, the magnetic needles will become aligned in one direction or another. But there is no telling which. At the critical point itself, each is equally probable, and which of the two alternatives will win out is purely a matter of chance – the result of one fluctuation happening to grow so big that it envelops the entire system. The critical point is like a needle on its tip, just about to fall one way or another.

So physicists tended to regard critical states as special but unstable. In the 1980s, however, they discovered that some systems seem able to adopt critical states that are robust. The canonical example was a pile of sand, onto which new grains are being slowly poured. Every so often, the new grains trigger an avalanche. But this could involve just a handful of grains tumbling down the slope, or it could entail a catastrophic sliding of the entire slope, or anything in between. These avalanches could occur on all size scales: they are scale-free. If you plot the probability distribution of avalanches of different sizes, it turns out to be a power law (see figure). This, incidentally, is the behaviour of a theoretical model of sand piles; it isn't clear that real sand piles behave exactly this way, although they do show something very like it.

So the sand pile is constantly undergoing scale-free fluctuations: it is in a critical state. But in this instance, the system is constantly returning to this critical state. After every avalanche, the constant addition of new grains returns the sand pile to the brink of a landslide. Instead of forever seeking to escape from the critical state, as a magnet does, the sand pile is constantly seeking to return to it. That's why this kind of behaviour is known as self-organized criticality (SOC).

Per Bak, a physicist who pioneered the study of SOC, was convinced that economic markets work in a state of self-organized criticality. He thought this because the market is constantly experiencing fluctuations that seem to be scale-free. 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 at least the

1960s that these fluctuations aren't random at all. That's to say, they have a non-gaussian probability distribution, which looks like this (see figure). 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. So if you try to make market forecasts based on the wrong statistical distribution, you can go badly astray.

To physicists, this power-law behaviour is a clue that the fluctuations arise from collective behaviour within the system: that is, to the interactions between market agents. There are now several theoretical models, inspired by physics, that reproduce this kind of probability distribution on the assumption that the agents' behaviour is influenced by that of their 'neighbours'. It turns out that, when the statistical properties of these fluctuations are studied very carefully, they don't exactly fit the model of self-organized criticality. But nevertheless, it provides a very general framework for understanding how scale-free fluctuations and power-law probability distributions (actually that's the same thing) can arise and be sustained in complex systems of interacting components.

Here, then are some of the key features of statistical physics that might prove useful when extending this science into the social sciences:

- inter-agent interactions and forces
- phase transitions
- collective and nonlinear behaviour
- critical points
- scale-free behaviour and power laws
- fluctuations and non-gaussian probability distributions
- self-organized criticality

And perhaps the most important consideration is that these phenomena are found in different physical systems that seem to share nothing in common when described at the level of the individual particles or components. The critical point of some magnets can be described mathematically in precisely the same way as that of a liquid/gas system. Phase transitions have universal characteristics. Self-organized criticality has been proposed in systems ranging from the mass extinctions of species in the geological record to the formation of solar flares to the statistics of earthquakes. The point is, then, that these are phenomena that don't depend at all on the specifics of a system – exactly what kind of forces exist between the constituent elements, how big or small these constituents are, what they are made from. Physicists don't *idly* speculate that such phenomena might also arise in social behaviour; they have the confidence to propose this because they have discovered how universal these things are in other systems, both non-living and living.

Applications in social sciences

Traffic flow

The flow of traffic along a road system can be modelled as a collection of particles all moving along a line – that is, in one dimension. They are not exactly like ordinary atoms moving through space, because they can speed up and slow down. But it isn't hard to prescribe precise rules by which they do this. People on an open road tend to accelerate until they reach the speed at which they are comfortable driving, and then they maintain that speed. If they come within a certain distance of a vehicle ahead, they will slow down to avoid collision. If necessary, this slowing will bring them to a standstill. These are relatively simple rules to implement in a computer model, and it turns out that they give rise to a kind of traffic fluid with properties somewhat reminiscent of the states of matter. That's to say, when the traffic is light, all the vehicles move freely according to their preferences: they barely interact. If the traffic gets denser, vehicles adjust their speed until they are moving more or less synchronously, all at much the same speed and with a roughly equal spacing between vehicles. And if the traffic is denser still, it forms a stationary jam. These three states are reminiscent of the gas, liquid and solid state, and they change from one to another abruptly, in what is essentially a kind of phase transition. A traffic model of this sort explains how 'phantom jams' occur, where there seems to be no visible cause, and they also predict observed phenomena such as travelling jams (where a jam moves steadily upstream of the flow) or stop-and-go oscillations, where a single perturbation to the flow can create a series of jams separated by intervals of moving traffic.

These models can be used to test out the effects of different driving regulations or road designs. For example, they might be used to test whether imposing speed limits on certain stretches of road might actually ease the flow of traffic by making the triggering of jams less likely; or where to position entries and exits on a motorway; or what the relative merits are of European and American lane rules (whether or not the lanes are used for a hierarchy of speeds). They can help to identify driving control measures that will reduce the chances of crashes.

So this is a simple yet perfect illustration of how a social physics might be used: as a test bed for exploring the consequences of structuring our rules or institutions one way or another. We have to decide for ourselves which of the various outcomes is the most desirable.

Alliance formation

Many of the choices we have to make are binary: we must do either one thing or the other. This makes us rather like the needles of magnetic atoms in a magnet, which can point either up or down. And indeed, physics-based models like this have been used to explore how decision-making is influenced by peer- and neighbour-pressure: recall how atomic magnets feel some obligation to align themselves with their neighbours.

This sort of bipartisanship is common in business and industry. It may arise, for example, in the setting of technical standards. Computer users have to decide whether to opt for a PC or a Mac system, and we might wonder under what conditions a minority product like Macs can persist indefinitely and when instead a market leader will inevitably come to

command the entire market. When two technical standards exist, manufacturers may be faced with the decision of backing one or the other, as they were in the early days of video with the choice between VHS and Betamax systems. 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.

This 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.

The evolution of technical standards for computer operating systems was a classic example. The Unix system, developed at Bell Labs and distributed at nominal cost, became the most popular. But by the 1980s there were around 250 different versions of Unix in use, each of them incompatible with the others. So there was an urgent need to standardize.

In 1987 Sun Microsystems and AT&T agreed that they would use the so-called Unix System V, and they formed an alliance which became Unix International Incorporated (UII). This forced seven of their rivals, including the Digital Equipment Corporation and IBM, to aggregate into the Open Software Foundation (OSF), which intended to develop a different standardized Unix operating system. So all other computer companies had to jump one way or the other: OSF or UII. Was there any way of a company predicting what others might do, and so making the best choice themselves?

The political scientist Robert Axelrod 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. Out of this push and pull emerge configurations in which the particle-like agents aggregate into alliances.

The force of attraction between two firms can be considered to be related to the size of the alliances in which they belong: the bigger the alliance to which firm A belongs, the greater the inducement for firm B to join. But the counteracting repulsion will depend on how much antipathy exists between the firms, which is likely to be related to the extent to which their products and markets overlap.

In the landscape model of alliance-formation developed by Axelrod and colleagues, each firm or ‘agent’ is therefore like a particle with an individually tailored force of interaction towards each other particle. The attractive component of the force that A exerts on B depends on how big A is (it makes sense, for example, for a tiny computer manufacturer

to align itself with a giant like Sun Microsystems). The repulsive force depends on whether B is a close or only a distant rival of A.

The model is like Maxwell's gas, except that each particle is unique and that there is typically only a handful of them. 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?

To find this, Axelrod and colleagues define a kind of 'total energy' for the group of firms, which is calculated by adding up all the forces of attraction and repulsion between each pair when the firms are aligned in various coalitions. Each possible configuration has an associated energy, and this defines an 'energy landscape' [pic]. 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.

There were nine principal US computer firms involved in the coalitions of the late 1980s, all of varying size and with different degrees of rivalry. There was no unique way of assigning the relative strength of the forces of repulsion between close and distant rivals, but these could be crudely chosen according to the degree of overlap between products. It turned out that the outcomes were not greatly affected by the precise strength of the forces.

There are 256 possible ways of dividing nine firms between two camps. Yet the calculations of the energy landscape showed that in general there were just two stable configurations—and the one with the lowest 'energy' was a very close match to the actual split between OSF and UII. In this configuration, only one company (IBM) was placed in the wrong camp ('UII' rather than 'OSF'). As expected, the two alliances had similar sizes in both of the stable configurations.

Since the probability of getting this close to the historical reality by pure chance is about one in fifteen, it looks as though the landscape model does a pretty good job.

One might imagine industrial decisions like this being made on the basis of all manner of long-term forecasts and cost-benefit analyses. But the landscape model invokes nothing of the sort. Instead, firms act on a decidedly myopic vision: in effect, they simply look at each competitor in turn and ask 'how do I feel about them?'

Axelrod has also applied the landscape model to the formation of alliances between seventeen nations in the approach to the Second World War. This is a more stringent test, since there are about 65,000 possible configurations (although admittedly most of them would be historically implausible) and it is far from obvious how to decide the degrees of attraction and repulsion between nations. Yet the model comes very close to predicting

exactly the right two camps – the Axis and Allied powers. There is just a 1 in 200 chance of getting a prediction this good simply by chance.

Firm formation

The question of what makes a firm successful is of course one of the most difficult and contentious in business. You can be sure that physics is not going to give you the answer. But I do suspect that some physics-based models have something useful to say about it.

The classical economic theory of the firm has great difficulty in dealing with the heterogeneous markets that exist in practice. It can deal with homogeneous markets under conditions of perfect competition, and it can deal with monopolies: under these conditions, theory can tell you how to maximize profits. By using game theory, it can make some headway in describing oligopolies. But none of these situations really describes how most markets are structured in terms of firm size, profitability or objectives. A large part of the problem is that there is still no good understanding of how the market gets to be structured that way. How do firms grow, and what limits this growth?

The first theory worthy of the name that attempted to answer these questions was formulated by Robert Gibrat in 1931. He decided, very sensibly, to start by looking at the statistics, which showed very evidently that there are many more small firms than large ones. In other words, it is a skew distribution. This distribution in firm sizes remains more or less invariant over time, in different countries and in different industrial sectors, and has been called “perhaps the most robust statistical regularity in all the social sciences.”

Gibrat proposed that a firm grows at a random rate, amplified by the existing size of the firm. This defines Gibrat’s now-celebrated Law of Proportionate Growth, which he envisaged as a kind of Newtonian law of the business world. It means the following. To predict how much a firm changes in size between now and some future time, one plucks a number at random from between -1 and 1 (it could be, say, 0.5 , or -0.3528 , or zero) and multiplies that by the current size of the firm. Thus bigger firms tend to change in size more markedly than small ones; but not inevitably so. There is an element of chance, since the various factors that influence growth are hard to predict. To put it another way, the bigger a firm is, the more able it is to capitalize on whatever opportunities come its way.

Gibrat’s ‘law’ leads to a size distribution with a mathematical form called ‘log normal’. (This means that the firm’s size fluctuates over time in such a way that the probability distribution of the *logarithm* of the size is gaussian.) Gibrat claimed that all kinds of data on firm sizes fitted his model rather well.

Nevertheless, the model is basically wrong—and everyone knows it. Whatever Gibrat’s model says about the distribution of firm sizes, it is apparent simply from observing how firms grow that they do not expand and shrink in randomly sized steps. Certainly, this is

hardly consistent with the idea of firms rationally maximizing their profits (which implies that firms should respond similarly, rather than independently at random, to changes in market conditions).

Moreover, the model is arbitrary. Gibrat's Law of Proportionate Growth is plucked out of a hat and is not backed up by microeconomic principles. Yet the standard economic theory of the firm can't offer anything better in its place.

And it now seems that Gibrat's model doesn't actually fit the data well after all. In 1996 physicists and economists at Boston University looked at the growth rates of all publicly traded US manufacturing companies between 1975 and 1991, encompassing around 8,000 firms. They found that the growth rates did not fit Gibrat's log normal distribution, but instead followed a power-law relationship (see figure).

So there does indeed seem to be a general mathematical law of firm growth – but a different one from that proposed by Gibrat. The question is: can we understand it?

The statistics identified by Stanley and coworkers have been supplemented by an even more extensive survey of 20 million US firms by Robert Axtell of the Brookings Institute in Washington DC, who found that firm sizes – not growth rates but absolute sizes – also follow a power law (see figure). And Axtell has offered an explanation of these distributions in terms of a kind of 'microeconomic' model of firm growth in which firms arise by the aggregation of many interacting agents each following their own agendas. Axtell's model is in the spirit of microeconomics in that it attempts to deduce the overall behaviour of a system from the motivations of the individual agents that constitute it. But unlike many such 'theories of the firm', it starts with no preconceptions about what firms are for or how they behave; indeed, the agents are not forced to aggregate into firms at all.

These agents are utility maximizers, each with a personal utility defined by a balance of two conflicting demands: money and leisure. The relative preferences for money and leisure vary throughout the agent population, and each agent tries to find the job that will allow it the preferred compromise. There is a mathematical equation in the model that relates an agent's efforts to the productivity of the group in which it belongs, and it is this that provides the incentive for forming firms.

This is not, however, quite the same as the standard idea that companies form and grow because of increasing returns of scale. An increasing return of scale is possible but not guaranteed. All the rule says is that it is in the interests of each agent to team up with others, no matter whether they are a workaholic or a slacker.

The agents join and leave firms in search of an optimal utility. Thus firms can grow and decline as agents come and go. The key point is that each agent has a choice about how hard to work, but that decision varies depends on the agent's situation. In a big firm, an agent can be a slacker and reap the rewards of other agents' efforts. On his own, on the other hand, he must do at least a little work, or face starvation.

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 (see figure). This means that its predictions are necessarily statistical. The probability distribution of firm sizes, for example, is a power law – just what is observed in practice. The model also generates the power-law distribution of growth rates. And these laws remain robust in the face of changes in the precise details of how agents interact and make decisions, so long as they remain utility maximizers.

Most of the firms that emerge in this model are ephemeral – as indeed they are in the real world. Of the largest 5,000 US firms operating in 1982, for example, only 35 percent existed as independent entities in 1996. There is a high ‘turnover’ of companies, which many economic theories of the firm do not acknowledge.

Why do firms fail? Axtell’s model shows that there is a typical trajectory. First, a new firm grows more or less exponentially in time as the increasing returns cause workers to flock to it. But at some point the firm reaches its peak, after which collapse is usually sudden and catastrophic. Reduced to a tiny fraction of its former size, the firm struggles on for a short time with a handful of determined workers before eventually vanishing (see figure).

This collapse is a consequence of a firm’s own success. Once it grows big enough, it becomes a haven for free-riders who capitalize on the efforts of others. So the firm becomes gradually riddled with slackers, until suddenly the other workers decide they have had enough and jump ship. Tellingly, just before a firm collapses, the average effort of its workers plummets to zero.

If we believe that such a simplified model can tell us anything at all about the real world, then we learn some revealing things about firms. First, they are not maximizers. Firms as a whole maximize neither profit nor overall utility (as conventional theories would have us believe). Individual agents *do* try to maximize their utilities, but this does not induce such behaviour in the group as a whole.

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.

Marriage

Finally, I want to look at what this kind of physical modelling can tell us about an institution not renowned for a pervasive influence of rationality: marriage.

George Bernard Shaw once said that ‘The common notion that the existing forms of marriage are not political contrivances, but sacred obligations... influences, or is believed to influence, so many votes, that no Government will touch the marriage question if it can possibly help it.’ Well, times change. Today some governments, alarmed at the declining

figures and, in supposed consequence, a diminution of ‘family values’, feel a duty to try to engineer a revitalization of marriage. But how do you do that? One can always pay people to marry, through the machinery of tax incentives. One can attempt to create a social climate that smiles upon families welded by marriage. One might even tamper with employment conditions to make families easier to rear, in the hope that raising a family increases the likelihood of marriage.

But in the end one will not get far without some understanding of *why* people choose to get married. Gary Becker brought this seemingly very personal question into the gamut of economic analysis in the late 1970s. He proposed that a cohabiting couple gains in efficiency over single people by specialization, just as in Adam Smith’s famous pin factory. One partner does the domestic duties, the other goes out to earn the wage. In this way the couple *maximizes their utility* like rational market traders.

Becker’s work, which helped to win him a Nobel prize, is very much in the spirit of a ‘physics of society’ in asserting that, to understand the reasons why things are the way they are, we are ill advised to rely on intuition and preconception. Rather, we should seek for models that illustrate how certain circumstances can arise if certain rules are followed.

The great value of Becker’s work is to identify the fallacy in conventional ‘neoclassical’ economic modelling that regards birth and marriage rates largely as ‘given’. Economists have not traditionally been interested in these things other than as background features of the economic landscape. In contrast, Becker has shown, economics both affects and is affected by these social factors.

On the other hand, his analysis is itself ‘neoclassical’ in its foundation, being based on rational agents independently maximizing their utility. This approach can only ever tell a partial truth, because again it neglects the crucial factor: interaction.

Paul Ormerod has addressed this shortcoming in a model of marriage demographics based on interacting agents. The model assumes that individuals can be regarded as more likely to make a certain choice if their fellow citizens do the same. One can then investigate how the various possible driving forces for marriage alter the proportions in an interactive society while keeping other factors constant.

In Ormerod’s model the population is divided into three groups: single, married and divorced. The state of being single is rather like virginity: once you’ve left it, there is no going back. But agents can switch back and forth between marriage and divorce as many times as they like. Two generalized factors are assumed to influence these choices: economic incentives (be they wage-earning potential, tax breaks, job opportunities or whatever) and social attitudes (public disapproval of unmarried cohabitation, unfashionableness of marriage).

This is what the model predicts, for example, about the effect of economic incentives, keeping other factors constant (see figure). We get a looped curve that provides two possible states of the system: a high and a low proportion of marriage in the population.

In other words, for a given set of socioeconomic conditions, two outcomes are possible. In fact, the two branches are linked by a continuous curve (see figure), which seems to imply that there are sometimes *three* possible states. However, one can show that beyond the ‘turning points’ of the upper and lower curves, the states represented by the dotted curve are unstable—they transform instantly into something else. This is exactly what emerges from van der Waals’s theory of the liquid-gas phase transition: a looped curve allowing for two possible stable states of the system.

This loop in the curve appears only if the strength of social attitudes is large enough. A three-dimensional graph which shows the dependence of marriage on both this and economic factors therefore looks like this (see figure). The possible states of the system lie on a *surface* which develops a kink.

This too is a familiar image to statistical physicists. It is the curve that shows the first-order liquid-gas transition vanishing into the critical point. The critical point itself corresponds to the place where the kink first starts (or ceases) to appear—where the upper surface is sufficiently twisted that it overhangs the lower surface. On this surface, ‘strength of social attitudes’ is replaced by temperature, ‘economic incentives’ by pressure, and ‘proportion of married population’ by density.

In other words, Ormerod’s marriage model has a critical point. At this point the distinction between a low-marriage and a high-marriage society disappears. So these interacting agents, influenced by the choices one another makes, display the whole gamut of behaviours that characterize the particles of a fluid, influenced by their mutual forces of attraction and repulsion.

This way of looking at apparently volitional human behaviour may seem strange. You might be inclined to agree with the implied scepticism of the writer Frederick Hunt in an 1850 issue of Charles Dickens’s periodical *Household Words*, commenting on the new fad for social statistics:

the savants are superseding the astrologers of old days, and the gipsies and wise women of modern ones, by finding out and revealing the hitherto hidden laws which rule that charming mystery of mysteries—that lode star of young maidens and gay bachelors—matrimony.

When wilful human deeds such as marriage and crime first entered the roster of social phenomena governed by statistical laws and regularities, the response was a mixture of amazement, delight and dismay. We might anticipate the same thing as physics makes its mark in social science. But I think it is worth remembering the words of William Newmarch, speaking to the Statistical Society of London in 1860, who pointed out that if social and political policy is to be just and effective, it must draw on something more than preconception and intuition:

The rain and the sun have long passed from 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.

This talk is based on material taken from Philip Ball's forthcoming book, which will be published by Heinemann/Farrar, Straus & Giroux in 2004.