

do the ends justify the means?

Suppose that some actions are right, and some are wrong. What's the difference between them? What makes some actions right, and others wrong?

Here is one very simple, but also very plausible, answer to this question:

Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would lead to the best outcome.

Consequentialism says, simply, that we should judge actions by their consequences. Whatever will lead to the best overall outcome is what one ought to do.

A slightly different way to get the general idea is this: if I am deciding between doing action A and action B, I should try to figure out what the world would be like if I did A, and what the world would be like if I did B; and I should do whichever action would lead to the better world.

Consequentialism


An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would lead to the best outcome.

This view raises two questions. The first is: what makes one outcome, or state of affairs, better than another?

To answer this question is to give a **theory of value**: a theory about what makes one state of the world better, or worse than, another.

Let us say that a **good** is something that makes a state of affairs better, and an **evil** is something that makes a state of affairs worse.

In the reading for today, John Stuart Mill gives the following statement of his theory of value — his view of which things are goods and evils.



The creed which accepts as the foundation of morals, Utility, or the Greatest Happiness Principle, holds that actions are right in proportion as they tend to promote happiness, wrong as they tend to produce the reverse of happiness. By happiness is intended pleasure, and the absence of pain; by unhappiness, pain, and the privation of pleasure. To give a clear view of the moral standard set up by the theory, much more requires to be said; in particular, what things it includes in the ideas of pain and pleasure; and to what extent this is left an open question. But these supplementary explanations do not affect the theory of life on which this theory of morality is grounded — namely, that pleasure, and freedom from pain, are the only things desirable as ends; and that all desirable things (which are as numerous in the utilitarian as in any other scheme) are desirable either for the pleasure inherent in themselves, or as means to the promotion of pleasure and the prevention of pain.

This view is sometimes called hedonism:

Hedonism

Pleasure is the only good
and pain is the only evil.

Suppose that this is true. Then how do we tell whether one outcome is better than another?

Here is a very natural answer. We 'add up' the pleasure, and 'subtract out' the pain. Whatever situation has the highest 'net pleasure' is the best.

In general, one might think, it is fairly straightforward to compare two different situations. One adds up the goods, subtracts out the evils, and determine the net good. On this view, one should always aim to maximize the **net good**.

Hedonism

Pleasure is the only good
and pain is the only evil.

In general, one might think, it is fairly straightforward to compare two different situations. One adds up the goods, subtracts out the evils, and determine the net good. On this view, one should always aim to maximize the **net good**.

This view can be stated as follows:

Maximizing Consequentialism

An action is the right thing to do
in certain circumstances if, of all
the actions available in those
circumstances, it produces the
highest net good.

(You might wonder: aren't Consequentialism and Maximizing Consequentialism pretty much the same thing? As we will see, they are not.)

Hedonism

Pleasure is the only good
and pain is the only evil.



Maximizing Consequentialism

An action is the right thing to do
in certain circumstances if, of all
the actions available in those
circumstances, it produces the
highest net good.

Utilitarianism

An action is the right thing to do
in certain circumstances if, of all
the actions available in those
circumstances, it would produce
the highest net pleasure.

Utilitarianism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would produce the highest net pleasure.

Utilitarianism is a very clear and plausible-sounding view about ethics.

This is the view which is often summed up with the slogan that one ought always to act to cause the greatest happiness for the greatest number. It is a paradigmatically unselfish theory: no one's pleasures and pains are more important than anyone else's.

Utilitarianism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would produce the highest net pleasure.


A historically influential objection to utilitarianism is that it is a 'doctrine fit for swine,' because it does not recognize the fact that, unlike pigs, human beings have goods other than mere pleasure.

Against this, Mill replies as follows:

When thus attacked, the Epicureans have always answered, that it is not they, but their accusers, who represent human nature in a degrading light; since the accusation supposes human beings to be capable of no pleasures except those of which swine are capable.

What is Mill's reply to the objection?

A more serious challenge to utilitarianism can be brought out by Robert Nozick's example of the experience machine.



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? If you are worried about missing out on desirable experiences, we can suppose that business enterprises have researched thoroughly the lives of many others. You can pick and choose from their large library or smorgasbord of such experiences, selecting your life's experiences for, say, the next two years. After two years have passed, you will have ten minutes or ten hours out of the tank, to select the experiences of your *next* two years. Of course, while in the tank you won't know that you're there; you'll think it's all actually happening. Others can also plug in to have the experiences they want, so there's no need to stay unplugged to serve them. (Ignore problems such as who will service the machines if everyone plugs in.) Would you plug in? *What else can matter to us, other than how our lives feel from the inside?* Nor should you refrain because of the few moments of distress between the moment you've decided and the moment you're plugged. What's a few moments of distress compared to a lifetime of bliss (if that's what you choose), and why feel any distress at all if your decision *is* the best one?

What must the utilitarian say about the relative goodness of the state of affairs in which everyone (or almost everyone) plugs in and the state of affairs in which no one does?

Suppose you face the decision whether to get into the experience machine. What would a utilitarian say about what you ought to do?

Suppose now that you face the decision of whether you should put **everyone** into the experience machine. (The machines are maintained by extremely reliable robots.) What would a utilitarian say about what you ought to do?


Does it matter if people ask you (or beg you) not to plug them in?

Is Nozick right that these consequences of utilitarianism are incorrect?

Here is a second challenge for the Utilitarian, which is based on another example due to Robert Nozick.

Imagine that there is a **utility monster** which gets more pleasure out of everything than any human does. No matter what things bring you pleasure, this thing gets more pleasure out of those things than you do.

Now suppose that you face a choice. You can either give some pleasure-causing thing to a friend of yours, or give it to the utility monster. Which course of action does the Utilitarian say you ought to pursue?



Recall that we presented Utilitarianism as the combination of two claims.

Hedonism

Pleasure is the only good
and pain is the only evil.

Maximizing Consequentialism

An action is the right thing to do
in certain circumstances if, of all
the actions available in those
circumstances, it produces the
highest net good.

You might think that the examples we have discussed — the experience machine and the utility monster — are problems for hedonism, but not for Maximizing Consequentialism. Couldn't the Maximizing Consequentialist just say that there are goods besides pleasure, and evils besides pain?

Maximizing Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it produces the highest net good.

Here are some other candidates for goods:

The extent to which the desires of agents are satisfied.

The extent to which the states of affairs contain beauty, or love, or friendship, or something else taken to be of objective value.

The extent to which the states of affairs maximize the well-being, or welfare, of agents.

Corresponding to each of these views about the good is a different version of Maximizing Consequentialism. For example, the first would yield the result that one should always act in such a way that maximizes the number of desires of people which are satisfied.


What would that view say about the experience machine?

Maximizing Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it produces the highest net good.

However, in the reading from John Rawls, we get a different sort of objection to Maximizing Consequentialism.



Rawls' objection is summed up with the concluding sentences of the passage we read:



This view of social cooperation is the consequence of extending to society the principle of choice for one man, and then, to make this extension work, conflating all persons into one through the imaginative acts of the impartial sympathetic spectator. Utilitarianism does not take seriously the distinction between persons.

To see what Rawls has in mind here, let's think about an example.

Suppose that we have a group of five people, whose 'goodness of life' — however we characterize goodness — is indicated by the numbers beside them.



Now imagine that I have the chance to bring about one of two states of affairs.

Situation A



0

0

0

0

942

Situation B



100

100

100

100

500

Which one, according to the Maximizing Consequentialist, should I bring about?

This is what Rawls means when he says that Maximizing Consequentialism fails to take account of the distinctness of persons. The Maximizing Consequentialist simply sums goods across persons, and thereby rules out the possibility that the goodness or badness of a situation can also depend on the distribution of goods across people.

Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would produce the best outcome.

Maximizing Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it produces the highest net good.

Does this sort of objection rule out Consequentialism generally?

It does not, because there is nothing to stop the Consequentialist from saying that what makes one situation better than another has to do with the distribution of the good, as well as the total net good. Many contemporary versions of Consequentialism are constructed in this way.

Consequentialism

An action is the right thing to do in certain circumstances if, of all the actions available in those circumstances, it would produce the best outcome.

Once one sees how flexible Consequentialism is, one might be tempted to think that some version of that view must just obviously be true.

One might also think that Consequentialism is so general that, without some explanation of what “best outcome” means, it does not tell us much at all about how we ought to act. The reading from Peter Singer, however, shows that this is a mistake. In particular, he argues that even very basic consequentialist assumptions imply that we owe much more to the poor than many people think.

This is the topic of Peter Singer's 1971 paper "Famine, Affluence, and Morality." Singer describes the contemporary example of refugees in Bengal and says the following:



What are the moral implications of a situation like this? In what follows, I shall argue that the way people in relatively affluent countries react to a situation like that in Bengal cannot be justified; indeed, the whole way we look at moral issues—our moral conceptual scheme—needs to be altered, and with it, the way of life that has come to be taken for granted in our society.

To understand Singer's position and argument, we need to do two things: (1) understand what sorts of situations he is talking about, and (2) understand what he thinks we are morally obliged to do in response to such situations.

Singer describes the situation in Bengal as follows:

As I write this, in November 1971, people are dying in East Bengal from lack of food, shelter, and medical care. The suffering and death that are occurring there now are not inevitable, not unavoidable in any fatalistic sense of the term. Constant poverty, a cyclone, and a civil war have turned at least nine million people into destitute refugees; nevertheless, it is not beyond the capacity of the richer nations to give enough assistance to reduce any further suffering to very small proportions. The decisions and actions of human beings can prevent this kind of suffering.

There seem to be two relevant aspects of the situation in Bengal: that it involves massive human suffering, and that it is, at least in large part, avoidable.

This leads to a natural question: are there today any situations of this sort — in other words, situations that both involve massive human suffering and are avoidable?

This leads to a natural question: are there today any situations of this sort — in other words, situations that both involve massive human suffering and are avoidable?

There are. About 25,000 people per day die of hunger. That is about one person every 3.5 seconds. This involves massive human suffering. And it is avoidable, because there is enough food on earth to feed everyone.

The next question is: what are our moral obligations, given this fact?

The next question is: what are our moral obligations, given this fact?

My next point is this: if it is in our power to prevent something bad from happening, without thereby sacrificing anything of comparable moral importance, we ought, morally, to do it. By “without sacrificing anything of comparable moral importance” I mean without causing anything else comparably bad to happen, or doing something that is wrong in itself, or failing to promote some moral good, comparable in significance to the bad thing that we can prevent. This principle seems almost as uncontroversial as the last one. It requires us only to prevent what is bad, and not to promote what is good, and it requires this of us only when we can do it without sacrificing anything that is, from the moral point of view, comparably important. I could even, as far as the application of my argument to the Bengal emergency is concerned, qualify the point so as to make it: if it is in our power to prevent something very bad from happening, without thereby sacrificing anything morally significant, we ought, morally, to do it. An application of this principle would be as follows: if I am walking past a shallow pond and see a child drowning in it, I ought to wade in and pull the child out. This will mean getting my clothes muddy, but this is insignificant, while the death of the child would presumably be a very bad thing.

In this passage, Singer states two different moral principles, which might be stated as follows:

The strong principle

One always ought to prevent something bad from happening if one can do so without sacrificing anything with moral importance comparable to the thing to be prevented.

The moderate principle

One always ought to prevent something bad from happening if one can do so without sacrificing anything of any moral importance.

Can you think of any examples where we seem to take for granted principles of this sort?

Let's look at a concrete example of what these principles imply, starting with the strong principle.

The strong principle

One always ought to prevent something bad from happening if one can do so without sacrificing anything with moral importance comparable to the thing to be prevented.

The importance of an ND education (vs an education at one's state university) is not of comparable importance to the lives of 30 people.

A Notre Dame education costs \$140,000 more than an average education in a state university

It costs roughly \$1 to feed one child in Africa for one day

The difference between an ND education and a state school education could feed 30 children in Africa, who would otherwise die of starvation, from age 5 to adulthood

1. One always ought to prevent something bad from happening if one can do so without sacrificing anything with moral importance comparable to the thing to be prevented. (Strong Principle)
2. A Notre Dame education costs \$140,000 more than an average education in a state university.
3. It costs roughly \$1 to feed one child in Africa for one day.
4. The difference between an ND education and a state school education could feed 30 children in Africa, who would otherwise die of starvation, from age 5 to adulthood. (2,3)
5. The importance of an ND education (vs an education at one's state university) is not of comparable importance to the lives of 30 people.

C. No one should attend Notre Dame. (1,4,5)

Is the argument valid?

It is difficult to reject premises 2 or 3. So if one wants to reject the conclusion of the argument, one must reject either premise 1 or premise 5.

It is natural to think of Singer's view as extremely radical. But, as Singer points out, many others throughout history would have regarded his suggestions as far from radical:

It may still be thought that my conclusions are so wildly out of line with what everyone else thinks and has always thought that there must be something wrong with the argument somewhere. In order to show that my conclusions, while certainly contrary to contemporary Western moral standards, would not have seemed so extraordinary at other times and in other places, I would like to quote a passage from a writer not normally thought of as a way-out radical, Thomas Aquinas.

Now, according to the natural order instituted by divine providence, material goods are provided for the satisfaction of human needs. Therefore the division and appropriation of property, which proceeds from human law, must not hinder the satisfaction of man's necessity from such goods. Equally, whatever a man has in superabundance is owed, of natural right, to the poor for their sustenance. So Ambrosius says, and it is also to be found in the *Decretum Gratiani*: "The bread which you withhold belongs to the hungry; the clothing you shut away, to the naked; and the money you bury in the earth is the redemption and freedom of the penniless."⁴

It is also worth pointing out that, for most of human history, moral opposition to slavery would have seemed extremely radical.

5. The importance of an ND education (vs an education at one's state university) is not of comparable importance to the lives of 30 people.

How might one argue against premise (5)?

Let's turn instead to the first premise: Singer's strong principle.

1. One always ought to prevent something bad from happening if one can do so without sacrificing anything with moral importance comparable to the thing to be prevented. (Strong Principle)

“If everyone gave to alleviate world hunger, it would only take very little money per person. So why should I give more?”

“Giving money to alleviate hunger only delays the problem, since doing so would only lead to further population growth, which in turn will just lead to more starvation.”

1. One always ought to prevent something bad from happening if one can do so without sacrificing anything with moral importance comparable to the thing to be prevented. (Strong Principle)

It is plausible that something like this principle will follow from most versions of consequentialism.

But it also has some surprising consequences. Imagine, for example, that killing one of my children will, for whatever reason, lead to 30 lives being saved. Is it clear that I must kill my child?

Let's turn to Singer's moderate principle.

This principle, though, can also be used to derive some surprising results.

The moderate principle

One always ought to prevent something bad from happening if one can do so without sacrificing anything of any moral importance.


It is bad for children to starve to death.

Starbucks coffee is of no moral importance.

A Starbucks coffee costs \$3.

It costs roughly \$1 to feed one child in Africa for one day

One can prevent three children from starving for a day by donating the amount of money you would have spent on a Starbucks coffee.



1. One always ought to prevent something bad from happening if one can do so without sacrificing anything of any moral importance (the moderate principle).
2. A Starbucks coffee costs \$3.
3. It costs roughly \$1 to feed one child in Africa for one day.
4. One can prevent three children from starving for a day by donating the amount of money you would have spent on a Starbucks coffee. (2,3)
5. Starbucks coffee is of no moral importance.
6. It is bad for children to starve to death.

C. No one should buy a Starbucks coffee. (1,4,5,6)

Is the argument valid?

Suppose that one were to argue that if no one drank Starbucks coffee, then the company would go out of business, and lots of people would lose their jobs, and that this would be of some moral importance. If all of this were true, would this falsify any premises in the argument?

We have now discussed a number of objections to specific forms of consequentialism — such as hedonistic consequentialism and maximizing consequentialism — and have discussed Singer's attempt to derive results about our obligations to the poor from very general consequentialist principles.

The last series of arguments which we will discuss is an attempt to show, not that some specific form of consequentialism fails, but rather that any consequentialist approach to morality should be rejected.

The last series of arguments which we will discuss is an attempt to show, not that some specific form of consequentialism fails, but rather that any consequentialist approach to morality should be rejected.

One general feature of consequentialism is its indifference to **how** consequences are brought about. What matters when deciding what to do is what one's various actions will bring about, not what those actions are.

One consequence of this general feature might be stated like this:

Act/omission indifference

Whether I bring about some state of affairs by doing something or failing to do it is morally irrelevant.

Some aspects of this principle are quite appealing. For example, the principle refuses to let people stand idly by as others suffer, on the grounds that one is not the **cause** of that suffering. One whose failure to act leads to suffering is, according to consequentialism, just as responsible for it as one whose action leads to that suffering.

But some troubling consequences of this principle are brought out by the following example, due to Judith Jarvis Thomson.



David is a great transplant surgeon. Five of his patients need new parts—one needs a heart, the others need, respectively, liver, stomach, spleen, and spinal cord—but all are of the same, relatively rare, blood-type. By chance, David learns of a healthy specimen with that very blood-type. David can take the healthy specimen's parts, killing him, and install them in his patients, saving them. Or he can refrain from taking the healthy specimen's parts, letting his patients die.

What does the consequentialist say that David ought to do in this case? What ought he to do?

This sort of case might lead you to think something like this: killing someone in order to save the lives of others is never morally permissible.

If this were true, this looks like it would be trouble for the Consequentialist, since it is hard to argue that killing someone, especially when it saves the lives of others, can never lead to an outcome which is, overall, the best of the available options.

But, as some of Thomson's other examples show, matters are not quite this simple.

Consider one of her examples involving a trolley car:

Edward is the driver of a trolley, whose brakes have just failed. On the track ahead of him are five people; the banks are so steep that they will not be able to get off the track in time. The track has a spur leading off to the right, and Edward can turn the trolley onto it. Unfortunately there is one person on the right-hand track. Edward can turn the trolley, killing the one; or he can refrain from turning the trolley, killing the five.

Is it permissible for Edward to turn the trolley? If so, wouldn't this be a case in which it is permissible — perhaps even obligatory — to kill one person in order to save five lives?

But then why might it be OK for Edward to turn the trolley, but clearly not permissible for the doctor to cut up his healthy specimen?

One might try to explain the difference here like this: Edward is choosing between **killing one and killing five**; either way, he is killing someone. David is choosing between **killing one and letting five die**, and this is something quite different. We have a stronger duty to avoid killing than to prevent people from dying.

But it is not clear that this is the right explanation of the difference between Edward and David, as is brought out by the example of Frank.

Frank is a passenger on a trolley whose driver has just shouted that the trolley's brakes have failed, and who then died of the shock. On the track ahead are five people; the banks are so steep that they will not be able to get off the track in time. The track has a spur leading off to the right, and Frank can turn the trolley onto it. Unfortunately there is one person on the right-hand track. Frank can turn the trolley, killing the one; or he can refrain from turning the trolley, letting the five die.

Here it seems as though Frank is faced with a choice between letting five die, and killing one — so his choice seems, in this respect, just like David's (the surgeon's). But it seems as though it **is** morally permissible for Frank to turn the trolley, even though it is **not** morally permissible for David to cut up the healthy specimen.

More complications are introduced by yet a third trolley example:

George is on a footbridge over the trolley tracks. He knows trolleys, and can see that the one approaching the bridge is out of control. On the track back of the bridge there are five people; the banks are so steep that they will not be able to get off the track in time. George knows that the only way to stop an out-of-control trolley is to drop a very heavy weight into its path. But the only available, sufficiently heavy weight is a fat man, also watching the trolley from the footbridge. George can shove the fat man onto the track in the path of the trolley, killing the fat man; or he can refrain from doing this, letting the five die.

Many people think that it is not permissible for George to push the fat man. But why is this any different from turning the trolley to kill the one on the right hand section of the trolley tracks? After all, in both cases, you are killing one rather than letting 5 die.

One thought is this: the fat man has a **right** not to be pushed onto the tracks in a way that people standing on trolley tracks **don't have a right** not to be run over by trolleys.

This sort of thought also promises to make sense of the example of David the surgeon; perhaps healthy specimens have a right not to be cut up, but that dying patients in need of transplants have no right to be saved.

This way of thinking about these cases is very different than the way of approaching them suggested by Consequentialism. According to this view, we should think about what we ought to do by **first thinking about the rights and obligations of the people involved** and not, at least in the first instance, about which action would bring about the best outcome.