

# PHI 202 | Notes 1

Michal Masny | 18 Sep 2019

## Readings:

Shaw, W. (2006). The Consequentialist Perspective.

Singer, P. (1972). Famine, Affluence, and Morality.

## A summary of what we talked about

In the first precept, we focused almost exclusively on the issues raised by Johann's second lecture. In particular, we talked about the role of intuitions in ethics, Singer's *Argument by Analogy*, supererogatory actions, effective altruism, and the ethics of giving more generally.

We did not talk about consequentialism, the attractions of this view, and the objections levelled against it. We also did not discuss about the similarities and differences between *Singer's Principle* and consequentialism. We will come back to this in future precepts.

## The purpose of this and other precepts

If you didn't leave the precept with a neat set of notes, don't despair! The point of this and other precepts was to introduce you to certain questions, to shake your belief system a little bit, and to give you the tools for thinking about these issues independently.

## Surveys and summaries

Johann's lectures will always be significantly more structured than my precepts, but neither he nor I are in the business of just presenting you with pros and cons of different answers to questions that ethicists have grappled with. Survey articles and summaries are easy to come by online. The Stanford Encyclopaedia of Philosophy and the Internet Encyclopaedia of Philosophy are two very helpful, regularly updated online resources which professional philosophers consult frequently. There, you can find comprehensive entries on all of the central topics. I can also recommend survey articles published in the journal *Philosophy Compass*. Whereas SEP and IEP *aim* to offer impartial surveys, those in *Philosophy Compass* tend to be more 'opinionated'.

## Intuitions

We started by talking about *intuitions*. Philosophers often appeal to hypothetical scenarios in their arguments which are usually called *cases*. For instance, in the lectures we talked about cases such as *Grenade*, *Trolley*, and *Organ Transplant*. Johann asked you what you think is the right action to perform in these circumstances. Importantly, he was not asking you to report whether you are consequentialists, Kantians, or whatever. Instead, he wanted to know what your intuitions are.

The nature of intuitions is a controversial issue in philosophy. You can find a long entry on this in SEP: <https://plato.stanford.edu/entries/intuition/>. We won't spend any more time discussing this, but you might find it interesting. I suggested that a helpful way to think about intuitions is as *considered pre-theoretic judgments*. When you reflect carefully on the details of a particular case, an answer to the question at hand will suggest itself to you. This is similar to what happens in mathematics and other sciences. You may be confronted with a problem or a puzzle and, even though you don't yet know how *exactly* to solve it or to prove it, you will have a sense of direction. My proposal is to use intuitions in ethics in a similar way: as a guide.

Of course, our intuitions might be mistaken, so we should not always rely on them unconditionally. Be prepared to revise your beliefs and to go against your intuitions in the process of inquiry.

Some intuitions are universally or almost universally shared. Singer, for instance, does not even want to consider the possibility that suffering and premature death are *not* bad. The intuition that it would be wrong not to save the child drowning in a shallow pond is also extremely common. But in other cases, people's intuitions about particular questions differ. We saw examples of this both in the lectures and in the precepts.

There is a branch of philosophy called *experimental philosophy* which is interested in these differences. Among other things, they conduct surveys to see whether intuitions about central philosophical puzzles differ across people from different cultures and socio-economic groups, speaking different languages, holding different religious beliefs, and levels of philosophical expertise, and so on. As many of you pointed out, these factors could influence what intuitions people have.

They also look at various *framing effects* which may be familiar to you from social sciences. As a number of you pointed out, the specific details of the case as well as the order in which cases are presented could influence what intuitions people have.

To illustrate these framing effects, we talked about the *Trolley* case. In *Trolley*, there is a runaway trolley, 5 people tied down to one track, and 1 person tied to another track. If you do nothing, the 5 will die. If you pull the lever and divert the trolley, the 1 will die. Many of you had the intuition that one should pull the lever and divert the trolley.

We contrasted this case with the *Large Man* case. (It was first discussed by Shelly Kagan in his 1989 book *The Limits of Morality*.) In *Large Man*, there is a runaway trolley and 5 people are tied down to one track. There is also a footbridge. You are on the footbridge and next to you is a very large man. The only way to prevent the death of the 5 people is to push the very large man off the footbridge. He will die but his body will stop the trolley. Should you push the large man? Fewer of you had the intuition that one should push the large man off the footbridge.

These cases are meant to be analogous, structurally similar. But you have correctly pointed out that there are differences between these two cases. One difference is that pushing a person off the footbridge feels different than pulling a lever. Another difference is that the large person is in some sense a by-stander to the whole situation, whereas the one person tied to the track is in some sense already involved in the situation.

Whenever we consider pairs of groups of cases like that, there will be such differences. The important question to ask then is whether these are *morally relevant differences*. If you think that they are not, you will regard them as distractions, as something that can confuse our intuitions about these cases. The *Trolley* case, you might think, is 'cleaner' and therefore a more reliable basis for our intuitions. This might prompt you to revise your intuitions about *Large Man*.

The order in which cases are presented is also likely to matter. It would be interesting to see what intuitions would people have if we considered *Large Man* first and *Trolley* second. Would the first response *prime* you to respond in a particular way in the second case? This is another known framing effect.

One immediate lesson from these considerations is that whenever you consider a case presented by someone else or one you came up with, you should play with the details and consider proximate cases.

Experimental philosophers want to reach further reaching conclusions on the basis of such considerations. Those involved in the *Negative Program of Experimental Philosophy* argue that the diversity of intuitions and how easily they can be influenced undermines their role in philosophy.

Those involved in the *Positive Program of Experimental Philosophy* are in the business of identifying factors that can influence our intuitions and studying the effects that they have. The thought here is that we can treat these influences as distortions and may be able to correct for them. You can read more about it at <https://plato.stanford.edu/entries/experimental-philosophy/>.

If you are interested in these issues, there will be an upper-level undergraduate course on experimental philosophy in the spring semester. It will be taught by Stephen Stich who is one of the pioneers of this branch of philosophy. He is a professor at Rutgers University, but he will be a visiting professor at Princeton's University Center for Human Values in the spring semester.

### **Singer' Argument by Analogy**

Next, we turned to Peter Singer's article. We observed that Singer offers two arguments. In addition to the more prominent *General Argument*, Singer offers an *Argument by Analogy*.

We broke this argument into premises and a conclusion. The first premise is roughly that it would be wrong not to save the child drowning in the shallow pond in the case he imagines. We noted that Singer explicitly appeals to intuitions to establish the truth of this premise. We could slightly re-phrase the premise to make it more explicit: *intuitively*, it would be wrong not to save ... .

The second premise is roughly that there are no *morally relevant* differences between (i) saving the child drowning in the shallow pond on one hand and (ii) saving lives and alleviating suffering of people in need in distant countries. We then discussed whether this premise is true. Many of you identified various differences between (i) and (ii). We then considered whether these differences are *morally relevant*.

One difference to which many of you pointed out is that the monetary sacrifice involved in saving the child drowning in a shallow pond is relatively small. You will ruin a pair of pants, perhaps a suit. Even the most expensive suits do not cost more than a few thousand dollars. By contrast, the conclusion of Singer's argument is that we should donate a significant portion of our income to aid organisations. It is not clear how much exactly it is and it will certainly vary from person to person. As Johann noted in the lecture, in his 2009 book *The Life You Can Save*, Singer defends a moderate position that, if we can, we should donate 10% of our income to aid organisations. In the USA, the mean income per capita is ~\$50,000. If an average American donates 10% of their income over, say, 40 years, this will amount to ~\$200,000. This is much more than the cost of the most expensive suit.

We then considered a question whether morality could require you to sacrifice *that much* to save someone's life. We considered a case of *Bob's Bugatti*. Bob is an average American who is about to retire. He invested almost all of his retirement savings in a rare expensive car: a Bugatti. He reasonably expects that the resale price of this car will increase over the coming years and he will be able to sell it to finance his retirement years. One day, Bob goes for a ride in his Bugatti and stops by the tracks. As it happens, there is a runaway trolley and a child on the track. If Bob does nothing, the child will die. He can, however, pull a lever which will divert the trolley onto his car; the car will be destroyed but the child will live. Should Bob divert the trolley?

(This case was first discussed by Peter Unger in his 1996 book *Living High and Letting Die: Our Illusions of Innocence*, which contains very many more variants on Singer's *Pond* case.)

You might also find the following variant on *Bob's Bugatti* interesting. The case is as earlier except that it is *you* and only you who can divert the trolley onto *Bob's* car. Would it be the right action?

## Supererogatory actions

We then discussed the importance of the following detail of Singer's case: the pond is shallow. This means that you would not be risking your life and health by embarking on this rescue mission.

But consider a different case: *Building On Fire*. Suppose there is a building on fire across the street. You know that there is a child trapped inside this building. The fire brigade isn't here and they won't arrive quickly enough. No one will hold it against you if you do nothing. If you choose to go inside, you will risk your life and, in the best case, suffer from burns. Is it morally required that you go in? Or is it permissible not to go in?

For consequentialists, the answer depends on weighing the actual consequences (will you die, how significant will the damage to your health be vs how good the life of the rescued child would be) or the expected consequences (what is the probability that you will die or suffer serious injuries and how significant they would be vs ...). But if we fill in the details suitably, there will be cases in which you should sacrifice your health or even life to save someone else's life.

Many non-consequentialists find this overly demanding. They say that there are certain actions which are *supererogatory*, actions which are *beyond the call of duty*. They say, for instance, that even if you could save 5 or 10 or 25 people by sacrificing your own life, you would not be morally required to do so. They think that there is a *threshold* for personal sacrifice which divides morally required actions from the supererogatory ones.

Where exactly is this threshold, if there is one? It is difficult to say. We can ask: would you be required to sacrifice your arm to save someone else's life? Likewise, we asked when we talked about *Bob's Bugatti*: would you be required to sacrifice your retirement savings to save someone else's life?

If you are interested in these questions, you can read more about it at:  
<https://plato.stanford.edu/entries/supererogation/>

Until this point we considered questions of *whether* and *how much to give*. But there are also important questions concerning *where* to give. Once again, we can wrap our heads around it not by thinking about donations directly but rather by considering other, structurally similar cases.

We discussed the following case from Shelly Kagan's 1989 book *The Limits of Morality*. Let's call it the *Child/Parrot Rescue* case. It goes like this: A building across the street is on fire. You know that there is a child trapped inside. You take the risk and enter the building. Inside, you find the child as well as a beautiful parrot. As it happens, you are a lover of birds and saving this parrot would bring you a lot of joy. Unfortunately, you can only save either the child or else the parrot. Is it permissible for you to save the parrot?

Many of you had the intuition that it would be wrong to save the parrot when you could save the child instead. Some of you had reservations about the idea of putting value on people's and animal's lives and about the idea that we compare such values. These are important issues which we will discuss later in the semester. For now, let's assume that the life a child is more valuable than the life of a parrot.

This case raises an important puzzle. Suppose that you think that going inside the building is supererogatory regardless of whether the child or the parrot are inside. Let's think about this more carefully. First, suppose the options are:

- (1) Do not enter the building.
- (2) Enter the building and save the parrot.

If you think that entering the building is supererogatory, then both (1) and (2) are permissible. Second, suppose that the options are:

- (1) Do not enter the building.
- (3) Enter the building and save the child.

Again, you think that entering the building is supererogatory, then both (1) and (3) are permissible. Third, consider the options in Kagan's *Child/Parrot Rescue*:

- (1) Do not enter the building.
- (2) Enter the building and save the parrot.
- (3) Enter the building and save the child.

Many of you thought that (2) is impermissible in these three-option choice; only (1) and (3) are permissible. One important question that philosophers have grappled with recently is why would (2) become impermissible in this context.

If you are interested in this issue, you can have a look at Theron Pummer's article "Whether and Where to Give". (Theron Pummer will be giving a talk at the Princeton Workshop in Normative Philosophy on 17 Oct 2019 at 4:30 pm in Marx 201. I don't know that it will be about yet, but you are welcome to attend). In fact, you might find it more interesting to read Johann Frick's precis and commentary of this very article on a prominent philosophy blog here: <http://peasoup.us/2017/04/philosophy-public-affairs-discussion-pea-soup-theron-pummers-whether-give-critical-precis-johann-frick/>. (Yes, philosophy blogs are a thing. Look at the comments section, many big names in philosophy participated in this discussion.)

In addition to raising a puzzle which is important on its own, this case can tell us something interesting about giving to aid organisations.

Suppose that you are not convinced by Singer's argument and think that donating to charities is beyond the call of duty. Many of you have nonetheless donated money to aid organisations in the past and will probably do it in the future. There may be issues that are close to your heart. Perhaps you lost a loved one to some rare version of cancer and you support an organisation which funds research into this disease. You might think that, because donating money to charities is beyond the call of duty, it is permissible for you to support any charity you see fit. However, it is plausible to think that some charities are *more effective* than others in alleviating suffering and preventing deaths. This could be because some organisations are more efficient, because the causes they support are more pressing, or both. An argument by analogy has it that if it is wrong to save the parrot rather than the child in the *Rescue* case, it is wrong to donate to less effective organisations rather than the more effective ones.

### **Effective Altruism**

We ended by talking about the effective altruism movement. Some of you have heard about it before and others did not.

Following up on the idea that some aid organisations are more effective than others, I asked you what you think are the most effective organisations. Many of you pointed to organisations which are very big and get a lot of publicity like UNICEF, Oxfam, and Amnesty International.

There is an organisation called Give Well which conducts research into the effectiveness of other aid organisations. I encourage you to go on their website and read their list of Top 10 most effective charities: <https://www.givewell.org/charities/top-charities>.

Of course, there are important questions about their methodology. What criteria do they use for evaluating the importance of causes and the efficiency of charities? You can read about it on their website as well.

You will see that many of these organisations focus on malaria. Malaria is a disease which affects over 200 million people each year. If you look at the percentage of deaths alone, you will reach the conclusion that malaria is not very deadly. It causes fever and dehydration. This doesn't sound too serious. But malaria nonetheless kills over 400,000 people every year, mostly in Sub-Saharan Africa. Many of these are children whose organisms are less resilient. These deaths are almost all preventable. All that is needed is access to fairly basic medical care.

Another organisation, GiveDirect, literally directly transfers your money to poor individuals in Kenya and Uganda. Three dollars in Princeton can get you a just a single cup of coffee, but it goes a long way in these countries.

If you are compelled by the arguments discussed earlier, and choose to support *some* charity, then you should perhaps support these charities rather than one which funds research into some rare variety of cancer which affects only a few or a few hundred people every year.

The last thing we talked about is your career choices. You will spend roughly 80,000 hours at work in your life. This will be the single most time-consuming activity in your life. In most cases, you will choose a career early on and stick to it until retirement. Thus, your career choice is really important.

I asked you what would you do if you had purely altruistic motivations and wanted to devote your life to helping others in need. Many of you suggested becoming a doctor or working for a charity. Yes, you would be most directly involved in doing good.

But *perhaps* this is not how you can do *the most good*. Maybe you have a unique set of skills. Maybe you are exceptionally good with numbers, could become an exceptionally successful investment banker, earn hundreds of thousands of dollars, and donate 50% or more of that to charity. Or maybe you are exceptionally good at baseball, could become a professional player, score a high-paying contract, and donate a large proportion of your money to charity. Or maybe you have a talent for acting, could become a Hollywood star, earn lots of money, and donate a large proportion of it to charity. *Perhaps*, you would actually do more good that way than a doctor would.

Of course, there is an important question whether you would *actually* keep donating a large portion of your income if you became an investment banker, a professional athlete, or a star actor. Maybe you would be 'corrupted' by your environment and change your priorities. In this respect, becoming a doctor seems like a much safer choice.

There is an organisation called *80,000 Hours* which provides advice on career choice for altruistically motivated people. They have a website. Check it out.