Kantian ethics

Last time, in our discussion of consequentialism, we discussed Williams' examples of George the chemist and Jim and the Indians. In each of these examples, Williams thinks that we should find the view of the Consequentialist implausible; and in each of these cases, it seems that what makes trouble for the Consequentialist is the fact that we are inclined to find the distinction between doing something and letting it happen morally relevant --- which is what the Strong Doctrine of Negative Responsibility denies.

At least two other worrying sorts of cases for consequentialism are worth considering:

- Cases which involve our intuitions about the *rights* of others. The example of the unwilling transplant.

- Cases in which, if Consequentialism is true, we seem to have a moral obligation to deceive ourselves about what we ought to do.

Consequentialism is one very general framework about how to think about what we ought to do. As the above makes clear, there are many different versions of Consequentialism.

But, as the above also makes clear, whether or not Consequentialism is true has very concrete consequences: for example, it seems to have the Strong Singer Principle as a consequence, and that Principle seems to have as a consequence that you are morally obliged to give almost all of your money to help suffering people around the world.

As we have seen, Consequentialism also faces some serious problems. One might wonder: if Consequentialism is false, what does that entail for Singer's argument? To answer this question, we need to understand how one might think about what we ought to do in a non-consequentialist way. We turn to that topic now.

One of the most important non-consequentialist ethical systems is due to Immanuel Kant, an 18th century German philosopher who is widely regarded as one of the most important, as well as one of the most difficult, philosophers in the history of Western philosophy.

Kant's central claim was that

It is impossible to conceive anything at all in the world, or even out of it, which can be taken as good without qualification, except a **good will.**

The question which his ethics tries to answer is: what is it for a will to be good?

We have already seen the consequentialist's answer to this question: according to a consequentialist, for a will to be good is for it to aim at acting so as to produce the best possible state of affairs. In this sense, the consequentialist thinks that what is "good without qualification" are states of affairs; good wills are defined in terms of the intention to produce good states of affairs.

According to Kant, this view gets things exactly backwards:

> A good will is not good because of what it effects or accomplishes—because of its fitness for attaining some proposed end: it is good through its willing alone—that is, good in itself.
>
> if by its utmost effort it still accomplishes nothing, and only good will is left (not, admittedly, as a mere wish, but as the straining of every means so far as they are in our control); even then it would still shine like a jewel for its own sake as something which has its full value in itself. Its use-

According to Kant, a good will is a will that chooses a certain action because it is the action dictated by duty. To understand his view of the good will, we therefore need to understand what duty is.

Kant calls rules which say what we ought to do *imperatives*.

There are different sorts of imperatives. Some of these are *hypothetical imperatives*: these say that we ought to do such-and-such if we are interested in reaching some end (which Kant calls a "condition"). An example of a hypothetical imperative might be that you ought to go to the dining hall tonight (given that you are interested in the end of eating dinner).

Kant does not think that the duties in terms of which the good will is defined are hypothetical imperatives. Rather, he says,

> there is an imperative which, without being based on, and conditioned by, any further purpose to be attained by a certain line of conduct, enjoins this conduct immediately. This imperative is **categorical.**

The categorical imperative simply says what you ought to do, period -- not what you ought to do, given some other interest.

So what are the commands of morality -- the categorical imperatives?

Somewhat surprisingly, Kant holds that there is just one:

> There is therefore only a single categorical imperative and it is this: *"Act only on that maxim through which you can at the same time will that it should become a universal law."*

Kant calls this the *formula of universal law.*

Your *maxim* is your reason for pursuing an action. What Kant is saying here is that the sole command of morality is to act only on maxim which are such that one could act from that maxim *while also willing that everyone else should also act from that maxim*. This is what it means for a maxim to "become a universal law."

The best way to see what Kant means is by considering an example; he discusses several, including the example of breaking a promise:

> 2. Another finds himself driven to borrowing money because of need. He well knows that he will not be able to pay it back; but he sees too that he will get no loan unless he gives a firm promise to pay it back within a fixed time. He is inclined to make such a promise; but he has still enough conscience to ask "Is it not unlawful and contrary to duty to get out of difficulties in this way?" Supposing, however, he did resolve to do so, the maxim of his action would run thus: "Whenever I believe myself short of money, I will borrow money and promise to pay it back, though I know that this will never be done." Now this principle of self-love or personal advantage is perhaps quite compatible with my own entire future welfare; only there remains the question "Is it right?" I therefore transform the demand of self-love into a universal law and frame my question thus: "How would things stand if my maxim became a universal law?" I then see straight away that this maxim can never rank as a universal law and be self-consistent, but must necessarily contradict itself. For the universality of a law that every one believing himself to be in need may make any promise he pleases with the intention not to keep it would make promising, and the very purpose of promising, itself impossible, since no one would believe he was being promised anything, but would laugh at utterances of this kind as empty shams.

Kant's line of reasoning here appears to be this: if I consider the maxim

Promise to get money whenever I need it with no intention of paying it back.

as a universal law, then I imagine a scenario in which everyone is constantly making false promises. But in this sort of scenario, the convention of promising would cease to exist: after all, no one would have any reason to lend money on the basis of promises if such promises are never kept. So in such a world it would be impossible to act on this maxim.

How might an analogous line of reasoning be used to show that lying is wrong?

For Kant, these are the paradigm cases of wrong action: actions proceeding from maxims whose universalization involves a contradiction.

This is already enough to bring out some important contrasts between consequentialist views of the sort we discussed and Kant's ethics.

1. Kant's ethics tells you, in the first instance, what morality forbids you from doing. But it does not tell you what you ought to do in every case; some actions might be morally praiseworthy even though not doing them would not be contrary to the Formula of Universal Law, and hence not morally forbidden. These actions are, therefore, neither morally required nor morally forbidden. For the consequentialist, on the other hand, one must always do what will bring about the best consequences: so (excluding ties) every action is either morally required or morally forbidden.

2. According to the consequentialist, the rightness or wrongness of a particular action depends on which action, in these particular circumstances, would lead to the best outcome. According to Kant, by contrast, the rightness or wrongness of acting from a particular maxim just depends on the *type* of maxim that it is. If making false promises, or lying, is sometimes morally forbidden, then it is always morally forbidden.

This last point -- that the rightness or wrongness of an action just depends on the type of maxim from which one is acting, rather than on the consequences of this particular action -- explains why Kantian ethics yields quite different results about what we ought to do than the sorts of consequentialist views we discussed.

Suppose, for example, that a judge knows that the defendant in a capital case is innocent, but also knows that not finding the defendant guilty and sentencing him to death will result in riots in which many will be killed. What would a consequentialist say about this sort of case? How about the Kantian?

In this sort of case, it might seem that the Kantian gets things right, and the consequentialist gets things wrong. But there are other cases where things might not seem to so clear. Here is one such example:

>You're living in Nazi Germany, and hiding a Jewish family in your basement. The authorities come to the door, and ask you whether you are hiding a Jewish family in your house. You know that they will believe you if you tell them that you are not; it is just a random check. What should you do?

What does the Kantian say about this sort of case? How about the consequentialist? What should we say?

Now recall the case of the unauthorized organ transplant, discussed above. What would the Kantian say about this case? Does the maxim from which you were acting lead to a contradiction if universalized?

It may not be obvious. However, in addition to the Formula of Universal Law, Kant gave another formulation of the categorical imperative, which he states as follows:

> Act in such a way that you always treat humanity, whether in your own person or in the person of any other, never simply as a means, but always at the same time as an end.

This is often called the *formula of humanity*. What does this principle say about the case of the unwilling transplant? How about the case of suicide?

How about cases of self-defense?

So we have two different moral rules: the formula of universal law and the formula of humanity.

One puzzling question is how these two principles are related. Kant seems to think that these two formulas are different formulations of the same idea; but it is a bit hard to see how this could be correct. Perhaps the two formulas are equivalent, in the sense that they yield the same result about which actions are right and which wrong; but it is not even clear that this is correct. In any case, there is no requirement that one's ethical theory be based on a single principle.

The choice between consequentialist and Kantian ethics is a difficult one, as there are many examples which are challenging to each sort of view.

Here is a final challenge to the Kantian perspective worth thinking about:

> According to the Kantian, what are really good or bad are not the consequences of our actions, but the actions themselves. So consider some bad actions, like acts of lying. Surely I ought to prevent such actions if I can, especially if I have the opportunity to prevent very many of them. But suppose that I can only prevent (say) the telling of 10 lies by myself telling a lie. Should I? Presumably the Kantian will say "no," since telling a lie is morally forbidden. But can this be right? Aren't 10 lies worse than 1?

What should the Kantian say? Is this a genuine problem?