

Utilitarianism

The Good and the Right (in that order)

An ideal moral theory

- What do we want from a moral theory? To know what is right, what is good, and what makes the good things good and the right things right. Also, we want to know how to put our theory of rightness and our theory of value together.
- The ideal theory would identify the fundamental principle(s) of morality. They would be *infallible* (they wouldn't tell you what to do and give you the wrong advice) and *universal* (they always offer advice when there is advice to give).
- If you had the fundamental principle(s) of morality, you would know everything there is to know about morality. I'll give you the fundamental principle...

The fundamental principle is:
(DtRT!) *Do the right thing!*

Is it infallible? Yes—if it says to jump, jumping is the right thing to do (not ducking).

Is it universal? Yes—if ducking is the right thing to do, it says to duck (not jump).

Are you now morally omniscient? Something is missing, but it isn't *infallibility* or *universality*.

One problem is that you can't apply the principle until you've worked out what the right thing is. We want a principle we can apply when we don't yet know what's right.

DtRT! has *universality* and *infallibility*, but it lacks *applicability*. It might be possible to translate DtRT and articulate a principle that's applicable. If it is a proper translation, it should preserve universality and infallibility.

How would that work?

First, could x be the right thing to do if x is not an option? No. So, the FMP should pick one (or more) option(s) and identify it as right. So, *Do the right thing* becomes *pick the right option*. (Assuming this is a proper translation, we retain U and I. *PtRO* isn't applicable, so we have work to do.)

Second, suppose you came to me for advice. I tell you $DtRT$ and you say that's not very helpful. I say $PtRO$ and you say that that's also not very helpful. I then ask you which option is best. Suppose you say this:

a is pretty good, but b is better than a . As for c , it's not as good as b . Those are all the options I have.

I'd ask, "So, b is best?" and you'd (better) say that that it is. Can there now be any doubt as to which option would be right to pick? That is, can you really ask:

Look, I realize that b is best, but should I pick it?

I submit to you that you *can't* rationally wonder whether it would be right to pick *b* if you insist that *b* is the best option.

If that's right, then we've made some progress:

Step 1: Do the right thing = Pick the right option.

Step 2: Pick the right option = Pick the best option.

Thus: Do the right thing = Pick the best option.

Pick the best option! is *infallible, universal, and applicable.*

Objection!

Objection: *Pick the best option* lacks applicability because we don't know what's best.

Response: Yes you do.

The best option = the option better than the others.

The option better than the others = the option gooder than the others.

The option gooder than the others = the option that contains more goodness than the others.

Thus, the best option = the option that contains more goodness than the other options.

Do the right thing = Do the thing that maximizes what's good.

Objection!!!

Objection: *Do the thing that maximizes what's good* is only applicable if we know what's good. You haven't told us what's good yet.

1st Response: You know this much—you know that what's right is determined entirely by what's good. You know that the right action is the action with the best consequences. So, you know *consequentialism* is true. Maybe you don't know what's good, but you know that you'll be morally omniscient once you have a theory of value.

2nd Response: Don't be fatuous. Of course you know what's good. I'm not some freaky goodness expert who knows better than you what's good. You're probably good enough to determine what's good. We'll see if you are in a second.

Where things stand

Here's what we have thus far. We have an argument for *consequentialism*:

C: You ought to perform an action iff its outcome would be better than the outcome of any other available action.

The basic idea is that you can't sensibly ask whether it would be right to do something once you've worked out that it's the best option there is.

We also know that we need to turn to value theory to fill in the details of our ideal moral theory. Once we know what's good, we know what's goodest, what's best, and what's right.

What's good?

G.E. Moore tells us how to think about the relation between rightness and value. John Stuart Mill tells us how to think about value:

The utilitarian doctrine is that happiness is desirable as an end, and is the only thing that is so; anything else that is desirable is only desirable as means to that end. What should be required regarding this doctrine—what conditions must it fulfill—to justify its claim. 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 similarly with the other sources of our experience. In like manner, I apprehend, the sole evidence it is possible to produce that anything is desirable is that people do actually desire it.

Intrinsic vs. Instrumental Value

Formally, we can distinguish between those things that are good for their own sake (intrinsically valuable) and those things that are good as a means to something else (instrumentally valuable). If Mill is right, what's desirable is good in some way. Perhaps what's desirable for its own sake is intrinsically good. What's desirable to the extent that it is a means to something else is instrumentally good.

Substantively, we can motivate hedonism as follows. It is clear that pleasure is desirable for its own sake and so intrinsically good if everything we've said thus far is correct. Is there a second kind of thing that's intrinsically good? Maybe not. Let x be that thing. Imagine a situation where x brings nobody any pleasure. Perhaps it brings about pain and misery. Would the presence of x count as a silver lining? If not, x might be good, but not on those occasions where it doesn't bring pleasure to anyone. So, perhaps x is good, but good as an instrument.

$$U = H + C$$

According to the *utilitarian*, the fundamental principle of morality is the greatest happiness principle:

GHP: An act is right to the extent that it promotes the greatest happiness for the greatest number.

Why think that? Because it follows from consequentialism and *hedonism*:

H: Pleasure is the only thing that is intrinsically good and pain is the only thing that is intrinsically bad.

According to GHP, right acts maximize happiness. According to CON, right acts maximize goodness. According to HED, pleasure is good and pain is bad. If you reject GHP and Mill's utilitarian view, you have to reject either hedonism or consequentialism.

Morality of swine

Objection: Because the hedonist thinks that only pleasures are valuable, they must think that the best life is that life of the pig.

Reply: The objection curiously assumes that we'd be happier as pigs than humans. As Mill puts it:

The accusation supposes human beings to be capable of no pleasures except those of which swine are capable. If this supposition were true, the charge could not be gainsaid, but would then be no longer an imputation; for if the sources of pleasure were precisely the same to human beings and to swine, the rule of life which is good enough for the one would be good enough for the other. The comparison of the Epicurean life to that of beasts is felt as degrading, precisely because a beast's pleasures do not satisfy a human being's conceptions of happiness.

The crib test

Imagine that you are filled with love as you look into the crib, checking on your newly arrived firstborn child. The infant is sleeping peacefully. You might think of various ways in which the baby's life could turn out. What schools will he attend? What career will he choose? What sort of personality and intellect will he have? Will he someday have children of his own? Your concern for the baby might express it self in the hope that, whatever he does, things will turn out well for him. You might hope that this baby gets a good life—a life good in itself for him.

Most of us, I think, would want our child to live a life filled with pleasure and free from pain. Even if we want our children to be good people and to do the right thing, we realize the difference between the morally good life and the life that's good for the child (e.g., the life of a saint might be filled with pain and hardship, so it scores well in one dimension but poorly when we focus just on welfare). Given the role that pleasure and pain seems to be play in our wishes for our children and our sense of what would be good *for them*, it is tempting to think that there is something right to hedonism.

Experience machine

Let's not be hasty. Consider Nozick's example:

Suppose there were an experience machine that would give you any experience you desired. Superduper neuropsychologists could stimulate your brain so that *you* would think and feel you were writing a great novel, or making a friend, or reading an interesting book. All the time you would be floating in a tank, with electrodes attached to your brain. Should you plug into this machine for life, preprogramming your life's experiences?

We can stipulate that nobody in the machine knows they are in the machine. And, we don't need to think about it from the first-person point of view. Even if we maximize the amount of pleasure the person receives, is the person's life really going all that well?

Alternatives to hedonism

Objection: Perhaps what's missing from life on the machine is this—you derive lots of pleasure but your desires are unsatisfied.

Reply: Not necessarily! It could be that many of your desires *are* satisfied. Compare two lives in two machines. In M1, the person desires that it is sunny and it is sunny outside. It also seems sunny to them, so they are happy. In M2, the person desires that it is raining when in fact it is sunny. It seems to them that it is raining.

Three points.

First, you're surely better off if you're happy about the weather. So, the pleasure should count for *something*. If that's right, the problem with hedonism can't be that it takes pleasure to be intrinsically good, but that it should recognize further goods.

Second, are you any better off if you're in M1 than M2? Isn't that an odd sort of double counting? If it is (and it is!), then desire satisfaction doesn't seem to make someone's life better.

Third, your desires can be about future states of the world that won't come about until after you are dead. It seems strange to think that how well your life goes depends (in part) upon how events turn out after you are dead.

[The issues concerning posthumous harms and double counting arise in connection with the view Nozick seems to like, which is that how well your life goes depends (in part) upon what you do and what you achieve.]

An argument from Parfit

When our acts can affect only one person, most of us accept maximization. We do not believe that we ought to give someone fewer happy days so as to be more fair in the way we spread them out over the parts of his life. There are, of course, arguments for spreading out enjoyments. We remain fresh, and have more to look forward to. But these arguments do not count against maximization; they remind us how to achieve it.

When our acts can affect several different people, Utilitarians make similar claims. They admit new arguments for spreading out enjoyments, such as that which appeals to the effect of relative deprivation, or do diminishing marginal utility. But Utilitarians treat equality as a mere means, not a separate aim.

Since their attitude to sets of lives is like ours to single lives, Utilitarians ignore the boundaries between lives. We may ask, “Why?”

Parfit doesn't think there's a good answer to that question. If not, then we can argue as follows. We ought to maximize pleasure and happiness if our actions affect only one person. If our actions affect two, things remain fundamentally the same. What's good or bad doesn't depend upon the identity of the benefactor. What's right or wrong doesn't depend upon anything but what's good or bad. (Parfit's way of filling this out is much more exciting than this, but it gives you a sense of how the utilitarian view works.)

Rawls

Rawls suggests that utilitarians ignore the boundaries between lives because they approach moral questions from the point of view of an impartial observer. When the utilitarian asks what would be right to do, she might identify with all of the affected people and this leads her to ignore that different people are affected.

Rawls thinks that this is objectionable, but is it? If you're not supposed to think of things from the point of view of an impartial observer, that must be because there's some other point of view to take that would lead you to classify cases differently. And, it better be that when you think about things from this point of view, you get things right in cases where utilitarianism gets things wrong. What would such a case look like?

Jim and the Indians

Jim finds himself in the central square of a small town in a small country ruled by a small despot. Tied up against the wall are twenty men, women, and children. In front of them are several armed men in uniform. A heavy man in a sweat-stained khaki shirt is the captain. He explains that his job is to round up twenty people off of the street and shoot them so that the people will learn not to protest these shootings. When he learns that you are visiting from abroad, he offers you the opportunity to shoot one of the protestors yourself. If you do, nineteen will be released to mark this special occasion. If not, the twenty will be summarily shot. What should you do?