



Behavioural Economics

Richard A Epstein

NEW ZEALAND BUSINESS ROUNDTABLE

MAY 2005

*This lecture, Behavioural Economics,
was given in Wellington on 3 August 2004
as a Treasury Guest Lecture.*

First published in 2005 by
New Zealand Business Roundtable,
PO Box 10-147, The Terrace,
Wellington, New Zealand
<http://www.nzbr.org.nz>

ISBN 1-877148-99-7

© Text: as acknowledged
© 2005 edition: New Zealand Business Roundtable

Printed and bound by *Astra Print Ltd, Wellington*

Contents

Richard A Epstein v

Introduction - Dr Veronica Jacobsen,
principal advisor, Labour Market and
Income, The Treasury vii

Behavioural Economics - Richard Epstein 1

Questions 23

Richard A Epstein

RICHARD A EPSTEIN is the James Parker Hall Distinguished Service Professor of Law at the University of Chicago, where he has taught since 1972. He has also been the Peter and Kirstin Bedford Senior Fellow at the Hoover Institution since 2000. Previously, he taught law at the University of Southern California from 1968 to 1972.

He has been a member of the American Academy of Arts and Sciences since 1985 and a Senior Fellow of the Center for Clinical Medical Ethics at the University of Chicago Medical School. He served as editor of the *Journal of Legal Studies* from 1981 to 1991, and as editor of the *Journal of Law and Economics* from 1991 to 2001.

His books include *Free Markets Under Siege: Cartels, Politics and Social Welfare* (New Zealand Business Roundtable and Institute of Public Affairs, 2004), *Skepticism and Freedom: A Modern Case for Classical Liberalism* (University of Chicago, 2003), *Cases and Materials on Torts* (Aspen Law and Business, 8th edition 2000), *Torts* (Aspen Law and Business, 1999), *Principles for a Free Society: Reconciling Individual Liberty with the Common Good* (Perseus Books, 1998), *Mortal Peril: Our Inalienable Right to Health Care* (Addison Wesley, 1997), *Simple Rules for a Complex World* (Harvard, 1995), *Bargaining with the State* (Princeton, 1993), *Forbidden Grounds: The Case Against Employment Discrimination Laws* (Harvard, 1992), *Takings: Private Property and the Power of Eminent Domain* (Harvard, 1985) and *Modern Products Liability Law* (Greenwood Press, 1980).

Professor Epstein has written numerous articles on a wide range of legal and interdisciplinary subjects. He has taught courses in civil procedure,

vi **Behavioural Economics**

communications, constitutional law, contracts, corporations, patents, individual, estate and corporate taxation, Roman law, criminal law, health law and policy, legal history, labour law, property, real estate development and finance, jurisprudence, land use planning, torts and workers' compensation.

Introduction

Dr Veronica Jacobsen,
principal advisor, Labour Market
and Income, The Treasury

It gives me great pleasure to introduce Professor Richard Epstein. I would like to thank the New Zealand Business Roundtable for organising Professor Epstein's visit to New Zealand.

Professor Epstein is the James Parker Hall Distinguished Service Professor of Law at the University of Chicago, where he has taught since 1972. He has also been the Peter and Kirstin Bedford Senior Fellow at the Hoover Institution since 2000. Although his training is in law, Professor Epstein's grasp of economics is phenomenal. He is an eminent scholar in law and economics, and has many other fields of interest. He has written innumerable articles and many books on a vast range of legal subjects and has taught courses in almost every branch of law.

Today, Professor Epstein will discuss the relatively new area of behavioural economics which questions some of the underlying assumptions of rationality in classical economics. This topic is also discussed in two chapters of his latest book, *Skepticism and Freedom: A Modern Case for Classical Liberalism*.¹

¹ University of Chicago, 2003.

Behavioural Economics

'We are all behaviouralists now'

There is a natural cycle in intellectual development. New ideas begin as isolated and idiosyncratic attacks from without. At first, the establishment treats them with scorn. But, as the ideas persist, it is forced to confront the pesky intruders on their merits, which the establishment does with ingenuity and determination. But then the tide of battle turns, and the former outsiders become the new orthodoxy.

One vivid illustration of this process took place in the arcane field of macroeconomics. Defending deficit spending against the conservative charge that it was 'Keynesian', President Richard Nixon is reported to have replied, "We are all Keynesians now". In recent times, a similar transformation has occurred in microeconomics. It seems there are no rational choice theorists anymore: we are all behaviouralists. Unfortunately, there is no clear vision of what behaviouralism consists of, nor any real agreement on how much of the traditional model of rational choice theory has survived, or why. In this lecture I hope to identify the Achilles heel of rational choice theory and discover if there is a theoretical structure that allows our understanding of human conduct to incorporate new insights into a behavioural programme that does not content itself with taking ad hoc potshots at the traditional rational choice theories.

The impulse for an improvement of economic theory sometimes comes from strange quarters, often in disciplines that seem to bear scant relationship to the problems at hand. In this instance, an understanding of human evolution – the field once known as sociobiology but now travelling under the sanitised label of evolutionary psychology – supplies the theoretical foundation

2 Behavioural Economics

for why we should expect consistent and predictable deviations from the standard models of rational choice theory.

The classical version of rational choice theory is vulnerable to attack on two fronts. First, the theory is weak in its ability to identify the choice of human ends, especially as these relate to the theory's strict assumption of strong individualism. Second, rational choice theory is vulnerable in its articulation of the means that people use to achieve their chosen ends; just how good at implementing our preferences are we ordinary human beings when we make calculations, probability judgments and decisions? A sizeable body of literature suggests individualistic assumptions and rational calculation assumptions do not hold water. But where do these assumptions break down, and how can we try to reconstruct the system? It is one matter to destroy a theory but another thing altogether to replace it with a richer and more serviceable framework.

How should we understand the assumption of radical individualism that is so central to rational choice theory? In other words, when an individual considers any action, the theory assumes the person asks one and only one question, "what is in this for me?". The emphasis is solely on the benefit to the individual actor. Taken to the extreme, rational choice can be associated with just one psychological approach to the welfare of others: complete indifference. In some cases, this seems to make pretty good sense. Buyers of goods, for example, know that any decrease in the price they pay will hurt the seller of those goods. Yet they push hard for a lower price. The same considerations apply in reverse to the sellers. But no one regards the behaviour on either side as worthy of reproach. The entire institutional framework that requires each to obtain the consent of the other before the deal goes forward protects against short-term opportunism. And the desire to preserve goodwill to grease the wheels for future transactions frequently leads both sides to be attentive to the needs of the other. The use of strong individualist assumptions about profit maximisation does not raise any hackles within this bounded environment.

The question is how well the same approach generalises when we move to non-market forms of behaviour that are not subject to these critical if implicit

institutional constraints. In this setting, the assumption of indifference to the welfare of others has chilling implications, because it makes it appear that each of us does not care whether others live or die. The only important matter is how we can benefit ourselves. Accordingly, negative consequences to third persons associated with an individual's actions can only be brought home through some form of sanction by others, for example, through private retaliation or social redress in the form of criminal prosecutions, tort actions or injunctions. This model leads us to expect that individuals would happily engage in murder against all sorts of people if they believed they would not be caught and punished. There is no sense that people tend to maximise welfare within a set of norms that include a respect for the physical integrity of others. There is no clear effort to say that people who bargain hard within the rules on price typically do not deviate from them, even if they thought they could get away with it.

Scholars who work with the unreconstructed rational choice theory have been forced to produce speculative and unlikely explanations of why all humans do not go around killing each other. The purported reasons include the indirect reputational effects of murder and the possibility of detection and conviction. But if that is the fear, then truly rational actors will also routinely lie and destroy evidence to cover their tracks, thereby compounding one type of wrong (force) with a second (deception). Most of us struggle to overcome our disbelief with these explanations, not least because the psychological literature describes any rare individuals indifferent to the fate of others as 'psychopaths', not as 'rational actors'. Economists can be unfairly caricatured as expecting everybody to behave in a fashion that, in real life, would see large portions of the population locked up and subjected to medical treatment, assuming we could find some benighted individuals whom we could entrust with the power to control the lives of others.

Rational choice theory runs into many of these difficulties when, as noted above, it uses the plausible assumptions for explaining the terms of voluntary exchanges in competitive markets to offer a comprehensive account of all aspects of human behaviour. But this tactic is a dangerous form of overgeneralisation. The psychological elements that do not seem very important

in the first context become absolutely critical in the second. So what can be done to salvage some rationality when it comes to the preference set of particular individuals? Many people want to find a way to soften the Hobbesian model of individual choice and replace it with something more in tune with the way that human beings really form their individual preferences.

The history of behavioural science

Viewed in the larger scheme of political thought, there is nothing new about behavioural science. David Hume wrote *A Treatise of Human Nature* in 1740. In it he tried to identify the condition of mankind that makes sociability possible. Hume seemed to understand intuitively that if the psychopathic model of rationality were correct then even Hobbesian manoeuvres intended to achieve a fragile social truce would quickly become unglued – intrigues and treacheries would undermine efforts at cooperative behaviour. Aware that he was deviating from tradition, Hume tried to build a system of rational choice that incorporated the reality that an individual's welfare is largely dependent upon the welfare of others. This led Hume to formulate the concept of “confined generosity”, which allows the welfare of other individuals to be part of the decision calculus of all persons. How this functions is a subject of some debate, but a couple of generalisations help pave the way. To those who are closer to us, “confined generosity” often requires that we provide them with financial assistance and moral support when they are down on their luck. For those who are further removed, positive assistance may be forthcoming in lesser amounts, or not at all. But there is still a strong sense that no one should kick, quite literally, those who are down on their luck or take advantage of their position. It is not good form to attack a defenceless individual, and even if some people will do it because they will not get hurt, many more will lend a helping hand, without any expectation of recompense.

These customary behaviours speak volumes about our accounts of human nature. More concretely, abandoning the radical individualist model means accepting that people avoid committing murder for reasons that are quite independent of the sanctions that might be imposed on them. Equally, people

in general are not indifferent as to whether they get \$100 through theft or trade. The welfare of others will be taken into account. Most people find this a much more comfortable model – after all, we see examples of Hume’s “confin’d generosity” in practice every day.

In making this general assertion, it is important not to become dogmatic in the opposite direction. Because most people will not take advantage of the helpless does not mean that everyone will follow that laudable path. In examining the role of egoism, or indeed any other human trait within a population, it is important to note that there is great variance around a median. An odd feature of rational choice theory is that it counter-intuitively assumes that the variance around the median in the model of self-interest propounded by Hobbes was precisely zero. We were expected to believe that each individual was exactly the same. In contrast, Hume did not assume that everybody has exactly the same amount of empathy as everybody else. Variation – higher or lower than the median – is an ever-present social feature, and that variation generates new possibilities for voluntary exchange and cooperation. Individual character will influence the kind of market, occupation and personal relations that each person will enter into. By assuming homogeneity on this single point, rational choice theory reduces its explanatory power. If the theory were in fact correct there would be no occasion for voluntary or involuntary sorting by personality type. If, however, there is some degree of variation with respect to a trait like egotism, then some sorting mechanism turns out to be thoroughly appropriate. In this case, we can start to explore which individuals move in particular directions as a function of where they lie on the distribution with respect to any given trait. Those with more empathy are more likely to become social workers or teachers. Loners are more likely to become hunters, and so on down the line.

Hume’s model, which predates modern behavioural economics by about 250 years, provides a fairly accurate rendition of what happens in real life. To this extent, behavioural economics should be welcomed; it relaxes an assumption that is so artificial and counter-factual that it discredits economics in the eyes of the public.

Evolutionary psychology

Accepting there is variation around a particular human trait is one thing. However, it is much harder to explain why the trait of sociability would emerge to begin with. I believe evolutionary psychology and the biological bases of cooperation may provide some answers. William D Hamilton wrote a famous article in 1964 in which he put forward a genetic basis for altruism. The principle he advanced was one of inclusive fitness. This principle fills a gap in the Darwinian accounts of evolution by holding that we should not be understood as sentient human beings with eyes, ears and noses but rather as vessels for the transmission of genes across generations. Human beings, like other living organisms, as it were, are unwitting participants of a fanatical plot by our genes to ensure that their distribution and frequency in the population are maximised across all generations.

There is only one way for a gene to replicate and that is for a person to have children. Unless you accept there is some empathy between parents and a child, you will never be able to explain the life cycle of human beings or any other kind of conscious organism. Parent-child interactions are not structured in accordance with the standard norms of rational choice theory. The care given to helpless infants cannot be explained by the use of this model.

Let me give two illustrations of this point. In the early 1980s, my wife Eileen and a Nobel laureate in economics, George Stigler, had a New Year's Eve discussion about why parents care for their children. Stigler, who worked within the individualistic model of rational choice theory, said that a bargain clearly existed: parents take care of young children, who later take care of elderly parents in return. Therefore there is a mutual exchange of benefits with a temporal dimension. When I later pondered this explanation, I came up with obvious objections: because of the time value of money, which requires benefits to be discounted, and because the amount of care in the two cases is uneven, there will be a vastly asymmetrical exchange of resources. My wife, however, had a better riposte. She simply asked: 'Why does any parent take care of a seriously disabled child?' It seemed Stigler had never considered that question, and it was one of the few occasions on which I saw him at a loss

for words. The empirical test is whether all parents of disabled children leave them to die because there is no reciprocity of benefits. That does not occur. What exists in practice is a complicated set of calculations regarding the level of investment parents make in their needy children relative to other demands on their resources, including those of caring for their other children and preparing for their own retirement. No exchange mechanism can explain what happens with the disabled child who can give no material support in return.

The second illustration concerns Gary Becker, another Nobel laureate in economics and a champion of rational choice theory. In a 1974 paper titled 'A Theory of Social Interactions', Becker developed a model of cooperation to explain why people care about the utility of others. He did not simply adopt the Hamilton model that focused on the level of cooperation between parents and children. Instead, Becker developed a model in which the parents cared about the children but the children did not care about the parents or each other. He suggested parents secure cooperation between children by removing resources from a misbehaving child and providing them to a better-behaved sibling. In this way, the parent establishes some sort of cooperative equilibrium between siblings who would otherwise not cooperate.

I remember the Center for Advanced Studies meeting in California in 1978 when this paper was discussed. Oliver Williamson pointed out that Becker's model was deeply flawed – the obvious response for a child would be to kill their siblings. Other children would no longer compete for care and resources, and the parent would lavish everything on the sole remaining child because the parent's well-being depended on that of the child. Yet fratricide is not a standard form of behaviour. There is not only cooperation at a genetic level between parents and children, but under Hamilton's view, also between siblings. Brothers and sisters share half their genes, so an individual would have to gain twice as much as they lost to make murdering a sibling individually rational. In ordinary circumstances, that is very unlikely to occur. However, if succession to a throne is at stake, fratricide may well be a rational response because the two-to-one ratio – or something much greater – could be satisfied, at least for some tiny fraction of individuals who are outliers in the distribution. The important thing to

understand about biological constraints is that they can be violated. They are not absolute, unmoving barriers against certain antisocial actions.

Under the Becker model, if parents die there is no regulator to adjudicate between a pair of siblings. Yet in practice, we rarely see orphaned siblings behaving like strangers towards one another. Rather, they often form a partnership and continue the same business that was previously run by a parent.

The point is confirmed historically. The Roman law of partnership did not begin with partnerships formed by strangers but with those of two children who had been subject to the authority of the same *pater familias* – a kind of family ruler. On the death of the father, the two children would form a partnership to continue the enterprise begun under his aegis. The term ‘going concern’ reflects the desire to keep a business going. To this day, the family partnership is seen as a highly important social institution that thrives because the level of cooperation between close relatives creates less need for sanctions or elaborate rules about control than are required when strangers pool resources. Conflicts of interest still exist but are likely to be less serious than in other settings.

Cooperation within an ethnic group

How far do bloodlines continue to matter? The theory of inclusive fitness does not stop with parents, children and siblings. It explains a certain level of consanguinity within members of a given clan or tribe. Some degree of cooperation and loyalty is to be expected because the genetic connection means relatives will find it easier to form social bonds. Reneging on a transaction is then more costly because the harm to another relative is a harm to yourself, and to others for whom you care.

Sometimes there will be uncertainty as to how far kinship extends, or what the connection with another individual might be. If two people have a general tendency toward cooperation, the prisoner’s dilemma game of mutual defection will not necessarily apply because intrinsic reflexes will incline both individuals to work together on at least a semi-cooperative basis, at least most of the time. This brings us back to what Hume talked about – confined generosity – or, in Adam Smith’s case, in *The Theory of Moral Sentiments*,

empathy. The level of interest in the welfare of other individuals within certain familial groups is high enough to encourage cooperation so that the rigid individualistic assumptions of the rational choice model do not hold.

In practice, there is a bad side to altruism that must also be taken into account. Let us suppose that I have strong connections with people in my own original tribal group. If I engage in hostile activities against another group, my altruism could lead me to be *more* destructive, not less, because, as a selective altruist, I am protecting not just myself, but also my clan. A sad implication is that conflict between competing groups with strong internal connections but few external links is likely to be fierce. It is amazing how many of today's wars and exterminations can be attributed to differences between ethnic groups. A large part of what civilisation turns out to be is the capacity to curb impulses that point in one direction and allow sociability to point you in another.

Implications for the marketplace

What implications does the model of limited generosity have for ordinary markets? It makes a system easier to operate successfully than if it were run on strictly Hobbesian lines. Exchange is facilitated. Sellers will be willing to take a little less and buyers to pay a little more if individuals care about the welfare of their trading partner. The movement need not be confined to the price term, but could easily extend to non-price terms where cooperation makes the deal easier, on such matters as date and place of inspection or delivery, for example. But no matter whether the cooperation moves in one dimension or many, overall the bargaining range will be greater, which means the class of successful trades will be larger. The risk of defection is reduced.

However, the model also has other implications that are harder to classify. With a strong rational choice theory, the social impulse in favour of redistribution through either voluntary means or state coercion does not make sense. But when the degree of commonality is taken into account, redistribution will be seen as less costly because people are giving to somebody whom they would like to see with more wealth. An individual's utility is now heavily and positively dependent upon the wealth that others will enjoy. The redistributive impulse may come at the expense of wealth creation.

Redistribution schemes also tend to work better in homogeneous populations than in situations where distinctive groups in the population regard themselves as relatively isolated from each other. Other factors must be taken into account, but no matter what complications are raised, relaxing the standard rational choice model means sociability – which originally seemed impossible – now turns out to be possible.

The model seems to explain many empirical phenomena. One is a game that is discussed *ad nauseam* in the economics and anthropological literature. I have \$10.00 that I will give to Veronica, provided she is able to buy off a second player, Steve, with some fraction of that sum. She must simply decide how much to give Steve, which could be anything from nothing to \$10.00. The standard rational choice theory would lead us to expect a 9:1 split every time. The outsider is better off with one than zero, so why should any rational person pay the outsider any more? Yet when we look at the data, we find that a 9:1 distribution only occurs among nomadic tribes with little cooperation. For everybody else, the median amount distributed turns out to be somewhere between \$3.00 and \$4.00.

Clearly, something more is at work, and a sense of self-respect may be part of it. But looking at this contrived example in the abstract makes it difficult to draw conclusions: there is nothing in this institutional setting that would place a lower-bound constraint on how much you must give somebody to secure their voluntary cooperation. The familiar maxim says nobody will participate in a venture unless the expected return exceeds their expected cost. But here we have an oddball example that deals with distribution without any implications for production, where most of our instincts about fair shares have been honed. So, when would most people in Veronica's position give either \$3.00 or \$4.00 to somebody else?

The better explanations are actually consistent with rational choice theory. Once you recognise the interdependence between two people and realise that the sentiments of sharing and reciprocity really do matter, it is clear that Veronica is likely to make certain calculations. She takes into account the personal esteem and self-respect of Steve, the person she needs to buy off. She realises that if she offers him \$1.00, there is only a slim chance that it will be taken. As she offers

more money, she stands to receive a smaller return, but secures a higher probability that the deal will proceed. Steve will usually not demand \$5.00 because that calls for the outsider to have equal gains from the venture. But backing off from that focal point, it is easy to see why the equilibrium position turns out to be \$3.00 or \$4.00. It is important to understand the embedded social constraints that define the problem, and then assume that, for the most part, people will get their calculations right, knowing what they know about the actual temperaments of the people with whom they deal.

Distribution of ability

This brings us to the second half of the question, concerning the connections between means and ends. Deep psychological insight is not required to understand that economists are wrong if the rational choice theory means everybody gets a perfect score in every examination they take, regardless of when they take it.

My colleagues Cass Sunstein and Richard Thaler are eager behaviouralists who invoke such terms as libertarian paternalism (which I regard as an oxymoron) in an effort to explain why the behaviouralist view upsets a normal understanding of how voluntary exchange works. They say that people are clearly not rational because everybody does not play chess as well as Gary Kasparov. In a way, this is an appealing insight because it means no one is rational enough to be a (perfect) government regulator, even if they are not rational enough to be anything else. This definition of rationality, associated with the ability to calculate, has a single-valued expression. The question you must ask is what to do when it turns out there is a distribution, either through nature or through training, in the relative abilities that people have? Relaxing assumptions of perfection makes it easier to explain the way markets work. In fact, easing the assumption of universal competence with respect to the ability to calculate and decide allows us to understand the emergence of many market institutions.

Let us look closer at the chess game. If you are a great player, you might offer your services to teach amateurs. This is perfectly sensible – a voluntary exchange would occur, in which the knowledgeable person is paid to train the novice. For this to work, novices simply need to know that they cannot play

chess at the level of masters. The inexperienced player only has to understand that there is a difference between their own abilities and those of the teacher, and believe that the teacher has the credentials and skills to teach. This model does not by definition preclude irrational behaviour. An inference of that sort would be defensible if we could find cases where somebody who has never played chess is able to charge money to teach a grand master. That improbable state of affairs would mean people who know nothing are being paid to spread ignorance by people who already have sufficient knowledge. Once you introduce the scope for *differential* competence, a rating system that establishes how much the differences matter becomes absolutely critical. It turns out that many social institutions deal with such problems.

Generally speaking, chess as an activity will not be fun if the master clobbers the amateur in every game. What is required for virtually every tournament is a rating system of sufficient reliability that each match will have maximum uncertainty because people are grouped according to ability. If you accepted an assumption of omniscience you would never play chess. The first time you sat down at the board you would be able to figure out all the permutations instantaneously and announce: 'I've looked at this board and a draw will occur once each of us follows our best strategies, as we will surely do. We had better try another game because this one is no more fun than tic-tac-toe'. But rationality means the following: if you know that you are in a state of relative disadvantage along a knowledge dimension, or indeed any other, and you have an opportunity to improve yourself, you will try to estimate whether the cost of improvement is worth it. This theory does not predict that you will not make any kind of mistake. A theory of rationality should embody differential competencies so that it can co-exist with theories of learning and sorting.

Heuristics

This seems to have enormous implications for the way one considers the entire body of law associated with the literature on heuristics. That literature is highly influential. Believers in behavioural economics fall into two camps. One group, led by Danny Kahneman and the late Amos Tversky, treats the phrase 'heuristic' (mental shortcuts that lead to decisions) as a synonym for the word

'bias'. Their articles always have titles like 'Heuristics and Bias'. Kahneman and Tversky say most heuristics work well but a few do not. Yet I cannot recall them providing a single illustration of a heuristic that works that they would be prepared to endorse. Rather, they prove that none of us is a Bayesian who iterates towards a single conclusion. We all get taken in by the availability heuristic or the representativeness heuristic. We all have hindsight bias. We do not know how to work out probabilities and we make a mess of things. Reading their work, one would be surprised that this building is still standing, and that you managed to arrive on time to hear this lecture.

The other side of this debate is headed by Gerd Gigerenzer, formerly of the University of Chicago and now at the Max Planck Institute for Human Development in Berlin. The title of his book is not 'Heuristics and Bias' but *Simple Heuristics That Make Us Smart*.² This represents the flipside. If the first group only showed the highlight reel in which the baseball players made error after error, Gigerenzer's film would show spectacular catches and pitches. We are supposed to believe that these two non-representative samples capture all of reality, yet clearly neither delivers on that promise. But, as the author of *Simple Rules for a Complex World*, I believe Gigerenzer does have a point, which is that many of these simple rules are pretty reliable. We have survived as a species and as individuals because we have relied on heuristics. We are better at what we do than we are at describing how we do it.

Let me offer a couple of simple examples. One of Gigerenzer's favourites is this: we see a ball in the air, and we do not have time to work out consciously all of the equations that allow us to tell where it will land. Yet humans have figured out that by keeping our eye on the ball, we can catch it or duck to avoid it. How do we do this? It is not a higher-level, conscious phenomenon. In nature there are many animals that cannot talk, including cats and dogs, with excellent skills of coordination when faced with this task. A dog can catch a frisbee just as well as humans can. In common with these animals we have a built-in perceptual network called the angular gaze. It relies

² Co-written with Peter M Todd and the ABC Research Group, published by Oxford University Press, 2000.

on a simple heuristic that tells us that if an object moves along a curvilinear path and we keep it at the exact same point in our field of vision, we will eventually meet the object. Aircraft pilots understand that if another plane is moving toward them and remains at the same spot on the windscreen, they will collide. Baseball players also demonstrate this phenomenon every time they catch a ball. Fielders do not move in a straight line to make a catch but in a curvilinear fashion because this makes it easier to maintain the angular gaze. This particular movement is called circling under a fly-ball.

What do we gain from a heuristic that we do not achieve by arriving at a fully worked-out solution? We catch the ball. But what do we lose? We do not know where we will catch the ball, because we are not solving the equations, just completing one particular task. Many baseball careers have been ruined by people circling under fly-balls and crashing into fences or each other because they followed this heuristic without paying attention to the physical constraints of their immediate environment.

How does this fit with the larger questions of institutional economics? Rationality tests that show people perform relatively poorly follow the same pattern: people uneducated and untutored in a particular specialty are run through a set of tests. In some senses their failure is foreordained even though they have every incentive to strive for the right answer and long before any issue of cognitive bias arises. With such a high level of individual failure, how do we manage to get the trains to run on time? Institutions have very different thresholds for success and failure. If all you are considering is some calculation bias, take solace in the observation that you do not have any motivation to deny the force of the truth once it is learned. Assume that a group of individuals work on a particular problem and one learns how to solve it. At that point, go back to what differential competence means: I may not be smart enough to prove Newton's fundamental theorem of calculus but I am smart enough to figure out that it is correct once somebody else has done so. The person who solves the problem must be able to persuade the others that he or she has gotten it right. At that particular point the cognitive bias disappears for the entire group because the good information has essentially driven out the bad information, in a kind of reverse Gresham's law.

Implications for organisational design: the case of contract defaults

One of the key principles of sound institutional design is to lodge the power of decision making with those individuals who are most likely to get the decision right. The implications of behavioural economics are that everyone we know is a bit less competent in this regard than we might have supposed. The question is whether this sombre news should influence our choice of organisational structure. It does. The general social preference for decentralised markets is in fact strengthened when one introduces the problem of imperfect knowledge associated with the standard cognitive bias theory. If a strong hierarchical system exists, what chance is there that the leader of the pack will get the answer right? Supposing the person is chosen by meritocratic principles, the chance is relatively high. If the average person gets it right 50 percent of the time, the leader may get it right 75 percent of the time. How many individuals are needed to work in parallel in order to reach that level of performance? Just two. If each of them has a 50 percent chance of getting it right then one of the two will get it right 75 percent of the time, and the other will acquiesce once shown the error of their ways. And if up to 10 or 12 people are involved, the probability of getting it wrong will be microscopic. We like redundancy in institutions precisely because we understand that the rational choice model, which assumes instantaneous and correct calculation ability, is an incorrect model of human behaviour.

This modestly cheery view of human nature helps address some of the practical issues that are raised by behavioural economics. Here are two contentious areas – one dealing with the term of service and the other with pensions – both of which relate to the employment context, and involve the role of default terms in dealing with human arrangements. Recall that any voluntary agreement must cover a large number of contingencies, some of which occur with only a very low probability. One easy way to reduce transaction costs while respecting the freedom of contract (and the prospects for gain that it holds out) is to set the default provisions in ways that anticipate what most parties would do if they addressed certain topics specifically. The fewer the terms that have to be flipped in negotiations, the easier it is to

complete a deal. And if the individuals in question have some confidence that the persons who set the defaults grasp what parties commonly desire, then the default provision need not function merely as a gap-filling device. It can also take on an additional educative function – thoughtful people who have pondered this issue tend to embrace this solution. As a beginner in this field, you might be wise to rely on the default. As a pro, you might decide to deviate from it, but only for some good practical reasons.

This sunny view of defaults is often at odds with behavioural economics. Thus, on the first topic of the contract at will, many behaviouralists worry that ordinary individuals will underestimate the probability of dismissal and therefore will sign on to an at-will contract when they would be better off getting additional protection for taking the job. The use of a for-cause default provision announces to these workers that they could enjoy some real measure of presumptive protection, while leaving open the ability to move to an at-will agreement.

Right at this point, we have to be careful to understand two different versions of this proposal, one of which turns out to be more interventionist than the other. The strongly interventionist view says that an employer's contract offer must come with the for-cause option in the default position, but leave the worker the option to flip it to an at-will agreement. That proposal is in fact a huge limitation on the freedom of contract because it necessarily prohibits the employer from insisting on an at-will agreement by an explicit term in any job offer.

Ironically, a proposal of this sort gives rise to its own cognitive difficulties, here in the form of another prisoner's dilemma game that should leave lawyers and economists of all stripes uneasy. Suppose that each worker is willing to accept the at-will term so long as others accept it too. The intuition in favour of this outcome is that any able worker would like to work at a place where co-workers who prove troublesome can be disciplined or dismissed. But if the deal is so structured that each person has to receive a for-cause offer, then who will flip over their individual contract, knowing that co-workers, especially ones less well-regarded by their employer, are likely to stick to the for-cause option? In effect, the employer who oversees the entire situation will know that

any favourable term that is supplied to one worker must be assessed in light of the implications, positive or negative, that it has for other workers. The decision by the employer, when appropriate, to impose a uniform at-will provision eliminates a lot of gamesmanship that would otherwise take place. The employer knows something about the full situation, has the right incentives to maximise worker welfare consistent with its own profits, and is less likely to fall into the familiar cognitive traps. So if the at-will deal promises higher productivity and higher wages across the board, why then oppose it?

The implicit mistake of the behaviouralist critique on for-cause employment contracts is to assume that the errors in probability estimation apply only to one risk, that of personal dismissal. But in fact, workers can make other errors as well, including the risk that inferior co-workers will have on their performance and prospects for advancement. The strong employer presence here goes a long way to dampen these various behavioural errors by substituting the judgment of a relatively experienced long-term player for those of the naive players. This strong version of the default rule – which allows only the worker to flip the for-cause into the at-will option – has to be a mistake.

What about the weaker form of the same rule, which says that in the event of joint silence, the employment contract should move forward on a for-cause basis? Here the option of the employer to insist on at-will contracts from day one is protected, so the only question is which default option leads us to the optimal contract via the shortest path. The practical experience in both the United States and New Zealand is that even the whiff of a for-cause protection brings forth a strong, bold-typed insistence on the old at-will arrangements. So why make the path more tortuous and the result less clear?

The standard employer's decision, moreover, is not taken idly, but emphatically. Recently, I had a conversation with an experienced director that went like this. Question: "Well, what would you do to hire a worker for a long-term project which may not work out." Answer: "Hire the person on an at-will basis where there is no penalty for dismissal during term." Question: "But would you pay the penalty amount, or some fraction of it, if things did not, in fact, work out?" Answer: "In most cases yes, because we care about our reputation. In extreme cases, the answer might be no." Question: "So why not

make the agreement track social expectations of presumptive behaviour by using for-cause language?” Answer: “Are you crazy? We would never put ourselves at the risk of the judicial system.”

There is nothing distinctive in that last response. The social equilibrium involves two tiers: a legal tier in which dismissal is without cause; and a social tier where it runs the other way. Indeed, behavioural economists should warm to this solution because it indicates that contracting parties are well aware of the rich underlay of what is sometimes called ‘social meaning’ in business relationships. Sophisticated employers are just emphatic in insisting that these two layers be kept apart. I see nothing in behavioural economics that upsets the systematic bias in favour of the contract at will.

Pensions raise other issues. Workers here have to worry about long time horizons, and decisions about making critical asset allocations for their future when it is quite likely, given their differential competence, that they are unclear about the difference between a stock and a bond. Plainly, no one in this context can argue that individual workers have full competence to make their investment decisions. They are overwhelmed by thousands of possibilities, without the time, equipment, and education to sort through them. But it hardly follows that they should be the subject of extensive state regulation. The usual experiments that highlight mistaken pension decisions involve individuals who are left to their own devices, often without third-party assistance.

Any employer who recognises this situation will supply that intermediary. Employers may well take weak paternalist positions by requiring workers to set aside a certain fraction of their income into pension funds, as is done without exceptions at the University of Chicago. And they could structure the defaults so as to make clear that some mix between stocks and bonds is optimal, while limiting the number of financial institutions with which they can transact. Both of these have been standard policy at the university as well. Employers need not insist on rigid default options, but could also send out sensible literature that explains the available choices so that individuals can tweak the basic programme further. And they could advise, and to an extent subsidise, some employees to get personal financial advice – a booming business – to integrate the financial planning for employee pensions with other

employee assets. In short, employers can, and typically do, engage in a wide range of practices to help their workers. The increased prominence of the behavioural critique is likely to spur private action long before any regulation, good or bad, could kick in.

There is, most critically, no conflict of interest on the pension funding issue. Any increase in the effective value of the pension is tantamount to an increase in salary, which works to the benefit of both sides. In fact, more firms are consciously using this approach, in sharp contrast with the debate over at-will versus for-cause contracts, where the firm has a decided stake in which arrangement is adopted. I see no reason for regulation to mandate the behaviours that higher levels of self-knowledge supply. Here is a case where contract and education fill an admitted void in the capacities for individual decision making.

Indeed, there are real dangers with regulation that are evident in the long-term protections that are supplied in the United States with both Social Security and Medicare. These programmes often speak about the dangers of 'voluntary' choice, but, by the same token, they offer no third-party help to the covered individuals. In addition, the dominant feature of these programmes is that they are not tailored to deal with any competence issues; in fact, they are often structured to promote forms of redistribution, frequently opaque, among plan members.

Here is the key difference. A simple form of paternalism that seeks to help workers over the humps would take the form of market pension and health plans. Individuals would have their own, and there would be no cross-subsidy between plan members. However, both Social Security and Medicare use the concerns with individual competence to shield a second agenda, which is to have subsidies between employees of the same generation – the rich underwrite the poor – and subsidies across generations – the younger underwrite the older. The net effect of this second agenda is that the cross-subsidies block any real substantive reform. For Social Security, private accounts will not work if the lion's share of the individual contributions are slated to pay other workers. For Medicare, privatisation would require the most elderly to pay much higher premiums than those enrollees who have just

turned 65. So long as markets can deal with the issue of limited competence, they should be preferred to state programmes that seek to do two tasks simultaneously. Here, as ever, the case for redistribution should be made out on its own merits, and not smuggled into programmes that have other purposes. If there is some degree of social empathy, it will find expression in the political system.

Modest conclusions on means and ends

Overall, I believe the correct approach is to understand the limitations of rational choice theory insofar as it deals with matters of ends and means. On ends, it is critical to recognise the vital role that empathy for the position of others plays in ordinary human affairs. Absent state coercion, the level of voluntary redistribution, both within and across tight-knit groups, is apt to be high. On means, we should openly acknowledge that individual calculations are often erroneous and unreliable, but take some comfort that self-knowledge leads people to that conclusion at least some of the time. At that point you will be able to explain how institutions are designed to exploit differential abilities in decision making and empathetic capacity. Occupational choice and everything else will follow.

This is a completely different conclusion from the standard view of behavioural economics, which is that once rational choice theory is found wanting, traditional market institutions will not work and some system of regulation must supplant them. I have tried to think of cases in which this might be true, and indicated two in the employment context where it is likely to prove false. Indeed, if my analysis of default terms can be generalised, I am hard pressed to find one in which the preference for a classical liberal approach is reversed by the insights associated with behavioural economics.

The first of the two failings of the rational choice model is met by introducing empathy, which essentially makes understanding behaviour a little easier. The second, imperfection with respect to calculation and judgment, makes it a little harder. Therefore we must find ways to exploit the first and remove the impediments associated with the second. We are generally able to do so. We do not travel with the speed of the gazelle or have the accuracy of

a hawk; we are not unerring in everything we do. Nor, however, are we confronted with a random, hit-or-miss situation. Powerful institutional arrangements tend to accentuate the positive and downplay the negative. The case for the traditional distribution of functions between the public and private sectors that I have always defended actually looks stronger with a more realistic set of assumptions than those associated with rational choice theory.

Hobbes does one thing for us, while Hume and Smith do something else. Hobbes explains the way the austere model works. When we introduce benevolence we depart from Hobbes, and Hume and Smith help us to appreciate the parameters of how imperfect calculation works. Together, these three giants assist us to create a basic framework.

Our main focus should be on designing institutions to ensure that those who tend to get the right answers are the decision makers. It is reassuring to observe that as the stakes get higher, the quality of selection processes and the level of performance also rise. Children trying to figure out mathematics problems may well get 50 percent of their answers wrong. But when it comes to people such as currency traders at Goldman Sachs, the ratios are different. They may make mistakes but I can assure you they are not making the errors that are revealed in Kahneman-Tversky rationality tests. Even I will not fall into the same trap twice.

Questions

Are there any areas where the limitations of rational choice theory might lead us to consider some form of regulation? One possibility that interests me is takeover regulation. In a takeover situation it is sometimes argued that the premium for control should be shared. I struggle with this idea, but might behavioural economics shed some light on its persistence?

I think that if this reflects anything, it is that utility functions are not perfectly independent. In a corporate venture in which we are both involved, you have to take account of my welfare and I have to take account of yours.

Leaving aside the takeover situation for a moment, let us consider a simple venture in which 10 people participate. There is a controlling block of three people. A new person enters the venture knowing of the existence of the control block. The new entrant might wish to bargain for certain kinds of minority protection. I do not believe there are any grounds to block such a transaction.

Consider the new entrant's point of view. This person might only make a 20 percent return, while the three members of the controlling block are making 25 percent. But the new entrant's next alternative, going into the open market, might only be a 15 percent return. Therefore, this individual will be willing to tolerate the control block situation, remembering that the controlling shareholders will be working to make the project a success. You must take into account both capital and labour in terms of what each person contributes to the firm, and given their efforts you may find that returns will roughly equalise. I believe that if you consider the situation from an *ex ante*

perspective there can be no objection to the arrangement, unless of course some kind of fraud or misrepresentation is involved.

What about simple mistakes? It can be argued that if there is a mistake (which means, in legal terms, an error in understanding facts, the meaning of words or the law) there is reason to believe the assumption of mutual gain from voluntary exchange is not going to be satisfied. In these circumstances re-contracting might take place. However, a mistake is different from a case of misrepresentation. In the latter case, the error is deliberately induced by one party. That changes the dynamics. It also reduces the number of cases one would wish to set aside, because there are at least 100 mistake cases for every case of misrepresentation.

Returning to the starting point, should a unilateral mistake (not induced by the other side) made by somebody with bad judgment mean that the transaction should be set aside? You have to consider the fact that there is no mutual gain in these cases, and weigh this up against the possibility that a person pleading a mistake is engaged in *ex post* opportunism. An investor may simply regret making a particular investment and wish to move money elsewhere. In such circumstances, claiming a mistake is simply a cheap way to back out of bad decisions. The higher the volatility in the market, the more one suspects that particular defence would be used.

Those actually involved in writing contracts typically take a narrow view of the sort of things that would allow a transaction to be set aside. In both England and the United States, a lot of discussion takes place around whether parole evidence should be allowed in looking at whether to vary the terms of a written agreement. Many commentators say it is unjust to allow an agreement with a collateral term to be enforced as though it did not exist. They favour open-ended inquiries. But evidence on contracts written in the market tells a different story with the so-called 'merger clauses'. These are clauses that say that a particular agreement shall be full and final and that nobody may introduce a collateral term to alter anything in it. The reason for this practice is that the contract-makers have made a strong judgment that the dark side of human nature – not the positive – would in fact undercut these transactions. They do not want that to occur. Courts then have to decide whether the stricter

view of the world should yield to the *bona fide* view, which would allow contracts to be unravelled.

There is a very interesting market illustration of this point. When English courts enforce international salvage and arbitration agreements, they take the following view: you write it; we enforce it. German courts at one time had a different approach: you write it; we will superimpose our own views. In other words, they would enforce the law in accordance with notions of good faith and unfair dealing as derived from the civil code. There was an outflow of business to English. What did German firms do? There was jurisdictional competition. They switched to the English model in order to get a greater share of the business.

I think control block issues are similar. There are many arguments for parity of treatment that are made on the basis of intuitive ideas of fairness. What often happens is that people who are not in the particular business can perceive the fairness argument but not the misallocation of resources that follows from the enforcement of their notions of fairness. When individuals get involved in these transactions over time, they realise that fairness is a very expensive good. As a result, they tend to prefer a Pareto improvement to a fair return. In other words, a return of 25 percent to three shareholders and 15 percent to a fourth is better than a 10 percent return across the board.

A shortcoming of the Kahneman-Tversky school is not to take into account expertise, specialisation and learning in conducting their experiments. An individual can be a whiz at corporate transactions but tone deaf when it comes to playing the violin. Competencies vary wildly across specialisations. As anyone in medicine or law could attest, an individual working in one field for a long period of time will have a huge advantage over an interloper who does not know the ropes. (Incidentally, the expression 'knowing the ropes' comes from old sailing ships. The person who knew the ropes was the sailor who knew exactly how to get to the top of the mast without killing himself. We have forgotten the origins of the term, but its meaning remains relevant.) University professors often think they are smarter than everybody else, which may be true, but in specialised areas experience really does matter, and on many issues we need to be humbler than we are.

You said that when the insights of behavioural economics are properly understood, they confirm that markets work well. What do you see as the implications of behavioural economics for government policy? Are there any?

There are relatively few. In effect, the model I proposed was one of confined generosity, which increases the range of bargains that are struck. It was not one of perfect benevolence. That means one should avoid falling into the socialist trap: from each according to their ability, to each according to their needs. Socialism simply does not work. It also means public choice issues remain important, not because every legislator turns out to be Hobbesian, but because one must guard against the dangers in a few that could subvert the well-being of the many.

On balance, I believe the importance of faction should not be underplayed. At the same time, the ridiculous public choice caricature that everybody who happens to work in a legislature or a bureaucracy has zero civic virtue should be discarded. Adopting this caricature simply offends people who know government officials who work very hard at their jobs. One must understand the difference between the Hobbesian view and the Humean view. A Hobbesian approach is to regard all humans as rogues and set up one devil to superintend the actions of another. This just does not work. The underlying assumptions cause resentment because they do not really match real-life experiences. The Humean view of confined generosity is more plausible. It means we can recognise a distribution starting with people who are totally selfless and ending with those who are completely roguish. We see this distribution in politics and elsewhere. The name of the game is to make sure that the enthusiasts do not kill you with their optimism and that the rogues do not kill you with their pessimism. That means you have to have checks and balances, but ones focused on constraining the tail of the distribution, not on everyone simultaneously.

Let us look at forms of intervention. In the case of voluntary exchanges in competitive markets, regulation that improves the security of exchanges is clearly beneficial. Therefore, laws governing property rights, contracts, frauds and recordation systems to create security of title have always been distinctive

government functions. These forms of regulation are so powerful and effective that nobody even debates them. Often people who pronounce themselves small-government supporters do not know what recordation statutes are or how they ought to function. That longstanding body of statutes remains broadly unchanged for traditional resources, as well as for newer forms of property, such as internet domain names, subject to a little fine tuning.

What about tort law and aggression against third parties? I do not see anything in behaviouralism that suggests that negative externalities are good. From behavioural economics we know that people will be somewhat more reluctant to inflict physical harms, for example, upon others because they may suffer the same harms themselves, and we know that when individuals make calculations they are likely to get them wrong. This does have some implications. It means that when we want to use a rule that enforces boundary conditions between neighbours – the first function of tort law – we want an outcome-based rule, not an input-based rule. A negligence rule is an input-based rule. It usually requires an examination of all the facts and circumstances of the case and a calculation of whether the cost of individual precautions at the margin exceeds the benefits. Every tort trial turns into a mini-administrative hearing. The problem is that people will make many mistakes, especially as jurors. Therefore, you want to look fundamentally after the fact at who was in the right and who was in the wrong. We do not want a case to turn on *ex ante* probabilities in part because behaviouralism tells us that in this area people are likely to make mistakes. If the right rules are in place outcomes will be more predictable. This will lead to higher investment in risky activities without investors needing to examine probabilistic factors.

With respect to property acquisition, sharing rules imposed between strangers are, generally speaking, not appropriate because trying to figure out how to divide the shares of ownership is too difficult. So with property and mortgages, for example, the right rule is one of strict priority unless there is a consensual agreement to set up a pooling arrangement. The implications are clear: keep things simple, let individuals figure out the probabilities for themselves, and take the pressure off the decision maker.

When it comes to antitrust, my view is that everything under the heading of market dominance as a ground for intervention ought to be scrapped, whilst everything under the heading of cartelisation ought to be retained. Every time you look at a market dominance case and find a clever explanation as to why it is anti-competitive you could propose an equally clever explanation as to why it is efficient. Frankly, I would rather not listen to either set of explanations. It is too costly, time-consuming and prone to error. With treble damages at issue, we need a level of confidence in the probabilities of anti-competitive conduct that is not attainable in the real world. Decision-making processes are imperfect and it is better not to head into those rocky waters.

Regarding taxation, I am a supporter of a flat tax scale because it is less complicated than other regimes. The moment you introduce elements of complexity, some people will be able to take advantage of the system better than others. This also applies to things like health care. Should we allow people to obtain a second doctor's opinion, if denied coverage the first time, when they are receiving medical benefits from the state? I am against the practice, because those who obtain a second opinion are the well-connected. Every time such a right is created, someone of lesser means and influence will suffer because they do not have the same access to the levers of power. When levels of ability and access are different, redistributive policies should be avoided.

One of my books is *Simple Rules for a Complex World*, but I am not so naive as to imagine that war or peace could be decided on the strength of simple rules. It is for epic decisions that public leaders are indispensable. However, there are many ways in which problems that look as though they could only be handled by gifted academics could be turned into routine practices. For example, in the world of medicine my aim would be to replace many of the specialist personnel with people who have the skills of bank tellers. I want to make sure the basic structure is right at the administrative level so more can be achieved with a high degree of reliability by lower-skilled individuals. The management lesson is clear: get the protocols right so that many things can be handled by way of delegation. If the protocols are wrong, uneducated people will make all sorts of wrong decisions at a decentralised level. With the right

protocols you do not want people romancing on their own unless they are extraordinarily skilled. Checks are required to ensure the right protocol is in place, but usually a unilateral deviation from standard protocol by line personnel is simply asking for malpractice.

Have people changed psychologically? Is behaviour different today from what it was in the past?

There is a variation in behaviour over time, just as there is variation at any given point in time. High levels of cooperation are required in a small and insular community. Consider the traditional rural practice of building a barn where neighbours help each other because the task is too big for anybody to do alone. In larger cities there is less interdependence and an increased opportunity for market-type transactions. The level of confined generosity will fall because the local monopoly situation is replaced by competitive supply. Instead of waiting for my neighbour to become free to help build my barn, I could hire somebody to do it when I want it done. In larger markets price becomes more relevant. In the absence of violence and fraud, the rational choice model would do a better job of explaining how various spot transactions take place.

As urbanisation increases, we see thicker markets. That means a little less generosity. That is one tendency. However, let us consider empathy toward the poor. There are two very different models to explore. One suggests that empathy has decreased while the other suggests it has increased.

Biblical texts tell us that redistribution essentially existed in the following form: when you ploughed your field you had to leave the gleanings at the edge for poor people. Let us imagine a scale of livelihood. You are at an income level of 10. The poverty-stricken are at a level of one and cannot survive unless they reach two. You, and others at your level, would sacrifice a unit of living standard and move down to nine, and the paupers would increase to a level of two. This is not simply a question of redistribution to get rid of a differential in wealth. Without assistance, people would die of malnutrition. If you read the work of Robert Fogel, a Nobel laureate in economics, you realise that over broad periods of history every calorie could mean the

difference between life and death for many people. One reason for the shorter life expectancy in the past was that caloric intake was often insufficient to enable humans to perform basic metabolic functions.

One assumes that survival motivates people more than a diminishing margin of utility. Therefore, one would expect that income redistribution would occur to a greater degree in a very poor society. By way of illustration, assume that one person has \$10 million and another person \$1 million. Looking at it as a logarithmic diminution of utility, there is the same difference of 10:1 mentioned above. Yet few people would argue that money should be redistributed from the multi-millionaire to the millionaire.

What occurs today is charity of monumental proportions. In the past, when caloric intake was insufficient, there were many contributions to alms and poverty. When income levels are much higher, the wealthy instead donate money to build classrooms, music halls and medical laboratories; poverty alleviation is no longer the sole aim of charity. Instead, a need is seen for things like cultural public goods or relief when great misfortune such as an earthquake or tsunami strikes.

Unlike Robert Putman, I would not be particularly concerned about a purported decline in socialisation. Putman's famous example was that fewer people go to bowling alleys. However, the number of people who play basketball has increased substantially with improvements in health. This shows that people have changed sports; it does not show that there is less sociability overall.

My understanding is that small changes in arrangements can lead to large changes in behaviour, and these changes cannot be predicted by the rational choice approach. Take the example of a savings scheme and default settings. Is the default arrangement that one opts out or opts in? That seems to make a big difference. Assume that people have a tendency to save too little. Individuals only have one opportunity to save for their retirement – there is no opportunity for learning. Why not take that insight from behavioural economics and use it to inform policy?

I believe that it is a mistake to over-generalise the impact of default rules or their stickiness.

Start with Thaler's work paper on this default issue, which is subject to one fairly serious criticism. Consider the situation where an employer is deciding whether to make a contribution to retirement savings part of a remuneration package. As explained earlier, the employer is willing to pay \$100 all up, and the choice is whether to pay the entire sum in salary, or \$90.00 in salary and \$10.00 into a pension fund. The employer would surely provide default terms that maximised the value of the package in the eyes of the employee. After all, the same amount is involved either way. To attract and retain staff the employer will set the opt-in or opt-out default in a way that is calculated to give the highest net compensation. So this is not an arbitrary setting; it is a considered judgment by the employer. If it turns out the employer is not making such a calculation, there is a serious disconnect. The appropriate response, that many employers have taken, is to spend a little more time considering the options.

I think it is simply bad management to assume that the only thing you do is to flip the opt-in or opt-out default in the opposite direction when there are information shortfalls. The right thing to do is to start educating your employees: send them a pamphlet or run a workshop to outline ways people could plan for their retirement, explain what the pitfalls are, and so forth. The University of Chicago gives employees a sum of \$3,000 to let faculty and staff hire a financial advisor. This individual approach makes sense because asset allocation becomes a much more difficult issue with age. At age 60 you would not want to rely completely on investments in equities, because there could be a 10-year stretch of negative returns.

In the case you mention, I would reverse the default if I believed there was a systematic mistake. However, in some cases that is not enough. Let us look at the University of Chicago again. They did not spend much time telling me how to allocate my savings. They advocated a 50-50 split, which is wrong for young people. But they also mandate the savings – there is no default rule as to whether you contribute or not. *Why?* The answer has nothing to do with the ideal form of pension savings. The stated reason is that if the employer does not require people to make savings, penniless employees might – after 40 years of devoted service – turn around and say 'you have a moral obligation to look after me'.

In the US context, I think that that is a strong argument to make Social Security a default option. I believe there would be flights to the exit at record speed if 30-year-olds knew that they were getting a negative rate of return from their investment in Social Security. This is now the case, taking inflation into account. So when the choices are fairly obvious people will tend to make the same ones regardless of the default setting. If it is harder to figure out the differences, switching is a costly way of getting information and the responses will not all be one way. An employer may provide some information but there may be a rational choice explanation as to why people stay or go because there is differential knowledge and people are not perfect calculators.

I have already indicated that the sticky defaults do not apply to the contract at will. No matter which way the default is set, every single employer of any size or consequence would reverse it within a day in just about every contract. The rule would be regarded as catastrophic and employers would want no part of it. The same is true of another hot issue in consumer law, that of liability for consequential damages. This problem arises, for example, when I sell you a computer and you lose all your data. You claim your business has now failed and you sue me. There is not a contract in the world that would allow you to recover for consequential damages, even though the default provision is set exactly the opposite under current law. Default predictions only hold in those cases where there is indifference on one side and uncertainty on the other. The defaults are utterly meaningless where there is a massive stake on one side and an institutional means of protecting it.

I have already indicated that the US Social Security system presents serious problems of long-term instability because of the mixing of paternalism and redistribution. The upshot is that we cannot allow for opting out because any dollar that is lost from the system has to be made up out of general revenues. Faced with this reality, the government does what we should expect. It obscures the situation by issuing statements of contributions and benefits that make it impossible for participants to have any sense of their projected rate of return. The US Social Security statement does not include the employer portion of an employee's contribution. The present value of a future entitlement is not shown in discounted form. The total amount of

contributions, grossed up for any standard return on investment, is never mentioned. The relative present values of the various payout options are never given. So we get these totally meaningless numbers that, if presented by private employers, would land them in jail for securities fraud.

The point I want to make is that the default issue matters. But the really important issue is the massive misinformation that the state disseminates to prop up a compulsory system that would never survive if it was a default option with full information.

What do you think about the endowment effect?

The endowment effect asserts that individuals will attach higher values to those items that they own relative to those that they do not. It postulates a systemic gap between what people are willing to pay to acquire something and what they are willing to accept by way of an offer for the same item.

The endowment effect is a funny kind of problem. First, the only way people discover it is through laboratory experiments. For the most part, it is not present in the field. One of the reasons is that, in most cases, the effect is overshadowed by perfectly normal judgments about subjective values.

For example, if the government is going to acquire my property forcibly, I do not need to be motivated by an instantaneous endowment effect to want to get more than market value. I have proved that its worth to me is greater than its market value by virtue of the fact that I have not put it up for sale on the market. Rational choice theory explains that market value is not equal to use value. When the former is greater you sell the property, and when the latter is greater, you keep it. No endowment effect is necessary to explain the decision.

I think this is extremely important because the only time one should care about an endowment effect is in a situation where there is no subjective value component. With respect to any transaction involving long-held properties, that will not be the case. In many public policy situations I believe the key question is subjective value, and that militates strongly in favour of property protection. The interesting thing about the endowment effect (to the extent that one perceives it to be an accurate reflection of what is going on) is that it cuts in

exactly the same direction as subjective value. It means that if the government decides to take your property, and is disrupting an endowment, it ought to pay above market value. It ought to pay the higher willingness-to-accept value rather than the lower willingness-to-pay value. If the phenomenon is true, it argues for stronger property protection than would otherwise be the case.

It is important to distinguish between two situations. One concerns transfer and the other the acquisition of property. We know the endowment effect cannot apply in a serious way to inventory for sale. If you look around any store the salesperson will not tell you: "This thing is worth \$1,000 to my employer and only \$500 to you so therefore we ought not to sell it because we will have to give up an endowment." The whole institutional arrangement of buying and selling simply dominates any kind of subjective effect. The number of times that somebody wants to buy the tie off your neck as opposed to a tie off the rack is so small that one is really worrying about casual, unorganised sales rather than organised markets. In practice, I think this phenomenon turns out, in general, not to be very important.

The one case where it may be important is in a state of nature. The only rule in nature for taking ownership of a thing is by capture or occupation. That turns out to be the dominant rule in virtually every system, particularly in every primitive system. Suppose there is a conflict between the person who takes something first and others who come after. That is extremely costly for both parties because the first thing to understand about any fight in nature is that even the winners lose. They may kill another individual but if they are wounded in the process they will be fair game for the next person who comes along.

Finding mechanisms to sort things out without actual conflict when there are multiple parties and repeat situations is an extremely important part of any legal system. If one party has an instantaneous endowment effect and the other has an instantaneous keep-off effect, that is a very powerful sorting mechanism. In the case of animals - everything from butterflies to antelope - where first possession is established the outsider will usually back off. What is relevant is not the inherent characteristics of the animal but the status of ownership, and the hormonal responses it stimulates.

I conclude that there probably is something to the endowment effect, but only in exceptional circumstances. I cannot think of any functional explanation as to why it is in one's interest to hold on to something when it could be exchanged for something of greater value. So I do not think this effect is a real impediment to exchange.

