

Utilitarianism

Recall that there are four questions one might ask an ethical theory to answer:

- i) Which acts are right and which are wrong? Which acts *ought* we to perform (understanding the "ought" as a moral "ought")?
- ii) What *makes* a particular action right or wrong? What is it about the action that determines its moral status?
- iii) How do we *know* what is right and wrong?
- iv) What, if anything, motivates us to do what is right?

There are a variety of strategies for answering (iii). One might hold e.g., that moral truths are revealed by God and encoded in holy texts; or that humans have a special moral faculty (something like conscience) that gives us access to moral truth; or that we can deduce moral principles from the structure of rationality; or even (with cultural relativists) that we come to know what is right and wrong simply by learning the moral code of our culture.

Mill (and other Utilitarians) use another strategy. They begin by asking what, if anything, is *good in and of itself*, something that is not merely instrumentally good--as a means to some other good--but is good *per se*. We answer this question by looking at the structure of human desire, what it is that (well-informed) people ultimately aim for in action. The idea is that once we understand what this good is, the right act is the one that promotes the maximum amount of it. Given that on this utilitarian account what is good is determined by what humans ultimately aim for, we'll have an answer to (ii) that points in the direction of an answer to (iv), i.e., what motivates humans to do what is right.

So what is the ultimate good? One can imagine several contenders: pleasure, satisfaction of desires, happiness. Pleasure in this context is meant to be just a kind of sensation; for instance, the way it feels to have your back scratched when it's itchy, the pleasure of warming up on a cold day, the pleasures of food and drink, or of being told you've just won the lottery, and so on.

Feelings like these are certainly good, but one could have a lot of pleasures and very little pain in life without your life being all that it could be. E.g., you may have various unsatisfied desires: the desire say to play a musical instrument, or speak a second language, or climb a mountain, or to devote your life to a cause. So maybe a better answer is: *satisfaction of your desires*. (Of course it's not just any old desires we're talking about, but the desires you really identify with, the ones endorsed by your "deeper self".)

Even satisfaction of "deep" desires isn't a completely satisfactory answer though, because there are lots of things that do or would make life valuable and worthwhile that don't correspond to any definite desires of yours. Perhaps you don't have the psychological insight to realize how good life would be for you if you, say, liked yourself better, or took more chances, or had greater enthusiasm for certain activities. How do we determine what constitutes human happiness if it isn't just a matter of satisfying our actual desires? Mill proposes the "higher pleasure test": "Of two pleasures, if there be one to which all or almost all who have experience

of both give a decided preference, irrespective of any feeling of moral obligation to prefer it, that is the more desirable pleasure." (598)

Happiness in Mill's sense is supposed to take in all aspects of a person's well-being, whether or not they are pleasures and whether or not they answer to the agent's desires. The happy person is someone who is thriving, whose life is full of deep satisfactions, who has what enables us to be our best as human beings.

Suppose that happiness/well-being in this sense is good in itself. It would seem then that the more happiness the better--if it is good for me to be happy, it's even better if I'm happy and others are too; and then it seems that if happiness is really what's important, then in order to promote the good we ought simply to aim to bring about the greatest amount of happiness. Our own happiness is a fine thing, but it's no more valuable than someone else's happiness--it's the total amount of happiness that matters. Thinking along these lines leads to the Utilitarian or "greatest-happiness" principle:

...actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. (597)

Or as contemporary moral theorists state it:

You ought always to act so as to maximize happiness, i.e., the right act is the act that results in the greatest amount of happiness overall.

Note that not all Utilitarians agree that happiness is the greatest good. Hedonistic Utilitarians, e.g., believe that pleasure is the greatest good. So you can substitute "pleasure" here if you are so inclined... Mill's version is also known as "Eudaimonistic Utilitarianism". Here is another contemporary version of the principle:

You ought always to act so as to maximize social utility, where "social utility" is simply another name for the general welfare.¹

For many people, this principle just seems obvious--if we can identify what's genuinely good, then clearly what we ought to do is maximize that. If everything we might do is good only insofar as it brings about happiness, then we ought always to judge our actions by this standard.

Questions:

1. *How do we know in a particular case what is going to bring about the most happiness?* In many cases we might not know for sure. According to Mill, our ordinary moral rules serve as a rough guide, e.g., in general, keeping promises, not telling lies, doing your duty, are likely to produce the greatest happiness. Of course there's no guarantee, and if we determine in advance of acting that normally good results wouldn't occur, then we shouldn't follow the moral rule.

2. *Suppose you can quantify the happiness of individuals; how should we calculate general welfare?* Simply *adding* amounts of happiness doesn't seem right, because then there would be a moral imperative to produce more people to get more total happiness. (Or to eliminate people to reduce the total unhappiness.) Plausibly we should aim to maximize average happiness. But then wouldn't wiping out the chronically unhappy people raise the overall average? What's wrong, then, from the Utilitarian point of view, in doing so?

¹ Gilbert Harman, "Utilitarianism" in *The Nature of Morality* (New York: Oxford Univ. Press, 1977), p. 152.

These questions are really just the tip of the iceberg, for Utilitarianism seems to pose quite systematic conflicts with our ordinary moral judgements. How can Utilitarians respond?

Utilitarianism flouts ordinary moral principles: Within limits we think that promise-keeping is obligatory, even when breaking the promise would lead to better consequences. Suppose I borrow \$10. from you and promise to pay it back. In the meantime your parents give you a ton of money and the \$10 means nothing to you. Ought I to pay it back? It seems that a Utilitarian would say no: I would be happier/better off if I could keep the money, and it won't make a substantial difference to you; so I maximize happiness by breaking my promise.

Response: there are more effects to be taken into account. There are *subtle direct* effects like disappointing your expectations, and making myself untrustworthy in people's eyes, and also *indirect* effects like weakening the institution of promise-keeping. When the indirect effects are considered the calculation favors returning the money.

What about the idea that death-bed promises are *especially* to be respected. Someone privately gets me to promise that his fortune will be used to buy a set of diamond-studded collars for his cat. If I break the promise no one is going to be disappointed, no one is going to think I am untrustworthy, the institution is not going to be weakened, etc. So the indirect effects are reduced to next to nothing. Shouldn't I still keep the promise?

Response: on further reflection, the Utilitarian answer is the correct one: you ought to disregard the deathbed promise and use the money for famine relief. But note: it's the correct answer about what is *right*, not necessarily the correct answer about *moral appraisal*. Our practices of praise and blame should themselves be subjects of Utilitarian evaluation; possibly it's better overall if we praise people for promise-keeping and blame them for promise-breaking regardless of the consequences. E.g., one might argue that the consequences of breaking a promise are almost always worse; the exceptions are very hard to recognize beforehand; and if you write exceptions into the rule people won't take it as seriously. Notice that the Utilitarian has two sorts of moral evaluation to work with:

Action X is *right* iff it maximizes happiness/utility.

Action X is *praiseworthy* iff the practice of praising actions like X maximizes happiness/utility.

So a Utilitarian could say that breaking the deathbed promise is *right* but *not praiseworthy*; it's in fact blameworthy.

Utilitarianism flouts ordinary moral distinctions Ordinarily we think there is an important moral distinction between harming and failing to help. E.g., should a doctor cut up one person to save five? Should one be willing to shoot one innocent person in a deal that would let twenty go free? No! But Utilitarianism says "yes". Can we accept this?

Response: Note that again the Utilitarian need not say that the actions in question are praiseworthy. The practice of praising actions of that sort leads to bad overall consequences, e.g., it is better if people are encouraged to have ultimate respect for human life. In general, causing harm to be regarded as worse than failing to help because the badness of harm is predictable; the goodness of help is much less predictable. So, for a Utilitarian it makes sense to train people to place greater weight on avoiding harm than on promoting good. In short, our intuition that it is wrong to kill one to save five is of the right sort: it will tend to promote

greatest happiness. Even so, killing one to save five in a particular case may be right because in fact doing so will produce the greatest happiness overall.

Utilitarianism asks too much If our actions must maximize happiness in order to be right, then we never do anything right (or at least most of us don't); because there are always things that would be more socially useful, i.e., would increase general welfare, than what we're actually doing. A moral view that says we never do anything right can't be the correct view.

Response: The Utilitarian can say: maybe we never do what's right, but it doesn't make sense to blame us. Again, separate right/wrong from what's praiseworthy/blameworthy. It doesn't make sense to blame people for falling short of ideals that they're going to find it difficult or impossible to meet. Only use praise and blame for those things that are correctable, and where our practice of blaming could make a difference. (Wouldn't it make a difference to expect the most?)

Utilitarianism asks too little The Utilitarian evaluates actions on the basis of their actual *consequences*. But intuition says that factors coming *before* the action are equally important: especially, the agent's *motivation* for acting in that way, the spirit in which the act is conducted. Some moral theorists even suggest that the inner decision is *all* that's important, the rest being just a matter of good or bad luck.

Response: Mill maintains that we must distinguish our evaluation of the person and the action; a good person may perform bad actions and a bad person may perform good actions. In considering the action, what matters is whether it *promotes the general welfare*. In considering persons, we evaluate their character, i.e., whether there is a "predominant tendency" to produce good or bad conduct. (704)

Is Utilitarianism self-effacing? At every turn, Utilitarianism tries to square itself with ordinary moral intuition by saying that these intuitions are although not actually *correct*, they are good intuitions for people to have: good *on Utilitarian grounds*. But of course one could turn this same strategy on Utilitarianism itself. Will it maximize happiness for everyone to accept Utilitarianism? Would be all be better off if we accepted a different moral theory instead? Is this a problem?

Broader Questions:

1. How do we go about determining what is the best moral theory? Should we be satisfied if the theory captures our ordinary intuitions about what is right and wrong?
2. If applying utilitarianism requires that we know the outcome of our actions, and if we can never be sure of the outcome, then, one might argue, utilitarianism is not a good guide to action. How important is it for a moral theory to provide a reliable guide to action?
3. What could possibly motivate one to act in accordance with utilitarianism, if it requires action contrary to self-interest? (See question 4 at the beginning of the handout).