

Consequentialism

Last time we were discussing the following view about what it is right and wrong to do:

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.

We noted one sort of objection to this view: namely that, as stated, it does not really give us much help in deciding which actions are the right actions to perform; and this is because it does not tell us what makes one outcome **better than** another.

The first question we will be interested in today is: what makes one outcome, or state of affairs, better than another?

One simple answer to this question might seem to emerge from Singer's discussion. He clearly thinks that what is bad about the sorts of situations he discusses is that they involve massive *suffering*. But what is suffering? Presumably, a certain amount and kind of a particular sensation, pain. This might suggest the following view the view that **one state of affairs is better than another if it contains less pain**.

If this sort of principle were true, it would support Singer's argument. But this is not a very plausible principle.

A better view takes into account pleasure, as well as pain. This sort of view about what makes an outcome good or bad might be stated as follows:

Hedonism

One state of affairs is better than another if and only if it involves the best overall distribution of pleasure and pain.

If we combine hedonism with consequentialism, we get a view about what we are morally obliged to do in every situation: we are morally obliged to pursue the course of action which will (in the long run) bring about the best overall distribution of pleasure and pain. This view might be called **hedonistic consequentialism**; an easier name for it is **utilitarianism**:

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 best overall distribution of pleasure and pain.

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 best overall distribution of pleasure and pain.

Utilitarianism is perhaps historically the most important form of consequentialism. It is a very simple, and very appealing, theory about our moral obligations.

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.

One of its strengths, as Singer's argument shows, is that it is a paradigmatically **unselfish** theory: no one's pleasures and pains are more important than anyone else's.

But the view also faces certain challenges. One forceful way of bringing this out is via 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?

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 best overall distribution of pleasure and pain.

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?



Nozick's example raises a few questions:

- What must the hedonist (and hence also 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 or not to have everyone plugged in to an experience machine. What must a utilitarian say about what you ought to do?
- Does it matter if people ask you not to plug them in?
- Is Nozick right that these consequences of utilitarianism, and hedonistic consequentialism, are incorrect?

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 best overall distribution of pleasure and pain.

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.

It is important to see that, even if you agree with Nozick, his example does not show that Consequentialism is false, but only that a particular version of that view - hedonistic consequentialism - is false. One might agree with Nozick about the experience machine, and still be a Consequentialist, if one holds that what makes one state of affairs better than another can sometimes depend on facts other than sensations of pleasure and pain.

What might make one state of affairs better than another, if not the overall distribution of pleasure and pain? This is a difficult question, to which many different answers have been given. Some relevant facts might include:

- 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 what makes one outcome better than another is a different version of consequentialism.

For example, “preference-satisfaction consequentialism” is the view that one should always act in such a way that maximizes the extent to which the desires of people are satisfied.

So let’s turn from our evaluation of Utilitarianism in particular to an evaluation of Consequentialism in general.

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 best overall distribution of pleasure and pain.

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.

So let's turn from our evaluation of Utilitarianism in particular to an evaluation of Consequentialism in general.

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 options will bring about, not what those options 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.

This principle seems to be a consequence of Consequentialism. And some aspects of this principle are quite appealing. For example, the principle -- as is again illustrated by the example of Singer -- 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 just as responsible for it as one whose action leads to that suffering.

But this principle of act/omission indifference also has some surprising consequences, as is brought out by the following example from 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.



Act/omission indifference

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

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.

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?

How might the consequentialist respond to this sort of case?

The consequentialist might respond to this case by saying that among the morally relevant outcomes of an action are whether that action leads to there having been a murder committed. So perhaps a situation which includes one death + one commission of murder is worse, by the relevant standards of what makes one outcome better or worse than another, than a situation which includes five deaths but no murders. (But: does it matter if the five healthy patients all need new parts because they were the victims of murder attempts?)

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 for the anti-consequentialist.

Act/omission indifference

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

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.

But, as some of Thomson's other examples show, matters are not quite this simple for the anti-consequentialist.

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.

This response to the problem of Edward and David clearly depends on the moral significance of the distinction between acts and omissions -- the moral significance, in this case, of the distinction between killing and letting die.

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.

Act/omission indifference

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

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.

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.

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.

This response to the problem of Edward and David clearly depends on the moral significance of the distinction between acts and omissions -- the moral significance, in this case, of the distinction between killing and letting die.

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. 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:

Act/omission indifference

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

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.

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.

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.

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.**

Act/omission indifference

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

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.

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.

Beginning next time, we will begin discussing this other, non-consequentialist approach to ethical questions.