

Week 2 Notes

Philo 101 Online | Hunter College | Fall 2019

Daniel W. Harris

Welcome to Week Two! For the next few weeks, our subject matter will be ethics. We've already encountered a bit of this last week: some of David Foster Wallace's central questions in *Consider the Lobster* are questions in ethics:

Is it all right to boil a sentient creature alive just for our gustatory pleasure? ...What does "all right" even mean in this context? Is it all just a matter of individual choice?

Over the next few weeks, we will take a closer look at some of the most influential ways that philosophers have attempted to answer ethical questions like these.

1 Ethical Theories

This week and next week, our main subject matter will be two of the most influential ethical theories that philosophers have articulated: utilitarianism and Kantianism. An ethical theory is a doctrine whose aim is to systematically answer ethical questions. It is an answer to the following question:

A CENTRAL QUESTION OF ETHICS

Which actions are morally permissible, which are morally obligatory, which are morally forbidden, and which, if any, are morally supererogatory (i.e., highly admirable but beyond what is morally required).

A completely general answer to this question would be an ethical theory—a theory that tells us how to evaluate the moral status of each action.

- An action is morally *permissible* if doing it is consistent with acting in an ethical way. For example: most of us would say that it is morally permissible to harm another person in self-defense, but morally impermissible to harm someone for no good reason.
- An action is morally *obligatory* if it's what you absolutely must do in order to behave in a moral way. For example, many of us would say that if you see a child who is about to wander out into a busy street and almost certainly be killed, and you can easily save them at no risk to yourself, then it is morally obligatory for you to do so. If you don't save them, you're doing something wrong. But we normally wouldn't tend to think that it is morally obligatory to sacrifice your own life to save another person's life.
- An action is morally *forbidden* if doing it is inconsistent with what is morally required. Murdering an innocent person is morally forbidden, if anything is. But most people wouldn't think of eating too many sweets as morally forbidden (even if it isn't good for you).
- An action that is morally *supererogatory* is one that is extremely morally praiseworthy—so much so that it goes beyond what morality requires. Supererogatory actions “go above and beyond the call of (moral) duty”. Many of us would think of donating a kidney to a stranger as a supererogatory action.

Notice that that these categories stand in certain relationships. If an action is obligatory or supererogatory, it must also be permissible: if you're required to do something, or if it's praiseworthy to do it, then it must be allowed. And if an action is forbidden, then it's not permissible, and you're obligated *not* to do it.

Asking how actions should be placed into these categories is a precise way of asking central ethical questions like “what is the right thing to do?” and “how should I live my life?”

2 What Ethics is Not

The first reading for this week is ‘About Ethics’ by Peter Singer. One of the central aims of this reading is to point out some bad ideas that newcomers to ethical theory tend to find tempting.

One common mistake is to confuse ethical questions with legal questions. For example, when asked whether it is morally permissible to use a racial slur, some find

it tempting to reply that of course it's permissible: after all, the American constitution protects freedom of speech. But this answer misses the point of the question, because it conflates moral permissibility with legal permissibility. The law doesn't always line up with morality. For example: there is no law against lying to your spouse, but it is at least sometimes immoral to do so. And it shouldn't be hard for us to think of immoral laws from the history (or present) of this country.

Some of the other moral theories that Singer debunks are often summed up in slogans that you should be familiar with. For example, the so-called golden rule: "Do unto others as you would have them do unto you". In modern English: treat people the way that you want to be treated. This sounds pretty good if you don't think about it too much. But after a bit of reflection it's clearly not a good principle. After all: some people like to be whipped. Does that mean that they should go around whipping everyone else? No! The golden rule is a bad rule, precisely because different people prefer to be treated in different ways.

Another tempting slogan is that you should "follow your conscience" or that you should "do what you think (or feel) is right". This slogan is suggestive of ethical *subjectivism*, which is the theory that each person is morally obligated to do whatever their own beliefs dictate is morally obligated. Again, this may sound plausible if you don't give it a second thought. But the problem is pretty clear upon further reflection: some people have false beliefs about how they should act! Some of America's founders believed that owning slaves was morally permissible, for example. According to subjectivism, their belief would have made it permissible for them to own slaves. (It's easy to think of similar examples involving terrorism, genocide, and so on.) Any ethical theory that has consequences like these must be wrong.

Finally, many of us are tempted by the idea that there is no universally right way to act, and that you should instead act in accordance with the customs and practices of the society in which you find yourself. Although it's wrong to push someone onto a subway car in New York for example, that's just how things work in some parts of Asia. Some are tempted by considerations like these into believing *cultural relativism*, which is the idea that what is right or wrong in a given society is just whatever the members of that society takes to be right or wrong. But again, as Singer points out, this position turns out to have bizarre and undesirable consequences when examined carefully. Just to take one of the points Singer makes: most societies don't agree on everything, and it often turns out that the majority is wrong. Slavery again provides a good example. Suppose that the majority of the people who lived in the Confederate States of America in 1862 believed that slavery

was morally permissible. According to cultural relativism, it would have followed that slavery was morally permissible, there and then. An ethical theory that tells us this should not be trusted.

3 Consequentialism and Utilitarianism

Singer's preferred ethical theory is a version of *utilitarianism*, which is in turn a version of *consequentialism*. Consequentialist theories are so-called because they classify actions into the different ethical categories on the basis of the action's consequences. What we are obligated to do, according to consequentialists, is whatever will have the best consequences.

Consequentialism is not a single ethical theory, but a family of them. The theories in this family differ from one another on the basis of what they take the best consequences to be. To arrive at a particular consequentialist theory, we need some way of ranking the possible outcomes of actions from best to worst. The theory will then tell us that we ought to take whatever action will have the best consequences, given the options available.

How should we evaluate the outcomes of actions? Put another way: what outcomes should we value most? For example, someone could argue that what we should value most is the preservation of the Earth's biodiversity. Or we could imagine someone arguing that what we should care about most is the production of as many geniuses as possible. And some people seem to act as though they believe that we should value outcomes solely according to the amount of economic prosperity. We could construct a version of consequentialism around any of these ranking systems. Quite generally, for any way of deciding that some possible futures are better than others, there is a corresponding version of consequentialism that tells us that we are obligated to aim for the best possible future according to that ranking.

Of all of these conceivable versions of consequentialism, the most influential versions fall into the subcategory called 'utilitarianism'. Peter Singer is probably the most famous utilitarian who is alive today, and John Stuart Mill is probably the most famous utilitarian who has ever lived, but there have been many others as well.

The central idea of utilitarianism is that the thing we should value most is happiness, and that the best possible outcomes are the ones in which happiness is maximized. In a slogan: we are morally obligated to act so as to maximize overall happiness and minimize overall unhappiness in the world. According to utilitarianism, things like biodiversity, genius, and economic prosperity may also be valuable, but

they are only instrumentally valuable, and only insofar as they lead to greater overall happiness. In many situations, it is likely that more overall economic prosperity will lead to more happiness, but this needn't always be the case. For example: we might be able to increase economic prosperity by forcing everyone to work 16 hours a day at gunpoint. But if we find ourselves with the option of increasing prosperity in this way, or in some other way that comes at the expense of overall happiness, a utilitarian would say that we should not do so.

This explanation of utilitarianism raises several further questions. First: whose happiness are we talking about here? Who must be taken into account, when deciding what would maximize happiness?

The utilitarian answer is that everyone's potential happiness and unhappiness matters equally. In particular: utilitarianism is *not* the idea that you should do whatever will make *you* most happy. (This is a common, and disastrous mistake that students sometimes make in introductory philosophy classes.) Rather, utilitarianism tells us that we must always weigh our own happiness against the happiness of everyone else, and do whatever will lead to the most overall happiness. If we define selfishness as acting in a way that benefits oneself at the expense of others, then utilitarianism forbids us to act selfishly.

This is a fairly intuitive ethical principle. Suppose you face a choice of either eating all of the ice cream in the freezer or eating only a small amount and leaving some for your two siblings, both of whom really like ice cream. You would be slightly happier if you get to eat it all, but your siblings would both be much less happy if they don't get any. In this case, the small amount of happiness at stake for you is outweighed by the large amount of happiness at stake for your two siblings. Utilitarianism says that you are morally obligated to share. Note that this is true even if you were the one who paid for the ice cream. Utilitarianism doesn't care about that! It only cares about maximizing happiness, even if that requires giving up some of your property to do so. We'll return to this point a bit later.

We still haven't fully answered the question of whose happiness should matter. After all: "everyone" is not a very informative answer. Does everyone include nonhuman animals? What about human fetuses? If we run into some aliens while exploring the Milky Way, how will we know if their happiness should count? Again, the utilitarian has a simple, if less than fully informative answer to these questions: if a being is capable of happiness or unhappiness, then we must take them into account when deciding what to do. If an animal or a fetus can experience happiness, then we should include them when we rank the potential outcomes of our actions.

Okay, but how can we tell whether animals can experience happiness?

This brings us to the hardest question that utilitarians have to answer: what is happiness, and how can we quantify it? Some of you will be tempted to say that these are stupid questions: happiness isn't any one thing, or isn't definable, you might be thinking, and it isn't quantifiable either. If you're right, then utilitarianism turns out to be an incoherent idea.

However, most utilitarians have tried to give definitions of happiness, and have thought that it is at least in principle quantifiable, even if it is very difficult to measure it in practice. But not all utilitarians define happiness in the same way. So, once again, utilitarianism can be thought of as a family of ethical theories, rather than a single ethical theory. Theories in the utilitarianism family differ according to how they define and quantify happiness.

An early and simple form of utilitarianism, which was formulated by the philosopher Jeremy Bentham, defines happiness as pleasure and the absence of pain. To be happy, according to Bentham, is to be in a state of physical or emotional pleasure, and to avoid suffering. His version of utilitarianism directs us to act in such a way as to increase the amount of pleasure in the world, and to decrease the amount of suffering, as much as we can.

Many have found this to be a highly intuitive moral principle. After all: suffering seems to be an inherently bad thing, and pleasure an inherently good one. And this theory seems to fit with some of our deepest-held moral convictions. The most morally repulsive actions we can think of—torture, assault, child abuse—are all acts that cause enormous amounts of suffering for no good reason. On the other hand, we tend to think that a little bit of suffering can be worth it if we can thereby avoid even more suffering (and create an opportunity for more pleasure) later on. This is how we're reasoning when we cause pain to babies by vaccinating them, for example. We are implicitly weighing their present pain against much more significant future suffering.

Bentham's theory also gives us a strategy for answering the question about which kinds of beings are morally relevant. We just need to figure out whether they are sentient—i.e., whether they are capable of pleasure and pain. As David Foster Wallace showed us, this is not always an easy question to answer in particular cases. But his example of the lobster also shows us that it is a question that science can make concrete progress on. And, since it is very plausible that many animals are capable of suffering, many utilitarians have argued that we are morally obligated to change the ways in which we treat farm animals. Peter Singer is not just the best-known

living utilitarian, he is also one of the most famous advocates for veganism, which is the practice of eliminating animal products from one's diet.

Our second reading for this week offers a refinement of Bentham's version of utilitarianism. John Stuart Mill agrees with Bentham that happiness is pleasure and the absence of pain. However, he doesn't think that all pleasures are created equal. Some forms of pleasure, he thinks, should be more highly weighted than others when we are ranking the possible outcomes of our actions. In particular, Mill thinks that intellectual pleasures (which he calls "higher pleasures") should be weighted more highly than mere bodily pleasures (which he calls "lower pleasures"). You should read Mill's argument carefully, and think about whether you agree, and why.

A third kind of utilitarianism has been defended by Peter Singer, who argues that it's wrong to think of pleasure itself as inherently valuable and pain as inherently bad. Pleasure is normally a good thing, he thinks, but not always. After all: some people like to feel a bit of pain every now and then. Some people like to "feel the burn" (a kind of suffering) that comes from running a marathon, for example. And, as mentioned earlier, some people like to be whipped, and this presumably causes them pain. Bentham and Mill might have said that these people are irrational, or might have said that the people in question choose these moments of suffering because of the greater forms of pleasure that tend to accompany or proceed them. But Singer thinks that we should not try to assume that there is anything wrong with choosing pain over pleasure. And so it isn't pleasure that we should try to maximize, according to Singer, but the *satisfaction of preferences*. To satisfy a preference is to get what you want. Often, we want things that will bring us pleasure, but maybe not always. According to Singer's version of utilitarianism, we are obligated to act in a way that will satisfy as many preferences as possible.

Notice that, in last week's reading, David Foster Wallace tried to convince us not only that lobsters have a capacity to suffer, but also that they genuinely prefer not to suffer. So, once again, there is a good case to be made for the idea that animals have preferences too. If so, then preference utilitarianism tells us that their preferences must be taken into account, just as humans' preferences must.

Finally, notice one more important difference between Bentham's and Mill's version of utilitarianism and Singer's version. It is sometimes possible to kill someone without causing them any pain. Of course, this may cause lots of suffering to their loved ones. But what about a hermit orphan who doesn't have any loved ones, and nobody would miss? What if you snuck up behind them, painlessly knocked them out with some chloroform, and painlessly killed them? Would it be wrong to do so?

According to hedonic utilitarianism, maybe not! After all: no pain is seemingly being caused, and maybe you're even getting some kind of psychopathic pleasure from your action. But this is a terrible consequence for an ethical theory to have! Singer's theory does a bit better: it tells us that since your victim had a very strong preference to stay alive, and presumably many other preferences beside, you're frustrating many strong preferences by ending their life. It is unlikely that your preference to kill could outweigh all of that. So, it's tempting to conclude that preference utilitarianism makes better sense of the fact that even painless murder is wrong.

4 Utilitarianism is Demanding

As I've already suggested, utilitarianism's recommendations often match up with our antecedent ethical beliefs. It's usually intuitive that we should do whatever will maximize happiness (whatever happiness amounts to). But this isn't always the case. Utilitarianism is sometimes very demanding in the obligations it imposes on us. Some have thought that this gives us a reason to reject utilitarianism, because it would require us all to be saints. But many utilitarians have embraced this consequence.

Peter Singer is an example: he argues that most of us in the developed world are morally obligated to donate large portions of our income to charity. And he practices what he preaches: donates about 70% of his income to charity. And although you might be thinking that that's fine for Peter Singer, who is a rich Princeton professor, his conclusions have nothing to do with you, I think he would disagree. Although you might not feel rich, you are almost certainly quite rich by global standards, and there is a good utilitarian argument that you should use more of what you have to help others than you probably currently do. This is the subject matter of the youtube video in this week's required readings.

Here's the argument, in brief. You often spend small amounts of money on things that give you a bit of pleasure. Maybe you buy a fancy coffee or a \$15 bowl of ramen, or some ice cream, or a new purse, or a slightly more expensive pair of shoes than you absolutely need. Singer points out that you could instead have saved that money (or bought a cheaper alternative) and given it to a charity that would have used it to reduce a great deal of suffering in another part of the world. So, for example, the cost of a bowl of fancy ramen could pay for several mosquito nets, one of which would likely prevent a child in Africa from dying a painful death from malaria. The amount of pleasure that you get from your ramen just doesn't com-

pare to the amount of suffering that would be inflicted on the malaria victim. It follows from utilitarianism that you are obligated to give away the money instead of spending it on the ramen (or the coffee, or whatever). That is what it would take to maximize the overall amount of happiness with your actions, after all. So, suppose you follow this reasoning and donate \$5 to charity. What about the next \$5 in your bank account? There are lots of children at risk of malaria, and so the same reasoning kicks in again! In the end, it seems to follow from utilitarianism that you should donate anything you can possibly spare to people who are in need—right up to the point where doing so would no longer maximize the overall level of happiness in the world. This is Peter Singer’s conclusion, and he has devoted a great deal of energy to what he calls the “effective altruism” movement, which aims to convince people to donate large portions of their income, and to identify the best places for them to donate it.