

# Deductive Arguments

Dr. Jake Wright<sup>1</sup>

University of Minnesota Rochester

## 1. Introduction

Philosophy is built on arguments. This isn't to say that no other subject is built on arguments; you will be expected to argue for your thesis when you write Sociology or History papers, for example. But when we say that philosophy is built on arguments, we mean something very particular; philosophy is especially interested in examining how good or bad a particular argument is by examining both the truth of one's claims and the way that those claims fit together. While this isn't all that philosophy does, it does capture the primary goal of one important area of philosophy—logic.

We're going to look at logic as part of our discussion of metaphilosophy because it serves as one of the foundations of good philosophy. If we want to discover truths about the world or even if we just want to examine our beliefs and determine whether those beliefs are justified, one of the things we need to know is whether or not our reasons fit together in the right sort of way.

As we'll see, there are many ways our reasons can fit together. Broadly speaking, these ways can be expressed in certain types of arguments. We can group these kinds of arguments into two main groups. The first, which we're going to spend most of our time discussing today, are called *deductive arguments*. The second, which we won't discuss in as much detail, are called *non-deductive arguments*. Not the most creative name, but there you go.

With each type of argument, I want to answer three rough questions. First, how does this type of argument work? In other words, how do we separate the good arguments from the bad? Second, what do these arguments try to show? Another way of asking this question is to ask what kinds of conclusions each type of argument tries to make. Finally, when would I want to use each type of argument? Since each type of argument has different criteria for goodness or badness, different types of conclusions you can reasonably reach, and so forth, the type of argument we use will depend a lot on the sort of question we're trying to answer.

I'll do this by first saying a bit about the role of reason in philosophy. You may take the role of reason as fairly obvious, but it's important to take a moment to realize that we often base our beliefs and actions on things other than reasons. This will not only help us better understand why reasons are important, but also the fact that reasons aren't always the best or most effective way to motivate people. After this, I'll discuss the two types of arguments in Section 3. Section 4 goes into more depth on two important features of deductive arguments—*validity* and *soundness*. Since these concepts are both

---

<sup>1</sup> Prepared for Introduction to Philosophy (Fall 2018) at the University of Minnesota Rochester. This document is licensed under a Creative Commons ([CC BY-NC-SA 4.0](https://creativecommons.org/licenses/by-nc-sa/4.0/)) license.

somewhat complicated and are the cornerstone of deductive reasoning, it's worthwhile to make sure we really understand them.

## 2. The role of reason in philosophy

You might ask yourself why we should care about how our reasons fit together. This is an excellent question! The ancient Greek philosopher Aristotle<sup>2</sup> gave us some persuasive arguments about why we should base our beliefs on reasons in his study of rhetoric, or how we use speech to convince people to believe us. Basically, Aristotle said that we try to convince each other to believe things in one of three ways. Of these three ways, Aristotle believed that using reasons was the best.

The first way we try to convince other people to believe certain things or act a certain way is by appealing to authority. For example, ask yourself why you brush your teeth. I bet your answer will be that brushing your teeth keeps them healthy and cavity free. But why do you believe that? My guess would be that your dentist, an authority on oral health, told you so and not because you've run a series of randomized control trials studying the connection between brushing and oral health.

There are many reasons to like appeals to authority. One of the biggest reasons to like appeals to authority is that such appeals make it a lot easier for us to make quick decisions. For example, if a doctor in the Emergency Room says that you need to be rushed into surgery immediately or else you'll die, I doubt you'll stop and demand to have the doctor explain their reasoning in detail.

Simply accepting the claims of an authority figure creates an interesting problem, though. How do we know which authorities are good ones? Presumably, we know because we trust the authority has good reasons for their beliefs. The ER doctor wants to rush me into surgery because of a medical problem and not out of a desire to be paid for an unnecessary surgery, for example. But if this is the case, then our trust in authority is only as good as the reasons used to back up that authority, which means that reasons are what are ultimately important. In fact, we can see this with our dentist example very clearly. My dentist has always told me to floss, and her authority as a dentist has given me very good reason to floss. But some recent research suggests that the evidence showing that flossing is beneficial is weak and of low quality.<sup>3</sup> When determining whether I should floss, it seems like reasons and evidence should carry more weight than the fact

---

<sup>2</sup> Aristotle lived in the 4<sup>th</sup> century BCE in what is today Greece. He studied philosophy at Plato's Academy and founded the Peripatetic School (often called the Lyceum, after the temple it was located in). His students included Alexander the Great. He wrote on topics ranging from biology to ethics to politics to art. The ideas we're discussing here are largely from his *Rhetoric*. None of his original writing remains; what we have is most likely a collection of notes prepared by Aristotle's students that have been translated from ancient Greek to Arabic (in the early Middle Ages) to Latin (mostly in the Late Middle Ages, though some weren't translated until the Renaissance) and eventually into modern languages. The Catholic Church banned many of Aristotle's books during a period where Greek thought in general (and Aristotle in particular) was very influential in the Islamic world, which is a major reason why the translation history is so convoluted. Essentially, without the intervention of medieval Muslim scholars, we would have lost Aristotle to history.

<sup>3</sup> Here's a link to one such study: <http://www.ncbi.nlm.nih.gov/pubmed/19138178>

that my dentist told me so, and if there was enough evidence that flossing was not beneficial, it seems like I would be totally justified in ignoring my dentist and not flossing.

Aristotle's second way that we try to convince each other involves appeals to emotions. For example, you've probably seen those commercials raising money to prevent animal abuse. It's a solid minute (sometimes two) of sad animal pictures with a weepy Sarah McLachlan soundtrack, followed by some celebrity asking you to send money now. These commercials are very sad. In fact, they're so sad that my wife always changes the channel or leaves the room when they come on.

Why do charities break out the sad puppy dogs and Sarah McLachlan when they want you to give? Because it's effective! Seeing all of that heartbreak makes us want to give and make the bad feelings go away. But once again, our desire to act based on something other than our reasons leads to some interesting problems.

Most notably, we recognize that it's not always best to act based on our emotions. For example, I like to think of the scene from *Bridesmaids* where Melissa McCarthy adopts a litter of puppies because they just looked so cute. Later, she shows up at Kristen Wiig's house and admits that maybe she let her emotions get the best of her.

But how can we know whether emotional decisions are good ideas? Just appealing to emotions leads to problems. For example, if we only appealed to emotions, then it was good for Melissa McCarthy to adopt all of those puppies when she felt emotionally compelled to do so, but a bad idea when her emotions changed. It seems like there's something wrong with that, if our actions can switch from justified to unjustified so easily. Alternatively, we could use reasons to figure out whether an emotional decision is the right one. Just get dumped? Maybe now isn't the best time to suddenly decide to shave your eyebrows. Just fail a test? Maybe now isn't the best time to decide to drop the class. But then again, maybe it is! Whether it is or not, though, seems to depend on your reasons. Just like appeals to authority, we might say that appeals to emotion motivate us to act a certain way or hold certain beliefs, but whether those actions or beliefs are the right ones ultimately depend on our reasons.

So, Aristotle thinks that the best way to motivate us to act is by appeal to reason. It's the only motivation that doesn't depend on something else for justification. In fact, it may be useful to remind ourselves that the other two ways, authority and emotion, ultimately rested on reason as their justification for action. More specifically, Aristotle thought that we try to convince each other by applying clear, meticulous reasoning. Brian Monty, one of the other UMR philosophers, puts this slightly differently when he describes philosophy as "thinking in slow motion." We usually think before we act (I hope), but those thoughts are often quick and not connected in especially clear ways.

For example, if I'm running late and haven't eaten breakfast, I might decide to get Dunkin' Donuts. If you asked me why I decided to do this, I could give you some reasons, like the fact that donuts are delicious, Dunkin' Donuts is close to U Square, and I feel hungry after three classes if I haven't eaten. All of this is true, but in the moment I decide to get a donut, I'm probably not connecting these reasons in a clear, meticulous

way. That's fine as far as donuts go, but when it comes to the truly important questions in life, it seems like I should take care to make sure that my reasons really do fit together in the right sort of way. As much as we can, we try to put emotions and other factors aside and decide based on nothing but the evidence.

There's a lot to recommend basing our beliefs on evidence and reasoning. For one thing, it seems like appealing to our reasons gives us a firmer foundation for our beliefs. For another, it seems like our other techniques for determining what to believe are ultimately supported by reasons. For these and other reasons, philosophers are especially interested in examining the reasons behind our beliefs, rather than the authorities that make claims or the emotions that spur us to action.

I should note that, for all of the benefits of appealing to clear, meticulous reasoning, this strategy isn't perfect. In fact, reasons and evidence often fail to motivate our actions and can serve as a powerful *disincentive* for action when reasons run counter to our closely held beliefs. This is a psychological phenomenon known as the *backfire effect*.<sup>4</sup> So when we actually try to convince people, it may be most productive rhetorically to appeal to authority or emotions or something other than cold, hard facts. But what's important for our purposes here is the recognition that these other rhetorical strategies ultimately get their justification from reasoning, and justification is what we're interested in at the end of the day in philosophy.

### 3. What is an argument?

There are some problems with relying on reason, including the fact that basing our beliefs solely on reason is incredibly hard. It might even be impossible. But at the very least, it gives us a goal to aspire to and a standard against which we can measure our arguments. This is why we're interested in examining what makes an argument a good argument.

In order to answer questions about how good an argument is, though, we're going to need to talk about some basic concepts. The first concept we need to talk about is the argument itself. When we talk about arguments in philosophy, we mean something very specific. We'll define an argument as follows:

An *argument* is a connected series of statements, called *premises*, that are intended to establish a claim, called a *conclusion*.

Basically, the goal of an argument is to use premises to show that the conclusion is either true or very likely to be true. Whether an argument is a good argument or not depends on what kind of argument it is how and well it achieves its goal.

Most of the arguments you'll encounter in your everyday life, along with most of the arguments we'll study in this class, are in what we might call the *vernacular form*.

---

<sup>4</sup> The backfire effect is fascinating, especially when you think about what we're trying to do in a philosophy class. If you want a quick overview, here's a clip from *Adam Ruins Everything*, one of my favorite TV shows, explaining the phenomenon. <https://www.youtube.com/watch?v=Q8NydsXl32s>

*Vernacular* is just a fancy word for how we communicate every day. For example, we usually don't say, "Smith was discovered by a faculty member to have been engaging in academic dishonesty." We say, "A professor caught Smith cheating." So, arguments in the vernacular form are just arguments that are written out as paragraphs, essays, and so forth, rather than being presented formally. This will be made clearer below.

While arguments in the vernacular form are easy to read and can be natural to write, we will often want to rewrite the argument in what we will call the *standard form*. Arguments in standard form list each of the argument's premises individually, along with individually listed, clearly identified conclusions. For example, we might take the following vernacular argument and present it in standard form as follows:

*Vernacular argument*

Today, I learned that if someone is a professor at UMR, then that person is also a state employee. As it happens, I know that Wright is a UMR professor. So it must be the case that Wright is also a state employee.

*Standard form argument*

P1: If someone is a UMR professor, then that person is a state employee.  
P2: Wright is a UMR professor.  
C: Therefore, Wright is a state employee.

Note that we have taken all of the reasons in the vernacular argument and presented them as individual premises in the standard-form argument. Note also that we have taken the conclusion and listed it separately underneath all of the premises that are supposed to give us good reason to accept the conclusion. Reframing arguments in standard form can be tedious at times, but as we'll see throughout the semester, it can make it much easier to separate good arguments from bad ones. More importantly, it can help us see how different claims relate to one another and help isolate claims that we think are suspect or in need of additional justification.

Philosophical arguments can broadly be divided into two types, which we call deductive and non-deductive arguments.

3.1. *Deductive arguments*

The first type of argument we're going to look at is called a *deductive argument*. Deductive arguments take premises and try to establish conclusions that are certainly true, which is to say that the argument's conclusion could not possibly be false. For example, the argument that Wright is a state employee, above, is an example of a deductive argument.

If the premises are true, then there just is no escaping the conclusion; Wright would certainly be a state employee. We measure the goodness of deductive arguments according to two criteria—*validity* and *soundness*. We'll talk about what each of these criteria means below.

### 3.2. *Non-deductive arguments*

The second type of argument we'll take a look at are *non-deductive arguments*. Unlike deductive arguments, non-deductive arguments don't try to show that a conclusion is certainly true. Instead, non-deductive arguments try to show that the conclusion is probably true. Another way of looking at non-deductive arguments is to say that the goal of such arguments is to give us good reason to believe the conclusion. We measure the goodness of non-deductive arguments in terms of *strength* and *weakness*. The more reason the argument gives us to accept the conclusion, the stronger the argument is.

While there's only one kind of deductive argument, there are a few kinds of non-deductive arguments.

*Inductive arguments* use evidence from a small group to make conclusions about a larger group. For example, randomized control trials are examples of inductive reasoning. They take a small number of cases and use the information from those cases to predict what the effect of a drug or procedure will be for everyone.

*Inference to the Best Explanation* compares two potential explanations for a particular observation and tries to ask which explanation would make the observation less surprising. Scientists often depend heavily on this kind of argument. For example, Galileo used Inference to the Best Explanation to demonstrate that the Earth revolves around the Sun.

*Argument from Analogy* takes two phenomena and compares them to determine if one phenomenon has an otherwise unobservable property. For example, if an archaeologist on a dig finds a long, curved knife with an intricately carved handle, she might conclude that it was used for a particular purpose (say, religious sacrifice) because other knives with similar properties were used for the same purpose.

There are two important features of non-deductive arguments that distinguish them from deductive arguments. First, the best we can say about non-deductive arguments is that their conclusions are probably or almost certainly true. Second, the conclusion of a non-deductive argument introduces new claims that aren't present in the premises. Let's take a minute to see why this is the case.

Consider an inductive argument you make every day. The sun will come out tomorrow. This is a claim I'm sure you are very confident about. In fact, I'm confident you'd bet your bottom dollar on this claim.<sup>5</sup> But let's consider the structure of the argument you'd make for this conclusion.

P1: The sun came up today.

P2: The sun came up yesterday.

---

<sup>5</sup> #dadjoke

P3: The sun came up two days ago.  
P<sub>n</sub>: The sun came up  $n$  days ago.<sup>6</sup>  
C: Therefore, the sun will come up tomorrow.

If we look at this argument closely, we notice that the claim made by the conclusion that the sun will come up tomorrow, isn't contained in any of the premises. As we'll see below, this is very different from the conclusion of a deductive argument. So, the conclusion introduces a new claim that isn't directly supported by any of the premises.

We also notice that, since the conclusion isn't directly supported by any of the premises, it's possible that the conclusion could be false, even if all of the premises are true. To see why this is the case, consider two famous examples.

The first example is sometimes known as Russell's Turkey, after the British philosopher Bertrand Russell,<sup>7</sup> who popularized it. Suppose that a farmer is raising a turkey to be butchered and served at Christmas dinner, as is customary in Brittan. Every morning, the turkey wakes up and sees the farmer enter the coop, and every morning, the farmer feeds the turkey. After many months, the turkey concludes that each morning, the farmer will enter the coop and feed him, and every morning, the turkey is correct. That is, the turkey is correct every morning except Christmas Eve, when the turkey gets a very rude introduction to the fact that non-deductive arguments don't guarantee the truth of their conclusion.

The second, somewhat less gruesome example, is known as the Black Swan Problem. For centuries, Europeans believed that all swans were white. This was a reasonable conclusion, since the only swans Europeans had ever seen were white, and these observations dated back centuries. However, as Europeans explored the globe, they made a startling discovery in Australia, where the swans were black. Despite centuries' worth of evidence, where every data point produced a white swan, the claim "All swans are white" was false.

At this point, you might wonder why anyone would use non-deductive arguments if their conclusions could be wrong like this. The answer has to do with the fact that their conclusions contain new information. Often, as we'll see below, we don't have enough evidence to present a good deductive argument that gets us the conclusion we're after.

---

<sup>6</sup> Here,  $n$  just means that the premises can be extended as far as we want, and  $n$  represents the number of premises we add (e.g., P365: The sun came out 365 days ago).

<sup>7</sup> Bertrand Russell was a late-19<sup>th</sup>/20<sup>th</sup> century philosopher famous for his work in logic and mathematics. In addition to living a very colorful personal life that saw him imprisoned for refusing military service during World War I (Russell was a pacifist, though he supported British involvement in World War II) and leading anti-nuclear marches, he is most well known for writing *Principia Mathematica* with Alfred North Whitehead. The *Principia* tried to prove all of mathematics without taking any axioms, or basic assumptions. This project was so complicated that Russell and Whitehead were only able to demonstrate that  $1+1=2$  on page 86 of *volume two*, with the cheeky footnote that "The above proposition is occasionally useful." Their entire project was overturned by German logician Kurt Gödel, who wrote a paper showing that mathematics is impossible without axioms.

For example, if I wanted to deductively prove that all swans were white, I would have to go observe every swan that ever existed. This would be quite impractical. Instead, I might observe a number of swans and use my observations about those swans to draw larger conclusions about all swans. I risk the possibility of being wrong, but I gain the ability to make claims that are likely to be true more efficiently. You might recognize that this is how most scientific disciplines work. Scientists take a small number of observations and use those observations to make general claims about much larger groups of individuals or events. You'll also use this kind of reasoning in statistics, which is basically a class in how to create high quality inductive arguments.

There's a lot more I could say about non-deductive arguments, but we need to make sure we have time to discuss validity and soundness in depth. So, let's move on.

#### **4. Analyzing deductive arguments**

Let's quickly recap where we've been. Philosophers, especially logicians, are very interested in understanding how our reasons fit together to give us justification for believing and acting certain ways. We examine our reasons by examining arguments for our beliefs, and these arguments come in two forms—deductive and non-deductive. Deductive arguments, which we're looking at closely, try to use premises to show that the conclusion must be true, no matter what. Now, let's take a look at deductive arguments and how we determine if they are good arguments.

We said above that there are two criteria for determining goodness in deductive arguments. The first is *validity*, and the second is *soundness*. This section is going to examine these concepts more closely.

##### *4.1. Validity*

We use *valid* in our everyday lives when we say things like, "That's a valid point." Usually, we mean something like, "That's a good point." In philosophy, though, we mean something more precise when it comes to validity. To see what we mean, consider the following argument:

- P1: If the sun is shining, then it is not night.
- P2: The sun is shining.
- C: Therefore, it is not night.

We want to know whether this is a good argument. First, I want you to notice something interesting about the premises; I have no idea whether they are true or not. This is because I'm writing the argument at a different time from when you are reading it. Maybe the sun is shining as you read this, maybe it's not. Maybe you're in a room with no windows, and you have no idea whether the sun is shining or not. As it turns out, none of this matters when it comes to validity.

All we need to do in order to determine whether an argument is valid or not is to determine how we would view the conclusion *if the premises happened to be true*. In other

words, all we have to do is ask whether the conclusion would have to be true if the premises were true. Suppose we do this for the argument above. In other words, suppose that if the sun is shining, then it can't possibly be night. Further suppose that the sun is shining right now, even if it's dark out or raining where you are right now. If we accepted both of those statements as true, what would we say about the conclusion? We would be forced to admit that the conclusion was true! It would be impossible for us to admit that the sun is shining, the sun shining meant that it was not night, and that it was actually night. This means that the argument above is valid. In this class, we will define validity as follows:

An argument is *valid* if and only if accepting all of the premises as true would require us to also accept the conclusion as true.

Arguments that are not valid are automatically *invalid*. Basically, invalid arguments are arguments where we could accept all of the premises as true and not be forced to accept the conclusion as true. Here's another way to think about invalidity. If you can think of a counterexample to the argument, then the argument is invalid. For example, consider the following argument:

P1: If something is a chicken, then it is a bird.  
P2: That is not a chicken.  
C: Therefore, that is not a bird.

Let's run our test for validity and see what happens. First, let's suppose that all of the premises are true. (Remember, the premises don't actually have to be true; we're just pretending!) If it were the case that all chickens are birds, and if it were the case that this particular mystery object, whatever it is, is not a chicken, does it follow that it cannot possibly be a bird? No! There are plenty of objects that are not chickens, but are still birds. For example, eagles, pigeons, and peregrine falcons are all birds, even though they are not chickens. In other words, eagles, pigeons, and peregrine falcons are all counterexamples to the claim that our mystery object cannot possibly be a bird. There is at least one circumstance where all of the premises are true, but the conclusion is false. So, the argument must be invalid.

So if an argument is valid, does that mean that the argument is good? In a certain sense, absolutely. If we have a valid argument, we at least know that the individual reasons for believing the conclusion fit together in the right sort of way, which is an important first step. But validity is only the first step when it comes to deductive arguments. This is because philosophers don't just want to know if reasons fit together in the right sort of way; we also want to know if we should actually accept the conclusion as true. In order to know the answer to that question, though, we need to examine the concept of soundness.

#### 4.2. *Soundness*

If all we were interested in was logic, we wouldn't need to worry about whether the premises in an argument were true or false. In fact, if all we were interested in was logic,

we would be so unconcerned with the content of our premises and conclusions that we would replace that content with variables, like in algebra. But since we are good philosophers who are interested in discovering truths about the world, it's not enough to look at an argument, figure out if it's valid, and move on with our lives; we need to know whether we really ought to accept the argument's conclusion as good. For example, consider the following two arguments:

P1: If someone is a UMR professor, then that person is a state employee.  
P2: Wright is a UMR professor.  
C: Therefore, Wright is a state employee.

P1: If someone is God-King of Earth, then everyone should obey that person.  
P2: Wright is God-King of Earth.  
C: Therefore, everyone should obey Wright.

Notice that both arguments are valid. If we accepted all of the premises as true, we would have to accept the conclusions as true, as well. But that doesn't mean that everyone should obey my every command. Why not? Because it turns out that in the real world, I am not God-King of Earth. In other words, (P2) of the God-King argument is false. And since (P2) is false, we don't have a good reason to accept the conclusion as being true.

On the other hand, we do have good reason to believe the conclusion of the UMR professor argument. The argument is valid, which means that all of the reasons fit together in the right sort of way; if the premises were actually true, then we would also have to accept the conclusion as true. Unlike the God-King argument, though, the premises of the UMR professor argument really are true. This means that we cannot escape the fact that I am actually a state employee. This conclusion must be true!

The fact that a sound argument's conclusion must be true is the defining feature of such arguments. But how do we know that the conclusion must be true, no matter what? There are two important features of sound arguments that force us to accept the conclusion as true.

First, the argument's premises have to fit together in the right sort of way, so that we would have to accept the conclusion as true if the premises were true. But, you'll notice, this is just the definition of validity. *This means that all sound arguments are valid arguments!* In other words, if you have an invalid argument, you can't possibly have a sound argument; that combination is impossible. This is the case even if all of the premises are actually true. In fact, this is the case even if the conclusion really is true, as well! To see why, consider the following argument:

P1: The most common element in the atmosphere is Nitrogen.  
P2: South America is south of Mexico.  
C: Therefore, Wright is less than ten feet tall.

All of the premises of this argument are true, but the argument is obviously invalid. There's nothing about how much Nitrogen is in the atmosphere or the location of South America relative to Mexico that requires me to be less than ten feet tall, even though this is the case. So, we would say that there's nothing about the argument that requires me to accept the conclusion as true. Whatever the reason is for my height, it has nothing to do with the reasons offered in the argument.

The second feature of a sound argument has already been discussed. Sound arguments aren't just valid arguments, they're arguments where the premises really are true. If the argument is valid, which means that we must accept the conclusion as true if the premises really were true, and if the premises actually are true, then the conclusion must be true, no matter what.

In the end, we can define soundness as follows:

An argument is *sound* if and only if the argument is valid and all of the argument's premises are true.

Is a sound argument a good argument? Absolutely! This is because sound arguments allow us to know, with certainty, that the argument's conclusion is true. If the ultimate goal of philosophy is coming up with good reasons for our beliefs and actions, what reason could possibly be better than certainty? This is why the goal of all deductive argumentation isn't just to come up with valid arguments, but to present sound arguments. This is also why we'll spend much of our time in this class when we discuss deductive arguments trying to figure out if particular premises are true or false; if you have a valid argument and you don't like the conclusion, you only have two choices. You can find a false premise, or you can admit that the conclusion must be true, even if you don't like it.

Before we end, I want to make a special note when it comes to validity, soundness, and truth: If you reread this section, you'll notice that I'm being very precise when I use certain terms like *truth*, *validity*, and *soundness*. In philosophy, concepts like validity and soundness are properties of arguments, and only arguments. Similarly, concepts like truth and falsity are properties of individual statements, like premises and conclusions, and only individual statements. Thus, we would never say that an argument is true, and we would never say that a premise is valid. This is a small but very important point that you'll need to master quickly!

## **5. Conclusion**

Before we can start examining philosophical issues like whether God exists or whether we have free will, we need to have some rules that let us figure out when our beliefs are justified. In philosophy, we think our beliefs are justified when we have good reasons to think that our beliefs are true, and the way we figure out whether our reasons are good reasons is by examining the arguments that support our beliefs.

Philosophical arguments come in two types—deductive and non-deductive arguments. Today, we looked at what makes these arguments good.

For non-deductive arguments, we saw that no matter how good the argument is, we can never guarantee that the conclusion will be true. We evaluate the strength of non-deductive arguments by judging how likely the conclusion is to be true, given the premises. Even though we can't guarantee the truth of a non-deductive argument's conclusion, there are still a lot of circumstances where we find such arguments useful. Most notably, we can use non-deductive arguments to draw broad conclusions without having to do the work required to make all of the observations that would guarantee the conclusion's truth, either because this would be too much work or because that sort of data gathering would be impossible.

Deductive arguments are powerful because, if they are sound, the conclusion must be true no matter what. Deductive arguments have two criteria for goodness. One criterion is validity. We defined valid arguments as arguments where we would have to accept the conclusion as true if we also (hypothetically) accepted the premises as true.

Valid arguments are important, but as philosophers, we're really searching for sound arguments when we use deductive reasoning. Sound arguments are valid arguments where the premises really are true. Sound arguments are important because they have conclusions that must be true, no matter what.