

Week 1 Notes

Philo 101 Online | Hunter College

Daniel W. Harris

1 Welcome to Philosophy 101 Online

Welcome to Philosophy 101 Online at Hunter College. Here are some things you should do this week:

- Familiarize yourself with the course website.
- Carefully read through the syllabus, which can be found via a link at the top of the website's home page.
- Sign up for our Slack Workspace, which you can do via a link on the website.
- Read this week's required reading, 'Consider the Lobster' by David Foster Wallace.
- Read through the rest of these notes.
- Answer this week's discussion questions, and submit your answers as per the instructions on the website's page for this week.
- Complete this week's writing assignment, which is explained on the website's page for this week.

And if you aren't sure how to do any of this, or if you have any questions about the course, you can email me at danielwharris@gmail.com.

2 This Week's Reading: 'Consider the Lobster' by David Foster Wallace

Our first reading of the semester is 'Consider the Lobster' by David Foster Wallace. We're starting with this essay because it works as a nice introduction to several of the central areas of philosophy, including ethics, epistemology, metaphysics, and the philosophy of mind, and because it demonstrates some of the connections between these areas.

You may have read something by David Foster Wallace before. He was a famous novelist and essayist who died in 2008. His best-known novel is called *Infinite Jest*. Before writing 'Consider the Lobster', he'd also become well known for the essays he'd written in magazines, including a couple of very funny travel pieces for *Harpers* about going on a Caribbean cruise ('A Supposedly Fun Thing I'll Never Do Again') and about going to the Illinois State Fair ('Getting Away from Already Being Pretty Much Away from It All'). So it seems that *Gourmet* commissioned him to write a piece about the Maine Lobster Festival hoping for another amusing, beautifully written travel piece of this kind.

It is tempting to find some irony in the fact that *Gourmet* (which stopped publishing in 2009) billed itself as "the magazine of good living". What *Gourmet* meant by this, it seems, is good living in the sense of the French term *bon vivant*, which, according to Merriam-Webster, means "a person having cultivated, refined, and sociable tastes especially with respect to food and drink". *Gourmet* was, essentially, a magazine that published photos, writing, and recipes that were designed to make readers' mouths water. But of course, the phrase 'good living' has another, ethical sense. In this sense, to live well is to be a good person, to do the right thing. And it is this latter sense of good living that David Foster Wallace wound up focusing on in *Consider the Lobster*. I enjoy imagining the reaction of the *Gourmet* editors when they got Wallace's piece, which is a somewhat unappetizing essay about a food that many people think of as a delicious luxury. Lucky for us, they printed it. And the essay is now recognized as a classic.

2.1 Ethics

After a bit of stage-setting, the essay takes its first philosophical turn with the following paragraph.

So then here is a question that's all but unavoidable at the World's Largest

Lobster Cooker, and may arise in kitchens across the U.S.: Is it all right to boil a sentient creature alive just for our gustatory pleasure? A related set of concerns: Is the previous question irksomely PC or sentimental? What does “all right” even mean in this context? Is it all just a matter of individual choice?

There are several philosophical questions here. First, the question of whether it’s “all right” to eat lobster is an instance of one of the central questions of ethics:

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, including the act of killing and eating lobster. In order to answer the question about lobster in a reasonable way, then, we need some principles to guide us. These principles should be consistent with one another and general in their application, which is to say that they should be organized into a theory.

Wallace’s question also raises the question of whether ethics is “just a matter of individual choice”, and the question of what “all right” even means in this case. These are what philosophers call questions of *metaethics*, because they’re questions about the nature of ethics itself. The possibility that ethical questions boil down to personal choice suggests either a version of *subjectivism* (the idea that what’s right and wrong is up to the beliefs of the person performing the action) or moral *anti-realism* (the idea that there’s really no objective matter of fact about what’s right and wrong). On the other hand, the most common sensical view, and one that is shared by many ethicists, is moral *realism*, which is the idea that there are objective moral truths waiting to be discovered.

Finally, Wallace raises the question of which beings matter for ethical purposes. To whom do we have moral obligations? Is it only other humans to which we are morally obligated? Isn’t it wrong to needlessly torture dogs? What about livestock, like pigs and cows? What about human fetuses? What if we encountered intelligent extra-terrestrials? It’s not enough, for philosophical purposes, merely to choose some cutoff between the creatures who matter and those who don’t. We also need to give reasons for including some beings within the sphere of those who matter,

and excluding others. Again, this is a question that an ethical theory should answer for us.

Although Wallace doesn't formulate a full ethical theory, the kind of reasoning that he engages in seems most closely related to a family of theories that philosophers call *consequentialism*. An ethical theory is consequentialist if it judges the moral worth of an action based solely on the action's results, or consequences. Any consequentialist ethical theory therefore needs a principle that tells us, in general, what kind of consequences are good. Wallace's version of this principle is that we should try, as much as possible, to avoid causing suffering (at least when the creature in question has a preference not to suffer). This seems like a highly intuitive principle: at least one thing that moves us to condemn actions as morally wrong is that they have caused unnecessary suffering. Moreover, it's not just humans to whom we intuitively have an obligation not to cause suffering. We become morally outraged when someone needlessly makes a dog suffer as well, for example. So, Wallace adopts a general version of this principle: it's wrong to cause more suffering than is necessary.

This version of consequentialism is usually called *utilitarianism*, which is the view that we should act so as to maximize happiness and minimize suffering. (There are many ways to make this idea more precise, some of which we'll consider later in the course.) This version of utilitarianism also comes with a ready-made answer to the question of which beings matter. Namely: in order to tell whether a being matters, we should find out whether it has a capacity for happiness and/or suffering. If it does, we need to take its potential suffering into account when deciding what we should do. This is what the influential utilitarian philosopher Jeremy Bentham meant when he said,

...the question is not, Can they *reason*? nor, Can they *talk*? but, Can they *suffer*? (Bentham, 1789, 122)

Much of Wallace's essay is devoted to answering Bentham's question, as applied to lobsters. If lobsters do have the capacity to suffer—and, in particular, if the suffering that we cause them when boiling them alive outweighs the happiness we get from eating them—then utilitarianism tells us that we shouldn't boil them alive.

2.2 Epistemology

An important question, then, is whether lobsters can suffer. More broadly: for any other being, how can we tell whether it has a capacity for suffering? This is an

epistemological question. Epistemology is the branch of philosophy dealing with questions about knowledge. It asks what knowledge is, and how we can get it. For example: how much of our knowledge comes via our senses, how much of it is innate, and how much of it can we attain by pure reasoning alone? This is something we'll consider later in the course.

Epistemology also deals with questions about particular kinds of knowledge. The question of whether lobsters can suffer is an instance of a more general question that epistemologists have long wrestled with:

THE PROBLEM OF OTHER MINDS

How can we tell whether some other being has a mind, or whether they are in some particular mental state, such as pain?

This is a tough epistemological question because mental states are apparently not the sorts of things we can observe directly. We seemingly have a kind of direct, introspective access to at least some of what happens in our own minds, but we can't directly experience the thoughts and feelings of others. This is a problem not just when it comes to animals, but also for other humans. The French philosopher René Descartes illustrates this point in his *Meditations* (which we'll read later in the semester):

If I look out of the window and see men crossing the square, as I have just done, I say that I see the men themselves...; yet do I see any more than hats and coats that could conceal robots? I judge that they are men. Something that I thought I saw with my eyes, therefore, was really grasped solely by my mind's faculty of judgment. (Descartes, 2017, 7)

We can generalize this point: all you ever see of other people is their body moving around. But we're not far away from a time when some Silicon Valley startup will be able to build convincing replicas of humans that behave roughly like us, but which function very differently on the inside. Indeed, maybe someone's already done this, and we just don't realize it. So how can we be sure, when encountering a new person, that they're not an unthinking automaton? And, for that matter: why should the fact that another *real* person is made of the same physical stuff as you convince you that they have inner lives like yours, given that you can't observe anybody's thoughts but your own?

We've now arrived at a kind of philosophical skepticism—specifically, *other-minds skepticism*, the view that other people may not actually be conscious, thinking

beings, but merely bodies that behave like our bodies. If you take it seriously, this kind of skepticism (like most other kinds) is a disturbing thought. Maybe you're the only one in the world with a mind? What's the best reason to think that this worry is unfounded?

An answer that many philosophers have given is called the ARGUMENT FROM ANALOGY, which goes like this. You know that *you* have a mind, and that your thoughts and feelings are systematically related to your own behavior. The reason you pull your hand away when it gets close to flame is that your *pain* causes you to do so, for example, and the reason you leave the house in the morning is that you *want* to go to class and *believe* that you need to leave in order to get there. Pains, wants, and beliefs are states of mind. But other people behave in roughly the same way as you, and we know from studying human anatomy and neuroscience that their bodies are hooked up in roughly the same way as well. So, this argument concludes, the best explanation of the fact that they behave in roughly the same way is that they are caused to do so by the same kinds of thoughts and feelings. Therefore, they must have minds too.

This is the same strategy that Wallace uses to argue that lobsters can suffer. Although lobsters' bodies and behavior are obviously quite different from our own, Wallace argues that they are similar enough in the relevant ways. For example: when you put a lobster in hot water, it becomes agitated, struggles, and tries to escape. (My favorite way of thinking about this: if you were dressed up in a big lobster costume and someone put you in a big pot of very hot water, you would do roughly the same thing that a lobster does.) Moreover, although lobsters' nervous systems are very different than ours, Wallace points out that they have some of the same components that we use to perceive pain. Of course, the analogy isn't perfect, and lobsters lack some of the neurological hardware that is involved in human suffering, so it's hard to be totally certain that lobsters feel the same kind of suffering that we do. But, Wallace argues, this uncertainty means that when we boil lobster alive, we're running a risk that we're causing a great deal of suffering.

2.3 Metaphysics and the Philosophy of Mind

The preceding style of argument seems to rest on some assumptions about the nature of pain—namely, that pain is something that happens in bodies rather than in minds or souls. One response that might tempt you is that, since lobsters don't have minds or souls, the whole question of whether they can suffer is just silly.

This line of thought takes us from epistemology into *metaphysics*. Whereas epistemology revolves around questions about how we can know about different aspects of reality, metaphysics revolves around questions about the nature of things in themselves. So, whereas an epistemologist might ask how we can know that another being is in pain, a metaphysician—and, more specifically, a philosopher of mind—would ask what the nature of pain is. More broadly: what is it to have a mind?

One traditional answer is that having a mind is having something in addition to one's body—that the mind and the body are two distinct ingredients that somehow combine to make up a person. This view is called *dualism*, because it is based on the idea that the mind and the body are two different things. Dualists sometimes also think of minds as “souls”, and dualism is seemingly presupposed by many religious traditions.

As we'll see in later readings, there are several good reasons to doubt dualism. And so, many philosophers now instead believe in some version of *physicalism*—the idea that everything, including the mind, is ultimately a part of the physical world. On this view, your thoughts and feelings are just processes happening in your body (in particular, perhaps, in your nervous system).

One popular theory of the mind that is compatible with physicalism is called ‘functionalism’. This is the idea that being in a mental state of a certain kind is a matter of having something going on in your body that plays the characteristic role of that mental state. A state of the body plays the role of pain, for example, if (roughly) it is normally caused by events that can damage the body, it brings about negative feelings, regrets, and fear of the cause, and it causes the person experiencing it to avoid the thing that caused it, to say ‘ouch’, to wince, and so on. (A full description of the functional role of pain would be much more elaborate.) The idea of functionalism is that although this role might be played by different kinds of bodily states in different kinds of organisms, all you need to be capable of feeling pain is to have *some* bodily state that plays this role. The benefit of functionalism is that it allows us to make sense of the idea that, for example, there could be a species of intelligent aliens who feel pain, but whose nervous systems are organized in very different ways from our own.

Wallace suggests that that's the situation with lobsters as well. They don't have centralized brains or cerebral cortices, as we do. But we can think of Wallace as arguing that lobsters do have bodily states that play the role of pain. In this sense, Wallace's argument seems to presuppose a kind of functionalism about the nature of the mind. Again, this is a topic we'll explore in greater detail later this semester.

3 Reading and Writing Philosophy

The two most important skills needed to do well in this course are philosophical reading and philosophical writing. I emphasize that I'm talking about *philosophical* reading and writing because philosophers read and write in a particular way. Part of the aim of this course is to get you to be better at that style of reading and writing—skills that are useful far beyond philosophy.

For these reasons, I have also included among the readings for this week two essays by the philosopher Jim Pryor about how to read and write philosophy. These are relatively informal pieces, designed to give newcomers to philosophy reading and writing advice.

The single most important thing to understand about philosophical reading and writing is that both activities center around *arguments*. By 'argument', I don't mean a verbal fight. Rather, I mean an attempt to persuade someone of a conclusion by structured reasoning from premises. Let's take that a bit more slowly.

Every argument is an argument for some particular *conclusion*. The conclusion is what the author of the argument is trying to persuade their audience of. The conclusion is the end of an argument. Usually, the author of an argument assumes that their audience doesn't already believe the conclusion of an argument before they read it. Otherwise, why would they bother making the argument? The point of an argument is to persuade its audience of something new.

If the end of an argument is its conclusion, then the starting points of an argument are its *premises*. These are the claims that the author begins with, and from which they plan to reason to the conclusion. In order to be rationally persuasive, an argument must have premises that the audience accepts, or is willing to accept once they hear it. After all: the premises of an argument are supposed to support the conclusion, and premises that the audience doesn't believe can't perform this function.

Reasoning is the process of moving from the premises to the conclusions of an argument. Even if we begin from true, acceptable premises, it is possible to reason badly and wind up with a false conclusion. In general, good reasoning is reasoning that, if beginning from acceptable premises, gives us good reason to accept the conclusion.

There are therefore two basic ways that an argument can go wrong: it can have false premises, or it can involve bad reasoning from its premises to its conclusion.

The main goal of philosophical writing is to construct good arguments for sur-

prising conclusions. An argument of this kind must have the following qualities:

- (1) It must have premises that the audience will accept as true. (Otherwise they can't support the conclusion.)
- (2) Its conclusion should be something that isn't already boringly obvious to the audience. (Otherwise what's the point?)
- (3) It must involve good reasoning to get from the premises to the conclusion. (Otherwise it won't be convincing.)
- (4) It must be written down in a way that makes it easy for the audience to understand. (Otherwise they won't get it.)

This may seem straightforward, but it is a very difficult skill to master. Jim Pryor gives some good advice about how to do it well in his *Guidelines on Writing a Philosophy Paper*. We will begin to practice philosophical writing in some of the later writing assignments in this course.

The purpose of this week's assignment is to get you to practice one of the basic philosophical *reading* skills. Getting good at reading philosophy requires two related skills: *reconstructing* and *critiquing* arguments. This week's assignment is all about reconstructing arguments. We'll practice critiquing arguments in a later week.

To reconstruct an argument is to understand and articulate its structure. This means figuring out what its conclusion is, figuring out what its premises are, and figuring out what sort of reasoning is supposed to connect them.

Reconstructing an argument may seem like a simple task, but it is often quite difficult. The main reason for this is that the arguments' authors don't always make the task easy for us. One problem is that there are no rules about the order in which we write the premises and conclusion of an argument, for example. The conclusion won't always be the first sentence or the last sentence. A second problem is that arguments are often accompanied by claims that are irrelevant to the argument itself, and that has to be ignored in order to understand the argument.

The third, and perhaps toughest obstacle to reconstructing arguments is that they may leave some of their premises or their conclusion unspoken. Instead, they might just assume that we can figure out what conclusion they're arguing for, or what their premises are supposed to be. Sometimes it is obvious what an unspoken premises is, but usually it's not. And there are some potentially sneaky reasons for

leaving premises and conclusions unspoken. By leaving a premise unspoken, the author of an argument can sometimes keep us from focusing on it, and thereby stop us from noticing that it's false, for example. And by failing to articulate the conclusion, sometimes an author is trying to get different people to supply different conclusions that they find more plausible. This is why reconstructing an argument is such an important philosophical skill: it forces us to pay close attention to the details of an attempt to persuade us, and to fully understand it. It's not possible to properly critique an argument before doing this.

So, your job for this week is to choose a paragraph from the reading and reconstruct the argument that you find in it by answering the questions on the website for this week. Here are a few tips:

- Choose your paragraph wisely. Not every paragraph will contain an argument, and some paragraphs' arguments will be easier to reconstruct than others.
- Once you've chosen a paragraph, I would recommend beginning by figuring out what that paragraph's argument's conclusion is. This will make the premises (including any unspoken premises) easier to spot.
- Think *hard* about whether there are unspoken premises. Ask yourself what assumptions, however obvious, we need to make in order for the author's premises to support their conclusion. It is very easy to miss something.
- Don't submit the first version of what you write down. Ask yourself repeatedly if there is any way to make it easier to understand, and edit accordingly. This is something that all serious writers (including college students who get good grades) do.

References

- Bentham, J. (1789). *An Introduction to the Principles of Morals and Legislation*. Clarendon Press, Oxford.
- Descartes, R. (2017). *Meditations on First Philosophy*. Early Modern Texts. Translation by Jonathan Bennett, <http://www.earlymoderntexts.com/assets/pdfs/descartes1641.pdf>, 2017 edition.