

Utilitarianism

Roger Crisp / Tim Chappell

Utilitarianism is a theory about rightness, according to which the only good thing is welfare (wellbeing or 'utility'). Welfare should, in some way, be maximized, and agents are to be neutral between their own welfare, and that of other people and of other sentient beings.

The roots of utilitarianism lie in ancient thought. Traditionally, welfare has been seen as the greatest balance of pleasure over pain, a view discussed in Plato. The notion of impartiality also has its roots in Plato, as well as in Stoicism and Christianity. In the modern period, utilitarianism grew out of the Enlightenment, its two major proponents being Jeremy Bentham and John Stuart Mill.

Hedonists, believing that pleasure is the good, have long been criticized for sensualism, a charge Mill attempted to answer with a distinction between higher and lower pleasures. He contended that welfare consists in the experiencing of pleasurable mental states, suggesting, in contrast to Bentham, that the quality, not simply the amount, of a pleasure is what matters. Others have doubted this conception, and developed desire accounts, according to which welfare lies in the satisfaction of desire. Ideal theorists suggest that certain things are just good or bad for people, independently of pleasure and desire.

Utilitarianism has usually focused on actions. The most common form is act-utilitarianism, according to which what makes an action right is its maximizing total or average utility. Some, however, have argued that constantly attempting to put utilitarianism into practice could be self-defeating, in that utility would not be maximized by so doing. Many utilitarians have therefore advocated non-utilitarian decision procedures, often based on common sense morality. Some have felt the appeal of common sense moral principles in themselves, and sought to reconcile utilitarianism with them. According to rule-utilitarianism, the right action is that which is consistent with those rules which would maximize utility if all accepted them.

There have been many arguments for utilitarianism, the most common being an appeal to reflective belief or 'intuition'. One of the most interesting is Henry Sidgwick's argument, which is ultimately intuitionist, and results from sustained reflection on common sense morality. The most

famous argument is Mill's 'proof'. In recent times, R.M. Hare has offered a logical argument for utilitarianism.

The main problems for utilitarianism emerge out of its conflict with common sense morality, in particular justice, and its impartial conception of practical reasoning.

1. Introduction and history

Defining utilitarianism is difficult, partly because of its many variations and complexities, but also because the utilitarian tradition has always seen itself as a broad church. But before offering a history, we must supply a working definition. First, utilitarianism is, usually, a version of welfarism, the view that the only good is welfare (see Welfare). Second, it assumes that we can compare welfare across different people's lives (see Economics and ethics). Third, it is a version of consequentialism (see Consequentialism). Consequentialists advocate the impartial maximization of certain values, which might include, say, equality. Utilitarianism is welfarist consequentialism, in its classical form, for instance, requiring that any action produce the greatest happiness (see Happiness).

The concern with welfare, its measurement and its maximization is found early, in Plato's Protagoras. In the process of attempting to prove that all virtues are one, Socrates advocates hedonism, the welfarist view that only pleasurable states of mind are valuable, and that they are valuable solely because of their pleasurableness (see Plato §9; Socrates §24; Hedonism).

The debate in the Protagoras is just one example of the many discussions of welfare in ancient ethics (see Eudaimonia). Some have seen Greek ethics as primarily egoistic, addressing the question of what each individual should do to further their own welfare (see Egoism and altruism §4). Utilitarianism, however, is impartial.

The Stoics, who followed Plato and Aristotle, began to develop a notion of impartiality according to which self-concern extended rationally to others, and eventually to the whole world (see Stoicism §18). This doctrine, allied to Christian conceptions of self-sacrifice, and conceptions of rationality with roots in Plato which emphasize the objective supra-individual point of view, could plausibly be said to be the source of utilitarian impartiality (see Impartiality).

In the modern period, the history of utilitarianism takes up again during the Enlightenment. The idea of impartial maximization is found in the work of the eighteenth-century Scottish philosopher Francis Hutcheson (1755) (see Hutcheson, F. §2). The work of his contemporary, David Hume

(1751), also stressed the importance to ethics of the notion of 'utility' (see Hume, D. §4.1). A little later, the so-called 'theological utilitarians', Joseph Priestley (1768) and William Paley (1785), argued that God requires us to promote the greatest happiness (see Priestley, J.; Paley, W.). Meanwhile, in France, Claude Helvétius (1758) advocated utilitarianism as a political theory, according to which the task of governments is to produce happiness for the people. He influenced one of the most extreme of all utilitarians, William Godwin (1793) (see Helvétius, C.; Godwin, W.).

It was Jeremy Bentham, however, who did most to systematize utilitarianism. Bentham's disciple, J.S. Mill, was the next great utilitarian, and he was followed by Henry Sidgwick. G.E. Moore (1903) distanced himself from Mill's hedonism, and offered an influential 'ideal' account of the good. One of the most important recent versions of utilitarianism is that of R.M. Hare (see Moore, G.E.; Hare, R.M.).

2. Conceptions of utility

Before you can maximize utility, you need to know what utility is. It is essential to note that the plausibility of utilitarianism as a theory of right action does not depend on any particular conception of welfare. An account of the good for a person is different from an account of right action (see Right and good).

Utilitarians have held many different views of utility. The 'classical' utilitarians – primarily Bentham (1789) and Mill (1861) – were hedonists. There are many objections to hedonism. What about masochists, for example, who seem to find pain desirable? Well, perhaps pain can be pleasurable. But is there really something common – pleasure – to all the experiences that go to make up a happy life? And would it be rational to plug oneself into a machine that gave one vast numbers of pleasurable sensations? Here there may be a move towards the more eclectic view of Sidgwick (1874), that utility consists in desirable consciousness of any kind. Some philosophers, however, such as Nietzsche (1888), have suggested that a life of mere enjoyment is inauthentic.

Hedonists have been criticized for sensualism for millennia. J.S. Mill sought to answer the charge, suggesting that hedonists do not have to accept that all pleasurable experiences – drinking lemonade and reading Wordsworth – are on a par, to be valued only according to the amount of pleasure they contain. Bentham and others had suggested that the value of a pleasure depends mainly on its intensity and its duration, but Mill insisted that the quality of a pleasure – its nature – also influences its pleasurableness and hence its value. But why must the effect on value of

the nature of an experience be filtered through pleasurable-ness? Why cannot its nature by itself add value?

Perhaps the most serious objection to any theory that welfare consists in mental states is the so-called 'experience machine'. This machine is better than the pleasure machine, and can give you the most desirable experiences you can imagine. Would it be best for you to be wired up to it throughout your life? Note that this is not the question whether it would be right to arrange for yourself to be wired up, leaving all your obligations in the real world unfulfilled. Even a utilitarian can argue that that would be immoral.

Some people think it makes sense to plug in, others that it would be a kind of death. If you are one of the latter, then you might consider moving to a desire theory of utility, according to which what makes life good for you is your desires' being maximally fulfilled. On the experience machine, many of your desires will remain unfulfilled. You want not just the experience of, say, bringing about world peace, but actually to bring it about. Desire theories have come to dominate contemporary thought because of economists' liking for the notion of 'revealed preferences' (see Rationality, practical). Pleasures and pains are hard to get at or measure, whereas people's preferences can be stated, and inferred objectively from their behaviour.

A simple desire theory fails immediately. I desire the glass of liquid, thinking it to be whisky. In fact it is poison, so satisfying my desire will not make me better off. What desire theorists should say here is that it is the satisfaction of intrinsic desires which counts for wellbeing. My intrinsic desire is for pleasure, the desire for the drink being merely derived.

The usual strategy adopted by desire theorists is to build constraints into the theory in response to such counterexamples: what makes me better off is not the fulfilment of my desires, but of my informed desires.

But why do desire theorists so respond to such counterexamples? It is probably because they already have a view of utility which guides them in the construction of their theories. This means that desire theories are themselves idle, which is to be expected once we realize that the fulfilment of a desire is in itself neither good nor bad for a person. What matters is whether what the person desires, and gets, is good or bad.

For reasons such as this, there is now a return to ancient ideal theories of utility, according to which certain things are good or bad for beings, independently in at least some cases of whether they are desired or whether they give rise to pleasurable experiences (see Perfectionism). Another interesting ancient view which has recently been revived is that certain nonhedonistic goods are valuable, but only when they are combined with pleasure or desire-fulfilment (see Plato, *Philebus* 21a-22b).

The nonhedonistic goods suggested include knowledge and friendship. Questions to ask of the ideal theorist include the following. What will go on your list of goods? How do you decide? How are the various items to be balanced?

3. Types of utilitarianism

Theories of right and wrong have to be about something, that is, have to have a focus. Usually, at least in recent centuries, they have focused on actions, attempting to answer the questions, 'Which actions are right?', and, 'What makes those actions right?'. The ancients also asked these questions, but were concerned also to focus on lives, characters, dispositions and virtues. Nearly all forms of utilitarianism have focused on actions, but in recent decades there has been some interest in utilitarianism as applied to motives, virtues and lives as a whole.

Utilitarianism is a form of consequentialism. But it is important to note that, since utilitarians can attach intrinsic moral importance to acts (especially, of course, the act of maximizing itself), there are problems in attempting to capture the nature of utilitarianism using the act/consequence distinction. A recent alternative has been to employ the 'agent-neutral'/agent-relative' distinction. Agent-neutral theories give every agent the same aim (for example, that utility be maximized), whereas agent-relative theories give agents different aims (say, that your children be looked after). Logically, however, there is nothing to prevent a utilitarian's insisting that your aim should be that you maximize utility. Though this theory would be practically equivalent to an agent-neutral theory, its possibility suggests there may be problems with attempting to use the agent-neutral/agent-relative distinction to capture the essence of utilitarianism.

What clearly distinguishes utilitarianism from other moral theories is what it requires and why, so we should now turn to that. The commonest, and most straightforward, version of utilitarianism is act-utilitarianism, according to which the criterion of an action's rightness is that it maximize utility.

Act-utilitarians might offer two accounts of rightness. The objectively right action would be that which actually does maximize utility, while the subjectively right action would be that which maximizes expected utility. Agents would usually be blamed for not doing what was subjectively right.

Another distinction is between total and average forms. According to the total view, the right act is the one that produces the largest overall total of utility. The average view says that the right action is that which

maximizes the average level of utility in a population. The theories are inconsistent only in cases in which the size of a population is under consideration. The most common such case occurs when one is thinking of having a child. Here, the average view has the absurd conclusion that I should not have a child, even if its life will be wonderful and there will be no detrimental effects from its existence, if its welfare will be lower than the existing average.

But the total view also runs into problems, most famously with Derek Parfit's 'repugnant conclusion' (1984), which commits the total view to the notion that if a population of people with lives barely worth living is large enough it is preferable to a smaller population with very good lives. One way out of this problem is to adopt a person-affecting version of utilitarianism, which restricts itself in scope to existing people. But there are problems with this view (see Parfit 1984: ch. 18). Recently, certain writers have suggested that one way to avoid the 'repugnant conclusion' would be to argue that there are discontinuities in value, such that once welfare drops below a certain level the loss cannot be compensated for by quantity. There is a link here with Mill's view of the relation of higher pleasures to lower.

Imagine being an act-utilitarian, brought up in an entirely act-utilitarian society. You will have to spend much time calculating the utility values of the various actions open to you. You are quite likely to make mistakes, and, being human, to cook the books in your own favour.

For these reasons, most act-utilitarians have argued that we should not attempt to put act-utilitarianism into practice wholesale, but stick by a lot of common sense morality (see Common-sense ethics). It will save a lot of valuable time, is based on long experience, and will keep us on the straight and narrow. Act-utilitarians who recommend sole and constant application of their theory as well as those who recommend that we never consult the theory and use common sense morality can both be called single-level theorists, since moral thinking will be carried on only at one level. Most utilitarians have adopted a two-level theory, according to which we consult utilitarianism only sometimes – in particular when the principles of ordinary morality conflict with one another.

The main problem with two-level views is their psychology. If I really accept utilitarianism, how can I abide by a common sense morality I know to be a fiction? And if I really do take that common sense morality seriously, how can I just forget it when I am supposed to think as a utilitarian? The two-level response here must be that this is indeed a messy compromise, but one made to deal with a messy reality.

Act-utilitarianism is an extremely demanding theory, since it requires you to be entirely impartial between your own interests, the interests of those you love, and the interests of all. The usual example offered is famine relief. By giving up all your time, money and energy to famine relief, you will save many lives and prevent much suffering. Utilitarians often claim at this point that there are limits to human capabilities, and utilitarianism requires us only to do what we can. But the sense of 'can' here is quite obscure, since in any ordinary sense I can give up my job and spend my life campaigning for Oxfam.

The demandingness objection seems particularly serious when taken in the context of widespread non-compliance with the demands of act-utilitarian morality. Most people do little or nothing for the developing world, and this is why the moral demands on me are so great. An argument such as this has been used to advocate rule-utilitarianism, according to which the right action is that which is in accord with that set of rules which, if generally or universally accepted, would maximize utility. (The version of the theory which speaks of the rules that are obeyed is likely to collapse into act-utilitarianism; see Lyons 1965.)

Unlike act-utilitarianism, which is a direct theory in that the rightness and wrongness of acts depends directly on whether they fit with the maximizing principle, rule-utilitarianism is an indirect theory, since rightness and wrongness depend on rules, the justification for which itself rests on the utilitarian principle.

The demandingness of act-utilitarianism has not been the main reason for adopting rule-utilitarianism. Rather, the latter theory has been thought to provide support for common sense moral principles, such as those speaking against killing or lying, which appear plausible in their own right.

Rule-utilitarianism has not received as much attention as act-utilitarianism, partly because it detaches itself from the attractiveness of maximization. According to rule-utilitarianism there may be times when the right action is to bring about less than the best possible world (such as when others are not complying). But if maximization is reasonable at the level of rules, why does it not apply straightforwardly to acts?

4. Arguments for utilitarianism

The most famous argument for utilitarianism is John Stuart Mill's 'proof' (1861). This has three stages:

1. Happiness is desirable.

2. The general happiness is desirable.
3. Nothing other than happiness is desirable.

Each stage has been subjected to much criticism, especially the first. Mill was an empiricist, who believed that matters of fact could be decided by appeal to the senses (see Empiricism). In his proof, he attempted to ground evaluative claims on an analogous appeal to desires, making unfortunate rhetorical use of 'visible' and 'desirable'. The first stage suggests to the reader that if they consult their own desires, they will see that they find happiness desirable.

The second stage is little more than assertion, since Mill did not see the vastness of the difference between egoistic and universalistic hedonism (utilitarianism). In an important footnote (1861: ch. 5, para. 36), we see the assumption that lies behind the proof: the more happiness one can promote by a certain action, the stronger the reason to perform it. Egoists will deny this, but it does put the ball back in their court.

The final stage again rests on introspection, the claim being that we desire, ultimately, only pleasurable states. Thus even a desire for virtue can be seen as a desire for happiness, since what we desire is the pleasure of acting virtuously or contemplating our virtue. One suspects that introspection by Mill's opponents would have had different results.

Perhaps the most common form of utilitarianism, as of any other moral theory, is, in a weak sense, intuitionist (see Intuitionism in ethics). To many, utilitarianism has just seemed, taken by itself, reasonable – so reasonable, indeed, that any attempt to prove it would probably rest on premises less secure than the conclusion. This view was expressed most powerfully by Henry Sidgwick (1874). Sidgwick supported his argument with a painstaking analysis of common sense morality. Sidgwick also believed that egoism was supported by intuition, so that practical reason was ultimately divided (see Egoism and altruism §§1, 3).

In the twentieth century, R.M. Hare wished to avoid appeal to moral intuition, which he saw as irrational. According to Hare (1981), if we are going to answer a moral question such as, 'What ought I to do?', we should first understand the logic of the words we are using. In the case of 'ought', we shall find that it has two properties: prescriptivity (it is action-guiding) and universalizability (I should be ready to assent to any moral judgment I make when it is applied to situations similar to the present one in their universal properties) (see Prescriptivism). Hare argues that putting yourself in another's position properly – 'universalizing' – involves taking on board their preferences. Once this has been done, the only rational

strategy is to maximize overall preference-satisfaction, which is equivalent to utilitarianism.

Hare's moral theory is one of the most sophisticated since Kant's, and he does indeed claim to incorporate elements of Kantianism into his theory (see Kantian ethics). Objectors have claimed, however, that, rather like Kant himself, Hare introduces 'intuitions' (that is, beliefs about morality or rationality) through the back door. For example, the logic of the word 'ought' may be said not to involve a commitment to the rationality of maximization even in one's own case.

5. Problems for utilitarianism

There are many technical problems with the various forms of utilitarianism. How are pleasure and pain to be measured? Which desires are to count? Is knowledge a good in itself? Should we take into account actual or probable effects on happiness? How do we characterize the possible world which is to guide us in our selection of rules? These are problems for the theorists themselves, and there has been a great deal said in attempts to resolve them.

More foundational, however, is a set of problems for any kind of utilitarian theory, emerging out of utilitarianism's peculiarly strict conception of impartiality. A famous utilitarian tag, from Bentham, is, 'Everybody to count for one, nobody for more than one'. This, however, as Mill implies (1861: ch. 5, para. 36), is slightly misleading. In a sense, according to utilitarianism, no one matters; all that matters is the level of utility. What are counted equally are not persons but pleasures or utilities.

This conception of impartiality has made it easy for opponents of utilitarianism to dream up examples in which utilitarianism seems to require something appalling. A famous such example requires a utilitarian sheriff to hang an innocent man, so as to prevent a riot and bring about the greatest overall happiness possible in the circumstances (see *Crime and punishment* §2).

Utilitarians can here respond that, in practice, they believe that people should abide by common sense morality, that people should accept practical principles of rights for utilitarian reasons (see §2). But this misses the serious point in many of these objections: that it matters not just how much utility there is, but how it is shared around. Imagine, for example, a case in which you can give a bundle of resources either to someone who is well-off and rich through no fault of their own, or to someone who is poor through no fault of their own. If the utility of giving the bundle to the rich person is only slightly higher than that of giving it to the poor person,

utilitarianism dictates giving it to the rich person. But many (including some consequentialists) would argue that it is reasonable to give some priority to the worse-off.

These are problems at the level of the social distribution of utility. But difficulties arise also because of the fact that human agents each have their own lives to live, and engage in their practical reasoning from their own personal point of view rather than from the imaginary point of view of an 'impartial spectator'. These problems have been stated influentially in recent years by Bernard Williams (Smart and Williams 1973), who puts them under the heading of what he calls 'integrity' (see Williams, B.A.O. §4).

In a famous example, Williams asks us to imagine the case of Jim, who is travelling in a South American jungle. He comes across a military firing squad, about to shoot twenty Indians from a nearby village where some insurrection has occurred. The captain in charge offers Jim a guest's privilege. Either Jim can choose to shoot one of the Indians himself, and the others will go free, or all twenty will be shot by the firing squad.

Williams' point here is not that utilitarianism gives the wrong answer; indeed he himself thinks that Jim should shoot. Rather, it is that utilitarianism reaches its answer too quickly, and cannot account for many of the thoughts we know that we should have ourselves in Jim's situation, such as, 'It is I who will be the killer'. Practical reasoning is not concerned only with arranging things so that the greatest utility is produced. Rather it matters to each agent what role they will be playing in the situation, and where the goods and bads occur. This point emerges even more starkly if we imagine a variation on the story about Jim, in which the captain asks Jim to commit suicide so as to set an example of courage and nobility to the local populace, on the condition that if he does so the twenty Indians will go free. The utility calculations are as clear, perhaps clearer, than in the original story. But it is only reasonable that Jim in this story should think it relevant that it is he who is going to die. To any individual, it matters not only how much happiness there is in the world, but who gets it.

References and further reading

Adams, R.M. (1976) 'Motive utilitarianism', *Journal of Philosophy* 73: 467-81. (Utilitarianism with focus on motives.)

Bentham, J. (1789) *An Introduction to the Principles of Morals and Legislation*, ed. J.H. Burns and H.L.A. Hart, revised F. Rosen, Oxford: Clarendon Press, 1996. (Highly influential statement of hedonistic act-utilitarianism.)

- Brandt, R.B. (1985) *A Theory of the Good and the Right*, Oxford: Clarendon Press. (Defends hedonistic rule-utilitarianism.)
- Brink, D. (1989) *Moral Realism and the Foundations of Ethics*, Cambridge: Cambridge University Press. (Defends two-level ideal utilitarianism.)
- Broome, J. (1991) *Weighing Goods*, Oxford: Blackwell, ch. 1. (Contains discussion of structure of moral theories. Moderately difficult.)
- Crisp, R. (1992) 'Utilitarianism and the life of virtue', *Philosophical Quarterly* 42: 139-60. (Discussion of two-level utilitarianism with focus on lives.)
- Frey, R. (ed.) (1985) *Utility and Rights*, Oxford: Blackwell. (Collection devoted to discussion of rights and utilitarianism.)
- Glover, J. (ed.) (1990) *Utilitarianism and Its Critics*, New York: Macmillan. (Contains essential readings, with helpful guidance by the editor. Good starting place; helpfully structured bibliography.)
- Godwin, W. (1793) *Enquiry Concerning Political Justice*, ed. K. Codell, Oxford: Clarendon Press, 1971. (Radical anarchist utilitarianism.)
- Goodin, R. (1991) 'Utility and the good', in P. Singer (ed.) *A Companion to Ethics*, Oxford: Blackwell, 241-8. (Introductory discussion of welfare.)
- Griffin, J. (1982) 'Modern utilitarianism', *Revue Internationale de Philosophie* 141: 331-75. (Good starting place for study of recent work on utilitarianism.)
- Griffin, J. (1986) *Well-Being: Its Meaning, Measurement, and Moral Importance*, Oxford: Clarendon Press. (Key discussion of different conceptions of utility. Defends ideal account with preferences entering at level of explanation of value.)
- Hare, R.M. (1981) *Moral Thinking: Its Methods, Levels and Point*, Oxford: Clarendon Press. (Sophisticated modern two-level act-utilitarian, influenced by Aristotle, Kant and Mill.)
- Helvétius, C. (1758) *De l'esprit (On the Mind)*, Paris: Arthème Fayard, 1988. (Enlightenment utilitarian political theorist.)

Hooker, B. (1990) 'Rule consequentialism', *Mind* 91: 67-77. (Contains defence of rule-consequentialism – also applicable to rule-utilitarianism – on grounds of undemandingness and fairness.)

Hume, D. (1751) *An Enquiry Concerning the Principles of Morals*, ed. L.A. Selby-Bigge, revised by P.H. Nidditch, Oxford: Clarendon Press, 3rd edn, 1975. (See esp. section 5, 'Why utility pleases', on the usefulness of the social virtues.)

Hutcheson, F. (1755) *A System of Moral Philosophy*, in *Collected Works*, vols 5-6, Hildesheim: Georg Olms, 1969. (Important Enlightenment work, anticipating Bentham and Mill.)

Kagan, S. (1989) *The Limits of Morality*, Oxford: Clarendon Press. (Defends demandingness of utilitarianism. Moderately difficult.)

Long, A.A. and Sedley, D. (1987) *The Hellenistic Philosophers*, vol. 1, Cambridge: Cambridge University Press, sections 56-67. (Contains excerpts from and discussions of Stoic ethics.)

Lyons, D. (1965) *Forms and Limits of Utilitarianism*, Oxford: Clarendon Press. (Classic discussion of act- and rule-utilitarianism.)

McCloskey, H.J. (1957) , *Philosophical Review* 66: 466-85. (Original version of the 'sheriff' case, in which a utilitarian sheriff is required to hang an innocent man, so as to prevent a riot and bring about the greatest overall happiness possible in the circumstances.)

McNaughton, D. and Rawling, P. (1995) 'Value and agent-relative reasons', *Utilitas* 7: 31-47. (Sophisticated attempt to capture consequentialist / nonconsequentialist distinction using deontic logic. Difficult.)

Mill, J.S. (1861) *Utilitarianism*, ed. R. Crisp, Oxford: Clarendon Press, 1998. (One of the most important and widely studied works in moral philosophy. Contains argument that pleasures can be seen as higher and lower.)

Miller, H.B. and Williams, W.H. (eds) (1982) *The Limits of Utilitarianism*, Minneapolis, MN: University of Minnesota Press. (Interesting and wide-ranging collection. Contains bibliography of works on utilitarianism from 1930 to 1980.)

Moore, G.E. (1903) *Principia Ethica*, ed. T. Baldwin, Cambridge: Cambridge University Press, revised edn, 1993. (Includes critique of Mill, and development of ideal account of the good.)

Nietzsche, F. (1888) *Die Götzen-Dämmerung*, trans. A.M. Ludovici, *The Twilight of the Idols*, London: Allen & Unwin, 1911, 94-6. (Criticizes the life consisting solely of happiness or pleasure.)

Nozick, R. (1984) *Anarchy, State, and Utopia*, Oxford: Blackwell. (Contains the 'experience machine' objection and defence of rights as constraints on maximization.)

Nussbaum, M. and Sen, A. (eds) (1993) *The Quality of Life*, Oxford: Clarendon Press. (Wide-ranging collection on welfare.)

Paley, W. (1785) *Principles of Moral and Political Philosophy*, New York: Garland, 1978. (Theological utilitarian, whose great success is said to have spurred Bentham to publish his *Introduction to the Principles*.)

Parfit, D. (1984) *Reasons and Persons*, Oxford: Clarendon Press. (Central modern work. Part IV contains discussion of population problems.)

Pettit, P. (1991) 'Consequentialism', in P. Singer (ed.) *A Companion to Ethics*, Oxford: Blackwell, 230-40. (Introductory discussion of consequentialism. Has brief bibliography.)

Plato (c.386-380 BC) *Protagoras*, trans. with notes by C.C.W. Taylor, Oxford: Clarendon Press, 2nd edn, 1991. (Develops early account of hedonism.)

Plato (c.360-347 BC) *Philebus*, trans. J.C.B. Gosling, *Plato: Philebus*, Oxford: Clarendon Press, 1975. (A discussion of the relative merits of pleasure and wisdom in the good life.)

Priestley, J. (1768) *Essay on the First Principles of Government; and on the Nature of Political, Civil, and Religious Liberty*, London: Dodsley, Cadell, Johnson; repr. in P. Miller (ed.) *Political Writings*, Cambridge: Cambridge University Press, 1993. (Influential theological utilitarian.)

Quinton, A. (1973) *Utilitarian Ethics*, London: Macmillan. (Historical introduction. Includes chapters on Bentham and Mill.)

Railton, P. (1984) 'Alienation, consequentialism, and the demands of morality', *Philosophy and Public Affairs* 13: 134-71. (On the everyday moral thinking recommended by utilitarianism.)

Rawls, J. (1971) *A Theory of Justice*, Cambridge, MA: Harvard University Press. (Significant defence of anti-utilitarian view of justice, based on 'separateness of persons'. Moderately difficult.)

Scarre, G. (1996) *Utilitarianism*, London: Routledge. (Useful introduction, including history. Contains bibliography.)

Scheffler, S. (ed.) (1988) *Consequentialism and Its Critics*, Oxford: Oxford University Press. (Collection of important recent critiques, with helpful introduction.)

Sen, A. (1981) 'Plural utility', *Proceedings of the Aristotelian Society* 81: 193-215. (Defends pluralistic conception of utility.)

Sen, A. and Williams, B. (eds) (1982) *Utilitarianism and Beyond*, Cambridge: Cambridge University Press. (Useful collection. Introduction provides general survey, including discussion of desire theory and rights. Contains bibliography. See essays by R.M. Hare for introduction to his views, and by T.M. Scanlon for contractualist discussion of utilitarianism.)

Sidgwick, H. (1874) *The Methods of Ethics*, London: Macmillan; 7th edn, 1907, books III and IV. (Intuitionist defence of utilitarianism embedded within perceptive discussion of common sense morality.)

Smart, J.J. and Williams, B. (1973) *Utilitarianism For and Against*, Cambridge: Cambridge University Press. (Smart discusses the 'experience machine', Williams integrity.)

Sumner, L.W. (1981) *Abortion and Moral Theory*, Princeton, NJ: Princeton University Press. (Utilitarian discussion of abortion.)

Williams, B. (1976) 'Persons, Character and Morality', in A.O. Rorty (ed.) *The Identities of Persons*, Berkeley, CA: University of California Press, 197-216; repr. in *Moral Luck: Philosophical Papers 1973-80*, Cambridge: Cambridge University Press, 1981, 1-19. (Critique of impartialist ethics.)

Williams, B. (1985) *Ethics and the Limits of Philosophy*, London: Fontana, ch. 6. (Contains sophisticated critique of split-level theories. Moderately difficult.)