Introduction to Moral Reasoning

This is book will provide a basic introduction to the logic behind moral reasoning and arguments about morality.

Site:	HCC Eagle Online
Course:	6142-PHIL-2389-Cooperative In Philosophy-RT-77609
Book:	Introduction to Moral Reasoning
Printed by:	Nathan Smith
Date:	Sunday, 1 June 2014, 10:54 PM

Table of contents

1 Moral Reasoning

The basic toolkit of the philosopher is logic. Logic defines the rules that govern good inferences. We make an inference any time we draw a conclusion about something based on information presented. Since philosophy discusses ideas and concepts that cannot be verified through experiment or observation and concerns topics that may not even have a defined method of resolution, we have to rely on our ways of reasoning in order to evaluate different philosophical positions.

Ethics is a special case of this more general philosophical problem. Consider a problem in ethics: 'Should I give a portion of my income to charity?' How can we resolve a question like that? It is obvious that there are no facts about the world that can decide whether or not we are obligated to give a portion of our income to charity. No matter what we learn about the good that charitable organizations do or the needs of people served by those organizations, we would still need a reason *why we are obligated* to give our money to help those people (why not some other institution, for instance?). So, we cannot decide moral questions just by appealing to facts. We also need to appeal to values or norms and we need to determine how far those values extend. Are they personal or cultural? Are they values that apply to all human beings? Are they obligations or do we just allow people to decide for themselves?

1.1 Normativity

Part of what makes moral reasoning unique is that it concerns questions of normativity, i.e., what we value (as good or bad, right or wrong). But normativity is not about what we, in fact, do value. Rather normative claims concern *what we ought to value*, or *what we ought to do*. So, moral arguments are not about matters of fact; they are about matters of prerogative. In ethics, we are not so much interested in the way the world is, but the way it ought to be.

As a result, moral arguments must include reference to a set of values or goods that are normative. It is not sufficient to make an argument that only refers to facts in order to draw a conclusion about what we ought to do.

Consider the following argument:

- In the United States, we consume 80 pounds of chicken per year.
- So, the average person in the US is responsible for the deaths of 20-25 chickens per year.
- Chickens raised for food live their lives in confined, dirty, and difficult conditions (often having their beaks clipped and being unable to associate with their offspring and perform other natural activities).
- So, on average, a person in the US contributes to a difficult life for 20-25 chickens per year.
- Therefore, you should stop eating chickens.

While all of these facts are perfectly true, they may not lead to the conclusion that we should stop eating chickens. After all, a person might agree with all of the premises and still reject the conclusion because, for instance, she might not think that it is a bad thing for chickens to suffer or die in order for humans to eat them. What the argument needs is an additional set of premises that show that it is wrong to treat chickens this way for the consumption of chicken meat.

1.2 Moral Foundations and Moral Theories

The issue of normativity naturally raises questions about the foundations of our ethical beliefs. While it might seem clear what foundations there are for our beliefs about the way the world is (e.g., that we can observe that the world is just this way), it is less clear what are the foundations of our ethical beliefs. On what basis does a person decide to give to charity or to stop eating chicken? There may be all sorts of reasons that a person might give for either of these actions, but we have the suspicion that, at some point, these questions have to come down to a matter of moral foundations: where do our ideas of right and wrong, good and bad, come from?

These are important questions that we will address elsewhere in the course. For the moment, what we need to appreciate is that the question of foundations is a difficult and controversial one. There are many different answers to the question of what provides the ultimate foundation for our moral beliefs. And these answers give rise to different moral theories. To give you an idea of what sorts of answers lead to what sorts of theories, I will outline a few here.

- The belief that morality is determined by God's commands leads to divine command theories of morality.
- The belief that cultural or personal preferences determine moral beliefs leads to relativism.
- The belief that the types of outcomes (whether they are beneficial our not) should determine moral beliefs leads to consequentialism.
- Kantianism and deontology derive from the notion that morality is determined by the nature of reason itself (and by our reflections on our own rational considerations when acting ethically).
- The belief that morality is entirely dependent on circumstances and situations and that it is impossible to derive general rules from these situations leads to the theory of particularism.

There are other theories and there are ways to differentiate them, but the important thing to see is that the question of the foundation of ethical beliefs is just as controversial as the question of which moral theory is true. As a result, even though it is natural to move from the issue of normativity to the issue of the foundations of normativity, we should resist making a direct appeal for ultimate foundations. Instead, we should seek methods of reasoning about morality

that allow us to evaluate moral arguments without resolving the bigger, metaethical debates.

1.3 Determining What is Morally Relevant

In any given moral argument, it is critical that we isolate the relevant moral features. In short, these are the features that focus on normativity. So, we want to ignore matters of fact unless they are directly tied to questions of value. Nevertheless, there do seem to be facts about a given situation that are directly morally relevant. Consider the case of a corrupt banker who stands trial for defrauding investors. Suppose that one of the witnesses in the trial is an elderly woman who was a family friend and lost her life's savings investing with this person. While it is a matter of fact that this elderly woman was led to invest with this man because he was a close family friend and that she lost her life's savings because of his behavior, we might think that this is a particularly morally relevant fact. After all, what kind of person takes advantage of an elderly family friend?

Some matters of fact have the potential to make us morally outraged or morally congratulatory, sympathetic or reverent. It seems to be part of human nature to interpret events in moral terms. So, how should we understand these facts when considering reasoning about morality? We might begin by asking ourselves what it is about the particular case that makes us outraged or congratulatory, sympathetic or reverent. For instance, what is it about the case of the corrupt banker described above that makes this fact seem morally relevant? We might identify a number of features, but here some general features that we might consider morally relevant:

- intention
- the extent to which the action violates accepted norms of behavior
- direct or indirect harm toward persons or animals
- deceitfulness or misrepresentation
- coercion (causing a person to act against their will)
- gratuitously bad consequences (harmful effects that could otherwise be avoided)

Not all of these features are uncontroversial. For instance, many philosophers believe that intention is not morally relevant or that deceitfulness is not always (or even generally) a bad thing. However, we can appeal to what are generally accepted standards of behavior and note what features of a given situation deviate from that behavior. This can help us to begin to get a grasp on the morally relevant considerations.

Another way of isolating the morally relevant features of a given situation is to make an

analogy between a controversial case and a case that is not controversial. Consider the case of online piracy (unauthorized downloads and sharing of copyrighted material). Many people argue that this is morally equivalent to stealing. Stealing is uncontroversially immoral. So, if piracy is analogous to stealing, then piracy is also immoral. Of course, the argument here hinges on whether and to what extent this analogy holds. But at the very least, we have identified the morally relevant features of the situation. Whether or not piracy is immoral will probably hinge on whether and to what extent it resembles stealing.

A further strategy is to ask about the set of beliefs that might lead to the judgment that a particular action is acceptable or unacceptable. What kind of beliefs would a person have to have in order to engage in online piracy or convincing an elderly family friend to risk her life's savings on questionable investments? This method asks us to consider a more general theory that might support a particular case. Sometimes we can get clear on a theory by examining the cases that it might be applied to; and sometimes we need to appeal to a theory in order to get clear on the particular cases.

1.4 Moral Principles

One of the ways that arguments about morality can find some basis in normativity is to appeal to principles of morality. Moral principles are prescriptive statements about what people ought to do in general. It is important to see that moral principles do not have to be absolute in order to stand as moral principles. For instance, we can take it as a general principle that 'you should help out those in need'. But we understand that there might be situations in which helping out someone in need does them further harm or violates some other duty that we have. So, it should be possible to affirm a general moral principle without being forced to defend it in all circumstances.

Examples of moral principles might include:

- do not intentionally harm another person
- do not gratuitously cause pain
- do not lie, mislead, or intentionally misinform
- help those who are in need
- develop your talents
- stand up for ideas you believe in
- thank those who have helped you
- give back what is owed

There are many others, of course. But the important thing is to see that moral principles can be used in moral arguments in order to appeal to the normative dimension of the argument. As an exercise, you should return the cases examined earlier of giving to charity or not eating chickens and try to determine which moral principles might be brought to the argument in order to lead us to the conclusion that we should or shouldn't perform some given behavior.

Moral principles can be based on a moral foundation (as we saw earlier). But we can also see moral principles as deriving from a common moral concept such as virtue. In other words, we can think of the virtuous person as one who embodies and acts on moral principles. The virtues, then, become a set of characteristics that lead to behaviors that conform to moral principles. As an exercise, consider the virtues of courage, generosity, helpfulness, or fairness and ask yourself which moral principle is embodied in each of these virtues.

2 Moral Arguments

Now that we have determined some of the background behind moral reasoning, we need to turn to the nuts and bolts of moral reasoning, i.e., the logic of constructing moral arguments.

Remember that moral arguments are just a specific kind of argument. In logic, we define an argument as a set of statements that provide support or justification for a belief. The set of supporting statements are called the premises and the statement that they support or justify is called the conclusion. We will see how the logic of arguments can be applied to ethics.

2.1 The Logic of Arguments

All arguments are composed of **sentences**. Sentences are statements that include a **subject** and a **predicate**: the predicate describes the subject. Sentences can be empirical observations, statements of belief, or statements of principle. In order for a sentence to be used in a philosophical argument it must have such a form that it could possibly be either true or false. We do not have to know whether it is true or false, but it has to be the kind of thing that could be true or false.

There are two kinds of sentences in philosophical arguments: **premises** and **conclusions**. Premises and conclusions can appear in any order, but when you are writing out arguments it is usually easiest to write the premises first and then the conclusion last. Conclusions can become the premises of further arguments. So, more complex arguments will contain multiple minor conclusions and a major conclusion. The conclusion is the claim that the argument is intended to support (this is similar to a thesis in a argumentative paper). The premises provide the support or evidence for the conclusion.

If the premises provide adequate support for the conclusion, then the argument is a **valid argument**. However, it is important to recognize that even a valid argument can lead to a false conclusion, since the premises might be false.

Sentences are either true or false. Arguments are valid or invalid.

2.2 Validity and Soundness

If an argument is valid and its premises are true, then it is unreasonable to reject the conclusion. This is the sense in which logic mirrors reason. We have an obligation either to accept the conclusion of an argument or to demonstrate either its invalidity or the falsehood of one or more of its premises.

This leads us to a two-step method for assessing arguments:

1. We must assess the validity of the argument.
2. We must assess the truth of the premises.

If the argument is valid and its premises are true, we say that the argument is **sound**. A valid argument can have a false conclusion, but a sound argument cannot.

2.3 Assessing Validity

Validity focuses only on the "form" of the argument, not the content. We are only interested in how the premises support the conclusion structurally. So, when you are assessing validity, you should ignore whether or not the premises are true. An easy way to do this is simply to assume that the premises are true and then ask yourself if the conclusion follows.

Sometimes we say that validity is "truth-preserving." In other words, if you start out with true premises, then a valid argument will preserve truth, so that you will wind up with a true conclusion. But if you start out with false premises, a valid argument might lead you to a false conclusion.

Consider the following two arguments:

Murder is wrong. Intentionally modified someone's automobile in order to cause a fatal crash is murder. ------------- So, intentionally modifying someone's automobile in order to cause a fatal crash is wrong.	Killing another human being is immoral. Soldiers in war kill other human beings. ------------ So, soldiers in war perform immoral actions.

Both of these arguments are valid. If you assumed the truth of the premises, you would have to accept the conclusion. In other words, if the premises were true, the conclusion would have to be true.

Of course, the second argument about soldiers in war is highly controversial. But you can see why: the premise that says 'killing another human being is immoral' leads to the conclusion that killing in war is immoral. So, this might make us question whether or not that premise is true. In fact, we might

suspect that it is not always wrong to kill another human being even though we accept that it is always wrong to commit murder. This might lead to a further question, 'what is the difference between murder and killing a human being?'.

2.4 Counterexamples

If an argument is invalid, then it is possible to generate what we call a **counterexample**. So, in order to show that an argument is invalid, you must provide an example that demonstrates the invalidity.

In order to think about counterexamples, sometimes it is useful to think of arguments in a slightly different way than we have up to this point. You can think about an argument in terms of an 'if, then' statement. The premises would be on the 'if' side, while the conclusion is on the 'then' side of the statement. From the example above, we can rewrite the argument as a statement:

> If killing human beings is immoral and soldiers in war kill human beings, then soldiers in war behave immorally!

A counterexample would be an example that shows the statement to be false. We can tell that this is a valid argument because there is no example that would make this statement false. That is, any instance of non-immoral behavior by soldiers in way cannot involve the killing of other human beings, since the statement says that the soldiers' behavior is only wrong when it involves killing other human beings. An if-then statement is made false only if the if-clause is true AND the then-clause is false. An if-then statement only says that *if* the if-clause is true, *then* the then-clause must be true.

Now, this is not necessarily a reason to accept the conclusion that soldiers behave immorally in war. If we reject the conclusion, however, we must be willing to revise the premises. So, if we think that when soldiers kill enemy combatants in war they are not doing anything immoral, then we cannot think that killing human beings is always wrong.

2.5 Moral Operators

An additional component to morality that differentiates it from straightforward, sentential logic is that moral reasoning involves claims that rely on **modal operators**. A **logical operator** is like an operator in mathematics (+, -, etc.) where the operator changes the meaning of the sentence. A traditional modal operator is like: 'it is necessary that...' or 'it is possible that...'. These statements clearly change the meaning of the sentence. Consider the sentence, 'it is raining outside'. Now, consider the sentence 'it is necessary that it is raining outside'. Clearly, even if it is raining outside, it is not necessary that it is raining outside, since sometimes it is sunny and no law of physics requires it to rain. So, modal operators change the meaning of sentences and thus change whether or not they are true or false. As a result, when we evaluate moral statements, we need to understand how moral operators can change the meaning of sentences.

I will introduce five moral operators:

- it is obligatory that ...
- it is permissible that ...
- it is impermissible that ...
- it is omissible that ...
- it is optional that ...

In order to define these operators and understand the effects they have on moral claims, I will describe the relationship between them.

- It is permissible that X if and only if it is *not* obligatory that *not* X.
- It is impermissible that X if and only if it is obligatory that *not* X.
- It is omissible that X if and only if it is *not* obligatory that X.
- It is optional that X if and only if it is *not* obligatory that X and it is *not* obligatory that *not* X.

So, something is permissible just in case the opposite is not obligatory. Something is impermissible just in case the opposite is obligatory. Something is omissible just in case it is not obligatory. And something is optional just in case neither it nor its opposite is obligatory.

Take a minute to try to digest that.

Furthermore, each of these statements has an influence on the others.

- If something is obligatory, then it is permissible.
- If something is obligatory, then it cannot be impermissible.
- If something is impermissible, then it cannot be obligatory and it must be omissible.
- If something is not permissible, then it must be omissible.
- And if something is not omissible, then it must be permissible.

Clearly, when we make moral statements about what is permissible, obligatory, optional, and so on, these statements have an influence on what other kinds of statements we must accept.

3 Moral Conflicts

The logic of moral reasoning leads to a very important result: we must be consistent in our moral views. If we hold to a certain moral principle, e.g., that killing human beings is always wrong, then we must be willing to accept the conclusions that logically follow from it. On the other hand, if we reject the conclusion, then we must be willing to revise the principle.

Many issues in ethics can be resolved by clarifying the morally relevant features, appealing to moral principles, and composing logically valid moral arguments. However, even when we do all of that, we may still find ourselves in a state of moral conflict. For instance, consider the following argument:

> P1 - It is immoral to inflict gratuitous harm on other animals.
> P2 - Eating meat inflicts harm on other animals.
> P3 - We, as human beings, do not have to eat meat.
> C1 - So, eating meat inflict gratuitous harm on other animals.
> C2 - So, eating meat is immoral.

Suppose that you want to deny the conclusion. What avenues are available to you? Well, P3 is certainly true, since there have been millions of vegetarians in the history of human civilization who have lived long, healthy lives. P2 is also undeniable, given that killing an animal before it dies of natural causes harms that animal. But it is even more true in the case of contemporary industrial agriculture, which inflicts a tremendous amount of pain and suffering on animals. The only alternative appears to be to reject P1. But, I suspect that people who eat meat are perfectly willing to reject P1, while people who are vegetarians hold that P1 is true.

The meat-eater might ask for further support for P1. But the vegetarian mights say that this principle is just as obvious as it is that you shouldn't lie, cheat, steal, or harm other human beings. So, here we have a clear case where a moral conflict remains even after we have clarified what is morally relevant, identified moral principles, and constructed a valid moral argument.

3.1 Moral Dilemmas

Moral dilemmas are a good way to test our moral beliefs and question our moral theories. Moral dilemmas place us in a circumstance where there is no clear best choice between alternatives. So, they may reveal to us the need for a clarifying moral principle or the need to revise a current moral belief.

Suppose that you are offered a job that requires you to do things that violate some of your moral beliefs. Suppose that this job offers you a lot of money and the chance to be well-respected by peers. Suppose, in fact, that in all respects (except for the issue of morality) it is preferable to other alternatives. But if you took the job, it would require that you compromise other moral beliefs you hold. Would you take that job? What considerations would inform your decision?

The case of moral dilemmas presents a slightly different scenario than overt moral conflicts, like the case of vegetarianism. In the case of moral conflicts, the problem arises because two (or more) people hold different sets of moral beliefs. Moral conflicts can be explained by the fact that different people hold different moral principles as true or that different people rank moral principles as having different importance in their decisions. The case of moral dilemmas is different, since the conflict is internal. Moral dilemmas reveal a tension with one person's set of moral beliefs. The problem of moral dilemmas cannot be explained by appealing to the differences between peoples' opinions on the issue of morality.

Nonetheless, both moral dilemmas and moral conflicts seem to be real things. How should we understand them?

3.2 Principles of Resolution

In order to resolve to resolve moral conflicts and moral dilemmas, we might hope to find some ultimate moral principle that can act as a final judge on these matters. The idea here is that given two competing values or moral principles, the only way to decide between them is to have a third value or moral principle that is more basic (has a stronger foundation) and more general (applies to all of the cases that each of the two other principles apply to). If this kind of reasoning is carried forward, we can easily see how it would lead to the conclusion that there is just one, final moral principle (a Supreme Court or final judge) that determines all others. This kind of reasoning supports the view of Bentham and Mill, that there is only one principle of morality (for them, it was to maximize the total amount of happiness).

However, we need not assume that all of the work of deciding between values and principles must be carried out by other values and principles. Many philosophers have concluded that it is just the job of human beings to make these sorts of decisions. The idea is that human beings should cultivate a kind of practical wisdom about actions and norms such that they develop a skill for deciding correctly in the case of moral conflicts and moral dilemmas. This is the view of Aristotle and (to an extent) Plato and the virtue tradition. And we can see that there is merit for this kind of thinking by observing that there are individuals in society who we admire as having been capable of making good moral judgments. It stands to reason that these individuals are exercising a capacity or skill that many people might be able to develop in similar (if not as dramatic or outstanding) ways. It is important, however, to see that Aristotle does not think there is one supreme principle in ethics. Instead, there are better and worse moral reasoners who are more or less capable of deciding between conflicting duties and values.

This sort of common-sense view on moral conflicts and moral dilemmas gives way to a more sophisticated version. Many philosophers, probably beginning with Immanuel Kant, have sought to take the ideal of practical reason and make it more formal. These philosophers consider the practical skill of good moral judgment as a formal procedure that can be considered in the abstract. One very illuminating instance of this abstraction of moral judgment can be seen in John Rawls with what is called the "veil of ignorance." In a book about the notion of justice, Rawls outlines a scenario that would allow us to determine whether a particular political arrangement (laws, constitution, etc.) is just. His reasoning is based on the idea that an objective rational observer would be able to determine whether or not a society is just. In a sense, he is appealing to an abstract, practically wise individual. The veil of ignorance

is the situation that would allow the individual to determine whether or not the society is just. In particular, this means that when you are considering whether a society is just, you should consider whether or not you would agree to live in that society from behind a veil of ignorance. The veil of ignorance is supposed to hide from your consideration any particular facts about you: your parents, age, ethnicity, region of the country, education, height, sex, religion, etc. Since your decision should be based only on the features of the society that apply equally to everyone, you should ignore all of those qualities that make you different. Rawls thought that from this point of view, anybody could be an objective rational observer and thus could determine whether or not the society is just.

We might imagine a similar principle in ethics. If we are confronted with an ethical conflict or an ethical dilemma, we might ask ourselves what scenario would enable us to view the situation as an objective, rational observer. Then we might try to reason through the issue from that perspective. As an exercise, try this out with the example of vegetarianism or killing in war. What perspective would enable you to be rational? What would the principles of good moral judgment alone lead you to conclude about these issues?

4 Review