

“UTILITY”

JOHN BROOME

University of Bristol

I

“Utility,” in plain English, means *usefulness*. In Australia, a ute is a useful vehicle. Jeremy Bentham specialized the meaning to a particular sort of usefulness. “By utility,” he said, “is meant that property in any object, whereby it tends to produce benefit, advantage, pleasure, good, or happiness (all this in the present case comes to the same thing) or (what comes again to the same thing) to prevent the happening of mischief, pain, evil, or unhappiness to the party whose interest is considered” (1823, p. 2). The “principle of utility” is the principle that actions are to be judged by their usefulness in this sense: their tendency to produce benefit, advantage, pleasure, good, or happiness. When John Stuart Mill (1969, p. 213) spoke of the “perfectly just conception of Utility or Happiness, considered as the directive rule of human conduct,” he was using “Utility” as a short name for this principle.¹ “The greatest happiness principle” was another name for it. People who subscribed to this principle came to be known as utilitarians.

Benthamism entered economics in 1873, with the publication of W. S. Jevons’s *Theory of Political Economy*. Jevons quoted Bentham’s definition of “utility” and announced: “This perfectly expresses the meaning of the term in Economy” (1871, p. 46).

But after Jevons’s time, the meaning of “utility” in economics shifted. The word came to refer, not to the tendency of an object to

My thanks to the editors and referees for helpful comments. This article was written while I was a visitor at Princeton University.

1. Mill is not using “utility” as a synonym for “happiness.” The initial capitals and the expression “directive rule of human conduct” make this clear. Mill never identifies utility with happiness or pleasure, though there are some loosely worded remarks in *Utilitarianism* that may suggest he does.

produce good, but to the good an object produces. By a person's "utility," economists came to mean, not the person's usefulness in promoting good around her, but her own good. "Utility" came to mean *good*. This meaning has since been overlaid by yet another, which I shall be describing later. But it still persists as one of the current meanings of "utility."

I cannot give an authoritative history of the shift in meaning. One difficulty is that the interpretation of an author's intentions is often debatable. For Jevons, "utility" definitely meant usefulness in Bentham's sense. Alfred Marshall, too, thought of utilities as useful properties of objects. He said, for instance: "As [man's] production of material products is really nothing more than a rearrangement of matter which gives it new utilities; so his consumption of them is nothing more than a disarrangement of matter, which diminishes or destroys its utilities" (1920, p. 64). But another remark of Marshall's illustrates the difficulty of interpretation: "The *total utility* of a thing to anyone (that is, the total pleasure or other benefit it yields him) increases with every increase of his stock of it, but not as fast as his stock increases" (1920, p. 93). I believe Marshall still meant usefulness by "utility" here. His parenthesis means, I think, that the *amount* of usefulness a thing has is equal to the *amount* of pleasure or other benefit it yields. But this remark could also be read – wrongly, I think – as *identifying* utility with pleasure or other benefit. Presumably it is ambiguities like this that allowed the shift of meaning to proceed unnoticed. And as early as 1881, F. Y. Edgeworth was occasionally using "utility" unambiguously in the shifted sense. He referred to "that quantity which alone the rational unionist is concerned to increase – *the labourer's utility*" (1881, p. 45). An employer might be concerned to increase the laborer's usefulness, but not a unionist. Edgeworth meant the laborer's good. But this was not his normal terminology. Where later economists would have used "utility," Edgeworth (like Jevons) normally used "pleasure." I do not think the shifted usage became common until much later.

Until recently it occurred exclusively in economics. I should be surprised to find an occurrence in philosophy from much before 1960. Henry Sidgwick's *Methods of Ethics*, the *locus classicus* of utilitarianism, hardly used the word "utility" at all. But it contains this footnote about its meaning:

I should point out that Hume uses "utility" in a narrower sense than that which Bentham gave it, and one more in accordance with the usage of ordinary language. He distinguishes the "useful" from the "immediately agreeable": so that while recognising "utility" as the main ground of our moral approbation of the more important virtues, he holds that there are other elements of personal merit which we approve because they are "immediately agreeable,"

either to the person possessed of them or to others. It appears, however, more convenient to use the word in the wider sense in which it has been current since Bentham. (1907, pp. 423–24)

Sidgwick said Bentham widened the sense of “utility,” whereas I say he narrowed it. I do not wish to quarrel about that. No doubt Sidgwick was right that to be immediately agreeable is not, in ordinary usage, a sort of usefulness. It is also true that some sorts of usefulness, according to ordinary usage, Bentham would not have included under “utility” (the usefulness of a thumb-screw, for instance). So Bentham widened the meaning in one way and narrowed it in another. But the point is that agreeableness is the tendency of an object to produce pleasure, not pleasure itself. So whether or not agreeableness is included in utility, utility is still a valuable tendency in an object, not a benefit derived from the object. Sidgwick, at this point in the book, was explaining that virtues have utility in that they “are directly or indirectly productive of pleasure to ourselves or to others” (1907, p. 424).

Recently, however, some philosophers have begun to adopt the economists’ usage. This is unwise. As used by economists, the term “utility” has become so ambiguous as to cause immense confusion. It should be used less, not more.

II

The confusion stems from a new meaning that was assigned to the word as axiomatic utility theory developed during the course of the twentieth century. (It was fully formed in Hicks and Allen, 1934.) The axiomatic theory sets out from a person’s *preferences*. It proves that, provided these preferences conform to some axioms, they can be *represented* by a “utility function.” The values taken by the function are called “utilities.” The sense in which the function represents the preferences is this: of any pair of alternatives, the function assigns a greater utility to the one that is preferred. So “utility” acquired the meaning: *the value of a function that represents a person’s preferences*. This is by now the official definition of utility in economics. For brevity, let us say: utility is *that which represents a person’s preferences*.

Now, let us ask this: of a pair of alternatives, is the one that a person prefers necessarily the one that is better for her? I mean nothing mysterious by this question. I use the word “good” and its cognates – bad, better, best, and so on – in exactly the sense they have in ordinary conversation. Mother is using this sense when she tells you it would be good for you to have a few days rest; an economist when she says inflation is worse for retired people than the unemployed; the politician when she tells you you would be better off dead than red. There is your good and my good: some things are good for you, and some are good

for me. Everybody knows what “good” means, though not many of us can define it. Of course, we endlessly disagree about what things are good and what things are bad. We argue with Mother, the economist, and the politician. One question we might argue over is whether it is necessarily good for a person to have what she prefers. That is the question I am asking now.

Again, then: of a pair of alternatives, is the one that a person prefers necessarily the one that is better for her? According to the official definition of “utility,” it has the greater utility. But a person’s utility, as officially defined, has no necessary connection with her good. So nothing in the definition suggests that the preferred alternative is necessarily better for her. However, many economists adopt the official definition of “utility,” while at the same time *also* using the word to stand for a person’s good. Because an alternative preferred by a person is defined to have a higher utility for her, they take it for granted that it must be better for her. They suppose, then, that a person always prefers what is better for her.

We may call a person who always prefers what is better for herself “self-interested.” This is using “self-interested” in a very strong sense.² It is saying, not only that the person pays no attention to the interests of other people, but also that she always prefers exactly what is in her own interest. It rules out, not only altruism, but also *imprudence*; a person is imprudent if, though concerned only for herself, she sometimes fails to do exactly what is best for herself. Simply by muddling the different meanings of a “utility,” many economists find themselves committed to the view that people are necessarily self-interested in this strong sense.

“The first principle of Economics,” said Edgeworth, “is that every agent is actuated only by self-interest” (1881, p. 16). That may have been true in Edgeworth’s day, but it was one of the achievements of modern utility theory to free economics from such a dubious first principle. The achievement was announced by Lionel Robbins: “So far as we are concerned,” he said, “our economic subjects can be pure egoists, pure altruists, pure ascetics, pure sensualists or – what is much more likely – bundles of all these impulses” (1935, p. 95). The first principle of economics is, I take it, utility theory. And modern, axiomatic utility theory makes no assumption that people are self-interested. All it assumes is that people’s preferences conform to a number of axioms: roughly, they simply need to be consistent. They can conform to the axioms without being self-interested. Yet the muddle over “utility” leads many economists to forget this important discovery.

It is certainly not very plausible that people’s preferences are always self-interested in the strong sense I described. It is a common opinion that many people – parents for instance – have preferences that are

2. As Derek Parfit pointed out to me.

partly altruistic: directed toward the good of others. And it is a common opinion that many people are imprudent: for instance, they prefer to take less exercise than is good for them. If either of these opinions is correct, people’s preferences are not self-interested.

It may turn out that common opinion is incorrect. Altruism and imprudence may not exist, and everyone may always prefer what is best for themselves. Utility, defined to represent a person’s preferences, may indeed turn out also to represent her good. All this is arguable. One argument, for instance, is that a person’s good actually *consists in* the satisfaction of her preferences, so that, of two alternatives, the one she prefers cannot fail to be better for her. But at least the argument needs to be made. It is a substantive question whether or not preferences are necessarily self-interested. If, though, you use “utility” to stand for a representation of a person’s preferences, and at the same time for the person’s good, you cannot even express the question. You will say: by definition, what a person prefers has more utility for her, so how can it fail to have more utility for her? The ambiguity is intolerable.

III

The confusion multiplies when it comes to *expected* utility theory: the branch of utility theory that takes account of uncertainty. Modern, axiomatic expected utility theory imposes axioms on a person’s preferences between uncertain prospects – more axioms than ordinary utility theory does. Granted these axioms, the theory demonstrates the existence of a utility function that has two properties. First, the function represents the preferences, just as before: of two prospects, the preferred one has the higher utility. And second, the function has the “expected utility form,” which means that the utility assigned to an uncertain prospect is the expectation (in probability theory’s sense) of the utilities assigned to the prospect’s possible outcomes. Axiomatic expected utility theory, then, is like ordinary axiomatic utility theory in that it defines utility to represent preferences. “Utility” still means *that which represents preferences*; the person still maximizes her utility. But since the utility of a prospect is also its expected utility – the expectation of the utility of its outcomes – we can also say the person maximizes her expected utility.

It happens that utility defined this way is unique up to increasing linear transformations. Utility with this degree of uniqueness is often said to be *cardinal*.

There is good evidence that people in practice do not conform to expected utility theory; they often violate its axioms (see, for instance, Tversky and Kahneman, 1986). But it can be argued (though this, too, is controversial) that fully rational people will conform to the theory. (My own argument for this point is in Broome, 1991, ch. 5.) For the sake of argument, let us take that for granted, and from now on consider

only fully rational people. So we can assume they are expected utility maximizers. For the sake of argument, too, I now want to set aside the issue of self-interest I mentioned in section II. So let us consider only people whose preferences happen to be self-interested, in the strong sense I described.

Take, then, a rational, self-interested person. When there is no uncertainty to worry about, this person prefers, of two alternatives, the one that is better for her. But what about her preferences between uncertain prospects? Of two prospects, will she necessarily prefer the one that gives her a greater *expectation of good*? Will she necessarily maximize her expected good?

A plausible answer is no, for two reasons. First, for an expectation of good even to exist, good must be an arithmetical quantity.³ And it is plausible that good is not such a precise notion as that. It makes clear sense to say that one prospect is better or worse than another, so goodness at least constitutes an ordering (though not necessarily a complete one). But it is plausible that there are no precise arithmetical quantities of good.

And second, even if there *are* arithmetical quantities of good, it is plausible that a rational, self-interested person might not maximize the expectation of her good. Suppose a person had a choice between 99 units of good for sure, on the one hand, and, on the other, a gamble at equal odds between 0 units and 200 units. The gamble has a higher expectation of good for her. Yet it seems perfectly rational for her to play safe and take the 99 units for sure. She would do this if she were risk-averse about her good. Maximizing the expectation of good implies *risk-neutrality* about good. And it seems perfectly rational to take some different attitude, such as risk-aversion, to risk about one's good.

This needs some more explanation. Our subject is self-interested. Therefore, of two prospects, we can take it that she will prefer the one that is better for her. But it does not follow that she will prefer the one with the greater *expectation of good* for her. The one that is better for her may actually have a lower expectation of good for her. The example shows this. Although the option of 99 units for sure has a lower expectation of good for her, it may nevertheless be better for her, because it is safe.

To be sure, since we are assuming the person conforms to expected utility theory, she maximizes the expectation of her *utility*. This means she is risk-neutral about utility. But it does *not* follow that she is risk-neutral about good. Axiomatic expected utility theory does not imply

3. More precisely, for it to be properly defined which of two alternatives has the greater expectation of good, good needs to be defined uniquely up to an increasing linear transformation.

risk-neutrality about good; it does not imply that a rational person maximizes the expectation of her good.

But once again the ambiguity of "utility" can get in the way of understanding this point. A rational person necessarily maximizes the expectation of utility. For an economist who also uses "utility" to mean good, this will make it seem as though a rational person necessarily maximizes the expectation of her good. But this is a mistaken deduction.

It may turn out, on further investigation, that actually a rational self-interested person *does* necessarily maximize the expectation of her good. This is a common view that began, I believe, with Daniel Bernoulli (1954). I call it "Bernoulli's hypothesis." There are arguments in its favor. (Some appear in Broome, 1991, chs. 10 and 11.) But it is not implied by axiomatic expected utility theory. And we certainly need to be able to ask whether or not it is true. If, however, you define utility, on the one hand, so that a person necessarily maximizes the expectation of it, and, on the other, you covertly identify utility with a person's good, then you cannot ask the question.

Let me express the question differently. Our self-interested subject always prefers, of two alternative prospects, the one that is better for her. Since she prefers it, by definition it has a higher utility. So her utility represents her good in the sense that, of any two prospects, the one that is better for her has the higher utility. This means that utility can properly be called an *ordinal* representation of her good. It has this property simply because the person is self-interested. But if she *also* maximizes the expectation of her good, then expected utility theory tells us that utility will represent her good more tightly than this: it will be an increasing linear transform of her good. In that case, it is said to represent her good *cardinally*. So Bernoulli's hypothesis is equivalent to the claim that utility represents good cardinally. Our question can be put this way, then: for a rational self-interested person, does her utility, as defined by expected utility theory, represent her good cardinally? We know already that utility itself is cardinal; it is unique up to increasing linear transformations. But that does not imply it represents good cardinally. The question is: actually, does it?

Put this way, it is a question that has very much interested welfare economists. A cardinal representation of good is a very useful thing to have if you are interested in evaluating distributions of income and wealth. It is essential if you are a utilitarian concerned to maximize the total of people's good. It was very much in demand in the early 1950s. At that time, welfare economists had been deprived by the "ordinalist revolution" of the 1930s of their right both to cardinal representations of good and to interpersonal comparisons of good. This had left them with almost nothing to say about distributions of income and wealth. The advent of expected utility theory in the 1940s appeared to offer them

back at least the cardinal representations. So they urgently needed to know whether the offer could be trusted. Could expected utility theory really supply cardinal representations of good?

There was great confusion about the question at the time, fueled by the ambiguity of "utility." The confusion showed up in a translation published by *Econometrica* in 1954 of Daniel Bernoulli's seminal article on expected utility theory. Bernoulli said that a rational person would maximize the expectation of her *emolumentum*. (Bernoulli wrote in Latin.) "Emolumentum" means benefit or advantage. So this is a statement of Bernoulli's hypothesis, as I call it: a rational person maximizes her expectation of good. But in *Econometrica*, "emolumentum" was translated as "utility." This lapse in scholarship prevented readers from seeing the crucial difference between Bernoulli's version of expected utility theory and the axiomatic version. The axiomatic version is not committed to the view that a rational person maximizes the expectation of her good. Bernoulli, on the other hand, was.

A clear-sighted article published by Daniel Ellsberg in 1954 ought to have sorted out of the muddle. Ellsberg, like me, was concerned about the ambiguity of "utility." Of von Neumann and Morgenstern, he said:

The operations that define their concepts are essentially new, and their results are neither intended nor suited to fill the main functions of the older, more familiar brands of "cardinal utility." It is unfortunate that old terms have been retained, for their associations arouse both hopes and antagonisms that have no real roots in the new context.

I am sorry to say, however, that confusion persists, still fueled by the same ambiguity. It is at work, for instance, in this more recent passage from John Harsanyi:

To be sure, the vNM utility function⁴ of any given individual is estimated from his choice behavior under risk and uncertainty. But this does not mean that his vNM utility function is *merely* an indication of his attitude towards risk taking. Rather, as its name shows, it is a utility function, and more specifically, it is what economists call a cardinal utility function. This means that the primary task of a vNM utility function is *not* to express a given individual's attitude toward risk taking; rather it is to indicate how much utility, i.e. how much subjective *importance*, he assigns to various goals. (Harsanyi, 1975, p. 600)

4. Von-Neumann–Morgenstern utility function. This is Harsanyi's term for a utility function that has the expected utility form.

IV

We cannot now return "utility" to its original meaning of usefulness. It is a technical term, thoroughly embedded in economics. But at least, as a technical term, we should confine it to one meaning. Which should we choose?

A natural place to look for leadership is the work of Amartya Sen. Let us examine how Sen uses this word "utility" in his recent work. I shall take as examples his books *The Standard of Living* and *On Ethics and Economics* (1987a, 1987b).

Sen's notion of utility plays an important role in his arguments. A major thesis of *The Standard of Living*, for instance, is that a person's standard of living cannot be identified with her utility. But Sen evidently assumes that his readers understand what he means by "utility" and does not explain it well. I do not find his meaning obvious, though. Two things are clear. Sen does not mean by "utility" what axiomatic utility theory means by it (see, e.g., 1987a, pp. 14–15). Nor does he mean a person's good; there are things he considers good for a person – for instance, her functionings and capabilities – that he does not include in her utility.

My best understanding of Sen's meaning for "utility" is: *that which utilitarians believe to constitute good*. His reason for adopting this meaning, I suppose, is this. "Utility" as a technical term was invented by utilitarians, so we ought to give it the meaning they give it. And utilitarians intend "utility" to refer to what they believe to constitute good, so we ought to use it to refer to that, too.

But this is a poor reason. First of all, it is false that, generally, utilitarians intend "utility" to refer to what they believe to constitute good. I said that in section I. Among the classical utilitarian philosophers – Bentham, Mill, Sidgwick – none of them used the term that way. They intended it to refer to the tendency to promote good. Among the classical utilitarian economists – Jevons, Marshall, Edgeworth, Pigou⁵ – only Edgeworth occasionally slipped into this usage. Even contemporary utilitarian philosophers rarely use "utility" that way. Generally, utilitarians refer to those things they believe to constitute good by their specific names: "pleasure," "happiness," "satisfaction," "wellbeing," "welfare," and so on.

Second, even if we suppose that utilitarians *do* intend "utility" to refer to what they believe to constitute good, and even if we want to give "utility" the meaning they give it, we should not use it to refer to the same thing as they do, unless we are utilitarians. (And Sen is not one.) In so far as utilitarians use this word in this way, they *mean* by it

5. Pigou (1932) did not use "utility" at all.

simply good. Of course, they intend it to *refer* to what they believe to constitute good. But if we are to give the word the same meaning as they do, we must use it to refer to what *does* constitute good. Only if utilitarians are right, will this be just what utilitarians believe to constitute good.

If, then, Sen means by "utility" that which utilitarians believe to constitute good, he is endowing the word with a new meaning. I believe this meaning may be unique to Sen. It is likely to be a useful one only if it is clear what utilitarians believe to constitute good; only then will its reference be unambiguous. But actually utilitarians are divided over what constitutes good. Some think good consists in good feelings, some in happiness, others in the satisfaction of desires, and so on.

Because of this ambiguity, Sen is constantly forced in these books to mention different "conceptions of utility" separately: the happiness conception, the desire-satisfaction conception, and so on. His arguments have to deal separately with each. In showing that the standard of living is not the same as utility, for instance, he has to show it for the various conceptions one by one (1987b, pp. 5–14). And since the conceptions differ radically, each demands a radically different argument. I see, therefore, little point in collecting them together under the one heading of "utility." It would be a harmless thing to do if "utility" were not already damagingly ambiguous. As it is, we can do without yet another meaning for it, and especially one that is itself ambiguous.

So I think Sen's lead points in the wrong direction.

V

As a meaning for "utility," we should choose either *that which represents preferences*, on the one hand, or *good*, on the other. Which should it be?

Both are economists' meanings. Philosophers have not used either until recently, and neither is well established in philosophy. So the needs of economics should have the first say in deciding between them. In any case, they will.

In economics the official meaning is the first: *that which represents preferences*. This is the meaning given in the major doctrinal texts (e.g., Debreu, 1959, p. 56) and in the best textbooks (e.g., Deaton and Muellbauer, 1980, p. 28). It is defined with great precision, as a technical term should be. Its use is universal in theoretical economics. There is no alternative term. Economics cannot do without it.

"Utility" in this sense need not be confined to a representation of a person's *actual* preferences. A function can also be called a utility function if it represents the preferences a person would have if she were rational and self-interested. But if she were rational and self-interested,

she would prefer, of two alternatives, the one that is better for her. So a function that represents the preferences she would have if she were rational and self-interested also represents, ordinally, her good. An ordinal representation of good can therefore be called a utility function. (I use “utility” this way myself in Broome, 1991.) We might also have a use for a notion of “social” preferences, suitably interpreted, and social preferences, too, could be represented by a utility function. All of this is within the scope of the official definition of “utility.”

Ellsberg thought it unfortunate that “utility” acquired its new technical meaning during the twentieth century (see the quotation in section III earlier). I disagree. Once divorced from usefulness, the word has no *natural* meaning. So it is ideally suited to perform the services of a technical term.

The second meaning, *good*, is also the property of economists. But it is an underground one. You will not often find it openly acknowledged. There is no classical warrant for it. Hiding underground, its main effect has been to cause confusion. And – a further sin – it is perfectly redundant. We already have an excellent word with the meaning of good: “good.”

I therefore propose that this second meaning for “utility” should be prohibited. “Utility” should be used only for a representation of preferences.

REFERENCES

- Bentham, Jeremy. 1823. *An Introduction to the Principles of Morals and Legislation*. London: Pickering.
- Bernoulli, Daniel. 1954. “Exposition of a New Theory on the Measurement of Risk.” Translated by Louise Sommer. *Econometrica* 22:23–36. Originally published in 1738.
- Broome, John. 1991. *Weighing Goods*. Oxford: Blackwell.
- Deaton, Angus, and John Muellbauer. 1980. *Economics and Consumer Behavior*. New York: Cambridge University Press.
- Debreu, Gerard. 1959. *Theory of Value*. New York: Wiley.
- Edgeworth, F. Y. 1881. *Mathematical Psychics*. London: Kegan Paul.
- Ellsberg, Daniel. 1954. “Classic and Current Notions of ‘Measurable Utility’.” *Economic Journal* 64:528–56.
- Harasanyi, John C. 1975. “Can the Maximin Principle Serve as a Basis for Morality? A Critique of John Rawls’s Theory.” *American Political Science Review* 69:594–606.
- Hicks, John, and R. G. D. Allen. 1934. “A Reconsideration of the Theory of Value.” *Economica* 1:52–76, 196–219.
- Jevons, W. Stanley. 1871. *The Theory of Political Economy*. London: Macmillan.
- Marshall, Alfred. 1920. *Principles of Economics*, 8th ed. London: Macmillan.
- Mill, John Stuart. 1969. *Utilitarianism*. In his *Collected Works*, Vol. 10. Toronto: Toronto University Press.
- Pigou, A. C. 1932. *The Economics of Welfare*, 4th ed. London: Macmillan.
- Robbins, Lionel. 1935. *An Essay on the Scope and Nature of Economic Science*, 2nd ed. London: Macmillan.
- Sen, Amartya. 1987a. *On Ethics and Economics*. Oxford: Blackwell.

- . 1987b. *The Standard of Living*, edited by Geoffrey Hawthorne. New York: Cambridge University Press.
- Sidgwick, Henry. 1907. *The Methods of Ethics*, 7th ed. London: Macmillan.
- Tversky, Amos, and Daniel Kahneman. 1986. "Rational Choice and the Framing of Decisions." *Journal of Business* 59:250–78. Reprinted in *The Limits of Rationality*, edited by Karen Cook and Margaret Levi. Chicago: Chicago University Press. 1991.