

Introduction to Philosophical Thinking

So you have decided to study philosophy. You may have a fairly clear idea of what studying philosophy involves, or you may have only a vague idea, or even no idea at all.¹ Since there are some interestingly different conceptions of philosophy (philosophers philosophize even about philosophy!) and since we need a single conception of philosophy to guide our work, we will begin our study of philosophy by first developing a fairly specific conception of philosophy. While we will rely on this conception of philosophy throughout the book, it is your responsibility (as a budding philosopher) to think carefully about it and have an opinion about its merits by the time you have finished the course. By then you may think that there should be changes or qualifications, small or even large, in the conception of philosophy that we have been using.

THE CONTENT OF PHILOSOPHY

Let's start building our conception of philosophy by distinguishing between the distinctive content involved in the discipline of philosophy and the distinctive method of philosophical thinking. The content concerns (obviously) what philosophers think about. For example, philosophers typically think about questions like these: What is knowledge? What is truth? What are minds? What is consciousness? Are we genuinely free? Does being morally responsible require being free? Are we, by our very natures, selfish? Is there a genuine difference between right and wrong or good and evil? What is justice? Does God exist? And even, as we have already seen, what is philosophy? In attempting to answer these questions, philosophers think about **claims** (specific, focused assertions that are put forth as being true or false), and also about more comprehensive views or positions (composed of many related claims) that purport to answer questions like the ones listed above.

To generalize from these examples, it would be reasonable to say that the content of philosophy concerns:

1. the fundamental nature of reality—the nature of space and time, of properties and universals, and especially but obviously not exclusively the part of reality that consists of **persons** (the branch of philosophy called **metaphysics**);
2. the fundamental nature of the cognitive relations between persons and other parts of reality—the relations of thinking about, knowing, and so on, (the branch of philosophy called **epistemology**); and

3. the fundamental nature of values, especially values pertaining to ethical or social relations between persons and between persons and other parts of reality, such as non-human animals, the environment, and so on (the branch of philosophy called **axiology**, which includes the more specific fields of *ethics*, *political philosophy*, and *aesthetics*).ⁱⁱ

THE METHOD OF PHILOSOPHY

We will forego for now any further explanation of the *content* of philosophy, since that is the main job of the rest of the book. But there are some implicit assumptions made by philosophers, and the clarification of these will require an explanation of the *method* of philosophical thinking.

The method of philosophical thinking requires a set of skills and some distinctive intellectual habits, what we call philosophical habits of mind. We will explain some of these skills and habits in this introductory essay, but the full appreciation of them requires exercising them on the philosophical views and arguments developed in the rest of the book. Two of the most basic skills involved in philosophical thinking are *clarifying* and *justifying* claims: as philosophers we are in the business of doing two main kinds of things, clarifying and justifying, in relation to a specific kind of object, a claim. What do we mean by clarifying and justifying claims? Let's break this phrase down.

First, what is a claim? As we have already seen briefly above, a claim is an assertion, something that is said with the intention of saying something that is either true or false. Here are some examples: there are cherry trees in the quad; Chicago is west of Washington, D.C.; $7 + 5 = 12$; grass is red; no dog has ever been lost; politicians are uniformly honest. Notice that claims can be false as well as true. Not everything that you say is a claim, since your intention is not always to state truths. For example, a question is not a claim, nor is an exclamation or a command.

Second, what is meant by clarification? When a philosopher *clarifies* a claim, he or she explains or spells out in detail the meaning of the claim. Clarification is often valuable or even urgently required, because the meaning of a claim as it is initially formulated can be seriously unclear in ways that make it difficult to discuss or evaluate it. Consider, for example, the claim that God is love. Presumably, the person who says that God is love intends to say something that is true, but some people find that claim very confusing. Does it mean merely that God is a loving person? No—it seems intended to mean something much more significant than that. But what? Since love is a kind of emotion, the literal meaning of the claim doesn't make clear sense (since God is surely not any kind of emotion).

So perhaps the claim is metaphorical instead of literal. It is much easier to clarify literal claims than it is to clarify metaphorical claims. But some important clarifying work is done even by saying that the claim is metaphorical.

Obviously, some claims need more clarification than others. Consider the following claims:

- (a) Money can't buy happiness.
- (b) God is love.
- (c) No bachelor is happy.
- (d) If a person is a mother, then that person is female.
- (e) Collies are dogs.
- (f) Studying philosophy is valuable.

These claims are not all equally clear. Which claim most clearly needs clarification? Claims (a) and (b) are both metaphorical, but one can more easily imagine explaining the metaphorical meaning of (a). Of course you cannot literally buy happiness, since it doesn't show up in any of the stores you go to, nor can you purchase it on the internet. But that is not the real point of the claim. Suppose, for example, that you were advising your younger sister against making what seemed to you to be a very bad marriage. The only good thing you can see about her prospective husband is that he is very wealthy. It would be very natural to say to her that money can't buy happiness, meaning that one can have lots of money and still be very unhappy. (You would, of course, be assuming that she wants to be happy.)

You can clarify a claim without thereby giving any reason to think that the claim is true. Think of how you might clarify claim (c), that no bachelor is happy. This is surely a false claim, but someone might still wonder about "happy." Before you think very hard about it, you might think that you surely understand what happiness is. But as soon as you try to clearly define it, all kinds of problems come up (see Chapters 5 and 8). Clarification sometimes requires explaining only one of the terms in the claim (as with (c)), while at other times it requires explaining the meaning of several of the terms (as with (f)). Sometimes a claim just needs to be made more precise. For example, someone might wonder if all collies are dogs, or just most collies are dogs, or only some collies are dogs. How much clarification a claim really needs might depend on the context.

Third, what is meant by **justification**? When philosophers justify a claim, they give *reasons* to believe the claim, and what better reason is there to believe a claim than a reason to think that it is true? Our conception of philosophy assumes that a reason to think that a claim is true is a good reason to believe it. Moreover, such a reason

seems at first glance to be the *only* kind of good reason to believe a claim—since to accept the claim is after all to accept it *as true*. In other words, if you don't have a reason to think that a claim is true, then you apparently have no reason at all to believe the claim. (Whether there might be some acceptable basis for believing a claim other than a reason to think that it is true is a question raised explicitly in Chapter 7.)

Let's suppose for the rest of this discussion that a reason for a claim will always be a reason for the truth of the claim. Another supposition that we'll make in explaining what is meant by justification is that the reasons advanced for the truth of a claim will themselves always be claims, assertions made in the attempt to say something true. And the assumption made in treating these claims as reasons is that the truth of the reasons provides evidence or warrant of some sort for the truth of the claim in question (the claim that we are trying to justify).

ARGUMENTS AND LOGIC

To put forth other claims in support of a claim you are defending is to give an **argument**. Thus, according to our account of justification, when a philosopher justifies a claim, he or she usually gives an argument. In philosophy, an argument is not a disagreement or a fight. According to the standard philosophical definition, an argument is a set of claims, one of which is the conclusion and the others of which are **premises** that are offered to support the conclusion: premises that are alleged by the person putting forth the argument to make it very likely or perhaps even guarantee that the conclusion is true. One of the first things you will learn when developing the skills that are important to the philosophical method is to become very sensitive to the difference between the conclusion and the premises of an argument: the claim being asserted by a philosopher, the claim that will be argued for is the conclusion, whereas the claims offered in support of the conclusion are the premises. One of the distinctively philosophical habits of mind is the one that clearly distinguishes between premises and conclusions, between what is being argued for and what is being offered as a reason.

One question that can be asked about the premises of an argument is whether they are true—or at least whether it is reasonable to think that they are true. But while the question of whether the premises are true is crucial to the strength of the argument, that shouldn't be the first question you ask when evaluating an argument. Before you worry about whether the claims offered as reasons are true, you should ask yourself whether, *if true*, they would genuinely support the claim. Reasons can support a claim more or less successfully, and when you ask how good the

reasons offered for the claim are (assuming that they are true), you are asking about the *strength* of the relation of support, the evidential relation, between the premises and the conclusion.

So the core idea of a philosophical argument is the idea of giving reasons for a claim: offering premises for the purpose of showing that the conclusion of the argument is true.

Some arguments are valid **deductive arguments**: arguments whose premises, if true, *guarantee* the truth of the conclusion. Consider the following argument for the claim that Mary took the car: either John took the car or Mary took the car, and I know that John didn't take the car, so Mary must have taken it. You can assess the strength of the relation of support of this argument without knowing John or Mary or knowing anything about the car. You just ask yourself whether the premises, if true, do support the conclusion. They do. If the premises of this argument are true, then the conclusion must be true.

But some arguments that are intended as deductive are **invalid**: it is possible for the premises to be true while the conclusion is false. Consider the following argument for the claim that Mary took the car: if Mary took the car, then John didn't take it; and I know that John didn't take it; so Mary must have taken it. Suppose that all of the premises are true. Does the truth of those premises guarantee (or even support at all) the truth of the conclusion? No, this argument commits the *fallacy of affirming the consequent*. (A fallacy is a mistake in reasoning.) It will be useful to digress just a bit to see clearly what this fallacy is, and why it is a mistake to reason this way.

As you will see, many philosophical arguments include **conditional** statements: statements of the form if *A*, then *B*. For example, if *Mary took the car*, then *John didn't take it*. The first part of the conditional, the *A* part, is called the *antecedent*, and the second part, the *B* part, is called the *consequent*. (A conditional statement can be confusing at first, before one thinks carefully about what exactly it is saying. Consider the claim that if George Washington was an octopus, then George Washington had eight legs. Is that claim true or false? Some students initially say that the claim is false, but once they look again, they see clearly that the claim is true.) The truth of a conditional statement does *not* require the truth of the antecedent (the claim that Mary took the car), because the conditional is claiming only that *if* Mary took the car, then John didn't take it. The *truth* of the conditional requires only that *if* the antecedent is true, then the consequent must be true (cannot be false); in other words, if the consequent is false, then the antecedent must also be false for the whole conditional statement to be true. Whereas if Mary and John both took the car, (in which case the antecedent is true and the consequent is false) then the conditional statement, *if* Mary took the car, then John didn't, is itself false. So if a conditional statement is true, and

the antecedent is true, then you know that the consequent must also be true; and if a conditional statement is true and the consequent is false, then you know that the antecedent must also be false. But *nothing* follows from the truth of a conditional and the truth of the consequent, and thus arguments that *claim* to draw a conclusion about the truth of the antecedent on the basis of the truth of the conditional and the truth of the consequent are making the mistake in reasoning called the fallacy of affirming the consequent.ⁱⁱⁱ

Since the premises in an instance of affirming the consequent don't actually support the conclusion at all, you might be tempted to say that this is no argument at all. But it doesn't seem right to say flatly that this isn't an argument: it seems clearer to say that it is a *bad* argument, and better still to say exactly what is bad about it. (When faced with an invalid argument, you need not worry about whether its premises are true, since even if they are, they don't provide even minimal support for the conclusion.)

There are other kinds of arguments whose premises provide good but not conclusive grounds for the truth of the conclusion: arguments that offer genuine support for their conclusions but where it is still possible, though unlikely, that the conclusion is false even though the premises are true. The arguments most standardly referred to as **inductive arguments** (or, more explicitly, enumerative inductive arguments) are like this, and many scientific arguments are of this kind. When, for example, someone reasons that all swans are white on the basis of many different observations of white swans, he or she is giving a very simple example of an inductive argument. You cannot reasonably conclude that all swans are white on the basis of one observation of a white swan, or even two or twenty, but if there are enough observations in sufficiently varied locations and circumstances, then you can reasonably conclude that *all* swans (and not just the ones you've observed so far) are white. When you reason that the sun will rise tomorrow on the basis of the claim that it has risen every morning so far for thousands of years, you are giving an inductive argument. Philosophers don't call good inductive arguments *valid*, because the definition of validity is that it is impossible for the conclusion to be false while the premises are true. Inductive arguments, by definition, have conclusions which might be false even though the premises are true. But the better the inductive argument, the more unlikely or improbable it is that the conclusion is false while the premises are true. Good inductive arguments, ones where the evidential or support relation between the premises and the conclusion is compelling, are usually described as strong. In such an argument, the truth of the premises provides a good reason to think that the conclusion is true.

A different kind of argument whose premises provide good but not conclusive reasons for the truth of its conclusion is what is referred to as an **explanatory argument** (also called an *inference to the best explanation* or an *abductive* argument—sometimes the term “induction” is used more broadly so as to also include arguments of this sort). The idea of an explanatory argument is that there is an alleged fact of some sort to be explained, other considerations that are relevant to the explanation of this fact, and some explanation that is claimed to be the best one in light of those considerations. Thus the premises of such an argument include both a statement of the alleged fact to be explained and statements of these other relevant considerations, and the conclusion is a statement of the explanation alleged to be best. And such an argument will be *strong* (never valid) if the explanation offered really is the best one, assuming that the fact in question really is a fact and the other supposedly relevant considerations are also true.

Here is a simple example. Suppose that the police call you at work to tell you that your car has been in an accident, and that the driver of the car left the scene. The question is how to explain the fact that your car has been in an accident (rather than still parked in the driveway where it was when you left to take the bus to work this morning). The following further considerations seem relevant: that only you and your sister Mary ever drive the car, even though she has recently been forbidden by you to drive it because of the many speeding tickets she has received; that there is only one key to this car (kept hanging on a hook by the back door); that Mary is the only person (other than you) with easy access to this key; and that the police found the car with the key still in it. Thus, it might be claimed, the most likely explanation of the fact that your car was in the accident, instead of still sitting in the driveway where you thought it was, is that Mary drove it (in spite of being forbidden to do so). Obviously you'll have even better evidence once you talk to her or find out from eyewitness reports what the driver of the car looked like, but don't you right now have a pretty good reason for the conclusion that Mary took the car? Your reason is not conclusive—the conclusion is not guaranteed to be true—because there are other possible explanations of the fact that your car is at the scene of an accident that might turn out to be better ones after all. (Maybe someone broke into your house, took the key, and drove the car away.)

We use explanatory arguments of this sort in everyday life, and scientists use explanatory arguments to draw conclusions about laws and theoretical entities. Such arguments also often play an important role in philosophical discussions.

Summing up, according to the standard definition of an argument that we started with, if premises are offered for the purpose of supporting the truth of a conclusion, then the set of claims consisting of those premises and the conclusion constitute an argument. The argument is **deductive** if the truth of the premises is intended to guarantee the truth of the conclusion; if the truth of the premises is merely intended to make the truth of the conclusion very likely or probable, but not guaranteed, the argument may be **inductive** or **explanatory**—though there are also other possibilities that we have not considered here, such as arguments that appeal to analogies. A deductive argument whose premises are related to its conclusion in the right way to accomplish its purpose is valid, while good inductive or explanatory arguments can only be strong. A further issue about any argument is whether the premises themselves are true.

AN ILLUSTRATION OF THE METHOD OF PHILOSOPHY

We have just thought through an initial explanation of two of the important skills involved in the method of philosophy: clarification and justification; let's now consider an illustration of those skills. Let's attempt to clarify and justify the claim that studying philosophy is valuable.

Clarification: Defining What We Mean

Let's first clarify the claim that studying philosophy is valuable. You reasonably wonder about the meaning of both parts of the claim: what is meant by "studying philosophy" and what is meant by "valuable"? Suppose that someone read Bertrand Russell's book *The Problems of Philosophy* one night—is that enough to have "studied philosophy"? Not according to what we mean when we make the claim. You have to do more than read one philosophy book to have studied philosophy. It takes more than one night to study philosophy. But there is no exact amount of study that can be said to be the precise meaning of "studying philosophy."

Sometimes what is needed to make an issue clear enough to be reasonably discussed is to replace the original claim with one that is clear and more precise, while still saying approximately the same thing. Thus, somewhat arbitrarily, we will understand "studying philosophy" to mean taking and passing at least four philosophy classes or doing something reasonably similar. (You could probably do the equivalent of taking and passing four philosophy courses on your own, if you were sufficiently motivated and had some resources for checking your understanding.)

Now what do we mean by the second part of the claim, “valuable”? All we mean by “valuable” here is that it is good for you, that you will be significantly better off for having done it. You might doubt that our claim is true, but you now have a pretty good sense of what we mean by it. We have clarified (though perhaps not yet enough) our claim that studying philosophy is valuable.

We might make one more clarification. Someone might wonder whether the claim is that studying philosophy is *always* valuable no matter who does it or just that it is *usually* valuable. For example, think about the following claims that have the form *doing A is B*:

Running a marathon is demanding.

Giving birth is done by females.

Watching TV is fun.

Regular exercise is important.

Getting a college degree is worthwhile.

The context, together sometimes with the content of the claim, usually determines whether someone who asserts one of these claims means “always” or “most of the time”—though this may sometimes be unclear. Let’s assume that what we mean when we say that studying philosophy is valuable is that it is *always* valuable.

So we now have a pretty good restatement of the original claim resulting from this initial effort at clarification: our claim that studying philosophy is valuable means that anyone who does the equivalent of taking and passing at least four philosophy courses will benefit significantly from it. One of the distinctively philosophical habits of mind is that of noticing when claims are more or less clear.

Justification: Giving an Argument

Now let’s move from clarification to justification, remembering that philosophical justification typically takes the form of an argument. Here is an argument for the claim we have just clarified:

1. Studying philosophy always makes you think more clearly.
2. Thinking more clearly is always valuable.

Therefore, studying philosophy is always valuable.

What makes this an argument? It is a set of claims, one of which is the conclusion and the others of which are premises that are put forward to support the conclusion. (When you give an argument for a claim, the claim is the conclusion of the argument.) So the first two sentences are the premises and the third sentence is the conclusion. Just pretend for a moment that the premises are true.

If they are true, does the conclusion follow? In other words, does the truth of the premises make the conclusion likely or even certain to be true? For this argument, as for many other arguments, you need to *think* about the premises and the conclusion in order to answer that question—the answer to that question is not the result of a clear mechanical test. As you acquire the intellectual habit of assessing arguments, you will get better and better at distinguishing the good ones from the bad ones. What you should do is suppose that the premises are true and then try to imagine whether it would be possible for the conclusion to be false, even given the truth of the premises.

We suggest that the truth of the premises (if they are true) provides a very good reason to think that the conclusion is true. In fact, the argument appears to be **valid**: it appears to be logically impossible for the premises to be true and the conclusion false. (We will consider further on a reason for questioning whether this is really so.) This argument therefore provides a good example of the relationship of *providing a good reason*, which is the core element of an argument.

In the argument above, if the premises are true, then the conclusion is apparently guaranteed to be true. But we obviously can't just assume that the premises are true. And since the conclusion of the argument has been justified only if the premises are true, our job of justifying the original claim is not finished until we have at least defended the premises (given reasons to think that the premises are true). Furthermore, we should also consider and respond to the most obvious objections, if there are any.

Giving Reasons for the Truth of the Premises

Let's begin with the first premise: studying philosophy always makes you think more clearly. In anticipation, think about what will be involved in studying philosophy: you will read many different philosophical texts from many different periods of history, learning what different philosophers have said about many different topics. Moreover, since the authors are philosophers, they will typically be arguing for their views, so you need to understand and critically evaluate those views and those arguments, in an attempt to figure out what you think about the philosophical topic at issue. In addition, since philosophers have to discuss many other sorts of things in explaining and clarifying their views, presenting arguments, and considering objections to other views (and, as we will see,

even to their own), a philosophical work is usually quite complicated, with all of these parts needing to be clearly sorted out. Studying philosophy involves carefully doing all these things.

What then is thinking clearly? Thinking clearly obviously involves being able to clarify various ideas and views that you encounter. But it also involves being logical: considering and sometimes discovering reasons for those views, together with being able to successfully evaluate when those reasons are good ones and when they aren't. One clear-thinking skill is the ability to juggle complicated combinations of ideas, keeping track of the different relations between the ideas, and like any skill, it takes time and practice to become good at it. You obviously need to think clearly in order even to understand philosophers, and you surely need to be able to think clearly in order to evaluate philosophical views. (When you evaluate a view, you decide whether it is a good view (likely to be true) or a bad view (likely to be false), and as a philosopher, you must have reasons for making that evaluation.) Therefore, someone who has studied philosophy (given how we have clarified that idea) has either learned to think really clearly for the first time, or else already knew how to think clearly to some extent but has now had lots of additional practice at it, and so presumably thinks even more clearly. There is a defense of the first premise: a reason to think that the first premise is true.

Now consider a defense for the second premise, that thinking more clearly is always valuable. Surely we can all agree that doing something that helps you get what you want is valuable, that you benefit significantly from doing something that enhances your ability to get what you want (unless, of course, what you want is not good for you). We contend that clear thinking always does that. Suppose, for example, that two people, call them Joe and Doug, each want to finish college as efficiently as possible, and suppose further that Joe thinks much more clearly than Doug. Suppose that Doug never quite thinks hard enough to keep the degree requirements straight or even realize that some of the requirements exist. Suppose that when an advisor suggests to Doug that he take a particular class, Doug never asks why or how that class will fit into his overall program. Doug probably doesn't have a clear plan, but he does have the desire to finish college as efficiently as possible. Suppose that Joe is constantly honing his thinking skills: he always asks the advisor to clarify her advice, always asks why this is a good class to take, and keeps the degree requirements clearly in mind. It is reasonable to assert that because of Joe's ability to think more clearly than Doug, Joe is more likely to realize his goal than Doug is—and this is not even the best example, is it? You don't have to think really clearly to realize the goal of finishing college as efficiently as possible. Imagine how clearly you have to think to be a responsible citizen or a loving friend or a terrific parent. Being able to distinguish

between believing something on the basis of wishful thinking as opposed to believing something on the basis of good evidence might make the difference between doing a good job and an inadequate job in many areas of human relations. So thinking clearly is always valuable because it helps you get what you want, whether you want to be a law clerk, a good parent, or a couch potato.

At this point, we have given reasons to think that both premises are true, and it seems initially clear that if the premises are true, then the conclusion must be true. So we have presented and defended an argument, but does that suffice to justify the claim? It surely is a justification, but it isn't yet the strongest justification we could give. An even better justification for a claim also includes considering and responding to objections to our own argument.

Objections: Considering Reasons Against the Truth of the Premises

Some students are reluctant to consider **objections** to an argument they are trying to defend, because it seems to them that they are weakening their own position. But an argument that considers and responds to objections is much stronger than an argument that considers no objections at all. Imagine that you are reading two editorials in the newspaper, one of which expresses your own political views, while the other expresses views contrary to yours. Suppose that each piece argues for its position without considering any alternative points of view that might lead to objections. When you read the one you agree with, it is—unfortunately—all too easy just to go along with the argument (you do, after all, already agree with the conclusion). But when you read the one you disagree with, you are probably thinking of objections along the way, and so perhaps you don't feel challenged in your own view because you think that you have good objections to the reasons given for the conclusion you disagree with. But imagine that the editorial you disagree with proceeded to consider objections similar to the ones you are thinking of as you read it, and imagine further that the responses it gives to those objections are pretty compelling. Wouldn't you feel more uncomfortable shrugging off the view it defends in that case? Wouldn't the challenge to your own view seem more serious? Analogously, wouldn't the piece arguing for the view you agree with be even stronger if it also considered and responded to objections? Any view argued from just one side, without considering alternative perspectives and the resulting objections, is not as compelling as a view that has considered the strongest objections and has also shown how those objections, however strong they seem, can be satisfactorily answered. This is an especially important philosophical habit of mind: see many sides of an issue—don't be satisfied with just one perspective.

We want to consider and respond to objections for the purpose of strengthening our argument, but presumably we find the argument pretty convincing (since we devised it), so how does one go about finding good objections? This is another one of the skills that you'll acquire as you develop philosophical habits of mind: you will need to be able to take up a critical attitude, criticizing the arguments of other philosophers. You can also take that same point of view in relation to your own argument, pretending that you hold the other point of view and looking for weaknesses in your original argument.

You might think at first that the way to object to an argument is to object to its conclusion—find reasons for thinking that the conclusion is false. But in fact this doesn't really work very well if you are trying to criticize the original argument. For if you offer reasons to think that the conclusion is false, then you have simply produced another argument for the opposite conclusion. You now have two opposed arguments, leading to opposed results, but not really engaging each other in any more substantial way. They can't both be valid arguments with true premises and it is unlikely (though not impossible) that they are both strong. But the mere conflict between them in itself gives no insight into which one is mistaken or—even more importantly—how it is mistaken. Thus what makes more sense, if the goal is to evaluate, criticize, or strengthen the original argument, is to consider objections to its premises or to the reasoning from the premises to the conclusion, rather than reasons to reject the conclusion. If there are good reasons to think that the premises of an argument are false or that the reasoning from the premises to the conclusion is faulty, then the argument fails to support its conclusion; but if those objections can be answered, then the overall case for the conclusion is strengthened.

Let's first consider an objection to each premise of our sample argument. An objection to a premise is a reason to think that the premise is false. Consider the first premise: studying philosophy makes you think more clearly. Someone might object by saying that studying philosophy is very confusing. To study philosophy, you have to read many different authors on many different questions, and many of the authors lived a long time ago, so their writing style is very different from ours and often hard to understand. Philosophical questions are hard to understand in the first place because they are so abstract and so remote from everyday concerns. Thus many people who study philosophy just end up very confused, and surely someone who is very confused is not someone who thinks clearly. Therefore studying philosophy does not make one think more clearly.

What about the second premise: that thinking more clearly is always valuable? Someone might object to this premise by pointing out that the more clearly one thinks, the more clearly one sees how vulnerable we puny humans

are. We have many grand desires, but our ability to “get what we want” is largely dependent upon conditions beyond our control, and thus all our planning and scheming is in the end just pathetic. The more clearly we think, the more clearly we realize this and the more paralyzed we become. We obviously do not benefit from being so paralyzed. On the contrary, ignorance and murky thinking are bliss. Therefore, thinking clearly is not always valuable.

Objections can also be made to the *reasoning* involved in arguing that the conclusion is likely to be true, given the truth of the premises. Consider the following objection to the reasoning of the sample argument. Suppose that someone grants the truth of both of the premises but argues that there are other, indeed much easier ways to learn to think more clearly than by studying philosophy. If the conclusion, that studying philosophy is always valuable, means that anyone should, all things considered, study philosophy, then that conclusion might very well be false, in which case the argument isn't really valid after all. Suppose that you are a math major; and suppose also that studying math teaches you to think more clearly; and suppose further that given your talents and interests, taking the time to study philosophy would take time away from other things that you enjoy without adding very much of benefit (since you are already learning to think clearly by studying math). So you could argue that even if it is true that studying philosophy would teach you to think more clearly and that thinking more clearly is always valuable, it is false that studying philosophy is valuable for you. It is false that you, given all of what is true about your life, should study philosophy. This objection challenges the reasoning involved in drawing the conclusion from the premises instead of challenging one of the premises.

At this point, we have formulated an objection to each of our premises and an objection to the reasoning of the argument, all with the ultimate goal of strengthening our case for the claim that studying philosophy is valuable by answering these objections.^{iv}

Responses: Showing Why the Objections Fail

Obviously we need to *respond* to these objections: we need to show why these objections are not strong enough to seriously affect the force of the original argument. In responding to objections, you will sometimes show that the reasoning of the objections is faulty, while at other times you may respond by showing that the original statement of the premise or the reasoning needs to be altered or qualified in order to accommodate the objection (while still being able to make the case for the claim you are defending).

The first objection claims that studying philosophy, rather than leading to clear thinking, is confusing. But while some people do find it confusing at first, that sense of confusion almost always goes away once one works at

it a bit harder. It is really not very easy to pass four philosophy classes without acquiring the thinking skills that clear up the confusion. This objection might have had serious weight if we had specified the idea of studying philosophy as meaning passing only one philosophy course, but it has no serious weight against the first premise when studying philosophy is understood as requiring one to have passed four philosophy courses (or done the equivalent).

The second objection claims that we are virtually powerless to get what we want no matter how carefully we plan and anticipate, so that clear thinking really isn't always valuable. But what evidence can be offered for such a claim? What reason can this objector offer for the view that our planning and careful thinking make no difference (or so little difference as to be irrelevant) to the outcome of our efforts? Suppose that you point out that you know of many examples of people who have often succeeded in getting what they want when they've planned carefully, and of other people who don't plan carefully and fail to get what they want. These examples are **counter-examples**^v to the objector's view: examples which give some reason to think that the view is false. Of course no one would claim that careful planning guarantees good results. The defense for the premise under consideration need only claim that careful, clear thinking makes getting what you want more likely—which is enough to make such thinking valuable. So the objector needs to give some reason to think that it is false that careful, clear thinking makes it more likely that you will get what you want.

The defender of the objection might respond that all of those people who planned carefully were just lucky, and the people who didn't plan well were just unlucky. Suppose that the objector continues to say the same thing (planners are just lucky) no matter how detailed your examples are and no matter how many examples you come up with. There is nothing that anyone can say that would conclusively prove that we have more control than the objection says we do, so the objector's view cannot be shown conclusively to be wrong. Is the objector's insistence on the view thereby legitimate? Is it intellectually respectable to insist on a view despite possible counter-examples, just because the view hasn't been *conclusively* proven to be false? This seems clearly unreasonable.

Our main response to this objection, then, is that the objection relies on a very controversial, ill-defended assumption, one that seems to us to be clearly wrong. Therefore, the objection fails to have any serious force against the premise it is aimed at.

What about the objection to the reasoning of the argument? The main response to that objection is that the objection misunderstood the content of the conclusion. The point of the argument was not to argue that everyone

ought to study philosophy, even though it was perhaps not entirely unreasonable for the objector to interpret the conclusion that way. One way to understand the claim that studying philosophy is always valuable is to think that studying philosophy will be valuable for *anyone*. And that strongly suggests that anyone ought to do it. But we could argue that the conclusion of the argument does not mean that anyone, no matter what else is true of him or her, ought to find some time to study philosophy. The conclusion instead just means that *if* you study philosophy, then that study will be valuable for you in the sense that you will benefit from it. This conclusion is completely consistent with the claim that for any particular person, he or she ought not to study philosophy because the study of philosophy, in spite of the benefits that it produces, would not be valuable, all things considered, given what would have to be sacrificed to engage in it and given the possibility of acquiring those same or very similar benefits in some other way. (Here you can see how considering objections also can help to clarify a position.)

Summary

We have clarified and justified the claim that studying philosophy is valuable. We have thereby illustrated many of the skills and philosophical habits of mind required by the method of philosophy. Can you see how the claim is clarified further in the process of justification? Can you see how much stronger the justification is because we have considered and responded to objections?

As you can see now, one of the first challenges of doing philosophy is to learn how to tell when a philosopher (whether it is you or someone you are reading or listening to) is arguing for a view, objecting to it, or responding to objections. As you practice thinking philosophically, you will get better at recognizing these different activities, and you will get better at clarifying claims, drawing distinctions, and making arguments, objections, and responses yourself.

READING PHILOSOPHY

With a little work, almost anyone can learn to think philosophically. But you might wonder what you need to do to work effectively at acquiring this ability. The first, most valuable resource on which you should practice thinking philosophically is the set of reading selections in this book (and any other philosophical texts you may read). But we've learned that students often find reading philosophy very difficult at first, and so we've included some advice about reading philosophy, together with some brief illustrations. And while the advice is directed primarily at

reading philosophy, the main points also apply to listening to a philosophical lecture or participating in a philosophical discussion.

Philosophical material is primarily argumentative and critical, almost never merely expository. You don't read philosophy in order to gather lots of facts that you will then memorize. You should read philosophy as if you were actively thinking along with the author of the text, as if you were having an intellectual conversation with him or her. Philosophers are arguing for a view or position. You should think of the author (or lecturer or discussant) as saying: "Look—this is *what* I think, and this is *why* I think it. What do you think?" Thus you should always keep two questions in mind while reading (or listening to lectures, or engaging in discussion with another philosopher): (1) What view or position is the philosopher advocating? The answer to this question may differ at different places in an article: an article may make more than one point, and you should ask how these points connect with each other. Notice also that what is being advocated at a particular point may be very simple and general (e.g., the claim that God exists) or very complicated and specific (e.g., the claim that a particular objection to a particular argument that God exists is mistaken). (2) The second question to ask is what reasons or arguments are being offered in support of the view being advocated. Try to sketch the answers to these two questions as you read—either in your head or, even better, on paper (perhaps in the margin of the book). You will then be in a good position to ask the next question: how *strong* are the reasons offered, and are there objections to them? It is virtually impossible to do this while relaxing in a passive frame of mind. Successfully reading philosophy requires an active, critical, imaginative mind-set (one of the distinctively philosophical habits of mind that you need to cultivate in order to learn to think philosophically).

There is one major source of confusion to watch out for. Because philosophy is essentially reflective and critical, a philosopher will typically discuss other positions and arguments besides the one defended in a particular article. These may include any of the following: (1) positions opposed to the one defended; (2) positions similar to, but still significantly different in some respect, from the one defended (where the difference helps to clarify the main view); (3) arguments in favor of positions opposed to the one defended, which will be criticized; (4) objections to the philosopher's own position, which will be responded to; and (5) sometimes even arguments for the view defended that the philosopher does *not* accept and wants to distinguish from ones that he does accept. It is obviously important to distinguish all of these from the philosopher's own positive views and arguments, and this is relatively easy to do if you are alert. (For example, arguments and positions other than his own will often be introduced by

noncommittal phrases like “it might be claimed that...” or “some writers argue that...” or perhaps by a reference to a particular person who holds the view or argument in question.)

It is usually a good idea first to look rather quickly through a reading to get an idea of the general “lay of the land” and then to read it again more carefully and critically. Otherwise it is too easy to lose sight of its main point as one works through the detailed arguments. Remember also that considerable rereading will certainly be necessary, especially but not only when reading a later selection that disagrees with an earlier one. Almost no one (including professional philosophers) is capable of adequately grasping a complicated argument in a single reading. And the most important piece of advice: don’t give up! The more you work at these skills and habits, the better you will become at them, and the more you will also come to see the point of philosophy. In this way, intellectual tenacity is one of the most essential philosophical habits of mind.

TWO SAMPLE PASSAGES

Sometimes a general characterization of how to do something, such as how to read philosophy, only goes so far, and what a student really wants and needs is a concrete example: a piece of philosophy with some discussion about how to implement some of that general advice. Here, then, are two sample passages: one from J.L. Mackie on the problem that evil poses to the belief in God’s existence and the other from John Locke on the problem of justifying our confidence that our senses provide reliable information about the world.

Example 1: a contemporary passage

Before we analyze the first passage (given below), we need to say something about its argumentative context. Mackie’s own view is that the problem of evil shows that God, conceived as omnipotent and wholly good, does not exist. But he recognizes that some people think that the free will response to the problem of evil undermines his view. According to the free will response to the problem of evil, God gave man free will in spite of the fact that such a gift would result in the occurrence of evil, because, according to Mackie’s interpretation of this response: “it is better on the whole that men should act freely, and sometimes err, than that they should be innocent automata, acting rightly in a wholly determined way.”

So in the passage we’ll analyze, Mackie is objecting to (giving reasons to reject) the free will response to the problem of evil. Here is the passage:

. . . if God has made men such that in their free choices they sometimes prefer what is good and sometimes what is evil, why could he not have made men such that they always freely choose the good? If there is no logical impossibility in a man's freely choosing the good on one, or on several occasions, there cannot be a logical impossibility in his freely choosing the good on every occasion. God was not, then, faced with a choice between making innocent automata and making beings who, in acting freely, would sometimes go wrong: there was open to him the obviously better possibility of making beings who would act freely but always go right. Clearly, his failure to avail himself of this possibility is inconsistent with his being both omnipotent and wholly good. (p. xxx)

Mackie starts this passage with a question, and that is often an effective rhetorical choice to get the reader thinking in the right direction. Notice, however, that he follows it up with a claim, one that is essential to his objection to the free will response:

If there is no logical impossibility in a man's freely choosing the good on one, or on several occasions, there cannot be a logical impossibility in his freely choosing the good on every occasion.

Immediately after that claim, Mackie draws the conclusion that God could have made free people who (as a result of the way he made them) always freely choose to do the good, but he formulates this in a more complicated and less immediately perspicuous way:

God was not, then, faced with a choice between making innocent automata and making beings who, in acting freely, would sometimes go wrong: there was open to him the obviously better possibility of making beings who would act freely but always go right.

The last part of this complicated sentence actually formulates the conclusion just stated, while the rest makes clear how it contrasts with the view he is opposing.

What other premises are required to make this argument fully explicit? Obviously Mackie is assuming, without explicitly saying so, that it is logically possible for someone to freely choose the good at least some of the time. It is hard to see how the defender of the free will solution could deny that. So now we have two premises, one explicit and one implicit:

1. If there is no logical impossibility in a man's freely choosing the good on one, or on several occasions, there cannot be a logical impossibility in his freely choosing the good on every occasion.
2. There is no logical impossibility in a man's freely choosing the good on one, or on several occasions.

From these two premises it follows that: there cannot be a logical impossibility in his freely choosing the good on every occasion, and it immediately follows from this intermediate conclusion that it is indeed logically possible that he freely chooses the good on every occasion.

Now what further thing is being assumed before it follows that God could have made free men who always choose to do the good? There is still one more implicit premise, one that Mackie is quite sure the defenders of the free will solution would agree with.

3. God can make actual whatever is logically possible.

So with just a little thought about logic and what it takes to fill in the argument, we can see what Mackie was up to when he moved from the following conditional statement: “if a person can freely choose to do the good at some times, then that person can freely choose to do the good at every time” to the conclusion: “God could have made people so that (as a result of the way he made them) they freely choose to do the good every time.” Here is the whole argument:

1. If there is no logical impossibility in a person’s freely choosing the good on one, or on several occasions, there cannot be a logical impossibility in his freely choosing the good on every occasion.
2. There is no logical impossibility in a person’s freely choosing the good on one, or on several occasions.

So, there cannot be a logical impossibility in his freely choosing the good on every occasion.

So, it is logically possible that he freely chooses the good on every occasion.

3. God can make actual whatever is logically possible.

So God could have made persons who (as a result of the way he made them) freely choose the good on every occasion.

Notice that you still don’t get the conclusion that God could have made people who did no evil, without adding the additional implicit premise that people who freely choose the good on every occasion do no evil. (But how explicit do we really have to be?—that’s a matter for you to decide in thinking about the argument.)

Mackie’s final sentence in the passage says more than he needs to say to object to the free will solution—do you see that? He is there reiterating his main view that no God who is wholly good and all-powerful could allow this

particular possibility, free men who perpetrate evil, when the men could be made just as free and while also being made so as to not perpetrate evil.

Now remember that your task is to read Mackie critically. To do that you must first see clearly what he is saying and why he thinks it is true, which is why we have clearly stated his conclusion as well as his premises. We have clearly laid out his argument, and now we should try to evaluate it: to determine how strong his reasons are. Remember that the first thing to do is to evaluate the form of the argument: is it strong—or perhaps even valid? Is it such that if all the premises are true, then the conclusion is likely (or even guaranteed to be) true? The form of the reasoning seems very strong, doesn't it? (In fact, it is valid.)

Now look to see which, if any, of the premises are most questionable. While there are some philosophers who would accept all of Mackie's premises and thus would accept his argument and his conclusion, there is one premise that many other philosophers, including many or most of those who are sympathetic to the view he is criticizing, would reject. Can you see which premise this is? Which premise would be easiest to object to? (This is not an easy question.)

The premise of this argument that is most likely to be challenged is in fact premise 3. While it sounds initially very reasonable to suppose that an omnipotent God can create anything that is logically possible, there is a subtle way in which this may be wrong. While it is surely possible that a free person might always make the morally best choice ("choose the good"), can God bring it about that a person always chooses in this way without making the person no longer free? The underlying idea here, explored at much greater length in the free will chapter, is that when a person *freely* chooses something, then it is always true that he (or she) could have chosen something else instead. But is this true of someone who makes the morally best choice because God has created him so as to bring about that he will make that very choice?

Now someone might worry that spending this much time on every single paragraph of your philosophy reading would take forever, and the response to that worry is that you obviously don't do this for every paragraph. We've chosen especially important, clearly argumentative passages to use as examples. But you should now be able to see why you are rarely assigned to read as many pages in your philosophy courses as you are for most of your other courses. Reading philosophy takes lots of time and thought and care and imagination. Notice that we could have gone on much longer, since we only began the evaluation phase. How far you go with the evaluation phase depends on your purposes and your level of experience. It is sufficient for just understanding the passage to give fairly brief

answers to the questions we framed above, but if you were going to write a paper evaluating Mackie's reasoning, you would have to do more.

Example 2: a historical passage

The other passage we will consider was written in a very different time, but it is still distinctively philosophical in that there is a philosophical claim being made and reasons being offered for the truth of that claim. In the *Essay Concerning Human Understanding*, John Locke is defending the very natural view that we have knowledge of the external (material) world and that our sense experience provides us with sufficient justification to ground that knowledge. In the early chapters of the book, Locke focuses on the source and nature of our ideas, because he believes that the reasonable starting point for responding to the skeptic is to show how we can trust that our senses give us the right (that is, largely true) ideas.

In Book IV, Chapter XI he offers four reasons for the conclusion that our senses “do not err in the information they give us of the existence of things without us.” We will focus here on the fourth of those reasons:

Our senses in many cases bear witness to the truth of each other's report concerning the existence of sensible things without us. He that sees a fire may, if he doubt whether it be anything more than a bare fancy, feel it too and be convinced by putting his hand in it. Which certainly could never be put into such exquisite pain by a bare idea or phantom, unless that the pain be a fancy too: which yet he cannot, when the burn is well, by raising the idea of it, bring upon himself again. (p. xxx)

First, what does Locke's conclusion mean? He is talking about ordinary sense perception. By “the existence of things without us” he means the existence of ordinary material objects—trees, buildings, mountains, rivers, and so on—outside of us, outside of both our bodies and our minds. By “the information they [the senses] give us” he just means the beliefs about such objects that we naturally form as a result of sense perception and that seem to reflect the content of perceptual experience. And by saying that our senses “do not err” in this information, he is saying that the beliefs in question are, at least for the most part, true.

So what does his reason for this conclusion amount to? He says: “Our senses in many cases bear witness to the truth of each other's report concerning the existence of sensible things without us.” Sensible things are just things that we can (apparently) sense—again material objects like tables and chairs. But what does he mean by “bear witness”?

Look at the example he gives. (Examples are often crucially important in understanding abstract philosophical claims.) As Locke would put it, your sense of sight reports to you that there is a fire, and then your sense of touch also reports the same thing, because it too will tell you that there is a fire—if you put your hand in (or near) to where the fire seems to be. If your visual perception of a fire was a “bare idea or phantom” (that is, was something like a mere mental illusion, not caused by and corresponding to something existing outside of you—something real), then when you tried to feel what you were seeing, you wouldn’t get the right feeling (that is, you wouldn’t experience pain—or heat)—unless, of course, the pain (or sensation of heat) was an illusion (“a fancy”) too. (But how likely is it that two illusions would fit together in that way?—this is part of Locke’s point.) Notice also that when the burn heals, one cannot, just with imagination, make oneself experience the same pain that one has when one really puts one’s hand in a fire—thus showing, he thinks, that the pain experienced in the actual case is more than just a bare idea.

So the example suggests that for the different senses to “bear witness to the truth of each other’s report” is for one sense to tell us what another sense also tells us. Clearly Locke thinks that this supports the conclusion that our senses do not err.” But we need to say more about how this conclusion is supposed to follow. How can we reason from the premise that our senses agree with each other in this way to the conclusion that what they tell us is true (or at least very likely to be true)? Something seems right about the thought, but just exactly how does the reasoning work?

Locke seems to be thinking that if what our senses tell us was not true, then it would at least be very unlikely that our senses would agree with each other in this way. But why think that is true? One way to make sense of this argument is to see it as an explanatory argument or an inference to the best explanation. What would explain the fact that our senses agree with each other?^{vi}

One possible explanation, the one that Locke seems to have in mind, is this: our perceptual experiences are systematically caused by external objects in a way that makes the perceptual experience accurately reflect the properties of those external objects. This is an explanation of the fact in question: if it were so, then the different senses, being affected by the same external objects (whichever ones are present where our bodies are located), should agree with each other in the way that they do. But is it the best explanation? Or are there other explanations that are at least equally good—equally good from the standpoint from which Locke is making this argument, one in which the accuracy of our perceptions and the very existence of the common-sense material world is in question?

In fact, there are a number of other possible explanations, though a full assessment of them is not possible in this discussion—see chapter 2 for more on this issue. Perhaps we are dreaming. (Do apparent perceptions agree with each others in dreams?) Perhaps a powerful being of some sort is systematically causing perceptions that do not correspond to any material reality, but which still agree. Perhaps instead your own subconscious mind is doing that. Or what about the possibility that the perceptions in question are systematically caused by external reality (which is why they agree) but in a way that distorts that reality very badly (which is why the resulting beliefs are not in fact true)? Which explanation is best: Locke's or one of these others, and why?

One other thing worth noticing is that we did not set out this argument in numbered steps, in the way that we did with Mackie's argument in the previous passage. This could have been done, but it would not, in our judgment, have been particularly helpful in this case. The structure of this argument is very simple: an alleged fact and the claim that a certain conclusion is the best explanation of that fact. What is complicated is the reasons for thinking that the alleged fact is a fact and the reasons for thinking that the alleged best explanation really has that status, and neither of these things lends itself very well to stepwise formulation. The moral here is that formulating an argument in numbered steps—or in any other way—is a *tool* for clarification and should be used where it is helpful and not otherwise.

One challenge to evaluating a philosopher's reasoning, not just here but in other places as well, is that philosophers are often defending views that we already believe to be true—as is surely the case with this passage from Locke. But be careful. Your job as a philosopher is to critically evaluate the cogency of the reasons offered for a claim, whether you independently believe the claim or not. So you need, in order to evaluate this passage, to scrutinize the reason Locke actually gives, asking yourself whether it are sufficient to show that the conclusion is true. Answering that question is hard, because evaluating explanatory arguments requires deciding which of many possible explanations is best—an issue for which the relevant standards are not entirely clear. But you can begin to do all of this when reading and trying to understand Locke's view even if you cannot see through to the end of such a line of reasoning. The harder you push yourself to understand exactly what he is saying and evaluating how cogent his reasons are, the more skilled you will become at reading philosophy and thinking philosophically.

Discussion Questions

1. Why do we need a single conception of philosophy for the purposes of this course? What if someone said, “We don’t need a single conception of philosophy. We can work together, doing philosophy, even if we each have different conceptions of philosophy.” Does that sound reasonable? Why or why not?
2. The area of philosophy called metaphysics studies the nature of reality. What do you think about reality? Is it clear what things are real and what things are not real? List some things that are real. Then list some other things that are not real. Does science study the nature of reality? What do you think is the difference between science and philosophy?
3. We claim that how much clarification a claim needs varies depending on the context. Consider, for example, the claim that there are trees in the quad. Now imagine the following context where the claim needs no clarification: suppose a landscaping firm has been hired to fertilize all the trees on campus, and they want to know if they need to go to the quad. You tell them that there are trees in the quad, and your claim is sufficiently clear and precise for the purposes at hand. But suppose there is another context: a landscaping firm has been hired to fertilize all the cherry trees on campus, and they want to know if they need to go to the quad. Now the claim that there are trees in the quad needs to be clarified, in this case made more precise: they need to be told whether there are *cherry* trees in the quad. Now you think of an example of a claim and two contexts, in one of which the claim needs clarification, and in the other of which the very same claim needs no clarification.
4. We claim that reasons for believing a claim should be reasons for thinking that the claim is true. Does that sound right? What other kinds of reasons might someone think are good ones for believing a claim? Suppose your best friend says that she believes she has won a \$10 million lottery. Should you believe her? Suppose that she offers no real reason for thinking that her claim is true, but wants you to believe her anyway (maybe she wants to celebrate). Is that a good reason to believe a claim (that someone *wants* you to believe it)? What about the claim that God exists? On what basis should someone believe that claim? Should you have some reason to think that the claim is true? Suppose that it makes you happy to believe that God exists. Is that a good reason to believe the claim? Think of some examples of claims that you think someone ought not to believe unless there is a good reason to think that they are true, and some examples of claims that you think can *perhaps* be believed on other grounds. You will obviously have to explain the context of the latter examples.

ⁱ **ST** What do you think philosophy is? Have you read anything that you would call philosophical? Have you had any discussions that you would regard as philosophical? Is there anyone in your family who is especially philosophical?

ⁱⁱ **ST** Place each of the questions in the previous paragraph into one of these three general categories.

ⁱⁱⁱ **ST** There is another, related fallacy called *denying the antecedent*. You should be able to figure out what that mistake is and why it is a mistake, given this explanation of the fallacy of affirming the consequent.

^{iv} **ST** Can you think of any other objections to this argument?

^v **C1** Another important philosophical skill is giving counter-examples to philosophical claims or theories. To create a counter-example, you need to first figure out just exactly what the claim or view says and then think carefully about how it applies to many situations, looking for examples that show that it is wrong.

^{vi} **ST** Think about this question on your own for a minute or two.