

# 1 Introduction: ethics out of economics

---

## 1.1 Economics and ethics

The traffic between economics and ethics travels in both directions. Each discipline can learn from the other. Economics is partly concerned with assessing the merits of economic arrangements, and with deciding how governments ought to conduct their economic affairs. It makes judgements of right and wrong, good and bad, in economic matters. It needs criteria for making these judgements, and the criteria must come from ethics. On the other hand, economists have developed for their own purposes sophisticated methods of analysis that turn out to be useful in philosophical ethics. They can help not only with questions of practical ethics, but also with fundamental issues in ethical theory.

This book concentrates on just one of the areas where the methods of economics can contribute to ethics. It is about those problems in ethics that require us to balance different interests or concerns against each other. For example, we have to balance the interests of future people against the interests of presently living people, fun in retirement against fun in youth, the wellbeing of the deprived against the wellbeing of the successful or lucky, the value of prolonging people's lives against the value of making people's lives better while they last. Many of these are problems of large-scale or public morality – problems for politics and the institutions of society. Weighing the prolonging of lives against the improving of lives is a pressing issue for the health service, for instance. How should the health service value hip replacements, which improve lives, compared with heart replacements, which prolong them? What importance should it give to palliative care, which does nothing to prolong lives, and may even shorten them? Other questions of weighing are problems for individuals rather than the public. Individual patients and their doctors sometimes have to choose between different treatments for a terminal disease: between treatments that will prolong life and treatments that will make life better while it lasts. All of us wonder whether to give up some of the nice things we eat for the

This chapter is an edited version of my inaugural lecture, delivered at the University of Bristol in February 1995.

2 Introduction: ethics out of economics

sake of living longer. In planning our lives, we all have to weigh good things in the future against good things now. Many of us have to weigh our own interests against those of others, such as our parents. Weighing and balancing are regular features of all our lives.

Many practical problems of weighing fall within the domain of both economists and philosophers. Economists are particularly concerned with some of the large-scale problems I mentioned. These are places where the scarcity of resources forces a society to weigh up alternative possible uses for these resources, and economics claims to be the science of scarcity. Indeed, there is no sharp division between economics and ethics in this area. Economists pursuing fundamental issues in welfare economics and philosophers pursuing practical issues in ethics have found themselves working on the same questions: is equality valuable?; how much should we care about the wellbeing of future people?; what value should we assign to preserving nature, or to human life?; and many others. Several of the chapters in this book lie in this borderline territory where practical questions of economics and ethics meet.

Others are concerned with more fundamental, theoretical issues involving weighing. From a theoretical point of view, problems of weighing are best handled within the department of ethics known as the theory of good or the theory of value. This is the department that is concerned, first, with what is good for individual people and, secondly, with what is good in general. Questions of weighing are questions about how different goods – for instance, goods coming at different times in a life or to different people – come together to determine how good a life is overall, or how good a society is. As I put it, they are questions about the structure of good.

The techniques of economics can help with understanding the structure of good. They can help to analyse the value of a life: whether a shorter exciting life is better than a longer boring one, for instance. They can help with the question of whether a more equal society is better than a less equal one, and so on. For its own purposes, economics has developed sophisticated formal theories that are intended to analyse the structure of people's preferences. It happens that these theories can be turned to analysing the structure of good. This is something ethics can take out of economics.

Ideas from economics have influenced the theory of good for a long time. Historically, there has been a long association between economists and utilitarian moral philosophers, and utilitarians have a particular interest in the nature and structure of good. Unfortunately, though, what moral philosophers have recently drawn out of economics in this area has generally been the wrong thing. I said the formal methods of economics can help analyse the structure of good. But when philosophers have adopted some-

thing from economics, it has often been the substance rather than the form. Implicit in economics, apart from the formal techniques, there is also a substantive theory of good. Few philosophers have taken over this theory wholeheartedly, but it often makes an appearance in ethical arguments. It is a mistaken theory, though. It is a pity it has penetrated ethics at all. In this introductory chapter, I shall argue that what ethics should take out of economics, at least in the area of the theory of good, is form and not substance. The substance should be left alone. Then I shall mention two examples of the useful formal lessons that ethics can learn from economics.

The arguments in this introduction will be brief, because many of them are taken up in more detail in later chapters.

## 1.2 The preference-satisfaction theory

What is this substantive theory implicit in economics? It is not a complete theory of good, but only a theory about what is good for a person. It is specifically:

*Preference-satisfaction theory of good.* One thing *A* is better for a person than another thing *B* if and only if the person prefers *A* to *B*.

(Actually, this biconditional does not express the preference-satisfaction theory completely. The theory also requires the determination to go from right to left: when a person prefers *A* to *B*, that makes it the case that *A* is better for her than *B*.)

Before assessing this theory, I must say something about the meanings of the terms ‘better for’ and ‘prefer’. In common usage, the expressions ‘good for you’ and ‘better for you’ sometimes carry a narrow connotation of prudence. Exercise is said to be good for you even if you hate it, because it will bring benefits in the future. Having mere fun may not count as good for you in this prudential sense. This sense would rule out the preference-satisfaction theory immediately, because people often prefer things just because they are fun or nice, even though they bring no long-term benefit. But I am using the expressions ‘good for’ and ‘better for’ more broadly to include anything that is worth your having. So fun is good for you just because it is good in itself, even if it does not lead to good things in the future. Your psychological makeup including your tastes helps to determine what is good for you in this broader sense, so it gives the preference-satisfaction theory a better chance. If you like kayaking more than mountain climbing – perhaps kayaking thrills you more than mountain climbing – then kayaking is better for you than mountain climbing. This broad sense of ‘is better for’ means the same as ‘is more in your interest’ or ‘makes you better off’. When I say *A* is better for a person than *B*, I mean the person is better off

4 Introduction: ethics out of economics

with *A* than with *B*. Having *A* rather than *B* would benefit her; it would improve her wellbeing.

Next the term 'prefers'. Preference is a slippery notion, and I cannot consider all its possible senses. Simply in order to make progress, I shall settle on a definition of 'prefer' that often appears in economics books:<sup>1</sup>

*Dispositional definition of preference.* 'A person prefers *A* to *B*' means the person would choose *A* rather than *B* if she were to have a choice between *A* and *B*.

This definition will certainly need some fixing. It will do for donkeys, but not for people. Jean Buridan tells us that if you stand a donkey exactly halfway between two equally delicious bales of hay, so that it is perfectly indifferent between them, it cannot decide which way to go. It will hesitate till it dies. But such extreme procrastination has never been seen in a person. When a person has a choice between two options, and she is perfectly indifferent between them, she chooses one or the other. According to this definition, then, she prefers one or the other. But that is not so; she is indifferent. So the definition definitely needs fixing to allow for indifference. This is not an important problem, though; the fixing can certainly be done. So let us adopt the dispositional definition of preference, subject to fixing.

After those explanatory remarks, let us go back to the substance of the preference-satisfaction theory of good. This theory is implicit in much of welfare economics. Welfare economists move, almost without noticing it, between saying a person prefers one thing to another and saying she is better off with the first than with the second. All the same, the preference-satisfaction theory is obviously false, and no one really believes it. It is obviously false because people often make mistakes and act on bad information. I brush my teeth thoroughly twice a day, thereby making it true, according to the dispositional definition, that I prefer brushing my teeth thoroughly to giving them a quick swish over. My reason is that I believe thorough brushing promotes my dental health. However, extraordinarily, my dentist has recently been hinting that I ought not to brush my teeth so thoroughly, because it wears them out. If he is right, thorough brushing, which I prefer to a quick swish, is worse for me than a quick swish.

So anyone who is attracted to the preference-satisfaction theory is really attracted to some more sophisticated version such as:

*Ideal-preference-satisfaction theory of good.* One thing *A* is better for a person than another *B* if and only if the person would prefer *A* to *B* were she in ideal conditions.

‘Ideal conditions’ include being well informed, and most people also include being in a rational frame of mind. Undoubtedly, this theory is nearer the truth than the unidealized version. Still, there are many objections to it, too. They have been thoroughly rehearsed by many authors, and I shall not dwell on them here.

I shall not dwell on the faults of the theory, but on its attractions. People, particularly economists, have been drawn to it for various reasons, and some of them are good ones. But the appeal of these reasons has been misdirected. They may give grounds for many of the conclusions welfare economists have drawn from them. For instance, they may give support to economic competition in particular circumstances. But they do not support the preference-satisfaction theory of good. My hope is that, once you see the good reasons that lead elsewhere, you will no longer be attracted by this theory.

**1.3 Liberalism versus a theory of good**

The good reasons I am thinking of are liberal ones. Economists are typically liberal, and typically they believe people should be left alone to manage their own lives. They should make their own choices about private things, and society should be democratically organized to put into effect the people’s choices about public things. These are nice ideas, but they do not lead to the preference-satisfaction theory of good. They lead to conclusions about the structure of society, and what should happen in society, not to any theory of good. I want to draw out this point by means of an example.<sup>2</sup>

A river runs through a fertile valley. Over the years, the forests upstream have been cleared and the dikes containing the river have been built higher and higher. By now, everyone knows that within the next few years the Great Flood will come. The river will breach the dikes. It may go left and destroy one side of the valley, bringing ruin to the farmers there. Or it may go right, and ruin the farmers on that side. (There are equal numbers on each side.) However, the surviving farmers will be relieved of competition from their colleagues, and grow rich. The government is wondering whether to do some expensive engineering works that will definitely prevent the flood. If it does these works, the farmers will be taxed to pay for them. Is it better to do them or not?

If the works are done, all the farmers in the valley will continue to live in modest prosperity, paying the tax for the engineering. If the works are not done, half the farmers will be ruined, but the remainder will rise to luxury, spared from competition and paying no tax. So the choice is between, on the one hand, a state of equality where everyone is modestly prosperous and, on the other, a state of inequality where half the people are ruined and the others are rich. The options are shown in table 1.1. Please adjust the

Table 1.1

Build flood-control works			Do not build flood-control works		
	River would go left	River would go right		River would go left	River would go right
Left-side farmers	Prosperity	Prosperity	Left-side farmers	Ruin	Great wealth
Right-side farmers	Prosperity	Prosperity	Right-side farmers	Great wealth	Ruin

conditions in your mind till you are sure the first alternative is better than the second. Make the state of ruin as dire and the state of modest prosperity as comfortable as they need to be in order to bring you to this conclusion.

Next imagine the farmers are optimists. They all know there will be a flood and that half the farms will be wiped out. But each one thinks it will probably be the ones on the other side of the valley who will be caught. Each farmer assigns a high probability to her own survival, then. This implies the farmers disagree about probabilities. Left-side farmers think the river will probably go right; right-side farmers think it will probably go left. Still, though they disagree, it does not follow that any of the farmers is irrational or ill informed, and indeed I assume they are not. Which dike the river breaches will depend on a complex conjunction of circumstances. Has the heavy construction in this town weakened the dikes? Is the river too tightly squeezed around that bend? The Great Flood is a one-of-a-kind event, and there is no objectively correct probability of the river's going one way or the other. Since each farmer is an optimist, she started by assigning a high prior probability to her own survival. She has taken account of all the information that has become available, and correctly updated her probabilities by applying Bayes's Rule. After that, she still assigns a high posterior probability to her own survival.

What are the farmers' preferences about the flood control works? If the works are done, each farmer will live in modest prosperity. She compares that outlook with her expectation of how things will go for her if the works are not done. If they are not, she will end up either ruined or rich. Please assume she is optimistic enough to prefer the risky option of not doing the works. She thinks riches are sufficiently likely and ruin sufficiently unlikely that she has this preference. So every farmer will prefer not to have the works done. Previously, I asked you to adjust your idea of the conditions till you were sure it was better to do the flood control works. So the situa-

tion is that the people are unanimously in favour of the worse option. Disagreement about probabilities has led to the result that everyone prefers the option that is worse.

What should happen? Since everyone is against them, I am inclined to think the works should not be done. This is presumably what democracy says; unanimous opinion determines what ought to be done. I insisted that doing the works is better than not doing them, so what ought to be done is the worse of the two alternatives. This may sound odd, but there is no contradiction in it. Indeed, if you believe in democracy, you are inevitably committed to believing, sometimes, that what ought to come about is not the best of the alternatives available. When a public decision is to be made, you will evaluate the options on offer, and form a belief about which is the best. At the same time, you may well believe the majority favours another option, and as a democrat you will therefore believe the other option is the one that should come about. So you will believe that what should come about is not the best option. Therefore, unless it is wrong to support democracy, there cannot be any contradiction in this.

When I described the farmers' unanimous preferences, you might have been tempted to go back and revise your view about the goodness of the options. But now you can see you have no need to do that. The preferences determine what ought to be done, but that is a different matter from what is best.

One thing that certainly cannot be true is that the option of not doing the works is better *for everyone* than doing them. It cannot be true, even though everyone prefers that option. Not doing the works will be better for the left-side farmers only if the river is likely to go right, but if the river is likely to go right, then it would be better for the right-side farmers to do the works.

You might think this argument presumes there is an objective probability that the river goes right. But that is not so; I am taking for granted there is no such thing. The basis of my argument is that any statement about goodness in an uncertain situation like this must be made relative to some probabilities. When I make such a statement, I am committing myself to some probability assessment of my own. I can say it is better for the left-side farmers not to do the works only if, like those farmers, I assign a high probability to the river's going right. I cannot then say it is better for the right-side farmers not to do the works. Even though the right-side farmers assign a low probability to the river's going right, I assign a high probability to it, so I must say it is better for them to do the works.

In practice, people differ in their judgements of the probability of many events, and their preferences will be based on their differing probabilities. But when we assess what is good for people, we must do so relative to some probabilities of our own. Therefore, what is good for people cannot always

8 Introduction: ethics out of economics

coincide with their preferences. This is one reason why the preference-satisfaction theory is false.

It constitutes an arcane and rather uninteresting objection to the theory. As I said, my intention was not to produce objections but to try and identify the theory's attraction. The attraction is liberalism and democracy. Correctly understood, however, liberalism and democracy lead to conclusions about what should happen in society, based on people's preferences. They do not imply any theory of good. This example shows we need to separate our account of what is good from our account of what should come about. I think there is a pervasive confusion in economics between what is good and what should come about for other reasons besides good. When an economist favours some economic policy, she needs to decide whether she favours it because it is the best policy, or because it is the liberal or the democratic thing to do.

Why does it matter? Why should economists not simply concern themselves with democratic processes and with preferences, and forget about good and bad altogether? Well, perhaps economics could go that way. It would require some changes in the theoretical structure of welfare economics, but I am not concerned here with what economics should do. I am concerned with what ethics should do. Ethics certainly needs a theory of good, whether economics needs one or not, and all I have said so far is that it should not be tempted to adopt the preference-satisfaction theory from economists. That is a bad theory.

#### 1.4 The formal techniques

But it should adopt some of the formal techniques of economics. Economic theory provides formal structures that correspond closely to the structure of good, and can be used to analyse the structure of good. I am thinking of the core theory of economics: 'preference theory' it is sometimes called, or 'consumer theory', or 'utility theory', or 'expected utility theory'. It is a theory that was originally designed as an account of the structure of a person's preferences. The basic primitive notion of the theory is a preference, which is a two-place relation

*Preference relation.* The person prefers \_\_\_ to \_\_\_.

The blanks are to be filled by things the person has preferences about: the objects of her preferences. The theory makes a number of assumptions – axioms – about the form of this relation. For instance, it assumes it is transitive:

*Transitivity.* If the person prefers *A* to *B* and prefers *B* to *C*, she prefers *A* to *C*.



The axioms provide the basis for many useful theorems about the structure of preferences. In more elaborate branches of the theory, the objects of preference are assumed to have some structure of their own. For instance, the theory accommodates uncertainty by taking the objects of preference to be 'prospects'. A prospect is a sort of portfolio of possible outcomes; it will lead to different outcomes depending on how the uncertainty resolves itself. This gives the basis for more complex axioms and more informative conclusions.

Formally, the theory is nothing more than the axioms and theorems. It can be reinterpreted by substituting in place of preference any other relation that happens to satisfy the axioms. Take these two relations:

*Betterness relation for a person.* \_\_\_ is better for the person than \_\_\_.  
*General betterness relation.* \_\_\_ is better than \_\_\_.

I think each of these betterness relations satisfies the axioms fairly well, though not perfectly. Each probably satisfies them better than people's preferences do in practice. So, provided we are careful, the whole of utility theory is available to provide an analysis of the structure of betterness.

What use is that? I think ethics has a lot to gain from formal analysis. I am going to give two examples of formal features of the structure of good, which it is important for ethics to recognize and understand. These are things that can be learnt from economics.

## 1.5 Think comparatively

The first is very basic: betterness, like preference, is a *relation*. It relates two things together: one thing is better than another. It is a matter of the comparative value of things. Since economists deal so regularly in preferences, they think naturally in comparative terms, and, when they come to think of good, they will naturally ask not what is good, but what is better than what. Philosophers seem not to have this same instinct to think comparatively.

For instance, although the preference-satisfaction theory of good has made some inroads into philosophical ethics, philosophers more often discuss the *want-satisfaction* theory:

*Want-satisfaction theory of good.* *A* is good for a person if and only if the person wants *A*.

This is really a fragment of the preference-satisfaction theory. It is no more true than the preference-satisfaction theory, but a lot more useless. To see its uselessness, let us suppose for a moment it is true. Imagine you meet a thirsty person. She wants water, she wants Coca-Cola, she wants beer, and

she has many other wants too. Suppose you know all her wants in great detail. You know she wants half a pint of water; she wants a pint of water; she wants a litre of beer; she does not want Coke and beer together; and so on. Given all that, what should you give her to drink? A pint of water? A half-pint of Coke? Coke and water together? Just from knowing everything she wants, you cannot tell.

To know what to give her, you need to know her *comparative* wants. You need to know what she wants more than what. You need to know her preferences, that is. Her preferences put all the options in an order: a pint of beer above a pint of water, a half-pint of Coke and half-pint of water above a half-pint of beer, and so on. If you know all her preferences, and if we grant for a moment the preference-satisfaction theory, you know what to give her; you should put her as high up her preference order as you can. But knowing just her wants is not enough.

Similarly, if you know only what is good for a person, even if you know everything that is good for her, that is useless. If you know everything that is good generally, that is useless too. You need to know what is better than what. Comparisons are the data you need for ethics. The lesson is: think comparatively. Economists know this instinctively; philosophers seem not to.

I shall give one example of a philosophical discussion that has been sent off the rails by a failure to think comparatively.<sup>3</sup> One of the principal doctrines of the Epicureans was that it does you no harm to die. Epicurus offers two arguments for this surprising claim. They are both in this quotation, but I am only going to discuss the first. He says:

Become accustomed to the belief that death is nothing to us. For all good and evil consists in sensation, but death is deprivation of sensation . . . So death, the most terrifying of ills, is nothing to us, since so long as we exist death is not with us; but when death comes, then we do not exist. It does not then concern either the living or the dead, since for the former it is not, and the latter are no more.<sup>4</sup>

The first argument is this. The only bad things that can afflict you are sensations. But death gives you no bad sensations because it deprives you of all sensations. So it cannot do you any harm.

This argument depends on the premise that ‘all good and evil consists in sensation’; the only things good for you or bad for you are good sensations or bad sensations. This is the theory of good known as ‘hedonism’. So one way of answering Epicurus’ argument is to deny hedonism. That is what Thomas Nagel does in his article ‘Death’, the most famous recent discussion of the evil of death. Nagel’s aim is to show that death is indeed bad for you, and his answer to Epicurus is to say that hedonism is false.

I agree that hedonism is false. But it has a lot of backers, and indeed there is quite a lot going for it. For instance, it is a better theory than one of its