

Introduction to Philosophy: Ethics

FRANK ARAGBONFOH ABUMERE, DOUGLAS GILES, YA-YUN (SHERRY) KAO, MICHAEL KLENK, JOSEPH KRANAK, KATHRYN MACKAY, JEFFREY MORGAN, PAUL REZKALLA, GEORGE MATTHEWS (BOOK EDITOR), AND CHRISTINA HENDRICKS (SERIES EDITOR)

Rebus Community

Introduction to Philosophy: Ethics by Frank Aragbonfoh Abumere, Douglas Giles, Ya-Yun (Sherry) Kao, Michael Klenk, Joseph Kranak, Kathryn MacKay, Jeffrey Morgan, Paul Rezkalla, George Matthews (Book Editor), and Christina Hendricks (Series Editor) is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

CHAPTER 5

Utilitarianism

Frank Aragbonfoh Abumere
Introduction

Let us start our introduction to utilitarianism with an example that shows how utilitarians answer the following question, “Can the ends justify the means?” Imagine that Peter is an unemployed poor man in New York. Although he has no money, his family still depends on him; his unemployed wife (Sandra) is sick and needs \$500 for treatment, and their little children (Ann and Sam) have been thrown out of school because they could not pay tuition fees (\$500 for both of them). Peter has no source of income and he cannot get a loan; even John (his friend and a millionaire) has refused to help him. From his perspective, there are only two alternatives: either he pays by stealing or he does not. So, he steals \$1000 from John in order to pay for Sandra’s treatment and to pay the tuition fees of Ann and Sam. One could say that stealing is morally wrong. Therefore, we will say that what Peter has done—stealing from John—is morally wrong.

Utilitarianism, however, will say what Peter has done is morally right. For utilitarians, stealing in itself is neither bad nor good; what makes it bad or good is the consequences it produces. In our example, Peter stole from one person who has less need for the money, and spent the money on three people who have more need for the money. Therefore, for utilitarians, Peter’s stealing from John (the “means”) can be justified by the fact that the money was used for the treatment of Sandra and the tuition fees of Ann and Sam (the “end”). This justification is based on the calculation that the benefits of the theft outweigh the losses caused by the theft. Peter’s act of stealing is morally right because it produced more good than bad. In other words, the action produced more pleasure or happiness than pain or unhappiness, that is, it increased net utility.

The aim of this chapter is to explain why utilitarianism reaches such a conclusion as described above, and then examine the strengths and weaknesses of utilitarianism. The discussion is divided into three parts: the first part explains what utilitarianism is, the second discusses some varieties (or types) of utilitarianism, and the third explores whether utilitarianism is persuasive and reasonable.

What is Utilitarianism?

Utilitarianism is a form of consequentialism. For consequentialism, the moral rightness or wrongness of an act depends on the consequences it produces. On consequentialist grounds, actions and inactions whose negative consequences outweigh the positive consequences will be deemed morally wrong while actions and inactions whose positive consequences outweigh the negative consequences will be deemed morally right. On utilitarian grounds, actions and inactions which benefit few people and harm more people will be deemed morally wrong while actions and inactions which harm fewer people and benefit more people will be deemed morally right.

Benefit and harm can be characterized in more than one way; for classical utilitarians such as Jeremy Bentham (1748-1832) and John Stuart Mill (1806-1873), they are defined in terms of happiness/unhappiness and pleasure/pain. On this view, actions and inactions that cause less pain or unhappiness and more pleasure or happiness than available alternative actions and inactions will be deemed morally right, while actions and inactions that cause more pain or unhappiness and less pleasure or happiness than available alternative actions and inactions will be deemed morally wrong. Although pleasure and happiness can have different meanings, in the context of this chapter they will be treated as synonymous.

Utilitarians' concern is how to increase net utility. Their moral theory is based on the principle of utility which states that "the morally right action is the action that produces the most good" (Driver 2014). The morally wrong action is the one that leads to the reduction of the maximum good. For instance, a utilitarian may argue that although some armed robbers robbed a bank in a heist, as long as there are more people who benefit from the robbery (say, in a Robin Hood-like manner the robbers generously shared the money with many people) than there are people who suffer from the robbery (say, only the billionaire who owns the bank will bear the cost of the loss), the heist will be morally right rather than morally wrong. And on this utilitarian premise, if more people suffer from the heist while fewer people benefit from it, the heist will be morally wrong.

From the above description of utilitarianism, it is noticeable that utilitarianism is opposed to deontology, which is a moral theory that says that as moral agents we have certain duties or obligations, and these duties or obligations are formalized in terms of rules (see Chapter 6). There is a variant of utilitarianism, namely rule utilitarianism, that provides rules for evaluating the utility of actions and inactions (see the next part of the chapter for a detailed explanation). The difference between a utilitarian rule and a deontological rule is that according to rule utilitarians, acting according to the rule is correct because the rule is one that, if widely accepted and followed, will produce the most good. According to deontologists, whether the consequences of our actions are positive or negative does not determine their moral rightness or moral wrongness. What determines their moral rightness or moral wrongness is whether we act or fail to act in accordance with our duty or duties (where our duty is based on rules that are not themselves justified by the consequences of their being widely accepted and followed).

Some Varieties (or Types) of Utilitarianism

The above description of utilitarianism is general. We can, however, distinguish between different types of utilitarianism. First, we can distinguish between "actual consequence utilitarians" and "foreseeable consequence utilitarians." The former base the evaluation of the moral rightness and moral wrongness of actions on the actual consequences of actions; while the latter base the evaluation of the moral rightness and moral wrongness of actions on the foreseeable consequences of actions. J. J. C. Smart (1920-2012) explains the rationale for this distinction with reference to the following example: imagine that you rescued someone from drowning. You were acting in good faith to save a drowning person, but it just so happens

that the person later became a mass murderer. Since the person became a mass murderer, actual consequence utilitarians would argue that in hindsight the act of rescuing the person was morally wrong. However, foreseeable consequence utilitarians would argue that—looking forward (i.e., in foresight)—it could not be foreseen that the person was going to be a mass murderer, hence the act of rescuing them was morally right (Smart 1973, 49). Moreover, they could have turned out to be a “saint” or Nelson Mandela or Martin Luther King Jr., in which case the action would be considered to be morally right by actual consequence utilitarians.

A second distinction we can make is that between act utilitarianism and rule utilitarianism. Act utilitarianism focuses on individual actions and says that we should apply the principle of utility in order to evaluate them. Therefore, act utilitarians argue that among possible actions, the action that produces the most utility would be the morally right action. But this may seem impossible to do in practice since, for every thing that we might do that has a potential effect on other people, we would thus be morally required to examine its consequences and pick the one with the best outcome. Rule utilitarianism responds to this problem by focusing on general types of actions and determining whether they typically lead to good or bad results. This, for them is the meaning of commonly held moral rules: they are generalizations of the typical consequences of our actions. For example, if stealing typically leads to bad consequences, stealing in general would be considered by a rule utilitarian to be wrong.¹

Hence rule utilitarians claim to be able to reinterpret talk of rights, justice, and fair treatment in terms of the principle of utility by claiming that the rationale behind any such rules is really that these rules generally lead to greater welfare for all concerned. We may wonder whether utilitarianism in general is capable of even addressing the notion that people have rights and deserve to be treated justly and fairly, because in critical situations the rights and wellbeing of persons can be sacrificed as long as this seems to lead to an increase overall utility.

For example, in a version of the famous “trolley problem,” imagine that you and an overweight stranger are standing next to each other on a footbridge above a rail track. You discover that there is a runaway trolley rolling down the track and the trolley is about to kill five people who cannot get off of the track quickly enough to avoid the accident. Being willing to sacrifice yourself to save the five persons, you consider

jumping off the bridge, in front of the trolley...but you realize that you are far too light to stop the trolley....The only way you can stop the trolley killing five people is

¹Of course, there may be exceptions to such a rule in particular, atypical cases if stealing might lead to better consequences. This raises a complication for rule utilitarians: if they were to argue that we should follow rules such as “do not steal” except in those cases where stealing would lead to better consequences, then this could mean rule utilitarianism wouldn’t be very different from act utilitarianism. One would still have to evaluate the consequences of each particular act to see if one should follow the rule or not. Hooker (2016) argues that rule utilitarianism need not collapse into act utilitarianism in this way, because it would be better to have a set of rules that are more clear and easily understood and followed than ones that require us to evaluate many possible exceptions.

by pushing this large stranger off the footbridge, in front of the trolley. If you push the stranger off, he will be killed, but you will save the other five. (Singer 2005, 340)

Utilitarianism, especially act utilitarianism, seems to suggest that the life of the overweight stranger should be sacrificed regardless of any purported right to life he may have. A rule utilitarian, however, may respond that since in general killing innocent people to save others is not what typically leads to the best outcomes, we should be very wary of making a decision to do so in this case. This is especially true in this scenario since everything rests on our calculation of what might possibly stop the trolley, while in fact there is really far too much uncertainty in the outcome to warrant such a serious decision. If nothing else, the emphasis placed on general principles by rule utilitarians can serve as a warning not to take too lightly the notion that the ends might justify the means.

Whether or not this response is adequate is something that has been extensively debated with reference to this famous example as well as countless variations. This brings us to our final question here about utilitarianism—whether it is ultimately a persuasive and reasonable approach to morality.

Is Utilitarianism Persuasive and Reasonable?

First of all, let us start by asking about the principle of utility as the foundational principle of morality, that is, about the claim that what is morally right is just what leads to the better outcome. John Stuart Mill's argument that it is based on his claim that "each person, so far as he believes it to be attainable, desires his own happiness" (Mill [1861] 1879, Ch. 4). Mill derives the principle of utility from this claim based on three considerations, namely desirability, exhaustiveness, and impartiality. That is, happiness is desirable as an end in itself; it is the only thing that is so desirable (exhaustiveness); and no one person's happiness is really any more desirable or less desirable than that of any other person (impartiality) (see Macleod 2017).

In defending desirability, Mill argues,

The only proof capable of being given that an object is visible, is that people actually see it. The only proof that a sound is audible, is that people hear it: and so of the other sources of our experience. In like manner...the sole evidence it is possible to produce that anything is desirable, is that people do actually desire it. (Mill [1861] 1879, Ch. 4)

In defending exhaustiveness, Mill does not argue that other things, apart from happiness, are not desired as such; but while other things *appear* to be desired, happiness is the only thing that is *really* desired since whatever else we may desire, we do so because attaining that thing would make us happy. Finally, in defending impartiality, Mill argues that equal amounts of happiness are equally desirable, whether the happiness is felt by the same person or by different persons. In Mill's words, "each person's happiness is a good to that person, and the general happiness, therefore, a good to the aggregate of all persons" (Mill [1861] 1879, Ch. 4). We may wonder, however, whether this last argument is truly adequate. Does

Mill really show here that we should treat everyone's happiness as equally worthy of pursuit, or does he simply assert this?

Let us grant that Mill's argument here is successful and the principle of utility is the basis of morality. Utilitarianism claims that we should thus calculate, to the best of our ability, the expected utility that will result from our actions and how it will affect us and others, and use that as the basis for the moral evaluation of our decisions. But then we may ask, how exactly do we quantify utility? Here there are two different but related problems: how can I come up with a way of comparing different types of pleasure and pain, benefit or harm that I myself might experience, and how can I compare my benefit and yours on some neutral scale of comparison? Bentham famously claimed that there was a single universal scale that could enable us to objectively compare all benefits and harms based on the following factors: intensity, duration, certainty/uncertainty, proximity, fecundity, purity, and extent. And he offered on this basis what he called a "felicific calculus" as a way of objectively comparing any two pleasures we might encounter (Bentham [1789] 1907).

For example, let us compare the pleasure of drinking a pint of beer to that of reading Shakespeare's *Hamlet*. Suppose the following are the case:

- The pleasure derived from drinking a pint of beer is more intense than the pleasure derived from reading *Hamlet* (intensity).
- The pleasure of drinking the beer lasts longer than that of reading *Hamlet* (duration).
- We are confident that drinking the beer is more pleasurable than reading *Hamlet* (certainty/uncertainty).
- The beer is closer to us than the play, and therefore it is easier for us to access the former than the latter (proximity).
- Drinking the beer is more likely to promote more pleasure in the future while reading *Hamlet* is less likely to promote more pleasure in the future (fecundity).
- Drinking the beer is pure pleasure while reading *Hamlet* is mixed with something else (purity).
- Finally, drinking the beer affects both myself and my friends in the bar and so has a greater extent than my solitary reading of *Hamlet* (extent).

Since, on all of these measures, drinking a pint of beer is more pleasurable than reading *Hamlet*, it follows according to Bentham that it is objectively better for you to drink the pint of beer and forget about reading *Hamlet*, and so you should. Of course, it is up to each individual to make such a calculation based on the intensity, duration, certainty, etc. of the pleasure resulting from each possible choice they may make in their eyes, but Bentham at least claims that such a comparison is possible.

This brings us back to the problem we mentioned before that, realistically, we cannot be expected to always engage in very difficult calculations every single time we want to make a decision. In an attempt to resolve this problem, utilitarians might claim that in the evaluation of the moral rightness and moral wrongness of actions, the application of the principle of utility can be backward-looking (based on hindsight) or forward-looking (based on foresight).

That is, we can use past experience of the results of our actions as a guide to estimating what the probable outcomes of our actions might be and save ourselves from the burden of having to make new estimates for each and every choice we may face.

In addition, we may wonder whether Bentham's approach misses something important about the different kinds of pleasurable outcomes we may pursue. Mill, for example, would respond to our claim that drinking beer is objectively more pleasurable than reading *Hamlet* by saying that it overlooks an important distinction between qualitatively different kinds of pleasure. In Mill's view Bentham's calculus misses the fact that not all pleasures are equal—there are “higher” and “lower” pleasures that make it “better to be a human being dissatisfied than a pig satisfied; better to be Socrates dissatisfied than a fool satisfied” (Mill [1861] 1879, Ch. 2). Mill justifies this claim by saying that between two pleasures, although one pleasure requires a greater amount of difficulty to attain than the other pleasure, if those who are competently acquainted with both pleasures prefer (or value) one over the other, then the one is a higher pleasure while the other is a lower pleasure. For Mill, although drinking a pint of beer may seem to be more pleasurable than reading *Hamlet*, if you are presented with these two options and you are to make a choice—each and every time or as a rule—you should still choose to read *Hamlet* and forego drinking the pint of beer. Reading *Hamlet* generates a higher quality (although perhaps a lower quantity) of pleasure, while drinking a pint of beer generates lower quality (although higher quantity) of pleasure.

In the end, these issues may be merely technical problems faced by utilitarianism—is there some neutral scale of comparison between pleasures? If there is, is it based on Bentham's scale which makes no distinctions between higher and lower pleasures, or Mill's which does? The more serious problem, however, remains, which is that utilitarianism seems willing in principle to sacrifice the interests and even perhaps lives of individuals for the sake of the benefit of a larger group. And this seems to conflict with our basic moral intuition that people have a *right* not to be used in this way. While Mill argued that the notion of rights could be accounted for on purely utilitarian terms, Bentham simply dismissed it. For him such “natural rights” are “simple nonsense, natural and imprescriptible rights, rhetorical nonsense—nonsense upon stilts” (Bentham [1796] 1843, 501).

Conclusion

Let us conclude by revisiting the question we started with: can the ends justify the means? I stated that as far as utilitarianism is concerned the answer to this question is in the affirmative. While the answer is plausible and right for utilitarians, it is implausible for many others, and notably wrong for deontologists. As we have seen in this chapter, on a close examination utilitarianism is less persuasive and less reasonable than it appears to be when it is far away.

References

- Bentham, Jeremy. (1789) 1907. *An Introduction to the Principles of Morals and Legislation*. Oxford: Clarendon Press.
- Bentham, Jeremy. (1796) 1843. *Anarchical Fallacies*. In *The Works of Jeremy Bentham*, ed. John Bowring. Vol 2. Edinburgh: William Tait.
- Driver, Julia. 2014. "The History of Utilitarianism." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. <https://plato.stanford.edu/archives/win2017/entries/utilitarianism-history/>
- Hooker, Brad. 2016. "Rule Consequentialism." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. <https://plato.stanford.edu/archives/win2016/entries/consequentialism-rule/>
- Macleod, Christopher. 2017. "John Stuart Mill." In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. <https://plato.stanford.edu/archives/spr2017/entries/mill/>
- Mill, John Stuart. *Utilitarianism*. (1861) 1879. 7th ed. London: Longmans, Green, and Co. <https://www.gutenberg.org/ebooks/11224>
- Singer, Peter. 2005. "Ethics and Intuitions." *The Journal of Ethics* 9(3/4): 331-352.
- Smart, J. J. C. 1973. "An Outline of a System of Utilitarian Ethics." In Smart, J. J. C. and Bernard Williams.
- Smart, J. J. C. and Bernard Williams. 1973. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.
- Further Reading
- Hare, R. M. 1981. *Moral Thinking*. Oxford: Oxford University Press.
- Hooker, Brad. 1990. "Rule-Consequentialism." *Mind* 99(393): 66-77.
- Scheffler, Samuel. 1988. *Consequentialism and Its Critics*. Oxford: Oxford University Press.
- Sen, Amartya and Bernard Williams, eds. 1982. *Utilitarianism and Beyond*. Cambridge: Cambridge University Press.
- Sidgwick, Henry. 1907. *The Methods of Ethics*. London: Macmillan.
- Singer, Peter. 2000. *Writings on an Ethical Life*. New York: HarperCollins.
- Smart, J. J. C. and Bernard Williams. 1973. *Utilitarianism: For and Against*. Cambridge: Cambridge University Press.
- Thomson, Judith Jarvis. 1985. "The Trolley Problem." *The Yale Law Journal* 94(6): 1395-1415.
- Williams, Bernard. 1973. "A Critique of Utilitarianism." In Smart, J. J. C and Bernard Williams.