The Skeptical Economist

Revealing the Ethics Inside Economics

JONATHAN ALDRED

The Skeptical Economist

The Skeptical Economist

Revealing the Ethics Inside Economics

JONATHAN ALDRED

earthscan
from Routledge

First published by Earthscan in the UK and USA in 2009
Paperback edition first published in 2010

For a full list of publications please contact:
Earthscan
2 Park Square, Milton Park, Abingdon, Oxon OX14 4RN
711 Third Avenue, New York, NY 10017

Earthscan is an imprint of the Taylor & Francis Group, an informa business

ISBN 13: 978-1-84407-705-2 (hbk)
ISBN 13: 978-1-84971-209-5 (pbk)

Typeset by FiSH Books
Cover design by Rob Watts

A catalogue record for this book is available from the British Library

Library of Congress Cataloging-in-Publication Data

Aldred, Jonathan,
The skeptical economist: revealing the ethics inside economics/Jonathon Aldred
p. cm.
Includes bibliographical references and index.
ISBN 978–1–84407–705–2 (hbk)
ISBN 978–1–84971–209–5 (pbk)
1. Economics– Moral and ethical aspects, 2. Economics. I. Title.
HB72.A345 2009
174–dc22

2008046911

Contents

Acknowledgements

The ideas in this book span a wide range of academic disciplines and I have benefited greatly from working in a multidisciplinary environment: Emmanuel College has been the site of many stimulating conversations which have helped develop my thinking. The Department of Land Economy is another multidisciplinary institution from which I have learnt a lot. Students in both Emmanuel College and Land Economy have challenged and helped sharpen my ideas. It has been a privilege to teach them.

This book has been influenced by discussions and correspondence over the years with many friends and colleagues, from whom I have learnt a great deal. There are too many to name, but I should at least mention Paul Anand, Geoffrey Browne, Anthony Gristwood, Geoff Harcourt, David Howarth, Paul Lewis, John McCombie and Serena Olsaretti.

Others who have helped make this book possible include Philip Arestis, Ian Hodge, Robert Macfarlane, Nick White and Anna Wierzbicka.

I have been lucky with Tim Hardwick, who has been a calm, perceptive and patient editor, and many others at Earthscan.

I am especially grateful to all those who have commented on drafts of chapters, often in time-consuming detail: Mike Gross, Clive Hamilton, Martin Heath, Michael Howard, Brent Kiernan, Simon Jacques, Rob Langham, David Lowen, John O'Neill, Cristiano Ristuccia, Derrick Robinson and Tom Warke.

I owe a particular debt of gratitude to John O'Neill, whose combination of philosophical insight and political passion continues to inspire me.

Finally, none of this would have been possible without two people. Hilary, for her faith in me and unwavering belief that 'the book' was worth writing. And I dedicate this book to my mother, Shirley Aldred.

Chapter One

Introduction: Ethical Economics?

Morality, it could be argued, represents the way that people would like the world to work – whereas economics represents how it actually *does* work. Economics is above all a science of measurement. (Levitt and Dubner, *Freakonomics*)

Few of us associate economics with ethics. Economics is a hard-nosed, pragmatic, 'dismal science'; ethics is philosophical daydreaming. And the quote seems to confirm the view that economics and ethics are strangers to each other. *Freakonomics* is perhaps the best-selling economics book, ever.[1] But its authors are wrong. Economics is not what it appears to be. Economics is an odd kind of science (if it is a science at all) and it is not just about measurement.

This book is not a conventional introduction to economics, although it will try to give an insight into how economists think. But a major obstacle to gaining that insight is that many of those who call themselves economists peddle a narrow or simplistic view of economics to serve vested interests and political ends. These people are better described as *policy entrepreneurs*.[2] Alongside the policy entrepreneurs stand others who are more naively confused in their misrepresentation of economics, but equally dangerous. Between them, these groups do a good job of misunderstanding, misrepresenting and misusing economics, with consequences from which we all suffer.

Policy entrepreneurs preserve their special claim to expertise by encouraging the myth that economics is a mysterious science whose workings are unintelligible to the uninitiated. Economics emerges in public debate as though out of a black box: we are supposed to accept various statements about economics as scientific facts, but the reasoning behind them remains hidden. 'Trust me, I'm an economist' seems to be the slogan. The effect of this black box presentation of economics is that its assertions become unchallengeable to outsiders. Skepticism is ruled out. We just have to accept on trust the 'inescapable economic logic', or similar threatening phrase, which leads to a particular conclusion, however unpalatable. Black box economics is not just the creation of policy entrepreneurs. Serious economists who make a virtue of their political neutrality can also unintentionally reinforce the black box image, because of their astonishing arrogance. For instance, Diane Coyle, formerly of *The Independent,* concludes her book with 'ten rules of economic thinking', one of which is 'where common sense and economics conflict, common sense is wrong'.[3] This imperious tone does not encourage people to embrace the wisdom of economists. People feel they are being told what to think, rather than encouraged to understand. Besides, common sense is sometimes wiser than economics. Traditional economic analysis recommends paying doctors according to the number of procedures or tasks they perform, in order to 'incentivize' them to increase productivity. Common sense points out that doctors will then stop doing anything for which they do not receive a financial incentive, and start claiming additional payment for activities they previously undertook freely out of a sense of professional duty. And traditional economic analysis recommends deciding what to do about climate change by adding up the costs and benefits of reducing carbon emissions, all measured in terms of money. But common sense suggests that not all costs and benefits of climate change can be measured in monetary terms, especially those costs which concern the loss of human lives, and the dramatic upheaval and dislocation of others. It is hard to escape the conclusion that economic analysis in these cases depends as much on value judgements and political and psychological assumptions as it does on neutral science. I do not object to value judgements and political beliefs creeping into economic arguments: I think they are inevitable. But then I do not claim economics is a science.

A close cousin of black box economics is veto economics, where assertions about economics are used as a kind of veto to rule out new ideas and proposals without further discussion. The veto is absolute because the assertions emerge from the black box: they are presented in a way that makes it extremely hard for the non-economist to dispute them. In its most extreme form, veto economics rejects ideas and proposals with just one word, offering no further explanation. Favourite veto words include 'inefficient', 'irrational', and 'anti-competitive'. For readers of the more specialist economics and business press, rejecting proposals on the grounds that they are 'suboptimal', 'time-inconsistent', or lack 'incentive-compatibility', has also become fashionable. And as a last resort there is always the plain but vacuous condemnation 'uneconomic'.

Veto economics serves to protect the economic orthodoxy. In some ways the orthodoxy has served us, in rich economies, quite well. For a brief period after the fall of the Berlin Wall, some even talked of 'The End of History'. Our economic problems were solved, and something called 'The New Economy' had arrived, promising endless prosperity – or at least an endlessly rising stock market. Although these fantasies are now largely forgotten, a more humble but still confidently optimistic orthodoxy prevails. The orthodoxy says that the role of governments is to maximize economic growth by providing a stable economic environment of low inflation and generally moderate levels of unemployment. And whatever the problem, markets are almost always the solution. This orthodoxy is not directly discussed in what follows. But a recurring theme is that the orthodoxy leaves something crucial out. Economic growth is not an end in itself. We should focus instead on our quality of life, our well-being, or to rehabilitate an embarrassing word, our *happiness.*

In rich countries the experience of recent years has been that increased material wealth has not led to improvements in quality of life. On the contrary, people report being less happy than 40 years ago. We seem to suffer from a kind of 'affluenza', a condition in which people become preoccupied with acquiring money and possessions, and gaining social status from them. A large body of psychological research suggests that these preoccupations are associated with increased rates of mental illness, including depression, substance abuse, anxiety and personality disorder.[4] Less dramatically, and anecdotally, we can live in a prosperous economic environment, and still discover that modern life is rubbish. So how can a better understanding of economics help?

Talk of improving quality of life might provoke expectations that the following chapters include detailed discussions of integrated transport policies, pension reforms and endogenous growth theory. In other words, the kind of policy wonkery which few of us care passionately about. But we *do* care about the principles at stake – whether it is fair to tax the rich more highly, whether environmental damage can be boiled down to a sum of money, whether surveys can really measure our quality of life or happiness. This book is about these principles. Black box economics obscures them. It takes certain views about these principles for granted, rarely mentioning them explicitly. In the two examples mentioned briefly above we have already seen that black box economics assumes:

- Employees are essentially selfish, and should be managed accordingly.
- The value of life can be measured in monetary terms.

Other presumptions include:

- What we buy always makes us better off.
- People respond predictably to financial incentives.
- Taxation damages the economy, and is morally wrong.
- Economic growth increases happiness.

Together, these assertions and many others form a web of beliefs which constrain and shape our economics and politics, affecting us in ways that extend beyond our economic lives as consumers and workers. These assertions are the subject of this book. Some might wonder whether they are about economics at all. It is true that these issues lie on the boundary between economics and philosophy – ethical economics – but that boundary is much more fragile than it seems.

Consider that favourite veto word of economists, inefficient. Efficiency is a core economic idea, one that to many seems like *the* core economic idea. The efficiency or otherwise of some proposal or outcome is often presented as a fact. In very rough terms, we do something more efficiently if we obtain the same desirable outcome using fewer valuable inputs of time, effort or resources more generally. Or instead, we can obtain a bigger or better outcome using the same inputs. So efficiency seems unambiguously a good thing, and also a scientific concept, referring just to the measurable relationship between inputs and outputs.

Unfortunately for this perspective, economics affects people. The scientific purity attributed to the efficiency idea is disrupted by people, because people are affected by actions taken in the name of efficiency. For example, economists and policy entrepreneurs use efficiency to make claims such as:

> Cutting the regulatory burden facing employers makes the labour market work more efficiently.

These regulations might concern working hours, overtime rules or minimum wages. The argument boils down to the idea that the same output can be produced using fewer workers, or produced by people who are paid less. Or perhaps instead, a greater output can be produced. Economists sometimes argue in support of reducing labour market regulations by pointing to the increased national output (economic growth) in countries which have adopted such policies. My aim here is neither to agree nor disagree with these policies. The crucial point is simply that, in almost all cases, there will be both winners and losers. If national output has risen then materially we are richer; conventional economics deduces that some consumers must be better off. And shareholders should benefit from firms making higher profits. But some workers may suffer from lower pay, or longer working hours, or an unpleasant increase in work intensity. Whether, on balance, the changes taken together represent an overall improvement depends on, among other things, ethical views about the relative priority we should give to impacts on the rich and poor, and ethical concepts of need, justice, entitlement and just deserts. So the 'efficiency gain' in this example is not an objective improvement; if it is regarded as an improvement, that is an ethical judgement. Efficiency is not an ethics-free concept.

Still, it seems there might be occasional exceptions to this rule. In the last paragraph I wrote that there will be both winners and losers *in almost all cases.* But what if everyone gains? Surely *that* would effectively count as an objective improvement – an unambiguous gain for everyone. Consider the common argument that free trade makes the economy more efficient. Some economists argue that, at least in the long run, *everyone* will be better off from freer global trade. Suppose this is true. A sympathetic reading of the evidence on recent trade liberalization suggests that, roughly, rich countries have gained moderately, recently industrialized countries such as China have gained the most, while very poor countries have gained the least. And within countries, the gains are very uneven too, with the poorest in India and China gaining much less than others in their countries. Even though 'everyone's a winner', it is clearly reasonable to regard this outcome as *not* representing an objective, unambiguous improvement, because inequality has increased. The poor may be literally better off, but not relative to the rest of us. Again, whether we view the 'efficiency gain' from freer trade as an overall improvement is an ethical judgement. Even if everyone wins, 'efficiency' still comes laden with ethical baggage.[5] So when veto economics invokes the word as though it were a neutral factual concept to which no one could reasonably object, we are misled.

It is worth stressing that the misrepresentation here is often unintentional. Few economists, if any, set out to deceive (even if the same cannot be said of policy entrepreneurs). Instead, many economists believe that economics can be ethics-free for the noblest of reasons – they are loathe to foist their own ethical views on other people. So they try to make their economic advice ethics-free. But as I have already indicated, ethics is unavoidable in economics. We should simply be open and explicit about it. Besides, these economists need not worry. They should follow the example set by their greatest forebears.

The great economists of the past saw economics and ethics as inextricably entwined; one example of this is their emphasis on the limits of economic prosperity in bringing about broader improvements in the human condition, and hence the relative insignificance they attached to economic prosperity alone. Adam Smith was dismissive of economic success as an end in itself: 'if the trappings of wealth are viewed philosophically, they will always appear in the highest degree contemptible and trifling'.[6] John Maynard Keynes, the most influential economist of the 20th century, shared similar concerns. In 'Economic possibilities for our grandchildren', Keynes described the 'love of money as a possession' as 'one of those semi-criminal, semi-pathological propensities which one hands over with a shudder to the specialists in mental disease'.[7] Keynes accurately predicted British average income per head in the early 21st century, and went on to anticipate the problems of affluenza, and the renewed importance of an ethical framework once society ceases to be focused solely on economic growth.[8] In all this, he was remarkably prescient: the essay was written in 1930.

Clearly, the early great economists would have rejected the belief that economics deals solely with material prosperity and financial wealth. But this belief, a contemporary variation on the illusion that 'economics and ethics don't mix', can be seen in action every day in the news media. Economics is discussed as if it were just concerned with the business and financial aspects of an issue, narrowly defined. If an airport expansion is being proposed, for example, the effects on the local economy and on the airline industry are included under the 'economic' heading, along with the construction cost, but all the other impacts are treated separately. Impacts such as carbon emissions, and the effect on quality of life for those living near the airport – and those able to go on holiday more easily and cheaply – are relegated to the 'non-economic' category. The same intellectual apartheid system continues in government.[9] 'Non-economic' impacts are all too easily regarded as less important, especially if they are not quantifiable in terms of money. If impacts cannot be measured in monetary terms, they are seen as soft, fuzzy, ephemeral… No wonder some economists begin to doubt if they are really there at all.

So those who believe that economics can be separated from ethics have forgotten its history. And quality of life impacts, such as those arising from a new airport, are just as much the subject of economics as inflation, unemployment and economic growth, even if they cannot be quantified in terms of money.

Many people, economists and non-economists alike, resist this conclusion because they fear it leads to an 'anything goes' dead end. In other words, they object to ethical economics because they believe that ethical debate is a waste of time. Some economists themselves have helped perpetuate this myth. Regarding disagreements about ethical values, Milton Friedman believed that 'men can ultimately only fight'.[10] Skepticism about ethical reflection has been popular throughout history and is likely to remain so, because it saves people from having to bother to defend their beliefs. But that does not mean it is right. Even if ethical judgements cannot be literally true or false, almost all of us regard some judgements as better than others. Although some ethical questions are plagued by doubt and disagreement, others suggest ethical agreement is possible. There is widespread condemnation of slavery or torture, and most of us will willingly defend our objections to these practices against those people who disagree. We do not regard this kind of ethical debate as pointless or meaningless. Finally, some people worry that it is intolerant or illiberal to advocate a particular set of ethical values, particularly when these values have implications for the conduct of others. But this worry is not compatible with the idea that ethical reflection is futile: tolerance and liberalism are themselves ethical values, requiring an ethical argument to justify them.

As well as avoiding undue pessimism about the usefulness of ethical reflection, we should, on the other hand, avoid undue optimism about what economics can tell us. As the joke goes:

> Economists don't answer questions because they know what the answer is. They answer because they are asked.

Once we recognize that economics does not fit our picture of a typical science, we must be more modest in what we expect of it. This does not mean that anything goes, that all the constraints of veto economics can be swept away with a single liberating stroke. But this book will suggest that the truth is often more complex than the policy entrepreneurs would have us believe. Often the truth is that economists don't know. Not knowing whether there is a constraint – whether, for instance, a lower rate of economic growth must lead to increases in unemployment – is not the same as knowing there is no constraint, but it cannot act as a veto on new policies either. This kind of modesty is not what many of us want to hear. We yearn for the comfort and security of definite answers. But an honest economic analysis can typically hope to do much less than that. At best, it will point out some inconsistencies in aims, methods or assumptions; it will rule out some options but leave several others open – and pass the decision back to whoever commissioned the analysis in the first place.

In case this sounds too pessimistic, let me sum up why you should read this book. You *can* read this book. As well as stripping out the mathematics, which often kills off public debate about economics, I have eliminated almost all the jargon too. As an economist, it is deeply satisfying to economize on economic jargon; for example, there is only one more passage in this book containing the word 'efficiency'. The concept of efficiency is elsewhere too, but not in a way which needs gratuitous labelling. In the following chapters, I discuss shopping and consuming, pay and taxation, economic growth and happiness. Then I look at the practice of putting prices on life and nature, and more generally the process of bringing money into new contexts: turning things and activities into commodities which can be bought and sold. The chapters are essentially free-standing and can be read in any order, although later chapters inevitably contain occasional references to earlier discussions. Throughout I show that economics draws on hidden ethical assumptions, both in its foundations and its practice. Whether we agree or disagree with these assumptions, we all need to know what they are. It is futile to try to rid economics of its ethical foundations, because all economic theories and policies must draw on particular views about how we ought to live, and what we value. We cannot avoid making choices between these views, so in what follows I set out to uncover them, with due skepticism about the associated economic arguments.

Chapter Two

The Sovereign Consumer

What is the basic principle of economics? As a wise elder once told me, 'People do what they get paid to do; what they don't get paid to do, they don't do.' People respond to incentives; all the rest is commentary. (William Easterly, former World Bank economist[1])

[When] self-interest and ethical values with wide verbal allegiance are in conflict, much of the time, most of the time in fact, self-interest-theory … will win. (George Stigler, Nobel Laureate economist[2])

The average human being is about 95 per cent selfish in the narrow sense of the term. (Gordon Tullock, Virginia economist[3])

Why is economics called 'the dismal science'? One reason is that many economists are profoundly cynical about human behaviour and the motivation that underlies it. Morality, they seem to suggest, is for losers: real people are almost always selfish. This is the key assumption behind the lead actor in modern economics, the 'sovereign consumer'.

Consumers go shopping, and shopping is clearly central to the economy – it is the fuel powering the engine of economic growth. And shopping is equally central to understanding how economists think. Modern economics is built on theories of 'rational choice', and the textbook example of rational choice is supposed to be the kind of choice made by consumers when they shop. All other choices – those made by workers in firms, or their managers or shareholders; choices made by bureaucrats and politicians; even choices made by private individuals about whom to marry or which moral values to adopt – are analysed by economists in essentially the same way, as variations on the basic story of consumer choice. The sovereign consumer is the actor at the heart of this story, a person who is fully informed, knows what they want, and never makes mistakes in getting it. In economics, the sovereign consumer is very much in control of their life.

As I suggested in Chapter 1, our anecdotal experience of shopping and consuming today is much less optimistic. We suffer from 'affluenza'.[4] But economists respond to these anecdotes with an awkward question: if our present patterns of shopping and consuming create problems and do not make us happy, why do we continue? This chapter begins to answer this question, with more to follow in Chapters 3 and 5. Before that, it is worth emphasizing what is at stake here. It is not just a matter of whether economists have an accurate picture of consumer choice. The story of the sovereign consumer lies behind some basic and influential ideas:

- economic growth is always beneficial
- people should get what they want
- more choice is always better.

These are widely held ethical views, which often masquerade as economic facts. In this chapter I will question them. I go on to argue that economic advice based on the assumption that people are selfish often has disastrous effects. The main problem is easy to summarize: assuming people are selfish becomes a self-fulfilling prophecy, with the result that altruism, trust and cooperation are all undermined. And the assumption of selfishness can even act to excuse, or indirectly justify, immoral behaviour. As we shall see below, for some economists, behaviour such as committing crime, or racial discrimination, can be both 'rational' and 'efficient'.

So there are two reasons to look in detail at the story of the sovereign consumer. It has influence far beyond the world of shopping. And at every stage it involves ethical judgements as well as economic analysis. The result is ethical economics.

The economist goes shopping

The economic analysis of how we choose is remarkably simple. It starts from the common-sense notion of choice arising from the combination of beliefs and desires. I *desire* to wash my clothes, and I *believe* that 'Razzle Dazzle' washing powder will help me wash them, so I buy some. The next step is to describe why I buy Razzle Dazzle rather than some other brand. Essential to making choice is not just desire, but *preference*: a desire to have or do this rather than that. I buy Razzle Dazzle because I prefer it to other brands. Obviously price and my budget play a part here too. I may prefer Razzle Dazzle to Washes Whiter, but nevertheless buy Washes Whiter because it is cheaper. There is a trade-off between my preference for Razzle Dazzle and its higher price. Economists are fascinated by these trade-offs, and have developed the theory in sophisticated ways to analyse how demand for a product may change in response to a change in its price, or the consumer's income, or the price of a rival product. But these complications need not concern us. The key components of the theory are just the preferences of the consumer, and the options available, such as different brands of washing powder. Together these determine what will be chosen: given her budget, the consumer will choose the most preferred option from among those available. But which will this be? In general terms, nothing can be said, because the preferences of different consumers will clearly vary for all sorts of reasons. I may prefer Razzle Dazzle just because I like the name, while you prefer Washes Whiter because you believe it does. And the pure economic theory stops there, making no assumptions about the content of people's preferences. But in practice, economists need to go further in order to draw practical conclusions from their analysis. They want to sidestep the details of consumer psychology, so instead they make a general assumption which is innocuous and self-evident to many – that consumers are largely self-interested. The idea is that in making choices between the available options, I pursue my own interests first and foremost. This assumption rules out, for instance, buying Washes Whiter because I think that the retailer or manufacturer makes more profit from it, and I want to help boost their profits. Also, it probably rules out buying Washes Whiter solely because my aunt will disapprove if I buy Razzle Dazzle, or buying Washes Whiter because I believe it is less harmful to the environment. 'Probably' because the definition of self-interest is opaque; we will be returning to this later.

The rest of the economic theory behind the sovereign consumer seems obvious enough. Indeed the theory is *so* obvious that it is easy to forget it is there at all, and to overlook the assumptions that were implicitly made in the outline above. To begin with, the theory takes both the consumer's preferences and the available options as given. Taking preferences as given implicitly assumes that the consumer already knows what they want – that they have had the requisite time and information to form a carefully considered set of preferences. Taking the options as given implies that the available alternatives are somehow fixed objectively and beyond the influence of the consumer. In reality, the process of making a choice is often very different. Preferences are hazy and hesitant because we lack the time and information to determine them. The decision-making process also leads to the discovery of new options, or the revision of existing ones to improve them; we do not act as if the available options are fixed in stone. Next, once the preferences and options are determined, the story of the sovereign consumer takes its conclusion for granted: the consumer chooses the most preferred option from among those available, implying that the consumer is infallible and never makes mistakes.

With these assumptions spelt out, you may be starting to have doubts. You certainly should be. This seemingly common-sense story of how we choose is far removed from reality. When the psychologist goes shopping, the world looks rather different.

The psychologist goes shopping

There is a rich body of research by psychologists on how we choose, much of it suggesting a sharply different picture of choice from that found in economics. For example, the options between which we choose are not as clear-cut or objective as economic theory implies. Essentially the same set of options can be perceived very differently, leading to completely different choices, due to *framing effects*, so-called because they concern the context or frame through which a person perceives a decision. A box of Razzle Dazzle can be perceived as expensive (if we remember that it cost less a month ago) or cheap (if the retailer displays it as being on 'special offer'). You may feel confident that you are not fooled by such simple retailing ploys (you always ignore 'special offer' claims when assessing price) but framing effects are pervasive and not always so obvious.

Framing effects: Two examples

1. Merely because an option is framed as the default option – one that will be selected unless the chooser actively decides otherwise – it is much more likely to be 'chosen', even when the decision is important. One study of enrolment in organ donation programmes in seven European countries found that on average 97 per cent of people were enrolled when this was the default option, but only 18 per cent otherwise.[5]
2. Suppose you have decided to see a show at the theatre, for which tickets cost $40. You arrive at the theatre, and as you queue to buy a ticket, you discover that you have just lost $40 in cash. Would you still go in? Most people say 'yes'. Now imagine instead that you had already purchased the $40 ticket, but as you arrive, discover that you have lost the ticket. The ticket is non-refundable and non-replaceable. Would you buy another ticket and still go in? In this situation most people say 'no'. Both situations involve a choice between seeing the show and being $80 worse off, or not seeing it and being $40 worse off: they are effectively identical choices, but framed differently.[6] The sovereign consumer is supposed to see through the surface descriptive differences between the two situations, treating them as irrelevant, and should therefore make the same choice on both occasions. Since most people in reality choose differently according to the way the choice is framed, they are irrational according to standard economic theory.

Another reason that the options should be thought of as subjectively defined, rather than as objective facts about the choice situation, is that our perception of them is heavily influenced by what psychologists term *availability*. Availability concerns how readily some piece of information can be brought to mind – how available it is to us – and depends on how vivid, striking or distinctive the information is. Availability often distorts our judgement. Many of us overestimate dramatic, vivid causes of death (murder, plane crash, lightning strike) but underestimate the much more likely causes, such as common illnesses. And watching one vivid interview with, say, a lazy recipient of state welfare benefits, will influence someone's judgement much more than overwhelming statistical evidence that most benefit recipients are not lazy. The distortionary effect persists even if people are told in advance that the interviewee is atypical.[7] The implications for consumer choice are clear. We pay more attention to available options, and may ignore others. We attach more weight to information about these options merely because it is more available. Advertisers of SUVs ('sports utility vehicles') can rest assured that one dramatic advert, showing a vehicle surviving a high-impact crash largely unscathed, will have more influence than well-publicized scientific research suggesting that SUVs are no safer for their occupants. Even if consumers know that the advert is unrepresentative, or downright manipulative, it still seems to exert its subconscious influence on choice.

Advertising also exerts its subconscious influence, via the availability effect, using brand recognition. Branding ensures an otherwise identical product becomes more familiar, and so more available to be brought to mind. This familiarity alone has been shown to breed approval of the product, again subconsciously. In one study, people were played different snippets of music, different numbers of times. On average, they preferred the snippets that were played more frequently, although their conscious verbal explanation of their preferences made no reference to frequency.[8]

Even without the distractions of availability or advertising, our predictions of the satisfaction from future experiences are unreliable, and so we struggle to choose the best option.[9] In particular, our predictions are biased by our current emotional state. We all know the danger of buying more food than intended when shopping on an empty stomach, a danger confirmed by psychological evidence.[10] Similarly, catalogue shoppers ordering by telephone seem overly influenced by the current weather: warm clothes ordered on cold days are more likely to be returned later.[11] And many people join gyms and health clubs which they subsequently rarely use, because at the time of joining they focus on the health benefits, rather than how they will feel in the future when visiting the gym.[12] A form of framing effect also distorts our predictions of future satisfaction. Suppose you are trying to decide between two equally expensive pairs of stereo speakers. In the audio shop, you will probably put great weight on which sounds better, and purchase on that basis, even though the differences are likely to be small. Unfortunately, at home you realize that you care more about the speakers' appearance, because the marginally superior sound quality is undetectable in isolation.[13] The context of choice in the shop leads us to put too much weight on an attribute which we later find less important.

The psychological phenomena mentioned so far all affect the framing of the choice situation, the options we consider, and the way we assess them. But even if these troublesome influences did not arise, there is a more fundamental worry about the consumer sovereignty story. So far we have assumed that the consumer's goal is clear – their own self-interest – and choice is just about pursuing that goal effectively. However, this assumption is questionable, even if we set aside for now the more philosophical doubts about the meaning of self-interest. Consider instead a simple consumer choice: choosing the best holiday destination. You may have noticed that you can find a holiday rather mixed at the time, with both good and bad experiences, but a few months later, the whole holiday seems good when looking back on it. In fact, even a short time later our memory provides an inaccurate guide to how pleasurable or otherwise an experience was, because of a form of availability effect known as 'Peak-End evaluation'. We remember any experience in terms of the extreme moments (peaks and troughs) and the final, end moment of the experience. This was vividly demonstrated in a famous experiment conducted by psychologist Daniel Kahneman.[14]

Kahneman studied a group of patients undergoing a colonoscopy, an uncomfortable medical procedure. For half the patients, the colonoscope was deliberately left stationary inside them for an extra minute, causing them to experience mild discomfort for this period, but less pain than earlier in the procedure. That is, this second group experienced the same procedure as the other patients, but with an additional one minute of mild discomfort at the end. And this is how the second group described their experience at the time. Nevertheless, a short time later, they rated the whole experience as *less* unpleasant than the first group, because they focused on the end moment, which was less painful. These Peak-End evaluations determined later choices. Over a five-year period after the colonoscopy procedure, patients in the second group agreed to more follow-up colonoscopies than members of the first group, presumably because they remembered the procedure as less unpleasant.

This tale casts doubt on the idea that pursuing our self-interest is as straightforward as it often appears to be. I may have an apparently clear goal – minimizing the unpleasantness of a medical procedure – but fail to attain it because of Peak-End evaluation. I favour the medical procedure involving not just a colonoscopy, but further discomfort on top. The problem is not just that we sometimes make mistakes in pursuing our self-interest. The deeper difficulty is identifying my self-interest, even after careful reflection, once 'mistakes' are pointed out to me. Suppose I had a choice: should I minimize the pain experienced during the colonoscopy procedure itself, or should I choose the option which, although it involves additional discomfort at the time, will make me feel happier about the whole experience shortly afterwards, and less fearful of follow-up operations?

Self-interest is also far from clear in the theatre example discussed earlier. Economic theory prescribes that we should reach the same decision about whether to see the show, regardless of whether we discover outside that we have lost $40 cash or a $40 ticket. We are supposed to choose solely on the basis of the material benefits and costs involved, and not be influenced by framing effects, the context of the choice. Does our self-interest demand this type of 'consistency'? Context matters. Even if the net financial impacts in two situations are identical, most of us distinguish them. Losing $40 cash is seen as one of life's minor misfortunes that inevitably befall us, with no implications for whether we see the show. But having to buy a replacement for a lost $40 ticket means that it will effectively cost $80 to go in, a price that may deter us. This kind of reasoning seems as legitimate as it is widespread, despite economists' advice that it runs counter to our interests.

Perhaps the most direct evidence that self-interest may be ill-defined comes from situations where we seem incapable of deciding what will satisfy it – where we lack the self-control at the heart of consumer sovereignty. Self-control problems are ubiquitous.[15] Often we anticipate them and take successful evasive action. I move my alarm clock out of reach, so that I will not be able to switch it off without getting out of bed. I join a pension scheme to commit myself to saving enough for retirement. But sometimes we struggle to decide where our self-interest lies, that is, whether to exercise self-control at all. I vow before dinner in a restaurant to forgo pudding, but succumb when the moment arrives. However, I never buy fattening puddings to eat at home. I buy cigarettes, but give them away, then run out to buy more. These changes of mind are not due to new information or hurried decision making. Many smokers and heavy eaters know that they are very likely to suffer from poorer health later in life, but continue to smoke and eat heavily nonetheless. They look back on years of similar behaviour, and have no regrets. Nor can we assume that they would resist if they had 'more will-power'– just ask a veteran smoker such as the painter David Hockney.[16] In short, people repeatedly have preferences pulling in opposite directions. They cannot decide what will best serve their self-interest, no matter how hard they scrutinize their preferences.

Whether due to framing effects, self-control problems, Peak-End evaluations or other common features of how we think, our 'self-interest' may not be clearly defined, even with the benefit of hindsight. And the bulk of psychological research more generally provides powerful evidence that consumer sovereignty is a myth.[17] But, you may be beginning to ask, so what?

Consumer sovereignty: An influential myth

While it might be embarrassing to economists that their description of how we choose is very often false, does it matter to the rest of us? Absolutely, because consumer sovereignty is taken for granted in all sorts of unlikely but crucial places. Consider advertising: if government policy were made by psychologists, then certain forms of advertising would surely be banned, because they lead people to buy too much, or things they later regret. In contrast, starting from the story of the sovereign consumer, economists have developed an entirely different perspective on advertising, one which has been very influential in justifying it. This is surprising, because the ubiquitousness of advertising poses a challenge for consumer sovereignty. If consumers are sovereign, they will not be influenced by advertising, but if they are not influenced, why do manufacturers and retailers bother to advertise? It seems that either consumers are influenced by advertising, implying consumer sovereignty is an illusion; or they are not, in which case it has no effect, and advertising spending is irrational for companies and an inefficient waste of society's resources.

The simplistic reply to this conundrum emphasizes that advertising can be informative: it tells you what is available, where and at what price. It does not *change* your preferences, but provides information which can help construct them. Thus informative advertising can influence the things we buy, without undermining consumer sovereignty. The problem with this reply is that most advertising seems persuasive rather than informative: it wants you to change your preferences between products – to buy Razzle

Dazzle rather than Washes Whiter, or this brand of car rather than that one – often on grounds which seem to have nothing to do with the quality or function of the product itself. The sovereign consumer would not be manipulated into changing preferences in this way. But economics offers a more sophisticated response to the conundrum.[18]

The essential idea is that our underlying preferences – the ones which the consumer sovereignty story holds sacrosanct – do not directly concern *products* at all. Rather, they concern the underlying *characteristics* of products.[19] For example, our fundamental preferences do not concern cars but characteristics of cars, such as speed, comfort, reliability, economy and image. We seek cars with particular bundles of characteristics. I prefer cars which are speedy, fashionable and comfortable to those which are reliable, economical but dull, while you may prefer the opposite. With this framework in place, it is argued that advertising serves a useful function. It provides information about the bundles of characteristics embodied in the products it promotes, information which may lead people to change the products they buy, even though their underlying preferences over characteristics are unchanged. An advert involving interviews with people who have owned a particular make of car for 20 years may convince you that it is a reliable brand. So advertising may *appear* to change consumer preferences – hence overturning consumer sovereignty – when in fact it simply provides the information to enable consumers better to fulfil the preferences over characteristics they had all along. In essence, the two assumptions mentioned above, that preferences are predetermined, and consumers never make mistakes in acting upon them, have merely shifted one step back. Since consumer sovereignty is indefensible for preferences over products themselves, it is invoked for preferences over characteristics instead.

Whatever the other drawbacks of this story, one lesson is clear. It is crucial to any defence of advertising. For example, the defenders of 'junk food' advertising have sought to convince the UK Food Standards Agency with an argument along exactly these lines.[20] And conversely, if consumer sovereignty is a myth, then the case for restricting some forms of advertising becomes much stronger. It is ironic that while one theory of choice is used to justify the value of advertising to society at large, another altogether is used by companies deciding how much to spend on advertising: they do not think in the narrow terms of preferences-over-characteristics. They draw on insights from psychology and sociology to create vague but emotive brand images to persuade people to want new and different goods; they are not just providing facts about their products.

Justifying advertising is not the only surprising use of the consumer sovereignty doctrine. If consumers are sovereign, then the choices we observe them making should be a reliable guide to what they actually prefer – their preferences are revealed by their choices. Economists use this *revealed preference* assumption frequently, in arguments which many other people find bizarre. Economic theories of 'rational addiction' assume that drug addicts' choices always reflect what they truly prefer. The possibility that their choices could be irrational is ruled out by definition.[21] And in Chapter 6 we shall see the revealed preference trick pulled again. Here's a brief glimpse. Policies to tackle climate change may save lives in the future, but involve large economic costs now. What should we do? Economists answer this question by trying to value the lives saved in terms of money. The first step in the calculation is to put a money value on an increased risk of death. Economists do *that* by measuring how much extra people are paid in risky occupations. This extra pay is supposed to reflect how much employees care about the risks, how much extra they must be paid in order to tolerate them. But this conclusion is only valid if employees' choices reveal their true preferences. In reality, workers are often unaware of the true magnitude of the risks they face, and they may end up in jobs because there seemed to be no alternative, not because they have consciously chosen extra risk in return for extra pay. Real people do not choose jobs as sovereign consumers choose washing powder.

The revealed preference device and the preferences-over-characteristics account of the role of advertising show the broad influence of the consumer sovereignty story. But the most important use of the story in modern economies is in justifying economic growth. The argument brings together the key assumptions behind consumer sovereignty: specifically, that consumers have clear-cut predetermined preferences which reflect the pursuit of self-interest, and that they never make mistakes in identifying their self-interest, or choosing in accordance with it. Since consumers are self-interested, they only prefer things that they believe will make them better off than any other available alternative. And since they are never mistaken, their beliefs are always correct, so they only prefer things that in fact *will* make them better off. Once this crucial link between preference satisfaction and being better off is in place, the conclusion that economic growth is a good thing follows quite easily. This is how:

> Growth is good, because it leads to more consumption.
> Consumption is good, because it leads to more preference satisfaction.
> Preference satisfaction is good, because it makes people better off.

Although this is clearly a simplified outline, it captures the essentials of an argument for economic growth in developed economies that is so widely taken for granted that it goes unmentioned. But I am *not* suggesting that economic growth is a bad thing (see Chapter 3), simply that this common argument for it begins with a very dubious premise, built on the myth of consumer sovereignty. Once the myth is abandoned, we cannot conclude that more material consumption is good just because more preferences are fulfilled, more wants are met. But what *can* we conclude? It is time to return to the question posed at the start of this chapter: if consumption creates problems and doesn't make us happy, why do we continue?

Addictive and competitive consumption

> Want is a growing giant whom the coat of Have was never large enough to cover.[22]

> A house may be large or small; as long as the neighbouring houses are likewise small, it satisfies all social requirement for a residence. But let there arise next to the little house a palace, and the little house shrinks to a hut.[23]

'Addictive' is a strong word. Very few of us are *compulsive* shoppers, in the sense of suffering from a psychological pathology. But many people's attitude to shopping is nonetheless akin to addiction. There is powerful psychological evidence that just having material goods, things, stuff, gear, is not enough. The satisfaction they provide depends on having more and more of them.[24]

When I first buy something, be it new clothes, electronic gadgets or even a car, it typically brings pleasure at first. The problem is that I rapidly become accustomed to owning and using the item, so my mood, or general