Positive Linking

How Networks Can Revolutionise the World

Contents

Preface

My previous book, *Why Most Things Fail*, was published in the mid-2000s. I addressed what is probably *the* most fundamental feature of both biological and human social and economic systems. Species fail and become extinct, brands fail, companies fail, public policies fail. Despite the rather gloomy title, the book did well. It was a *Business Week* US Business Book of the Year.

The book built on general themes which I had already explored in my two previous books. In *The Death of Economics* in the mid-1990s I argued that conventional economics views the economy and society as machines, whose behaviour, no matter how complicated, is ultimately predictable and controllable. The financial crisis of the late 2000s showed only too clearly how deeply flawed is this view of the world, embraced enthusiastically by mainstream economists, international bodies such as the International Monetary Fund, central banks and politicians the world over. On the contrary, the economy is much more like a living organism.

In *Butterfly Economics* in the late 1990s, I developed this theme. I analysed a wide and seemingly disparate range of economic and social questions, seeing them as analogous to living creatures whose behaviour can be understood only by looking at the complex interactions of their individual parts.

Why Most Things Fail drew inspiration from the biological sciences even more, demonstrating close parallels between, for example, the extinction of biological species in the fossil record and the extinction of companies. The idea that economics should look to biology for intellectual inspiration is a long and distinguished one. Alfred Marshall, who founded the faculty of economics at Cambridge University around 1900, was the first major scholar to articulate this view.

Vernon Smith, in his economics Nobel Prize lecture in 2003, stated bluntly: ‘I urge students to read narrowly within economics, but widely in science. Within economics there is essentially only one model to be adapted to

every application . . . The economic literature is not the best place to find new inspiration beyond these traditional technical methods of modelling.' I have followed this precept in this book. In addition to biology, I draw on powerful insights from, amongst other disciplines, psychology and anthropology.

Both *Butterfly Economics* and *Why Most Things Fail* were fundamentally based on the concept of networks, the idea that individuals do not operate, as conventional economics assumes, in isolation, but are connected together in society. Both the theory and practice of networks is a rapidly developing area, and I make use of results at the forefront of knowledge in this book.

Viewed from a network perspective, many aspects of our social and economic world look completely different than they do from the conventional view of mainstream economics. The world is not by any means a machine whose behaviour is predictable and controllable by pulling a lever here, by pressing a button there. The individual components – people, firms, regulators, governments – interact with each other and each component has the capacity to change directly how other components behave.

In many ways, this makes successful policy making, whether in the public or private sectors – much harder. So much is contingent on who influences whom on a network and when. Simple causal relationships between a change in policy and any given outcome no longer exist – if they ever did!

At the same time, far more effective policy making becomes possible. It requires a fundamental change of mindset by policy makers. This does not mean no government, but it certainly means much more thoughtful government instead of the complacent, tick box mentality which currently dominates the public sector in the West. The positive aspects of the huge recent increase in knowledge about social and economic networks open up new possibilities for solving many long-standing problems. A lighter, smarter touch, one which exploits the positive linking aspects of our modern, networked world.

I am grateful to a large number of people for encouragement and discussions which have helped to develop the ideas of this book, and in particular to Alex Bentley, Greg Fisher and Bridget Rosewell. Julian Loose of Faber and Faber has once again proved to be a very helpful and inspirational editor.

Paul Ormerod

London, Wiltshire and Red Lumb, October 2011

Introduction

Modern economic theory was first set out on a formal basis in the late nineteenth century. There have certainly been developments since then, but at heart the basic view in economics of how the world operates remains the same. Economics is essentially a theory of how decisions are made by individuals, of what information is gathered and how it is used by the decision maker.

All scientific theories, even quantum physics, are approximations to reality. Theories involve making assumptions, simplifications, to enable us to understand problems better. A key feature of a good theory is that its assumptions are a reasonable description of the real world.

In the early twenty-first century, just as it did in the late nineteenth, economics in general makes the assumption that individuals operate autonomously, isolated from the direct influences of others. A person has a fixed set of tastes and preferences. When choosing amongst a set of alternatives, he or she compares the attributes of these alternatives and selects the one which most closely corresponds to his or her preferences.

At first sight, this may seem quite reasonable, indeed even 'rational', as economists choose to describe this theory of behaviour. But there is a serious problem with the assumption that individuals operate in isolation from each other, that their preferences are not affected directly by the decisions of others. The social and economic worlds of the twenty-first century are simply not like this at all. We are far more aware than ever before of the choices, decisions, behaviours and opinions of other people. In 1900, not much more than 10 per cent of the world's population lived in cities. Now, for the first time in human history, more than half of us live in cities, in close, everyday proximity to large numbers of other people. In the last decade or so, the internet has revolutionised communications in a manner not experienced since the invention of the printing press in the mid-fifteenth century.

The assumption that people make choices in isolation, that they do not adopt different tastes or opinions simply because other people have them, is no longer sustainable. Perhaps – perhaps, and it is a big 'perhaps' – over a hundred years ago this might not have been a bad assumption to make. But no longer.

The choices people make, their attitudes, their opinions, are influenced directly by other people. The medium via which this influence spreads is the social network. Often, social networks are thought of as purely a web-based phenomenon: sites such as Facebook. These can indeed influence behaviour. But it is real-life social networks – family, friends, colleagues – that are even more important in helping us shape our preferences and beliefs, what we like and what we do not like.

Network effects, the fact that a person can and often does decide to change his or her preferences simply on the basis of what others do, pervade the modern world. Throughout history, a crucial feature of human behaviour has been our propensity to copy or imitate the behaviours, choices, opinions of others. We can see it in the fashions in pottery in the Middle Eastern Hittite Empire of three and a half millennia ago. And we can see it today in the behaviour of traders on financial markets, where the propensity to follow the herd can lead all too easily to the booms and crashes we have lately experienced. Scientists such as Robin Dunbar have argued that our anomalously large brain (compared to other mammals) evolved precisely because, from an evolutionary perspective, copying is a very successful strategy to follow.

This concept is just as crucial for companies and markets as it is for people. In September 2008 Lehman Brothers went bankrupt, precipitating a crisis which almost led to a total collapse of the world economy and a repeat of the Great Depression of the 1930s. It was precisely because Lehman was connected via a network to other banks that made the situation so serious. Lehman's failure could easily have led to a cascade of bankruptcies across the world financial network, first in those institutions to which Lehman owed money, then spreading wider and wider from these across the entire network. Incredibly, neither the systems of financial

regulations which were in place, nor the thinking of mainstream economics which influenced policy so strongly, took any account of the possibility of such a network effect.

A world in which network effects are a driving force of behaviour is completely different from the world of conventional economics, in which isolated individuals carefully weigh up the costs and benefits of any particular course of action. A world in which network effects are important is a much more realistic description of the human social and economic realities which exist in the twenty-first century. It is the implications of this world which I explore in this book.

Incentives have not disappeared as a driver of human behaviour. It is still the case that if, say, Pepsi raises its price compared to Coke, more Coke and less Pepsi will be sold. This is the world which economic theory describes. It is not wrong. But it is often misleading, for it offers only a very partial account of how decisions are made in reality. Network effects can be far more powerful than incentives, and we will see many examples in which network effects have completely swamped the impact of incentives, leading to outcomes completely different from those intended by policy makers.

Network effects require policy makers, whether in the public or corporate spheres, to change radically their view of how the world operates. In part, they make policy much harder to implement successfully, and they help explain many of the failures of policies based on the assumption that incentives and not network effects are the key drivers of behaviour. But they open up the possibility of much more effective and successful policies, ones which harness our knowledge of network effects and how they work in practice. Hence the main title of this book: *Positive Linking.*

Chapter 1

Unintended Consequences

On Wednesday 16 October 1555, two of the leading members of the reformed English Protestant Church, Hugh Latimer and Nicholas Ridley, were chained to a stake in the city of Oxford. They were then burned to death. By what amounted to a series of historical accidents a Catholic, Mary, had become Queen, the ruler of all England, scarcely two decades after the Church of England made its historic break with the papacy. She was attempting to impose Catholicism by a policy of publicly burning leading Protestants. If they renounced their faith, their lives would be spared and they might even continue to enjoy the power and trappings of high office. If not, they faced the fire.

But far from quailing at a terrible fate, Latimer and Ridley embraced it cheerfully. 'Be of good comfort, Master Ridley, and play the man; we shall this day light such a candle, by God's grace, in England, as I trust shall never be put out,' Latimer allegedly pronounced. They, along with other condemned Protestants, had formed a deliberate policy of facing death with equanimity, in order to make a positive impression on those who witnessed the burnings.

They believed that the story of their end would spread by word of mouth far beyond those present at the executions. Existing Protestants might be encouraged by their example to be steadfast in their faith, and new converts gained. And on this occasion, the martyrs were ultimately proved to be correct. On Mary's death, Protestantism was restored as England's religion.

Flash forward over 400 years to another event in English history, far less momentous, but one which offers a vignette of popular culture, not of the mid-sixteenth century but the late twentieth.

In Sardinia during the 1990 soccer World Cup, the English supporters were feared for their violent reputation. One evening in Cagliari, a large number gathered in the streets. Facing them were the police. As Bill Buford relates in his excellent book *Among the Thugs*, various individuals made attempts to stir the fans into collective action without success. Making himself conspicuous so that others could see his actions, one threw a metal object at the police. Another charged the police and yelled for others to follow. Further attempts were made by isolated fans to encourage the crowd into collective action, but none joined in.

Tiring of the whole situation, in response to the actions of one particular youth, a police captain fired his pistol into the air in a signal for the potential mob to disperse. The reaction was unexpected. At the sound of live ammunition being discharged, the English supporters immediately began to destroy property and attack the police. The very action intended to subdue the fans into quiescence provoked exactly the opposite reaction. The individual supporters suddenly turned into a mob.

These two stories, disparate though they may seem, have a great deal in common. They illustrate the seemingly perverse and apparently irrational ways in which people can behave. Rather than sullenly dispersing back to the safety of their hotels or into bars when a firearm was discharged, the soccer fans ran at the police. Latimer and Ridley were offered the choice not just of their lives but their freedom if they embraced Catholicism. Rather than meekly agreeing, whatever private reservations might have remained, they chose to suffer an appalling death. Not just history but contemporary life is replete with examples of people behaving in seemingly inexplicable ways.

A key theme of this book is that these widespread forms of behaviour are explicable. They are illustrations of the power of social networks. Today, the phrase ‘social networks’ is often synonymous with networking across the web on sites such as Facebook. But this is just one, albeit new and important, aspect of a phenomenon that has existed for centuries. People do not live in isolation, but in society. Their lives are filled with interactions

across social networks. The network of their families, the network of their work colleagues, the networks of their hobbies. Real-life social networks in which people meet, gossip, chat, argue. Networks in which people's choices, behaviour, opinions can be influenced, shaped, even altered dramatically by the process of social exchange with other people.

Within these social networks, people often copy or imitate what others do or think, for a variety of motives. An individual might have formed a private view on a matter, but might believe that others with a different opinion are better informed and so change his or her mind as a result. Or someone may accept the behaviour of a particular social group simply from a desire to conform. More subtly, peer acceptance might give an individual permission to behave in a way that, in a different social context, would be unacceptable.

The idea that copying is an important aspect of behaviour does not mean that individuals operate as automatons, that they have surrendered control over their decision to others. People can copy and still retain a clear sense of agency, of purpose and intent over their own actions. So in a strange city, you may consciously decide to copy others, to go to the restaurant where there are lots of customers rather than to the one in the same street where there are few. Lacking any other reliable information, lacking local knowledge yourself, you decide to be influenced by the choices made by others. Even in the highly connected world of the twenty-first century, networks are not everything. People still retain their individuality, their capacity to decide actions and beliefs for themselves, despite what is popular, either in society as a whole or amongst their particular group of friends, family or work colleagues.

Most public policy on social and economic matters is based on the premise that people, or indeed companies, behave as individuals when they are making decisions. Like so many Robinson Crusoes, people exist in splendid isolation. And it is this view of the world that is epitomised by mainstream economic theory.

We explore in this book the connection between the impact of incentives, of the assessment of costs and benefits of different actions, on individuals, and the effect of social interaction across networks. When the power of the network takes over, people are no longer acting autonomously, but as part of a social group, and their behaviour and decisions are driven by the process of copying, of imitation.

Sometimes the initial impact of changes to incentives on the behaviour and decisions of a few individuals will be seen to be enhanced as the power of the social network takes over, and this effect can on occasion be dramatic. But, equally, there are times when the impact of copying behaviour across a social network, of imitating the behaviour of others, does not just offset the effect of incentives, but takes the system in the completely opposite direction to what was intended.

In recent decades, the discipline of economics has exhibited powerful tendencies of intellectual imperialism. Not content merely to analyse the familiar areas of firms, consumers, prices and markets, economists have turned their attention to a wide range of social issues, seemingly far away from the original scope of economics: the study of the allocation of scarce resources. The institution of marriage, crime, piracy, drug addiction – economists now focus on all of these and more.

Indeed, the two historical vignettes which opened this chapter can be translated into the context of economic theory. The popular image of economics is that it deals with ‘big’ things, national output (GDP), unemployment, inflation, interest rates. This is macroeconomics and these are the topics which appear in the newspapers and on our television screens and on which economists are regularly seen to pronounce.

But, in essence, economics is a theory about how individuals make decisions. About decisions made at the microeconomic level. The measure of the relative importance of microeconomics is indicated by the fact that, over the past two decades, the Nobel Prize in economics has been awarded for work which has been

unequivocally ‘macro’ in character on only four occasions. Not all the others have gone to micro, for some of the Laureates have made advances in techniques of statistical analysis, but micro distinctly outweighs macro in these awards.

At the core of microeconomics is a series of theoretical postulates about how the so-called ‘rational’ individual makes decisions. For example he or she has a well-defined and fixed set of preferences concerning the choices on offer. He or she gathers all available information when making a decision, matches it against his or her preferences, and then makes the best possible decision – the ‘optimal’ decision, as economists like to say to give it a more scientific air – given the information and the preferences.*

* I have discussed this model in more detail in previous books, such as *The Death of Economics*, a critique of free-market economic theory written in the mid-1990s.

So the agent – the jargon phrase in economics for the person making the decision – may, if the products have the same price, prefer Pepsi to Coke. (It is in fact rather useful to use ‘agent’, rather than ‘he or she’, since the word subsumes the two genders and avoids having to repeat the two.) But if the price of Pepsi rises relative to Coke, at some point any given agent will switch and buy Coke instead. This is not because the agent’s preferences have fundamentally altered, but because the money saved in buying Coke rather than Pepsi in this illustrative example can be used to buy more of other products. So, overall, the preferences of the agent might be more closely matched by switching to Coke.

Another example is the British savoury spread Marmite. The only other countries where I believe either it or a close variant are on general sale are Britain’s closest cultural neighbours, Australia and New Zealand. Based as it is on the scrapings of the fermented residue at the bottom of beer barrels, agents’ preferences on this tend to be sharply divided. I cannot abide it. My wife adores it. But there is some price at which I could be persuaded to eat it, probably a negative one in which the producer paid me rather than the other way round.

When making choices between fairly straightforward, inexpensive, well-established consumer products, the economist's view of 'rational' choice may be reasonable. Coke, Pepsi, Marmite have all been around for a long time and agents have formed their preferences. They are unlikely to suddenly alter them, at least in any appreciable numbers.

In making a decision about the choices on offer, only a relatively small amount of information needs to be gathered, mainly concerning price. This latter factor might not be completely obvious, because the price per unit of weight or volume may vary from store to store, by pack size, or because of special offers such as 'buy one, get one free' or 'three for the price of two'. The mathematical capabilities of many people are known to be low. For example on the very day I write these words a TV advert to recruit teachers created by the Training and Development Agency for Schools has been exposed by a fifteen-year-old schoolboy as containing the wrong answer to a fairly straightforward question.* But even making allowances for this, there is a limited amount of information for agents to gather before matching it to their preferences.

* The question is: what is 'x' in the expression $(g^2)^7 = g^x$?

This model of rational behaviour is no longer relevant in many circumstances in the world of the twenty-first century. Agents face a vast proliferation of choice, massive information overload. Many of the products on offer are highly sophisticated, difficult to evaluate in terms of their attributes. And we live in a world which is far more connected, in which we are far more aware of the opinions and behaviour of others, than we were a hundred years ago when standard economic theory was first being formalised. In 1900, the clear majority of the population of the world lived in relatively isolated villages. In the twenty-first century, the majority lives in cities, in close proximity to large numbers of other people. And the revolution in communication technology brought about by the internet makes us dramatically more aware of the behaviour of others than at any time in the whole of human history.

We need a new model of rational behaviour, one which is empirically consistent with the real world, the world of the twenty-first century. The economist's definition of rational behaviour is only one possible way to define the concept of rationality. Behaviour which does not follow the precepts of economic rationality is *not* irrational, as economists would have us believe. Indeed, in the modern world in many contexts it is the economic definition of rationality which has become irrational!

The development of such a view of the world, a more realistic view of how agents actually behave in the social and economic contexts of the twenty-first century, is a main theme of the book. It has radical implications for the conduct of policy, both corporate and public. Potentially, its impact is very positive. Our knowledge of how networks influence behaviour in the social and economic worlds is growing rapidly, both theoretically and empirically. The opportunity both to exploit this knowledge and to develop it even further, for there is much still to be done, over the coming decades is enormous. Successful policy making in the highly connected, networks world of the 21st century will be impossible without understanding positive linking, how we can use what are often abstract and difficult concepts to help shape a better world.

*

How, then, might we think about the two historical episodes described above in the context of standard economic theory? The theory makes claim to be a general description of human behaviour, a general theory of how people make choices. These examples may seem outside the conventional areas of economics, but if a theory is claimed to be general, it ought to be able to illuminate these events. Besides, it is economists themselves who have pushed the theory into areas such as marriage and crime and claimed that it has strong explanatory power in what might more usually be thought of as social rather than economic settings. The Chicago economist Gary Becker received the Nobel Prize for exactly this kind of work.

The individual preferences of the soccer fans were to have some sort of riot in which property would be vandalised and innocent passers-by made to cower in fear or, even better, injured in some way or other. This is why they had assembled as they did. A message had been passed to meet at a particular time, six o'clock as it happens, in a particular square. The colloquial phrase used was 'it's going to go off', meaning that, for those interested, the gathering would offer an opportunity to participate in creating mayhem in the city of Cagliari.

The fans would derive 'utility', again using the jargon of economics, from rioting. But what were the related costs to set against these 'benefits', using the phrase in inverted commas on this occasion to emphasise that these were, of course, benefits only to those involved in trashing the city, not to those unfortunates on the receiving end. The most obvious cost was the phalanx of police standing in front of them. Heavily equipped with helmets, shields, truncheons and guns, they were clearly capable of inflicting costs, such as a beating or arrest and prison, on anyone foolish enough to provoke them.

Economic theory usually allows individuals to differ in their preferences. Incredibly, as we shall see later in the book, the trends in macro theory in recent decades have been to suppress this, trying to explain the economy as a whole in terms of a single 'representative agent'. But more of this later. For now, we remain firmly in the terrain of microeconomics, where agents can have different preferences. The youths gathered in the square would undoubtedly differ in the benefit each individual believed he would gain from having a riot compared to carrying out other activities, such as having a beer or reading a book on Einstein's theory of general relativity or Shakespearean sonnets. They would differ in the evaluation of the costs of any police action inflicted on them. And, according to standard economic theory, they would even be allowed to differ in their assessment of the probability of being the recipient of such action themselves (with the strict proviso, and I am not making this up, that over the course of a series of such riot events, each fan on average assesses the probability correctly).

On this view of the world, every single one of the fans who responded to the verbal message 'it's going to go off' derived utility from participating in hooliganism. And on arriving in the square and seeing the police, at some point they all formed the view that these benefits to them outweighed the likely costs. It is possible that a few crept discreetly away, having come to the opposite view, but Bill Buford's description of the events certainly suggests that almost all the English fans present participated in the subsequent vandalism and general criminal behaviour which occurred.

But this does not take us very far in understanding why the riot started when it did. The fans stand, confronting the police on a hot late afternoon. But at first, they are a collection of individuals and not a collective mob. They have the potential to become a mob, but nothing happens. Several fans try to incentivise them all to start behaving badly by carrying out prominent acts of bravery, or lunacy as most people would see it, against the police. We can in fact readily understand the behaviour of these particular individuals from the point of view of economic theory. Considerable status would be attached to being seen as a leader by the other fans, being seen as a Top Boy, to use the British colloquial expression for the leader of a gang of thugs or hooligans. The benefit from this would be perceived as outweighing the undoubted increase in the potential cost to the individual by identifying himself so prominently to the forces of law and order.

This whole rationale for the event, as described by conventional economic theory, may already seem somewhat convoluted, but it now becomes even more so. Why did the fans as a whole not respond to the actions of the individuals who deliberately tried to incite a riot by their provocative actions? According this theory of how agents behave, we have to suppose that the responses to these by the police were such as to temporarily tip the balance between costs and benefits in the minds of all the would-be rioters. In other words, when a youth came forward and threw a metal object at the police, they perhaps brandished their truncheons more fiercely, and that such a signal increased the likely costs of a riot in the minds of the fans.

But then the police captain, tiring of confronting this unpleasant group of badly dressed, smelly individuals,* fired live ammunition into the air. He clearly believed that signalling this potential cost – the possibility they would be fired upon – would be a sufficiently large incentive to make them eschew the pleasures of a riot on this occasion. The cost of being the recipient of a bullet surely outweighs that of even a savage truncheoning, but the response of the fans suggests otherwise.

* *Among the Thugs* graphically illustrates these qualities in a variety of contexts.

The fans immediately charged the police. A possible reconciliation with the core model of individual behaviour in economic theory is that the shot fired into the air was a sign of weakness on the part of the police, a sign that they would not actually open fire on the English, regardless of what they did. But this argument is now getting pretty tenuous. Ex post, economic theory can rationalise almost anything which has ever happened, but these attempts often amount to no more than a Just So story, as is certainly the case here. Their credibility gets stretched well beyond breaking point.

A much simpler explanation can be given in terms of networks. When the fans first gathered in the square, it is not implausible to interpret their behaviour in terms of individualistic economic theory. Each of them enjoyed a riot, they gained 'utility' from it, in the jargon of economics. As noted above, the strength of their individual preferences for participating in a riot compared to other activities undoubtedly varied, as did their assessments of the costs. The delay between the fans assembling and the riot starting, and the lack of collective response to the efforts of reckless individuals to incite them, suggest that almost all of the supporters had formed the view, given the serried ranks of the forces of law and order confronting them, that the potential costs involved outweighed the benefits.

A shot was fired. The collection of individuals immediately became a mob. They lost their individual identities. And their preferences altered dramatically, so that when the gun went off they charged the police, acting as a

single unit. They had arrived as individuals as part of a social network with a shared interest in hooliganism. Information about a potential outlet for this activity had been passed across this network. But the action of the unfortunate police captain altered qualitatively the structure of this network. The individuals became fused as one, with an overwhelming preference to riot almost regardless of costs to themselves as individuals.

What scientists call a 'phase transition' had taken place. When water is gradually cooled, it remains water as the temperature drops from ten to nine to eight degrees and so on. Then, suddenly, as it passes through zero, a phase transition occurs. Water becomes ice.

A simple example of this phenomenon which is almost certainly more familiar to most readers of this book than taking part in a public riot is a social gathering with friends, in a bar or perhaps at a party. Each individual present enjoys alcohol in moderation and dislikes hangovers. But the company is delightful, the wine flows. The collective mood temporarily overcomes the preferences of individuals. And the effect of the social network present in this particular milieu is that almost everyone is induced to drink more than he or she intended at the start of the evening, or would drink on their own or in a smaller or less congenial group. The next day, operating once more as individuals with their individual preferences restored, some will undoubtedly regret the collective set of preferences – to consume yet more alcohol – which spread across the social network. One or two may go so far as to pledge to themselves never to drink again. (Until the next time, of course!)

James Surowiecki wrote a very interesting book in 2004, *The Wisdom of Crowds*. This is essentially about the process of answering a question by taking into account the opinions of a large number of individuals rather than relying on just a few, no matter how expert these people might be. There are many practical examples where the 'crowd' certainly gives a more accurate estimate, such as the classic ones of guessing the number of sweets in a jar or the weight of a prize bull at a country fair. The word 'crowd' is put in inverted commas here, because the

process only really works when the individuals participating in the process remain as individuals and not part of a crowd in the way the soccer vandals were.

The crucial assumption needed for the average of the collection of individual opinions to be more accurate than a single expert is that they do indeed form their views independently, without reference to those of others. Once this independence vanishes, once the agents become fused into a single whole, often the outcome is not so much the wisdom, more the 'madness of crowds', as it was described by Charles Mackay as long ago as 1841.

*

There is another very general point, another key theme of this book, to take from the example of the English soccer fans and their rampage through the streets of Cagliari. The individuals received several attempts to incite them. But none of these worked. Then, completely unexpectedly, the one event which any detached observer might think would offer a clear deterrent, the firing of the gun, turned the group of individuals into an enraged crowd.

And this is the point. Most attempts to spread a choice, an opinion, a type of behaviour, across a network of individuals fail. The events in Sardinia are simply an example of this general point. And it is why we have to be very careful when designing public policy. Like the police captain, policy makers will often have little idea about the likely consequences of their attempts at nudging groups towards particular decisions or opinions. Duncan Watts, formerly Professor of Mathematical Sociology at Columbia, now director of the Human Social Dynamics group at Yahoo! and someone we will meet in much more detail later, coined a phrase for this fundamental property of networks. They are 'robust yet fragile'.

The collection of individuals who make up a network will, most of the time, exhibit stability with respect to most of the 'shocks' which this particular system receives. The shock could be a piece of news in the context of financial markets, an advertising campaign in a consumer market, or, as here, attempts to incite a group of fans

to alter their preferences and attack the police. The system is stable in the sense that most shocks make very little difference, they are absorbed, shrugged off, and few people change either their minds or their behaviour as a result. So the network is 'robust'.

But, every so often, a particular shock may have a dramatic effect. So the network is also 'fragile'. The behaviour of individuals across the whole, or almost the whole, of any particular network might be altered. Before the event, it can be very difficult if not impossible to discover what the eventual impact is going to be. A big shock, almost by definition, will have big consequences. So if the Italian police had opened up with machine-gun fire directly at the fans, we could reasonably conclude that in this particular instance a riot would not have taken place. But most events, most attempts to change behaviour, do not fall into this category. Most have very little impact. But occasionally, one does.

From these sordid events in Sardinia, we can now return to the altogether more dramatic happenings in Oxford over 450 years ago. Again, we can offer a partial explanation in terms of incentives, of costs and benefits, when people are acting as if they were isolated individuals. But, again, this kind of rationale soon becomes incomplete. Networks are needed to complete the picture.

Since time immemorial religion has been, and continues to be outside Western Europe, a major presence in human society. Yet mainstream economics has virtually ignored the topic, certainly in comparison to the enormous amount of work carried out in sociology, anthropology, psychology, history – disciplines considered 'soft' by most economists.

There has been some work on religion in economics. But when Laurence Iannaccone of George Mason University in Fairfax, Virginia, probably the leading modern economic scholar in this area, wrote an 'Introduction to the Economics of Religion' in the prestigious *Journal of Economic Literature*, he noted that

'the study of religion does not yet warrant a *JEL* classification number'. This simple observation is significant in revealing the amount of attention paid to religion by economists up to the late 1990s.

The example of the Oxford Martyrs is specifically religious, but the arguments being considered are relevant much more generally to all human belief systems where faith or ideology is important. In the decades around the middle of the twentieth century, why did highly placed individuals in both America and Britain decide to give their loyalties to the ideology of communism and betray their countries by revealing secrets to the Soviet Union? Neither Kim Philby not Alger Hiss, two of the most notorious spies, appears to have been motivated by money, by the set of standard incentives in the economist's toolkit. They were motivated by faith, the wholly misplaced faith that the Soviets would create a better future for all humanity. They were utterly and completely wrong. But they believed.

History is replete with examples of ideological differences which cannot be accounted for on the basis of 'rational' economic decision making. Given the historical importance of religion, many such disputes involve this topic. But thinking still of the Soviet Union, after the collapse of tsarist rule in 1917, a vast proliferation of competing political ideologies bubbled to the surface. The Western liberalism of Kerensky, who formed a government for a few brief months in 1917. The several varieties of Whites, believers in monarchy, against whom the Bolsheviks fought a brutal and debilitating civil war. Within the revolutionaries themselves were differing ideological tendencies: anarchists, social revolutionaries, Mensheviks, Bolsheviks, to name but a few. And Bolshevism, the ultimately victorious faction embodied in the Communist Party of the Soviet Union, was notorious for vicious internal ideological disputes even when the key authority figure of Lenin was still alive.

There were undoubtedly many motives at work in each of the arguments and struggles. Personal ambition mixed with genuine belief that your opinions and those of your faction were the correct ones, the ones which would bring about Paradise on Earth. But all of them involved faith and ideology rather than rational, incentive-

based decisions. So, although we now resume the discussion focused on religion, we should keep in mind that the points are relevant to any faith- or ideology-based dispute in human affairs.

*

All sciences classify the various aspects of their discipline. We earlier came across the basic distinction between micro- and macroeconomics. But the scientific classification goes into much finer detail than this. The *JEL* (*Journal of Economic Literature*) system is the one used by all economists. It divides the subject into well over 500 – five hundred – sub-categories. And in the late 1990s religion, one of the most fundamental features of human society, did not warrant a category of its own, so little work had been done on it. The situation has now changed. Religion does have its own economics sub-category. But, revealingly, it is allocated in section Z, 'Z12' no less, coming right at the end of the very long list, lower down even than the ten sub-categories in category Y, all of them 'Miscellaneous Categories'.

In some ways this is surprising, given that Adam Smith, the founding father of modern economics, wrote about the topic extensively in one of his two great books, the *Theory of Moral Sentiment*. He even analysed religious issues from the perspective of agents responding to incentives in his *Wealth of Nations*. He discussed how the clergy could be motivated by self-interest, how monopoly is as bad for religion as it is in other areas of human activity, arguing that competition – being able to choose between competing religions – is good.

Economics had to wait almost exactly 200 years before Smith's analysis was extended, in a model developed by Corry Azzi and Ronald Ehrenberg and published in the top-ranking *Journal of Political Economy*, based in the free market-oriented University of Chicago. A short summary gives a flavour of both this particular model and subsequent work by economists on the topic of religion.

Individuals allocate their time and money amongst religious and secular commodities with the aim of maximising lifetime *and* afterlife utility. 'Afterlife consumption', as Azzi and Ehrenberg describe it, is the

primary goal of religious participation. Secular utility depends in the standard way on inputs of time (work) and the products which are purchased. Afterlife utility depends upon the entire effort devoted by the individual to religious activities over his or her lifetime.

The article is not a spoof, though it would be quite difficult to invent a more effective satire of the model of utility-maximising Rational Economic Man which dominates the entire literature of economics. So many points spring to mind. For example, on a purely technical point within the spirit of the literature itself, but one which is important empirically in some main religions, 'afterlife utility' does not vary continuously with the amount of effort devoted to religion during your lifetime. It is a simple binary outcome: either you are in Heaven, with boundless pleasure, or in Hell, with endless pain. And the outcome might very well not depend on the amount of time and effort which you devote to religious activity. Who will be saved in the well-known parable, the self-righteous Pharisee, obsessed not only with his own virtue but with the constant public display of it, or the sinful but repentant publican who devotes very little time and effort to religion? To be fair to the economics of religion, it has moved on to consider participation more as a group activity and to focus on institutions and their behaviour. But it has very little to say about the most fundamental question: why believe at all?

Despite all this, incentives were certainly at work in the religious world of England in the 1550s. Although there were many nuances within each religion, individuals faced a basic choice between being Catholic or Protestant. Queen Mary's father, Henry VIII, had broken with the Pope in the 1530s and established the Church of England. The institution had gradually come under the control of hard-line Protestants, a trend which accelerated during the short reign of his young son, Edward VI, in the years around 1550. Following Edward's premature death from tuberculosis, Mary – forever known in English iconography as 'Bloody Mary' because of her burning of the martyrs – had come to the throne determined to restore Catholicism.

There were some important directly economic issues to settle. Henry had carried out the biggest seizure of private property in English history, when the monasteries were dissolved and their lands confiscated by the Crown. Under Edward, the church leaders had gone even further, stripping and looting churches of the elaborate trappings and ornaments of Catholicism. How far should these measures be reversed? Mary's main adviser, Cardinal Pole, an Englishman who had almost become Pope in 1550, advocated a complete restoration. Mary had to balance the immediate benefits to her Church against the potentially destabilising political consequences and costs of expropriating the property which Henry had sold on almost immediately to wealthy noblemen and merchants.

But the main question facing her was how to restore the old religion of Catholicism, how to persuade people to re-embrace what she regarded as the true faith. It appears to be the case that the clear majority of the population in fact still adhered to Catholicism. Protestantism was the new brand, as it were, and was still some considerable distance from displacing the market leader. But it had achieved a strong market presence in London and its immediately surrounding areas, then as now the key focus of English political and economic power. Even more pertinently, the leaders of the Protestant Church were, in general, militants.

The bishops and other prominent churchmen could be, and were, removed from their formal positions by simple administrative acts and put in prison. A few were willing to adapt, presumably because they were attracted by the benefits of office and attached little weight to the potential afterlife costs of displaying devotion to possibly erroneous doctrines. The most notorious of these was the remarkable Anthony Kitchin, Bishop of Llandaff in Wales. He was the only person to serve as a bishop under all the various forms of religion embraced by Henry, Edward, Mary and her successor Elizabeth I and who would, in the words of one prominent historian of the times, have doubtless become a Hindu provided he could continue to remain Bishop of Llandaff. Behaviour such as this was satirised immortally in the eighteenth century popular song 'The Vicar of Bray', recounting the

contortion of its eponymous subject in remaining in ecclesiastical office through the religious changes brought about by successive English monarchs. The chorus is a monument to placemen and timeservers everywhere:

> And this is Law, I will maintain,
> Until my Dying Day, Sir,
> That whatsoever King may reign,
> I will be the Vicar of Bray, Sir!

But the removal from office of most of Edward's leading clergy altered neither their beliefs nor those of lay believers in Protestantism. Mary and her advisers soon settled on a policy of terror to deal with this problem. Well-known Protestants would be given every opportunity to recant, but if they continued to refuse they would face the flames.

On the face of it, the strategy was a sensible one to follow from Mary's perspective. There had been many previous examples in human history of terror being successful in achieving its aim. And specifically in England, only 150 years previously, the Lollard heresy, an early form of Protestantism, had been suppressed effectively by a few selective burnings.

*

We might usefully pause to ask why this might be the case. The question of religious or ideological belief is enormously complicated, and one which is ultimately not susceptible to explanation by the model of 'Rational' Economic Man. In terms of this set of behavioural postulates, the agent has first of all to gather available information. But in this context, what is the relevant set of information? By definition, the existence of the afterlife can never be proved, no matter how much information we might gather. The information than has to be processed to come up with the best – sorry, the 'optimal' – choice.

The seventeenth-century French philosopher Blaise Pascal came up with his famous wager to claim that the best strategy is in fact to believe in God. Essentially, he argued that since we are incapable of knowing whether God exists or not, we have to wager on the outcome. In terms of the agent's overall happiness, the gains and losses of belief or non-belief have to be taken into account. In Pascal's own words, 'Let us weigh the gain and the loss in wagering that God is . . . If you gain, you gain all; if you lose, you lose nothing. Wager, then, without hesitation that He is.' In other words, if God exists and you believe, you gain an infinite amount of happiness but if He does not, you lose nothing. But equally if God does not exist, you lose very little either way.

Pascal was a highly original thinker, and his wager is one of the seminal contributions to modern theories of probability and decision making. Not surprisingly, there is a very large academic literature on his wager, a good introduction being in the online Stanford Encyclopedia of Philosophy.* The details need not concern us here, but suffice to say that even after thousands of academic articles, the outcome is unclear. We cannot establish an agreed basis on which an agent might use rational behaviour to believe in God, or a supreme being, or not.

* http://plato.stanford.edu/entries/pascal-wager/.

Ultimately, religious or ideological belief for the individual is a matter of faith and not rational analysis. The social networks in which the person is embedded are also crucial in terms of cultural norms and peer pressure and acceptance or otherwise of belief. 'Embedded' can mean far more than the current position. It can embrace the networks in which they grew up, for example, or which they have been part of previously and which helped shape their current beliefs. And in mid-sixteenth-century Europe, the question was: what variety of religious ideology to believe?

Even the most devout believer experiences doubt from time to time. Within Christianity, even the Apostles themselves experienced crises of faith in the immediate aftermath of the Resurrection. The two men on the road to Emmaus with Jesus were unable to recognise him, a scene given modern vibrancy in T. S. Eliot's memorable

phrase in *The Waste Land*: ‘who is the third who walks always beside you?’ The phrase ‘a doubting Thomas’, meaning a sceptic, has its origin in the Apostle Thomas, who refused to believed until he had placed his hands in the Crucifixion wounds.

In any event, in the rapidly changing circumstances of sixteenth-century England, who could be really sure what was the true faith? Was the Pope the true Head of the Universal Church, or was he merely the Bishop of Rome, or even the Antichrist? The Bishop of Llandaff, whom we met above, would not have been alone in being able, if required, to subscribe to any one of these three distinct propositions.

In contrast, the prospect of being burned alive was only too real and certain. Sometimes, if the fire took hold well, death could be reasonably quick owing to oxygen deprivation, but it was an appalling end nonetheless. Equally, however, contemporary documents record examples of victims dying in prolonged agony, pleading for ‘more fire’ as a combination of damp wood and perverse winds slowly roasted them alive.

So Mary’s policy had a definite incentive element, and agents did react. Some fled abroad. Others openly renounced Protestantism. The most famous of all, Thomas Cranmer, Archbishop of Canterbury and head of the Anglican Church, recanted in prison no fewer than six times before finally summoning the courage to be led to the stake.

But we know from history that her policy failed. Even as Mary lay dying in the summer and autumn of 1558 after four years of terror, there was still a persistent supply of martyrs willing to be burned. And when her sister Elizabeth restored the Protestant faith, the nation embraced it with enthusiasm.

Networks were the reason for this failure. The negative incentive of the fear of the stake took her so far, but not far enough. It was overwhelmed in completely the opposite way by the power of networks. And networks were present in two separate but closely related ways.

The first was the very close network between the militant Protestants themselves, maintained even when they were in jail awaiting interrogation or execution itself. They sent messages, exchanged letters, a veritable torrent of encouragement to keep the faith and set an example to the population as a whole.

This latter, the entire people of England, constituted the second network. Much more loosely structured than the tight-knit religious one, information even in those days did pass pretty rapidly across communities. Most of the population of the country lived within three days' ride of London. The internet it was not, but news certainly travelled. And the burnings themselves were major public events, often attended by thousands of people. The Marian regime sometimes unwittingly contributed to the potential number of favourable message bearers. For the execution of Latimer and Ridley, for example, every household in Oxford was compelled to send at least one member to witness the event. The burnings were not just a deterrent. The authorities were in fact aware that they might also be a source of inspiration. So, leading Catholic figures preached on pulpits specially constructed near the pyre, explaining the heresies of the condemned and expounding what they saw as the true faith. People could be either educated or frightened by the spectacle.

Much of our information on these events comes from a remarkable book, John Foxe's *Book of Martyrs*, first published very early in the reign of Queen Elizabeth I. The Protestant clerical elite were convinced from the outset that their deaths could cause the policy to rebound on the persecutors. They were well aware of the necessity of creating a good impression at the stake. Accounts of their behaviour would spread by word of mouth. Existing Protestants would be encouraged to keep the faith, and waverers influenced by their steadfastness.

To this end, then, they encouraged one another, On 8 February 1555, on the morning of his execution, Laurence Saunders, a noted Protestant preacher in London and the Midlands, wrote to his wife and supporters: 'God's people shall prevayle: yea our blood shal be their perdition, who do most triumphantly spill it.' He actively

encouraged them to attend and enjoy the event: ‘Make haste my deare brethren, to come unto me that we may be mery.’

In terms of dramatic impact, Saunders died embracing the stake. John Rogers, the first of the 300 or so martyrs to be burned, was seen to be washing his hands in the flames and Archbishop Cranmer signalled his adherence to the Protestant faith by thrusting into the fire the right hand which had previously signed his humiliating recantations. John Hooper, Bishop of Gloucester and a notoriously grumpy man, took the opportunity on his way to execution to bless a blind child and greet local dignitaries – for all the world like a modern politician out on the campaign trail for votes. Foxe put the crowd assembled to watch him burn at 7,000, ‘for it was market-day and many also came to see his behaviour towards death’. Seven thousand may not seem a lot, but relative to the size of the population, it was the equivalent of around three quarters of a million in present-day America. In short, a massive crowd.

We know the outcome of this particular historical event. Incentives were put in place to persuade people to adopt one particular set of beliefs, and these had a certain amount of success. But their impact was completely offset and indeed dwarfed by the impact made by the martyrs across the network of the population as a whole. The deliberate policy of calm and even joyous acceptance of death made an impression across the land and people were influenced directly by this behaviour.

But there was no guarantee in advance that this would happen. Indeed, during the years of the Marian terror, there was evidence that the policy was working. As mentioned earlier, some leading Protestants left the country and others, from all walks and levels of life, recanted.

*

A contemporary example both of the potentially huge effect of networks and of the inherent uncertainty of outcome they create is the momentous events taking place in North Africa and the Middle East. I am writing

these words in early April 2011, when neither I nor anyone else knows how they will unfold even in the (historically) short time between the writing and the publication of this book. And I am leaving the original words unchanged as the book itself is being revised during the summer and autumn of 2011. I want to capture on record how things stood in April of that year, to illustrate the uncertainties involved in such situations.

As I write, protests in some countries have already been followed by changes in regime. In Algeria, the government appears to have been able to defuse the tension. In Syria, the protests seem to be in the process of being ruthlessly suppressed. And the current situation in Libya is, to say the least, chaotic and uncertain.

The immediate catalyst to the events was Mohamed Bouazizi, a twenty-six-year-old Tunisian. Bouazizi was a university graduate living in a provincial town where he was unemployed and trying to find, but unable to get, work. He started selling fruit and vegetables in the street without a licence. The authorities put a stop to his activity, confiscated his goods and humiliated him. In response, he set himself alight and died in hospital on 4 January 2011. This sparked the riots which forced the President of Tunisia to flee the country, and was followed by similar uprisings in Egypt, Libya, Jordan, Yemen, Bahrain and Syria.

Clearly, the potential for social unrest was already in place in all these countries, ruled for decades by undemocratic regimes of various degrees of corruption and brutality, with large numbers of discontented young people. The incentives to replace the regimes were there. And doubtless, there had been many individual protests against the regimes, any one of which might have spread like wildfire across the latent network of the desire for change. But these protests have vanished into the mists of history. The network proved robust with respect to these now-unknown shocks. In the case of Bouazizi, the network responded in a fragile way.

Bouazizi's act of defiance was spectacular. Is this why it succeeded and others did not? Perhaps. Readers of a certain age will undoubtedly recall Jan Palach and his role in the events following the so-called 'Czech Spring'

of 1968. Under the leadership of Alexander Dubček, the government of the then-Czechoslovakian state had carried out a series of liberalising measures which alarmed the Soviet leadership in the Kremlin. Czechoslovakia was part of the Warsaw Pact group of nominally independent countries, controlled in practice by the Soviet Union. In August 1968, Soviet military forces occupied the country and Dubček's government was removed from office.

A group of students made a pact to burn themselves to death in public as an act of protest against the Soviet invasion. Jan Palach actually carried it out, not in some remote provincial town but in the principal square of the Czech capital, Prague, on 16 January 1969. The event attracted widespread publicity world-wide, in contrast to the self-immolation of Bouazizi, with a leading British newspaper feeling able to write just four months after his death that 'the name of Mohamed Bouazizi has largely been lost in the unfolding story of the Arab Spring'.* Palach's self-immolation did trigger demonstrations, but these proved to be far from sufficient to bring about change and were suppressed by the security services. So in a network context, we cannot even say in advance whether or not a truly dramatic gesture such as being burned to death for your beliefs will be sufficient to persuade others and bring about change.

* *Daily Telegraph*, 2 April 2011.

*

These rather disparate historical events have introduced the main themes of the book. Most policy, certainly public policy, is based upon the idea that people respond as rational individuals , in the sense in which economics uses this word, to incentives. If you fire a shot into the air as a sign to a group of football supporters to disperse, you believe they will be induced to regard the costs of a riot as being too high, and they will disperse. If you threaten to burn someone to death unless he or she changes his religious opinion or ideological belief, you think the negative incentive which such a death entails will be sufficient to achieve your aim, at least

as far as most individuals are concerned. And if someone is sufficiently stubborn as to ignore the incentive and to undergo the dreadful ordeal, you think the public spectacle will act as a serious deterrent to others.

In the main examples in this chapter, incentives have not worked in a way which has achieved the desired outcome. But quite often, they do. They even worked up to a point in the religious turmoil of mid-sixteenth-century England. So this approach to policy is not always without merit.

But the impact of networks can be considerably greater than that of incentives alone. If the two operate in conjunction, if we experience the phenomenon of positive linking, the changes to behaviour of a relatively small number of people which incentives might induce can spread across a larger group because of network effects. Equally, however, as we have seen, incentives can be swamped if behaviour across a network surges in a manner different from that intended by the policy makers applying the incentives.

Networks introduce an entirely different dimension into the policy picture. This argument is based on the truism that humans are social creatures. As mentioned earlier, in economic theory, individuals operate like multiple Robinson Crusoes, taking independent, autonomous decisions that are not directly influenced by the decisions or opinions of others. Network theory allows the social dimension of human activity to be taken into account when trying to understand how agents behave, and when thinking through the policy implications of their behaviour.

The examples used to illustrate this so far, those of rioting and being burned at the stake, are rather extreme events which relatively few people will ever encounter. But networks on which people copy or imitate the behaviour of others matter in many, more homely, everyday situations. Think of Crocs. For the few readers who have never come across the term, these are shoes which became very fashionable during the 2000s. Shaped like clogs, the upper material of the shoe is studded with holes. Shoes with holes? These may be perfectly functional

in the arid climates of Adelaide or Arizona. But in rainy Seattle or Scotland? Yet the brand proliferated everywhere. And a key reason for the success of Crocs was precisely that they became fashionable in the first place.

A difficult problem, to which we will keep returning during the course of the book, is why something starts to become fashionable at all. Fashion seems to 'emerge' from nowhere in particular. Certainly, many companies, as we will later in the book, are attempting to make use of the concept of positive linking to try to make their brands fashionable, to promote a cascade of sales. But for the moment we will park this issue safely, and simply note that once a brand, a concept, a way of behaving, starts to become fashionable, it becomes even more so simply because it is fashionable. People decide to buy something, to adopt an opinion, simply because others are doing so. It is this social dimension to choice and decision making which network theory captures. In turn, the increased demand sets up positive feedback on the supply side. More shops stock the item because it is becoming popular, which means that even more people become aware of it.

When network effects are present, people are using rules of behaviour which are quite different from those of the behavioural postulates of economics, in which individuals carefully gather information about alternative choices or courses of action and match them against their own fixed preferences. Instead, they may copy, they may imitate the behaviour of others.

And a key implication of basing your decision, at least in part, simply on what other people are doing or thinking is that your preferences are no longer fixed. Instead, they change and evolve over time, as the impact of people on your various social networks alters your own behaviour. Indeed, the preference of the English soccer followers changed in a matter of seconds once the shot was fired. In terms of fashion, you may be influenced by what you see other people wearing in the street – people you may never see again. In deciding whether to buy into a particular pension scheme, your network here is likely to be a very small number of individuals whose

opinion in these matters you trust and value. But in each case, your own private opinion may be influenced and changed by what others are doing.

This is a fundamentally different view of the world from the one in which people are assumed to operate in isolation and to base their decisions on a fixed set of preferences. Even in the latter type of world, the world of standard economic theory, successful policy making might still be difficult. Discovering what people really do want might itself be a considerable challenge. But if their wants, their needs, their desires are fixed, in principle a smart policy maker, whether in the public or corporate sector, can discover them. And the right levers can be pulled, the correct buttons pressed, to change incentives in such a way that the desired outcome of the policy maker can be reliably achieved. In a world in which tastes change in response to the choices and actions of others, this model of policy can simply no longer be relied upon.

The problem for policy makers is even more complicated. Networks can appear in a variety of guises and be activated in a variety of ways. The soccer hooligans were connected by a network of shared interests such as violence and, in case we forget, the game of soccer itself. This is why they bonded together in this particular social network. The choices of individuals in this group on matters such as which brand of trainers to wear, or their opinions on which players should be selected for the team, would be influenced by the choices of others – exactly as they are in the wider, innocuous world of fashion and the purchase of items such as Crocs.

But with the vandals, the shot fired into the air was the signal for a qualitative change in the nature of this network to take place. From being a network of individuals connected in terms of shared interests but nevertheless capable of making decisions as individuals, it became fused into one. A particular mode of behaviour, involving the destruction of property and physical violence, flashed across the entire network. Everyone adopted and became consumed by the same mode of behaviour.

In the example of religious faith, a network of awareness was already in place across England. Many people knew of Cranmer, the Archbishop of Canterbury, and other nationally prominent church leaders such as Latimer and Ridley. At a more parochial level, people would be likely to know of the diocesan bishop, say, or a leading local preacher, even though people on the other side of the country might not have heard of them. So, the population of England was connected on a network which transmitted awareness of the existence of a range of Protestant leaders. The martyrs took the gamble that this pre-existing network, which essentially consisted of one in which information about them was exchanged, could be used by them to influence behaviour and opinions.

But the feature which all these examples have in common, in their widely different contexts and impacts, is that behaviour and opinions can be altered not just by individuals reacting to changes in incentives but directly by what others think, believe or do. The challenge for policy makers in the interconnected, networked world of the twenty-first century is to harness this positive power of networks and to use them in conjunction with incentives. We need networks both to make sense of the policy terrain and to design more effective policies, We need positive linking.

Chapter 2

‘Up to a point, Lord Copper!’

We have seen a range of examples, from harmless choices in footwear fashion to altogether more dramatic decisions on ideology, in which the effects of networks were much more powerful than those of incentives. One purpose of this chapter is to redress the balance somewhat by providing diverse examples of incentives having clear and identifiable impacts on behaviour. Successful policy making in the twenty-first century requires an understanding of *both* networks and incentives, a point which is illustrated again as we move through the chapter.

Incentives do matter. This is the one great insight into human behaviour which economics provides, an insight which is supported by an enormous amount of empirical evidence.

I should say immediately that this does not mean that free-market, equilibrium economics is correct. This latter sentence is so important that I will repeat it in bold type. **The fact that agents respond to incentives does not mean that free-market, equilibrium economics is correct.** It is not at all necessary to believe in the whole of the standard behavioural paradigm in economics in order to recognise that incentives matter. Indeed, in the centrally planned economies of the old Soviet bloc, incentives could take the form of social acclaim for meeting the production norm, or for being awarded a medal as a Socialist Hero of Labour.

We will see during the course of the chapter that the use of incentives to achieve aims or targets, perhaps changing tax rates or giving out subsidies, is by no means a panacea for the policy maker. This is the case even when network effects are either weak or absent more or less altogether. Humans are inventive, innovative, and they may very well respond to changes in incentives in ways which are very hard to anticipate.

But incentives work only ‘Up to a point, Lord Copper’, as the editor of the *Daily Beast* in Evelyn Waugh’s novel *Scoop* used to say when his titled proprietor asserted something which was at best only partly true and was often unequivocally wrong.

Sometimes, changing incentives does work out more or less as expected. This is certainly true in a qualitative sense, even if the exact quantitative predictions are not borne out. If Coca-Cola, say, puts the price of its products up and nothing else changes at the same time, to state that the sales of Pepsi will probably go up is perhaps merely stating the obvious. Agents who like this sort of fizzy drink are given an incentive to buy less Coca-Cola: its price has gone up. But incentives have been used in less expected ways, sometimes by the most unexpected people.

For example in 2003 Ken Livingstone, as mayor of London, introduced the ‘congestion charge’, a tax on vehicles entering central London during the day, in an effort to solve the problem of traffic congestion. Even the mayor’s worst enemies could scarcely accuse him of being a gung-ho free-market economist. His political stance has always been firmly on the left.

Nevertheless, in a politically bold move, Livingstone attempted to deal with the traffic problem in a major world city by the use of incentives. There was great uncertainty in advance, which persisted even during the early months of the actual operation of the scheme, about how agents – motorists in this case – would respond. Many different forecasts were made. But the tax has worked *qualitatively* exactly as one would expect. Traffic flows into Central London are lower than they would have been without this charge. In other words, faced with an additional cost of driving into central London, some motorists have decided either to reduce their visits to the area, or to use alternative means of transport.

Livingstone made use of the reaction of agents to incentives to achieve a desirable social goal. It does not mean he had become a convert to the political ideology of Mrs Thatcher. Nor does it necessarily mean that motorists carried out a rational analysis of the costs and benefits of the scheme. But it worked.

Much more generally, the economics of the mainstream works up to a point exactly *because* it incorporates this fundamental insight into human behaviour, that changes to incentives alter it. But such insight is perceived as through a glass darkly. Except in certain limited circumstances, the rest of its theoretical constructs are at best shaky and often plain wrong. The core model of standard economics assumes that an agent gathers all the available information relevant to a decision, and is then able to process it in a way which enables the agent to arrive at the very best decision possible, given the tastes and preferences of the agent. These latter are assumed to be fixed and cannot be influenced by what other agents do. All theories are approximations to reality. The question is always: how good are the approximations?

Policy based on the use of incentives is mistrusted by many people, precisely because it has the image of being derived from the highly mathematical abstractions of economic theory. So, before going on to give examples of unintended consequences of changes in incentives, it is worth considering briefly whether the use of abstraction and maths can be justified in analysing social and economic problems. We are, after all, dealing with human behaviour.

I have no problem at all with abstraction, though I am mindful of the fact that this needs to be justified. Many scientific theories are highly abstract, far removed from everyday life. But this abstract quality is the very nature of the beast. From the myriad complicated details which surround many situations in reality, we are trying to distil a few key factors and to describe how their interactions help us to understand what is going on. Ideally, we would like this theoretical account not just to shed light on one particular problem, but to be capable of generalisation across many situations.

Ultimately of course, any theory, if it is to be regarded as being truly scientific, has to be tested by empirical evidence. It must be judged not by its abstract beauty, but by its ability to explain messy reality. This is how, for example, we are able to dismiss astrology as not being a science. For a number of years at the start of my career I was an economic forecaster, attempting to predict the future course of the British economy. Despite being equipped with the latest red-hot developments in both economic and statistical theory, the ability to achieve any degree of systematic accuracy proved elusive. Growing disillusioned with the whole process, for a short time I carried out a comparison between the accuracy of my macroeconomic forecasts and those of my daily horoscope. If anything, the latter had the edge!

But theoretical abstraction is both desirable and necessary for any sort of progress to be made in understanding the world, whether in the natural, biological or social sciences. So the abstract nature of much economic theory does not of itself make it a legitimate object of criticism. Even if the theory is misleading, or even plain wrong in many situations, this is not because of its abstract nature. It is because the assumptions, the simplifications which are needed to have any sort of theory, are not supported by the evidence.

*

It is rather harder to justify the use of mathematics in economics, especially the specific sort favoured by economists. Maths is pervasive in economics. It is an integral part of the self-image of the discipline. And it serves as a very distinct barrier to entry, a very clear 'Keep Out!' message to anyone who does not feel comfortable in this area. Given its importance, a fairly lengthy detour is warranted before we move back to an explicit discussion of incentives, the main theme of the chapter.

Maths can in fact be very useful, provided that we think of it as just a tool. Economists often make clever maths an end in itself, and in doing so overlook the fact that we are trying to understand and explain what happens in the real world. Elegant maths also often leads economists to make the mistake of confusing the model with

reality. It is a tool which can assist logical thinking. It's like another language. It can help us find our way around and serve as a medium of communication amongst people discussing the same subject.

Here is a cautionary tale of how maths has become fetishised at the very highest levels of the discipline of economics. A friend of mine teaches economics at the University of Cambridge in England. Fairly recently, she had a first-year student who was very good indeed at maths. So much so that he complained there simply was not enough of it in his course. For his second year, he was sent on an exchange to the other Cambridge, to the Massachusetts Institute of Technology. Emails of an increasingly desperate nature began to whizz back to my friend across the Atlantic. The final one said simply: 'Help! Please let me back home. There isn't any economics in this course. It's *all* maths.'

Things are not quite as bad as that in most places, but the use of maths has become pervasive in economics. Just for the record, at the right time and amongst consenting adults, I, too, use maths extensively, albeit of a different kind from that which pervades economics. You can tell I am an economist myself when I say, on the one hand there are good reasons for the use of maths in economics, and on the other hand there are bad ones. So far, mainly the bad ones have prevailed.

It would not matter very much if economics was not taken so seriously by policy makers. Hardly anyone bothers about some of the lunacies in literary theory, for example. But economics matters. Why is it that maths came to be so pervasive in economics, when so much was achieved without it? The worst reason is that the use of maths makes economists feel that they are proper scientists. They suffer from deep physics-envy. Ironically, economists seem to envy the classical physics of a century and more ago. In many ways, physics itself has moved on to incorporate network-based ideas.

Physicists have to use maths – try doing quantum physics in words. And they are real scientists, who really have explained how lots of things really do work. So if we use maths, that makes us real scientists, does it not? The logical error in this last sentence is pretty obvious. But it does not stop the inner glow of satisfaction that most economists feel when they cover the page in mathematical symbols.

There is a more serious and more damaging reason why maths, or at least a particular kind of maths, is used in economics. This is inextricably linked with the concept of Rational Economic Man. In essence, as noted in the previous chapter, economics is a theory about how individuals behave. And in the standard theory, it is not just that people are assumed to be self-interested. Rational Economic Man acts like some sort of super-computer, always gathering every single bit of information which is relevant to a decision, and then making the best possible decision out of all the available options. Not just a good decision, or even a very good one, but the best. The 'optimal' decision.

Now there is a whole branch of maths devoted to 'optimal' solutions. This is differential calculus, which many readers will have come across at school. It is the ideal tool for a theory which says that individuals behave in a way which is optimal for them, given their tastes and preferences. So if you eat junk food and weigh 300 pounds as a result, or if you drink heavily and destroy your liver, or if you smoke and get cancer, if you riot when the police open fire, that is your choice. You must have been making what you believed to be the best possible lifestyle choice for you, and calculus can prove this.

This is still the basis for a lot of the economics which is taught today. Yet, paradoxically, it has been precisely the use of maths within economics itself which has exposed fundamental problems at the very heart of the model of the Rational Economic Man theory of behaviour.

Working out the full implications of these behavioural postulates proved an exceptionally demanding scientific task, which took a century to complete. By the mid-1970s, this programme of research was eventually finalised. There is nothing left to discover. It is a marvellous intellectual construct, but it turns out to be a scientifically empty box. It has no testable implications. In other words, there is no empirical test we can use with which the theory could be refuted. Such tests might be very hard to devise, but any true science has to be capable of being refuted empirically.

*

We might pause and offer some postulates which might be capable of being refuted by evidence. For example the familiar diagrams of basic Economics 1.01 show downward-sloping demand curves. This simply means that a lower amount of a product is demanded if its price goes up. A higher price means less sales by volume, and if we plot such a relationship in a simple chart, the relationship between sales and price will slope downwards. Figure 5.2, on page 00, illustrates the point. But we cannot deduce logically from the theory of Rational Economic Man behaviour that this key statement, widely used by economists, is true.

This result was established through pages of intricate maths. A translation into English could be carried out, but a full explanation would take many chapters. However, a useful insight into the proposition that we cannot deduce theoretically that market demand curves slope downwards, even if the demand curves for every single agent do slope down, is as follows. Suppose the price of a product is increased. The volume of sales will go down. The value, which is just price multiplied by quantity, may either rise or fall depending on how much the volume changes compared to the increase in price. If a 10 per cent price increase only reduces volume sales by 1 per cent, the value of sales will rise, but if it cuts volume by 20 per cent, value will fall as well. Either way, the income of the company making the product is changed. If the market for the product is tiny compared to the economy as a whole, we can reasonably stop our analysis of the implications of the price rise here.

But suppose we are thinking about the demand for labour, about how many workers firms want to employ. This market – or collection of markets – is typically enormous relative to any national economy. A persistent theme in economic discourse, and indeed policy, is the need to 'price people back into work'. In other words, to reduce wages so that more people are employed. A cynic might note in passing that the argument is usually applied to less skilled people, and few in the financial services sector have been suggesting that bankers take a salary cut in order to price them back into work after the financial crisis. However, this is indeed a point made in passing.

Suppose further that, by some means, a government succeeded in cutting wages to increase the demand for labour. The potential problem here is that wages are not just the 'price' of employing someone to the company, they represent spending power across the whole range of goods and services. So a reduction in the price of labour might lead to *fewer* people being employed because consumption by the workers, the total amount they spend, may decline too. On the one hand, labour has become cheaper. On the other, workers have less to spend throughout the economy as a whole, so the demand for many products may fall, and less workers be needed to produce them. Of course, there could be an offsetting effect if firms spent any increase in profits resulting from the wage cut. This account, however, does give some insight into why it is not possible theoretically to prove that 'demand curves slope downwards'.

Despite the theoretical indeterminacy of the core model of conventional economics on this point, the idea that labour is too expensive, that it should be priced back into work, that real wages should be reduced, was a key ideological theme of the 1990s and 2000s in many Western countries.

Of course, the reality is that if the price of a product or service is increased then usually, but not always, its sales fall. In practice, and certainly when added up across any particular company, salaries cannot be significantly higher for a long period of time than the value of the contribution of the workforce. If they are, the firm will eventually fail. But these are empirical observations, which seem generally valid but which do not obtain on

every occasion. The point about the theoretical model is that even if we observed the opposite, that, say, when the price of a product rises its sales increase, this would not be a refutation of the theory, because the theory allows the relationship between price and sales to take any shape whatsoever.

So, paradoxically, the use of maths in economic theory, and more precisely the type of maths preferred by economists themselves, has provided some very powerful results which do much to undermine many of the policy-related claims it makes.

*

In practice, of course, returning directly to the main topic of the chapter, incentives often matter. To stress again, this statement does not necessarily imply that the theory of Rational Economic Man is correct. Further, when we speak of 'incentives', the usual concepts associated with the word are things like pay and prices. You may be offered a bonus at work in order to try and incentivise you, either to work hard or to stay loyal. A government may increase tax on fuel, say, to try to get people to use less of it, to respond to the negative incentive of the higher price. Yet as we have seen, incentives may often appear in an unconventional form, such as the 'price' of being burned to death.

It is not just that incentives may appear in odd or unexpected guises. Their effects may be hard to anticipate. This does not mean that they are not working at all. Nor does it mean that we have simply just not done enough careful statistical analysis of behaviour to be able to say that if, for example, we put the price of Coca-Cola up by 10p or 10 cents, sales will fall by *x* or *y* per cent. Rather, it can mean that agents *do* respond to changes in incentives, but the ways in which they respond may not be anticipated by the policy maker. Or it can mean that the incentives work in the way they were intended to, but that the wider consequences of this are not foreseen.

On the latter point, a few years ago, my wife telephoned to make an appointment with a doctor at our local National Health Service surgery. It was not urgent. Previously, she had been able to make such an appointment

at her convenience in several days' time. But the rules had changed. Appointments could now only be made on the same day that the request was made. The government had brought in a target that a high percentage of patients had to be seen on the same day that they contacted the doctor. On the face of it, a perfectly desirable aim. Doctors received payments which were conditional on the target being met, so they were incentivised to do so.

But this led to great inconvenience. Working people cannot always guarantee to be free on any particular day, and even when they can, others may get in before them and take the slots. When she finally managed to get to see the doctor, many fruitless phone calls and over a week later, she raised this policy change with him. The doctor was most apologetic. It was not his fault. His funding depended in part on him meeting the government's new target. The doctor was responding to incentives in a perfectly sensible way from his point of view. The unintended consequence was that his unfortunate patients experienced considerable inconvenience in no longer being able to make appointments to suit their schedules.

The limitations of what we might term the 'clever regulator' approach to policy, using incentives to achieve specific targets, were also present during the UN Climate Change Conference in Copenhagen in 2009. This time, the authorities failed to grasp how agents might respond innovatively and imaginatively to the incentives put in place. The city's mayor, Ritt Bjerregaard, and the city council wanted to curb prostitution during the conference. They sent postcards to hotels and delegates to the conference urging them not to patronise the city's sex workers. The delegates were exhorted to 'Be sustainable – don't buy sex'. The hotels themselves were admonished: 'Dear hotel owner, we would like to urge you not to arrange contacts between hotel guests and prostitutes'. Here was a clear set of incentives placed in front of agents. The incentives were perhaps particularly strong for the hotels, for it is not in the interests of service providers such as these to incur the

potential wrath of the local authority, equipped as it is with all kinds of powers which can make their lives difficult.

The response of the prostitutes shows how inventive humans can be. Members of the Sex Workers Interest Group simply offered free sex to anyone who could produce both their delegate credentials to the UN conference and one of the notices sent out by the mayor. Their incentive to maintain their business was sufficiently strong for them to introduce this innovative marketing arrangement. The choice they faced was between a much reduced income if the mayor's strategy was complied with, and a normal income reduced by the occasional free service.

A much more detailed example of unintended consequences, or, more precisely, consequences which are very hard to foresee because of the innovative responses of agents, is as follows. Jérôme Adda and Francesca Cornaglia of University College London published a study in 2006 in the top-ranking *American Economic Review*. The potentially detrimental effect of nicotine on smokers' health is well established. In recent decades, most Western governments have attempted to reduce cigarette consumption, and an important way of trying to achieve this aim has been by increasing excise duties – taxes – on cigarettes to make them more expensive.

The policies have undoubtedly been successful. A number of detailed academic studies have shown a distinct correlation between higher prices and reduced cigarette consumption. The initial impact of a tax increase tends to be diluted over time because of the addictive nature of the product, but it nevertheless persists.

In this context, we might usefully note that social networks have reinforced the impact of incentives, and in particular by their influence in persuading people to stop smoking altogether. The Framingham Heart Study is a unique database, monitoring the health of individuals over many decades in the eponymous town in Massachusetts. It is a rich source for medical research. But it also provides material for social scientists.

Unusually for such medical surveys, the study contains information not just on the individuals but on their family and friends.

Nicholas Christakis and James Fowler of Harvard University analysed the data from a network perspective, publishing their findings in the *New England Journal of Medicine* in 2008. Their results were striking. The cessation of smoking by a co-worker in a small company decreased a person's chances of smoking by 34 per cent. If a friend gave up, the person was 36 per cent less likely to smoke, and the chances were 59 per cent less if a spouse stopped smoking. Their study does not identify the separate impact of incentives on individuals, such as price increase or the public information provided on the health risks of smoking. But once an individual makes the decision to stop, for whatever reason, the effect is potentially transmitted through his or her social networks – family, friends, work colleagues – to others, who might stop simply because of the example which has been set to them. So networks can operate *with* incentives, to reinforce and magnify the initial impact of the latter.

To return, however, to Adda and Cornaglia and the potential difficulties of anticipating the effects of changes in incentives. Their article is based on the so-called 'rational theory of addiction'. Full of heavy-duty maths, it is replete with phrases such as 'we assume a quadratic utility function' and 'the proof requires a second-order Taylor approximation'. But, to reflect on points made earlier in the chapter, we really do not need the assumption of rational behaviour at all. The empirical results of the study, based upon careful statistical analysis of the data, are clear.

They used the American National Health and Nutrition Examination Survey, a database of some 20,000 people across the United States, which contains information on the number of cigarettes smoked and their nicotine, tar and carbon monoxide concentration. Tax rates on cigarettes vary across states, providing plenty of variation with which to estimate their impact on behaviour. In common with many other studies, Adda and Cornaglia

found that the higher the tax rate, the fewer cigarettes were smoked. So far, good and entirely expected news for the health-promoting policy maker. But they discovered that higher tax rates led smokers to switch to brands with a higher tar and nicotine yield. This was not in itself a novel finding, though it increased the credibility of the two previous research papers which has previously reported it. A large number of papers had been written on the impact of taxes and prices on the number of cigarettes smoked, but only two on this switching behaviour, so it was valuable confirmation of this effect, though nonetheless worrying.

The real originality of the research was the discovery that smokers also increased the intensity of their smoking by extracting more nicotine per cigarette, regardless of the brand which was consumed. Smokers become more inclined to smoke the cigarette right to the end, behaviour which not only increases tar and nicotine consumption, but also leads the smoker to inhale more dangerous chemicals, which in turn has been shown to cause cancer deeper into the lung. So, yes, higher taxes do reduce the sales of cigarettes. Incentives work as expected. But at the same time, smokers compensate by both switching to brands containing more tar and nicotine, and by consuming cigarettes in ways which are more dangerous to their health.

So, traffic in central London, the response of professional doctors to changes in payment structures, the supply of sexual services, nicotine consumption – a disparate range of circumstances in which incentives have altered behaviour. Sometimes in ways which the policy maker did not foresee and did not, in hindsight, desire. But incentives undoubtedly mattered.

Network effects have also been demonstrated to be important, reinforcing the impact of incentives, in at least one of these examples. And in general, as we saw in the opening chapter, in most real-world social and economic situations, we need to understand the potentially subtle and powerful interplay between the effects of incentives and the effects of networks.

*

So far in this chapter we have focused deliberately on incentives, on some of the strengths and weaknesses of the traditional tool of policy makers. It is time for a shift of gear. Time to discuss at some length an important policy area where both network effects and incentives are at work.

The complex relationship between incentives and networks to which I refer is that within the criminal justice system. A perennial question for policy makers is: does prison work? It is at this point that we put on our detached philosopher's hats and ask in turn what precisely this question means.

Even in the United States, where the rate of incarceration per 100,000 inhabitants is five or six times the average in the rest of the developed world, it is hard to get sent to prison. There are occasional highly publicised stories in which individuals leading hitherto blameless lives suddenly in a fit of rage murder their spouse, or even go on a killing spree. But the vast majority of people who get sent to prison have already had fairly extensive experience of crime and the criminal justice system. They are steeped in the culture of crime. Superb television series such as *The Sopranos* or *The Wire* hold our attention through the quality of their scripts and acting. But they also succeed because they are realistic. Almost like Dante's Circles of Hell, there are concentric rings of individuals, from the hard-core gang leaders out through people only peripherally engaged with them. But they are all engaged with the process of crime. And not surprisingly, once they become engaged in this way through their social networks, some of them end up in prison.

One thing we have observed empirically about crime is that once a person has been in prison, he (or, far less frequently, she) is very likely to commit a crime again. Recidivism rates are high, even though a vast range of policies have been tried in an attempt to get the rate down. In this sense, prison does not work.

Despite stories in the popular press about how criminals live a life of luxury in jail, in reality a prison term remains an unpleasant experience for most of those incarcerated. Apart from the loss of liberty, crime within

prisons is often rampant, and many inmates live in fear of physical violence. Tom Wolfe's description in *A Man in Full* of the Californian prison in which one of the book's more sympathetic characters, Conrad Hensley, is incarcerated is awful, brutal – and entirely accurate. Yet the experience seems to provide scant deterrence against reoffending. Once someone is sent to prison, the influence of social networks on individual behaviour appears to dominate that of incentives. The social and cultural milieus of hard-core criminals, the social networks in which they are embedded, are ones in which crime is itself the norm.

A key fact is that most crime is committed by young men possessing little money or intelligence and few skills. And there does appear to be something inherently implausible about the idea that such individuals assess all the available information and choose the 'optimal' decision when they are contemplating breaking into a car or thinking about punching someone in a bar.

The standard response by economists to such points is to invoke the 'as if' argument. In other words, whilst it may not appear that agents go through the process of finding optimal decisions, they behave 'as if' they do. There are layers of subtleties to this argument which need not delay us. But even the simple statement of the point is not as foolish as it might first appear. Very few of us know how to solve the difficult non-linear differential equations which describe the flight of a thrown ball, yet most of us can predict its path well enough to catch one. It is 'as if' we had done the maths.

But, to repeat the point made at the start of the chapter, we do not need to invoke the idea that agents are responding in some optimal sense to incentives. People may be short-sighted in terms of the decisions they make, they may consistently make decisions whose outcomes go against their own self-interest. Yet they are still changing their behaviour in response to changes in incentives.

This is a point which most economists find hard to accept. Surely, the mainstream 'rational' agent argument goes, people do not necessarily always make the optimal choice, but rather over time they gradually learn to avoid decisions which are not in their own interests. This raises an issue of great importance to which the whole of the next chapter is devoted. In essence, in many situations, the best choice can *never* be identified, even after the event, no matter how smart we may be.

For now, however, in the spirit of empiricism, we simply note that, for most people, crime does not pay. Most of the young men who spend their days steeped in petty crime would actually be better off in straight money terms in low-paid legitimate employment. The proceeds from most crimes are very small. Criminals often act impulsively, paying less regard to the potential costs to them of committing a crime than the objective evidence indicates they should. For example, as known criminals in a locality, they have to endure the stress of frequent visits from the police, their own records making them natural suspects. So the benefits from crime are not as high as those from a regular job, and the costs – the stress of being known criminals, the court appearances, the fines, the frequent prison sentences – are much higher. Yet individuals persistently choose to follow a criminal life. In part, this can be explained by the culture in which they become involved, the social network of crime. But in part it can only be said that they are making a constant stream of decisions which, from the point of view of economic theory, are irrational, which go against their own self-interest.

However, this certainly does not mean that criminals fail to respond to incentives. True, they make decisions which, in terms of their own self-interest, are often not very sensible. True, they do not necessarily respond exactly along the lines of the theory of Rational Economic Man. But their behaviour can nevertheless be influenced by the various positive and negative incentives which criminals face.

An example of positive incentives occurred in April 1999, when Britain introduced for the first time a minimum wage, which provided pay increases for a large number of low-paid workers. Two London-based economists,

Kirstine Hansen and Steve Machin, carried out a careful and very sophisticated statistical analysis of its impact across the forty-three police-force areas of the UK. They concluded unequivocally that 'altering wage incentives can affect crime and therefore that there exists a link between crime and the low wage labour market'. By making the alternative option of regular employment, albeit at the minimum wage, more lucrative, some potential criminals were incentivised to choose this rather than to 'earn' their living from crime. This does not mean that they had suddenly become rational agents in the economic sense of the term, able to assess costs and benefits more effectively. It simply means that incentives had changed and, however imperfectly, some agents responded.

*

Steve Levitt is famous for his blockbuster book *Freakonomics*. But he is also an extremely distinguished economist, winner in 2003 of the John Bates Clark Medal awarded to 'that American economist under the age of forty who is adjudged to have made a significant contribution to economic thought and knowledge'. One of his areas of interest is crime. And the discussion of why crime has fallen in America in his book is based in turn on an article he wrote in the prestigious *Journal of Economic Perspectives*, which in turn is based on a large number of technical academic articles.

One of his conclusions has struck a notable chord with many people, quite possibly because of the rather startling and unexpected nature of the topic: that one reason for crime falling sharply is a rise in the number of abortions. As noted, most crime is committed by poor, unskilled young men, and more abortions in their social group means that there are fewer of them around to commit crime.

But Levitt also concludes that prison is in part responsible for the dramatic reductions in crime, a finding often conveniently forgotten. Both Britain and America saw large increases in the prison population starting around twenty years ago, and in both countries there have been sharp subsequent falls in crime. Of course, simple

correlation such as this does not prove causation, but Levitt's conclusions are based upon highly sophisticated statistical studies which readily encompass such issues.

One obvious factor is that people in prison cannot commit further crimes – at least against society in general. So, by simple arithmetic, a bigger prison population means smaller crime figures. But the more important impact arises from the deterrent effect, not on those who are actually serving sentences, but on those at liberty who are contemplating breaking the law.

A key step seems to be moving from a situation in which a young man has not committed a crime, to one in which he has. Although most crime is committed by poor, unskilled men, most such individuals remain law-abiding citizens. They will probably know who the criminals are in their neighbourhood, and may even socialise with them. But the main influence of their particular social networks, the main impact on what they regard as normal behaviour, remains people like themselves who live their lives within the law.

Of course, carrying out a single criminal act does not immediately lead to the destruction of an individual's existing social networks and to his re-embedding into a group of hardened career criminals. But it may begin to alter these relationships.

An elusive goal for criminologists is to identify individuals who are more likely to become prolific criminals (at some point in their lives) than others. The group of prolific criminals is usually thought to be some 5 or 10 per cent of the total population of offenders. Some progress has been made. Being born into a family where most members are criminals increases this probability substantially. And it is now clear that boys raised by single-parent, never-married mothers also exhibit a higher probability of being involved in crime than others with different family backgrounds. I should stress that it is not the case that boys from such backgrounds automatically become criminals. Indeed, most grow up to be perfectly respectable members of society.

Nonetheless, the probability of them turning to crime is distinctly higher than it is for the population of boys as a whole, even taking into account factors such as family income. In general, the vast majority of offenders share the characteristics of the persistent offenders, making prior identification very difficult. Many such offenders abandon crime during their twenties and become productive, taxpaying citizens.

It is, of course, neither practical nor acceptable to incarcerate boys from criminal families as soon as they reach puberty, still less every boy from a poor, single-parent family. Could we not instead attempt to identify the much smaller number of those who are likely to commit large numbers of crimes and devote resources diverting them from such a path before it is too late?

If we could do so, there would be a double impact: first, fewer criminals would mean fewer crimes; second, the influence of criminality as a social norm amongst their peers would be weakened, since individuals known to have committed large numbers of crimes undoubtedly attract attention and gain influence in such social circles.

*

Surprisingly little systematic work has been done on the number of crimes committed by individuals, but there are two well-established databases which record criminality amongst a group of individuals over time. The first, the Cambridge Study in Delinquent Development, is a prospective longitudinal survey of 411 males in a working-class area of north London. Data collection began in 1961–2. The second, the Pittsburgh Youth Study, began in 1986 with a random sample of boys in the first, fourth, and seventh grades of the Pittsburgh public school system. The sample contains approximately 500 boys at each grade level, for a total of 1,517 boys.

The Cambridge data relates to the number of convictions for each boy over a period spanning the mid-1960s and 1970s. The Pittsburgh data describes self-reported acts of delinquency over short time intervals beginning in the late 1980s. In other words, the studies differ both in their time coverage and in the fact that the Cambridge study is based on convictions whilst the Pittsburgh one utilises self-reporting.

Despite these differences, there is a remarkable similarity between the two in the statistical distribution of the number of crimes committed. The 'statistical distribution' in this context describes how many individuals in each database commit (or record) zero crime, how many commit just one, how many commit two, and so on.*

* The actual analysis relies on a number of mathematical concepts which would take considerable time to describe in words. For those interested in the details, I have published the analysis as 'Scaling Behaviour in the Number of Criminal Acts Committed by Individuals', *Journal of Statistical Mechanics: Theory and Experiment*, July 2004. It may be thought unusual that a statistical physics journal would be interested in this analysis, but the statistical distribution which is identified is one of general interest to this particular research community.

There are two striking features of the results. First, a much better description of the number of crimes committed by individuals is given if we segment the number into two separate groups than if we analyse them all together. Specifically, the groups are 'the numbers who commit zero crimes' and 'the numbers who commit any crime'. In other words, the description of the data when the number of boys committing or reporting zero crimes are excluded is different from that when they are included.

Second, once this distinction is made, there is no 'typical' number of crimes which an individual commits. Once a boy has moved from committing no crime to committing just one crime, the total number of criminal acts he might commit can take place on all scales. Moreover, the number of crimes which any individual does in fact commit can be thought of as the outcome of a purely random process.

We see again, incidentally, the concept of 'robust and fragile' networks introduced in the opening chapter, albeit in a slightly different guise. Here we have a network, in this case a population of young men, living on the same public housing scheme perhaps, who have not yet committed a crime. For whatever reason, one of them carries out a criminal act. Most of the time, he will never go on to commit more than a handful of crimes. Occasionally, he will graduate to a life involving numerous criminal acts spread over a period of years. The network is robust in the sense that most of the people in it carry out either no crimes at all or just a small number. And it is fragile because a small number do go on to be career criminals.

These abstract concepts have two important practical implications. First, the fact that the number of crimes committed by an individual is compatible with the outcome of a purely random process means that it is not possible to identify in advance, once a crime has been committed, how many crimes that individual will go on to commit. So we cannot hope to target *in advance* those boys who will have a highly prolific career in crime, and who may therefore exercise a strong influence over the behaviour of their peers. We may, as discussed above, be able to go some way in identifying those who are more likely to make the first crucial step from zero to one crime, but we cannot then go on to separate those who will commit many more crimes from those whose criminal career will involve only a small number.

The second is that the crucial step is indeed to make the transition from being law abiding to carrying out the first criminal act. In terms of the numbers of individuals committing different numbers of crimes, more commit just one crime than commit two, more commit two than three, and so on. But the largest and most important distinction by far is between zero and one.

Here is where we see an interplay between incentives and networks. Once a young man makes the initial transition to crime, his perception of himself starts to alter, as does the perception others have of him. He becomes potentially less acceptable to the members of his various social networks who do not commit crime, and, conversely, more in tune with the social values of those who do. He may himself accelerate the process of potential change, depending on how much his own self-image is altered as a result of his actions. His identity changes. Once a young man has carried out a few crimes, he is by no means predestined to become a career criminal. Indeed, most do not: the effect of the non-criminal social networks in which he has been involved as a non-criminal often draw a young man back into a law-abiding life.

Incentives, whether positive or negative, may influence either the crucial initial decision to commit the first crime, or, later, the decision to withdraw from a life of crime. But the influence of peer pressure, peer

acceptance, the gradual increase in the relative importance of copying or imitative behaviour compared to that of incentives, increases the more a young man associates with criminals, and there is a chance that a criminal career has been born.

It is this which provides the intellectual basis for successful policies of containing crime. It is the positive use of networks of attitudes amongst the 'at risk' group, the search for the triggers which will generate positive linking across these networks, which will keep crime down. There is no check list of policies, each 'rationally' evaluated by teams of economists, which will guarantee success. Rather, it is the much more subtle concepts of social norms, of what constitutes reasonable behaviour in the relevant peer groups, which is the key. This is hard to achieve, not least because there is no readily specified tick-box approach to the problem. But positive linking has the potential to create massive changes for the better.

Yet a nagging question runs through this whole discussion of crime. Most criminals have backgrounds of poverty, they have low levels of conventional skills, and are often barely literate or numerate. It seems implausible that they behave as 'rational' economic agents, gathering information and meticulously processing it in order to arrive at the best possible decision, given their own tastes and preferences. Indeed, we have argued that this is *not* how we need to see them as behaving. Agents can still react to incentives even though they are not following the behavioural precepts of conventional economic theory.

The question is: does economics have anything to say about behavioural patterns which do not square with its core theoretical assumptions? Do we simply dismiss such behaviour as 'irrational', or is there something more useful we can say?

The concept of the rational agent does indeed remain very much alive and well within economics. But, as it happens, in recent decades, a whole new empirically driven field has developed within the subject itself, one which poses challenges to the mainstream view of how the world operates. This is known as 'behavioural economics', and it is to this which we now turn.