Philosophical Arguments

An introduction to logic and philosophical reasoning.

Nathan D. Smith, PhD.

Houston Community College

Table of contents

    e. Begging the Question
    f. False Dichotomy
7. Induction
    a. Generalizations
    b. Weak Generalizations
    c. Strong Generalizations
    d. The Problem of Induction
8. Arguments about Causes
    a. Correlation is not Causation
    b. Advice in Making Causal Arguments
9. Arguments from Analogy

Philosophical Arguments

In philosophy, we evaluate philosophical claims based on the arguments that can be made in support of them. Sometimes arguments appeal to empirical or objective evidence, sometimes intuitions, principles, or beliefs, but philosophical arguments always use some sort logical reasoning to support the claims they make.

In the following chapters, we will focus on the nature of logical reasoning and how it can be used in philosophical arguments.

The Logic of Arguments

All arguments are composed of **sentences**. Sentences are statements that include a **subject** and a **predicate**: the predicate describes the subject, while subjects refer to things or ideas. Sentences can be empirical observations, statements of belief, or statements of principle. In order for a sentence to be used in a philosophical argument it must have such a form that it could possibly be either true or false. We do not have to know whether it is true or false, but it has to be the kind of thing that could be true or false. This property of sentences is what allows them to be used in philosophical arguments; it is central to the underlying structure of logic.

There are two kinds of sentences in philosophical arguments: **premises** and **conclusions**. Premises and conclusions can appear in any order, but when we write out arguments in what is called "cannonical form," we always write the premises first and then the conclusion last. Conclusions can become the premises of further arguments. So, more complex arguments may contain multiple preliminary conclusions before they reach their final conclusion. The conclusion is the claim that the argument is intended to support (this is similar to a thesis in a argumentative paper). The premises provide the support or evidence for the conclusion.

If the premises provide adequate support for the conclusion, then the argument is either **strong** or **valid**. However, it is important to recognize that even strong or valid arguments can lead to a false conclusions if the premises are not true. In other words, no matter how good your reasoning is, if the evidence you start with is faulty, then the conclusion will be faulty, too.

Sentences are either true or false; arguments are valid or invalid.

Validity and Soundness

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

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

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

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

Assessing Validity

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

Sometimes we say that validity is "truth-preserving." In other words, if you start out with true premises, then a valid argument will preserve truth, so that you will wind up with a true conclusion. But if you start out with false premises, a valid argument may preserve that falsity conclusion, so you can’t be sure if the conclusion is true only based on the premises.

Consider the following two arguments:

All human beings are mortal. Socrates is a human being. -------------- So, Socrates is mortal.	All apples are delicious. This is an apple. ------------- So, this is delicious.

Both of these arguments are valid. If you assumed the truth of the premises, you would have to accept the conclusion. In other words, if the premises were true, the conclusion would have to be true. In fact, they have the exact same form or structure, so one cannot be valid while the other is invalid.

Of course, in the case of the second argument about apples, the first premise is not true (i.e., all apples are not delicious). And, as a consequence, the conclusion does not necessarily follow. But the argument is valid because if the premise were true, the conclusion would necessarily follow.

Counterexamples

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

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

> If all apples are delicious and this is an apple, then this is delicious!

A counterexample would be an example that shows the statement to be false. We can tell that this is a valid argument because there is no example that would make this statement false. This is because any non-delicious apple would make the first part of the sentence (the 'if-clause') false. But an if-then statement is made false only if the if-clause is true AND the then-clause is false. An if-then statement only says that *if* the if-clause is true, *then* the then-clause must be true.

Consider the following argument:

> When it rains, the streets are wet.
> The streets are wet this morning.
> ------------------
> So, it rained last night.

Now consider it as a sentence:

> If it is true that when it rains the streets are wet and the streets are wet, then it rained.

Can you think of a counterexample? Try to think of some scenario that makes the if-clause true, but the then-clause false.

In order to do this, assume that the if-clause is true and the then-clause is false. What scenario would make the if-clause true but the then-clause false?

Assessing Soundness

Now that we have discussed validity, we need to be able to test for soundness.

Remember, validity is a claim about the structure or form of the argument: a valid argument is such that *if* the premises are true, then the conclusion *must* be true.

Soundness, on the other hand, requires that the argument is valid *and* that the premises are true. So, in order to assess for soundness, we need first to assess for validity and then to assess the truth of the premises.

So, how do we test the truth of philosophical claims?

Two Theories of Truth

The question of what makes a philosophical claim true is a serious philosophical issue that we can't spend adequate time on here. There are many different theories of truth, but I will present two historically prominent theories of truth and a bit of a rationale for why you might think that this theory of truth is the right one.

First, remember that in arguments, sentences can be true or false while arguments are valid or invalid, sound or unsound.

> *Correspondence*: A sentence is true if and only if it accurately represents some state of affairs.

This is perhaps the most natural theory of truth. It seems right to say that something is true if and only if it corresponds to some actual state of affairs. But sometimes it is difficult to tell which states of affairs a sentence is supposed to correspond to. For instance, what is the state of affairs corresponding to the statement 'physical objects are the cause of my sensations'? Or, what is the state of affairs corresponding to the statement 'murder is wrong'?

> *Coherence*: A sentence is true if and only if it is logically consistent with a set of other sentences.

This view of truth is surprisingly strong. In fact, we can imagine coherence substituting for correspondence simply by recognizing that in order for us to know about states of affairs, we have to make statements about them. More importantly, coherence allows us to determine the truth of things like mathematics or morality that do not have a clear reference to states of affairs.

Assessing Truth

Whatever may be said for what makes a sentence true--be it correspondence or coherence, or something else--it's important when making claims of fact that are supposed to support the conclusion of an argument, that one is prepared to support the premises too. That is, arguments should be built from relatively well-founded premises: claims for which there is evidence, of one sort or another. When formulating arguments on premises drawn from your own beliefs, it will sometimes help to ask yourself: How do I know this is true? Can I support this claim if I am asked to? When you can do this, you are well on your way to doing philosophy.

Logical Operators Natural-Language Sentences Defined

So far, we have considered logic from the perspective of English-language sentences. Of course, we could translate any of these sentences into another language, such as French or Mandarin Chinese, and we would expect that the basic logic of the sentences would be preserved. (Anthropologists have discovered some aboriginal language for which this is not possible, but we will set these aside for the moment.) The ability to translate the logic (and meaning) of one sentence into another sentence is a key component in the universality of claims in philosophy. Untranslatable concepts are rare and interesting, but even when a particular term cannot be translated in exacltly the same way in another language, it is almost always possible to convey the meaning from one language into another. This is a central feature of all translation.

However, when we translate from one language to another language we can be careful to preserve logic and meaning, but there are inevitable shortcomings that are the result of the vagueness and ambiguity inherent in natural languages. Natural languages are those languages (like French, English, or Mandarin Chinese) that have evolved naturally within a particular linguistic community. Though these languages have an underlying grammar and their words have meaning (we can see this in grammar books and dictionaries), the grammar and meaning is not designed, but rather emerges naturally. Like most structures that emerge naturally, natural languages can be beautiful, subtle, complex, and evocative, but they are also often messy and unclear. In fact, the very same property of language that makes poetry possible is what makes logic difficult.

Consider the following poem by Theodore Roethke:

My Papa's Waltz

The whiskey on your breath
Could make a small boy dizzy;
But I hung on like death;
Such waltzing was not easy.

We romped until the pans
Slid from the kitchen shelf;
My mother's countenance
Could not unfrown itself.

The hand that held my wrist
Was battered on one knuckle;
At every step you missed
My right ear scraped a buckle.

You beat time on my head

With a palm caked hard by dirt,
Then waltzed me off to bed
Still clinging to your shirt.

Is this a poem recalling a nostalgic memory of dancing with a hard working father who had a bit too much to drink? Is it a terrified and troubling poem about a drunken father who would beat his son when he returned from the bar after work? The poem is intentionally ambiguous. Roethke uses words like "beat time," which could indicate the *beats* of music, but could also mean *beating* in terms of corporal punishment. Again, the image of "My right ear scraped a buckle" is unclear: was it intentional or the result of a whipping? There is a deep truth here: after all, even an abusive father is still a dad. Despite the terror of possible anger, even an abused child probably has fond memories of his or her parent. The poem manages to evoke all of these truths, these emotions in a few simple stanzas, using the fluidity, ambiguity, and vagueness of the English language.

While poetry is a great example of how an author can exploit the ambiguities of language in order to bring out subtle truths, it is necessary to remove ambiguity when we do logic. Exactly this same feature of lanaguage – it's fluid capacity to refer to multiple different scenarios at the same time – is a reason why we have to move away from natural language in order to do logic in the most precise way. In this chapeter, we will begin to gain some precision in our logic of sentences by defining the components of sentences and their operators. This is a precursor to actually symbolizing sentences using symbolic form.

Sentential Components

When we begin to symbolize sentences, the first thing we need to do is to analyze the sentence into its component parts. Imagine that sentences are themselves composites of many different component parts. These components work together to generate a meaning, but the components themselves have a meaning on their own. You get the sense of this in natural language when you study vocabulary and grammar. But what is the underlying logic?

For the purposes of this discussion, I will assume that we are doing sentential logic. First order predicate logic will allow us to consider different types of sentential components and other types of logic (for instance, the logic of propositional attitudes, modal logic, or second order predicate logic) will allow still other types of sentential components. Different types of logic have the capacity to assign meaning to parts of sentences that have no meaning in other types of logic.

In sentential logic, all sentences have the property of being either true or false. The basic meaning of any sentence in sentential logic is its truth or falsity, its "truth value." When we analyze a sentence into its parts in sentential logic, we get three different kinds of components: atomic sentences, sentential connectives, and other stuff. Atomic sentences, are the smallest components of a sentence that have the property of being either true or false. Connectives link atomic sentences in such a way that the combination of atomic sentences with sentential connectives alters the meaning (truth or falsity) of the whole sentence. The “other stuff” can be discarded or it can be lumped with the atomic sentences. Basically, “other stuff” are the words that function as transitions, intensifiers, or some other rhetorical function that doesn’t alter the underlying meaning of the sentence. Consider the following sentence:

Sometimes when the weather is dreary, I get a headache and feel tired all day.

Try to analyze this sentence into its atomic components, sentential connectives, and other stuff. Start with atomic sentences. What are the smallest components of this sentence that are either true or false?

1. The weather is dreary.
2. I get a headache.
3. I feel tired all day.

Notice that I had to slightly alter #3 from its form in the original sentece, but not in a way that altered its meaning. How are these three components connected?

1. When _________, then ___________.
2. _________ and ___________.

Again, I slightly altered the sentence in #1 in order to make its meaning clearer, but I haven’t altered the underlying meaning. The important thing to see is that there are two different sentential connectives at work here (the ‘when’ and the ‘and’). Is there anything left over?

Sometimes

Right. So, this last word is not a meaningful part of the sentence in sentential logic. What ‘sometimes’ means is that there is a certain set of times for which the sentence is true and a certain set of times for which the sentence is false. ‘Sometimes’ tells us that the sentence is true ‘not all the time’ or ‘not at any given time’. ‘Sometimes’ is a quantifier (it tells us how much or how often). Quantifiers are part of predicate logic, so we can just ignore them for the purposes of sentential logic.

Logical Operators

Sentential connectives can be thought of as operators in the logic of sentences. An operator performs a function on given inputs. That is, it takes those inputs from some domain and maps them onto some range. Logical operators in sentential logic take sentential components (that are either true or false) and maps those components to either the true or the false. Let’s use our sentence above and see how the logical operators are used. In order to see how the logical operators function in the sentence, we have to assume truth values for our atomic sentences.

Let’s assume that #1 is False, #2 is True, and #3 is True. Now, let’s plug these values into the logical operators:

> When (False), then (True) and (True).

How should we understand the entire sentence? Is it true or false? It turns out to be true, but in order to understand this answer, we need to look at the function of each of the logical operators in the sentence. There are five logical operators in sentential logic and we will look at each in turn: negation, conjunction, disjunction, conditional, and biconditional.

Negation

One of the simplest of the logical operators to understand is the negation operator. Negation takes a sentential component and reverses its truth value: if the sentence is true, then negation makes it false; if the sentence is false, then negation makes it true. This sounds simple enough, but it’s important to understand what negation is and what it is not. Consider the following sentences:

> N1- I am going to the party on Friday.
>
> N2- I have a family dinner on Friday.
>
> N3- I am not going to the party on Friday.

In casual conversation, you might think that N2 is the negation of N3. Imagine a conversation between two friends: “Are you going to the party on Friday?” / “I have a family dinner on Friday.” Obviously, having a family dinner is a reason not to go to the party. But it’s not inconsistent with going to the party. It is still possible for N1 and N2 both to be true. The case is quite different for N1 and N3; there is no possible way for both sentence N1 and N3 to be true at the same time. N3 is the negation of N1; N2 is not the negation of N1.

To make this even more clear, let’s imagine a simple argument using the sentences above.

> N2- I have a family dinner on Friday.
>
> ------------------------
>
> N3- Therefore, I am not going to the party on Friday.

Is this a valid argument?

No. We agreed that it was possible for N1 to be true even though N2 is true. We also agreed that if N1 is true, N3 must be false. So, it follows that it is possible for N2 to be true, while N3 is false. Just imagine a case where I go to the party after I have dinner with my family. As long as it is possible to do both of these things on Friday night, then it is possible for N2 to be true while N3 is false. So, the argument is invalid. This shows that N2 does not entail (and therefore is not logically equivalent) to the negation of N3. So, N2 is not the same as the negation of N1.

Conjunction

Whereas negation operates on a single sentential component (or atomic sentence), all of our other operators will operate on two or more sentential components. The ‘conjunction’ operator is most commonly associated with the word ‘and’, but it can be associated with many other words that serve the same function, such as ‘but’, ‘however’, ‘although’, and so on. What the conjunction operator does is it says that all sentence components joined by the conjunction are true. Here are some examples:

> C1- It’s raining outside and the streets are wet.
>
> C2- I have a car, but I don’t drive very much.
>
> C3- Many Americans follow politics, although not many vote in every election.

Conjunctions are true if and only if all of the conjuncts are true. So, whenever any part of the conjunct is false, the entire conjunction is false. This should actually be pretty intuitive. For instance, if the streets are wet, but it’s not raining outside, then C1 is false. The only case where C1 is true is when it is both raining and the streets are wet. Similarly, C2 is only true when it is true that I have a car and it is also true that I don’t drive very much.

Let’s consider some arguments that employ conjunctions:

> Many Americans follow politics.
> Not many Americans vote in every election.
> -------------------------------
> Therefore, many Americans follow politics, but not many Americans vote in every election.
>
> I have a car, but I don’t drive very much.
> -------------------------------
> Therefore, I have a car.

These are valid arguments and they illustrate the key property of conjunctions: whenever a conjunction is true, it follows that both of the conjuncts are true. So, if we are given two propositions, like in the first example, then we can conclude that the conjunction of the two propositions is also true. Conversely, if we are given a conjunction, as in the second example, then we can conclude that each of the two conjuncts is true.

Disjunction

Disjunction is another logical operator that connects two or more sentence components. In order to get a handle on disjunction, let's take some exmaples from everyday life:

D1- This entree comes with soup or salad.

D2- This semester I can take Math, English, or Philosophy.

D3- An Associates Degree is 60 hours or more of college credit.

These examples are intended to bring out an important ambiguity in our English-language use of the word 'or'. We typically use 'or' to connect a list of possible alternatives. However, it's not always clear if the or is *inclusive* or *exclusive*. In other words, it's not clear whether the statement could be true if both of the disjuncts (both sentences connected by a disjunction) are true. D1 is probably exclusive. When the server says, "This entree comes with soup or salad," this does not mean that the entree can be served with both soup and salad. However, in D3 it is clear that you will receive your Assoiates Degree if you complete 60 hours and more. So, D3 is an inclusive or.

The 'or' operator in English is ambiguous. For logic, we need to remove that ambiguity. In order to do that, let's consider the circumstances under which each of the statements is true. Given a sentence like

D4- The hens are in the coop or the cows are in the field.

there are four possible scenarios:

D4-1- The hens are in the coop and the cows are in the field.

D4-2- The hens are in the coop, but the cows are not in the field.

D4-3- The hens are not in the coop, but the cows are in the field.

D4-4- The hens are not in the coop and the cows are not in the field.

Under which circumstances is D4 true? Let's start by ruling out the last scenario. We can agree that when the hens are not in the coop and the cows are not in the field, D4 is false. Then, let's rule out D4-2 and D4-3 because any way we conceive of the meaning of the 'or' statement, both scenarios will be true. This leaves us with the first scenario. What happens when both the hens are in the coop and the cows are in the field? Is the sentence 'The hens are in the coop or the cows are in the field' true or false?

Let's return to our three examples from the start in reverse order because I think this will be the easiest way to understand the problem.

D3: Suppose I complete 75 hours of college credit. It follows that I have completed 60 hours (on my way to 75) and that I have completed more than 60 hours. Do I have the hours to get an Associates Degree (let's assume – big assumption – that I have fulfilled all my requirements!)? Obviously, yes. So, when both of the disjuncts are true, D3 is true. D3 is an inclusive or.

D2: Suppose that I am taking Math, English, and Philosophy this semester. Is D2 true or false? This is not as clear as the case of D3, but I think it is still obvious that, under this scenario, D2 is true. In effect, D2 is saying that my options for this semester are Math, English, and Philosophy. It is saying that I don't have to select them all, but if I do, then that doesn't negate the fact that they were still my options. So, it looks like D2 is also an inclusive or.

D1: Now, suppose that I order my entree with both soup and salad. Does this make D1 true or false? This is somewhat tricky. I think the easy answer is yes. After all, much like the case of D2, soup and salad are options that come with the entree. If I get both of them, then surely I had both options. So, the sentence remains true. However, I think someone could disagree with me here and that's what I want to talk about now.

Consider the possibility that I get both soup and salad with my entree. I'm loving life because I have lots of food. Yum, yum! But now the store manager comes over to my table and says, "Excuse me, sir, we are going to have to charge you extra for the soup, since our entrees come with either soup or salad but not both. If you choose both soup and salad, then we have to charge you for an extra side dish." Well, that seems reasonable, but notice what just happened. The manager has clarified what D1 *really* means. In effect, D1 should be written like this:

> D1´: This entree comes with soup or salad but not both soup and salad.

The last part of the sentence is *implied* in the original sentence and this is what the manager is telling me. But now I think we can see that D1´ is not really like D2, D3, or D4. It's really a much more complex sentence. D1´ has the following form:

> A or B and not (A and B)

So, D1´ is not a simple disjunction; it's actually a conjunction of a disjunction with the negation of another conjunction. What this shows us is that the exclusive or is actually much more than a simple 'or' sentence. So, when we encounter disjunctions in logic, we will always assume that they are *inclusive* disjunctions.

Now let's look at some simple arguments that use the disjunction.

> My entree comes with soup.
> **---------------------------**
> Therefore, my entree comes with soup or salad.

This semester I can take Math, English, or Philosophy.
I can't take Math this semester.

Therefore, I can take English or Philosophy this semester.

Both of these arguments are valid. And they demonstrate important features of the 'or' operator. The 'or' operator makes two or more sentence components true if at least one of the sentence components is true. Moreover, the 'or' operator says that if two or more sentence components are joined with the 'or' operator, then at least one of the components must be true. The first case is sometimes called addition and the second case is called process of elimination or disjunctive syllogism (if you want to get fancy).

Conditionals

Conditionals are probably the most difficult of the logical operators for students to understand. So, expect to struggle with this one a little bit. However, if you take the process step by step – as we have for the other operators – I hope that you will discover you understand a great deal about conditionals.

So, like all of our operators (except negation), conditionals connect two sentential components. The conditional operator expresses a conditional relation between two propositions. That is, each proposition is conditional on the other. The conditional can usually be expressed as an if-then statement. Consider the following example:

T1- If you drop the china teapot on a hard floor, then it will break.

This expression says that there is a conditional relation between dropping the china teapot on a hard floor and the china teapot breaking. But what is that conditional relation? In other words, when is this statement true and when is it false? In order to assess this, let's imagine some different scenarios:

T1-1- I drop the china teapot and it breaks.

T1-2- I drop the china teapot and it doesn't break.

T1-3- I don't drop the china teapot, but it still breaks.

T1-4- I don't drop the china teapot and it doesn't break.

The first two cases are the clearest. In the first case, the conditional holds true. I said that if I dropped the china teapot, it would break; I dropped it; and it broke. So the sentence is true. In the second case, the conditional does not hold true: I said that if I dropped the teapot, it would break; but I dropped it and it didn't break. So, the second case clearly shows the conditional to be false.

The latter two cases are a bit more tricky. Considering T1-3, there are many ways to break a china teapot without dropping it on a hard floor. I could hit it with a hammer, for instance. Suppose that I do. Does this make the conditional false? T1 didn't say that dropping the teapot on the floor was the *only way* to break it. All it said was that *if* I did drop it on the hard floor, then it would break. In fact, it looks like the conditional leaves open the possibility that there are other ways to break china teapots. So, case T1-3 does not clearly make T1 false. Similarly, T1-4 does not clearly make T1 false. Indeed, we could imagine a person who sincerely believes that T1 is true, that is, they are seriously worried that if they drop the china teapot on the hard floor, it will break. So, they take considerable care never to drop the china teapot on the floor. And, as a result, the teapot never breaks. This clearly does not falsify T1. Indeed, it may be a consequence of believing that T1 is true.

Let's consider another example:

T2- Where there is smoke, there is a fire.

This is also a conditional (which we could easily rewrite as an if-then statement). Now, we can consider under what circumstances this statement might be false. One possibility is to consider a scenario where there is fire, but no smoke. In fact, this happens with very hot burning fires. If a fire is very hot, it consumes all of the large particles and so does not let off any visible smoke. So, there are fires without smoke. Does this falsify T2?

No. After all, T2 doesn't say that the only way there could be a fire is if there is smoke. It just says that if there is smoke, then there will be a fire. Again, if there is no smoke and no fire, the statement still holds true. The only way to falsify T2 is to provide a case where there is smoke, but no fire. In our world, this doesn't happen (unless you think that smoldering hot material is not a fire, but then I think you're being overly restrictive in your definition of a fire). But what about a world (say, in outer space) where there is no smoke. Could this sentence be true in that world? Well, this is a very interesting question. You might think that in such a scenario, we just can't know. And we shouldn't say that something is true or false if we can't know. Some philosophers have taken that position, but for our purposes, we will say that whenever the *if-clause* is false, the conditional must be true. This is the standard interpretation of the conditional in sentential logic. And I think it highlights the most important feature of the conditional, namely, that it asserts a relation that could be falsified. If that relation is not falsified, then we have to assume that it is true. (It's sort of like 'innocent until proven guilty'.)

But what kind of relation is expressed by the conditional? This is the interesting question and I think that if you understand this, you will be able to master the conditional. Based on the examples above, you might think that conditionals express a cause-and-effect relation. But this would be mistaken. Consider the following true conditional:

T3- If I am the present king of France, then groundhogs rule the world.

This statement is always true because the *if-clause* is always false: that is, there is no present king of France and so I cannot be the present king of France. Consider another true conditional:

T4- If the earth is 4.5 billion years old, then 2 + 2 = 4.

This statement is always true because the *then-clause* can never be false. Here, the *then-clause* is necessarily true (that is, it is impossible for it to be false) and so the conditional can never be false, even if it is false that the earth is 4.5 billion years old.

At this point, you might think that I am being a philosopher in the worst sense of the word – I'm just making things up that have no bearing on reality! And I can understand that viewpoint. But what I want you to appreciate is that there are an infinite number of statements like this, some that might be more realistic than others. And, while they may seem irrelevant to you, they demonstrate a very important property of conditionals: *there does not have to be any **real** relation between the if-clause and the then-clause*. In particular, there does not have to be any causal connection between the if-clause and the then-clause. Now, interestingly enough, it is very difficult to generate realistic examples (that is, examples where the if-clause and then-clause appear to be related) that do not have some underlying causal connection. However, there are a great many causal connections that do not satisfy sufficient conditions. So, causality is somewhat more complex, though it appears to be related to conditionals. Nevertheless, this is a metaphysical issue, not a purely logical one. So, let's just set it aside for now.

In fact, we don't even need to talk about causality in order to understand the relation expressed by conditionals. That is because we can understand it perfectly as a *conditional relation*. In particular, we should try to understand the if-clause of a conditional as the **sufficient condition** for the then-clause. And we should understand the then-clause as a **necessary condition** for the if-clause. What this means is that if we can satisfy the if-clause, then the then-clause will obtain (this expresses the sense of the sufficient condition). Put otherwise, whenever the if-clause is true, the then-clause must be true (this expresses the sense of the necessary condition). Let's look at some examples that make the necessary and sufficient conditions of conditionals apparent:

T5- If the weather is pleasant, then we will have a picnic.

T5 expresses the idea that pleasant weather is a sufficient condition for having a picnic. You might express this sentiment on your day off, when you want to get outside for the day. So, the only thing you are waiting on is the weather. And as long as the weather is pleasant, you will have a picnic.

T6- If you can make yourself dinner, then you know how to cook.

This sentence expresses the idea that knowing how to cook is a necessary condition for making yourself dinner. The idea would be that you can't make yourself dinner without knowing how to cook, so as long as you can make yourself dinner, then you must know how to cook.

Each of these statements obeys the general rules of truth and falsity that we outlined above. Considering T5, we can say that whenever there is bad weather, the statement is not falsified and therefore (by standard sentential logic) considered to be true. However, if the weather is pleasant and you do not have a picnic, then it is clear that the statement is false (that is, that there must have been other factors that determined whether you would have a picnic and pleasant weather was not sufficient). Considering T6, we can say that if you never make yourself dinner, then we can't say whether or not you can cook. And if this is the case, then I think the basic idea of the sentence still holds true. So, according to sentential logic, we will count the statement as true. However, if you can somehow make yourself dinner without having any knowledge of how to cook, then this would show the sentence to be false. So, under that scenario, we would have to say the sentence is false.

Conditional statements are very powerful and we will encounter them frequently in logical arguments. So, let's look at some basic argument we could construct, using the sentences above.

> If the weather is pleasant, then we will have a picnic.
> The weather is pleasant.
> --
> Therefore, we will have a picnic.

This perhaps the clearest expression of the sufficient condition in argument form. What this argument shows is that satisfying the if-clause (making the second premise true) results in the then-clause (making the conclusion true). This is called modus ponens.

> If I am the present king of France, then groundhogs rule the world.
> Groundhogs do not rule the world.
> ---------------------------------------
> Therefore, I am not the present king of France.

It is important to see that even though the conclusion here states that part of the conditional is false, arriving at this conclusion requires us to assume that the conditional is true. In other words, we are only able to infer that I am not the present king of France *because* the first premise says that groundhogs ruling the world is a necessary condition for my being the present king of France *and* the second premise says that groundhogs do not rule the world. So, the inference to the conclusion is dependent on the truth of the first premise. This argument is called modus tolens.

> If there is smoke, there is a fire.
> If there is a fire, then something is burning.
> --------------------------------------
> Therefore, if there is smoke, then something is burning

Here, I added another conditional expression in order to show how you can link two conditionals together. This argument is called a hypothetical syllogism.

Biconditionals

Similar to the conditional operator, the biconditional operator connects two sentential components in a way that expresses necessary and sufficient conditions. The difference between the conditional and the biconditional is (just as you might expect and please forgive the pun) that the biconditional goes both ways. Whereas the conditional expression shows a difference between the if-clause and the then-clause – the if-clause is a sufficient but not necessary condition, while the then-clause is a necessary but not sufficient condition – the biconditional shows that the relation is mutual, that both clauses are necessary and sufficient conditions. This is usually expressed using the language of “if and only if.” Consider the difference.

First we have a conditional expression from above.

T5- If the weather is pleasant, then we will have a picnic.

This expression is not the same as

T5´- If we have a picnic, then the weather is pleasant.

T5 says that pleasant weather is sufficient for us to have a picnic, so having a picnic is necessary whenever the weather is pleasant. However, T5´says something different, namely, that having a picnic is sufficient for the weather being pleasant, so the weather being pleasant is necessary to having a picnic.

To see the difference more clearly, consider the following two cases:

T5-1- The weather is pleasant and we do not have a picnic.

T5-2- The weather is not pleasant, but we still have a picnic.

Now the difference should be clear: the scenario presented in T5-1 makes T5 false, but T5´ remains true; the scenario in T5-2 makes T5´ false, but T5 remains true. So, T5 and T5´ clearly do not mean the same thing: since they have different truth values for these two scenarios.

Biconditionals are not like this. You can change the order of the biconditional without altering its meaning. Consider these examples:

B1- A person has a beard if and only if there is hair on his face.

B2- You are an astronaut if and only if you are prepared to go into space.

B3- Something is water soluble if and only if it dissolves in water.

You can reverse any one of these expressions without changing its meaning:

B1´- A person has hair on his face if and only if he has a beard.

B2´- You are prepared to go into space if and only if you are an astronaut.

B3´- Something dissolves in water if and only if it is water soluble.

B1´, B2´, and B3´ do not change the meaning of B1, B2, or B3. You can see that they are, in essence, the same. The reason for this is because the biconditional expresses the fact that each sentence component is a necessary and sufficient condition for the other. So, switching the order does not change the conditional relation. We can see this by considering the various scenarios under which these biconditionals might be true of false. Consider B3:

B3-1- Something is water soluble and it dissolves in water.

B3-2- Something is water soluble, but it does not dissolve in water.

B3-3- Something is not water soluble, but it still dissolves in water.

B3-4- Something is not water soluble and it does not dissolve in water.

B3 should be true in both the cases of B3-1 and B3-4, but it should be false in cases B3-2 and B3-3. In fact, it is. So B3 is true. What the biconditional asserts is that the truth value of one of the sentence component tracks the truth value of the other: whenever one is true, the other is true and whenever one is false, the other is false.

And now that we have established this, we can see that B1 is false. That is, it is possible for someone to have hair on his face without having a beard. He could have stubble, peach fuzz, a mustache, or a goat-tee! And, today, B2 is true, but some time in the near future, it will be false because we will have commercial space travel that takes people who are not astronauts into space. So, being an astronaut will, someday, not be necessary for being prepared to go into space.

We can now see that biconditionals express a very important relation. In the real world, they frequently express a kind of definitional relation, when we say 'if and only if' we are asserting that there is a very tight relationship between the two sentence components – for instance, one is a definition of the other. Of course, we can imagine more fanciful (or trivial) biconditionals that remain true even though one is not the definition of the other.

B4- The sun revolves around the earth if and only if six is greater than 8.

Since both of these sentence components are always false, the biconditional holds. But there is obiously no close relation between the two sentence components. Similarly,

B5- There are an infinite number of odd numbers if and only if all matter has mass.

Here, the biconditional connects two sentences that are always true and so the biconditional is always true even though there is no relation between the sentence components.

Now, let's consider some arguments using biconditionals.

You are an astronaut if and only if you are prepared to go into space.
You are an astronaut.

Therefore, you are prepared to go into space.

Simlarly,

You are an astronaut if and only if you are prepared to go into space.
You are prepared to go into space.

Therefore, you are an astronaut.

These two arguments show that each side of the conditional is a sufficient condition for the other side of the conditional. Likewise,

Something is water soluble if and only if it dissolves in water.
This substance does not dissolve in water.

Therefore, this substance is not water soluble.

And

Something is water soluble if and only if it dissolves in water.
This substance is not water soluble.

Therefore, this substance will not dissolve in water.

Finally, we can decompose a biconditional into a conjunction of two conditionals. Here we are showing how a biconditional, in effect, expresses the fact that the two sentence components are conditionally related in both directions.

Someone has a beard if and only if he has hair on his face.

Therefore, if someone has a beard, then he has hair on his face; and if someone has hair on his face, then he has a beard.

In fact, this decomposition shows us exactly why the biconditional (B1) is false. This biconditional is false because it does not hold in the right-to-left direction. In other words, it is not true that if someone has hari on his face, then he has a beard. This is where the cases of stubble, peach fuzz, and mustaches come in. However, it is still true in the left-to-right direction, that is, if someone has a beard, then he must have hair on his face.

So, you can see that there is a very natural relationship between biconditionals and conditionals such that many of the same logical moves that hold for biconditionals also hold for conditionals. However, there are important differences between the two and they, fundamentally, express different relations between sentence components.

Valid Deductive Arguments

Logic, however, does not necessarily have anything to do with truth. Again, the logic of an argument is determined by its form: whether the premises, *if true*, logically entail the conclusion. If so, the argument is valid (whether the premises are true or not).

In order for us to appreciate what makes an argument valid, it is useful to look at some examples. Below, I provide several examples of valid deductive arguments along with the names that are traditionally associated with these arguments. There are, in principle, an infinite number of valid deductive arguments, but these common arguments will give you a sense of the sorts of rules and reasonings that make deductive arguments valid.

Modus Ponens

This is a very natural form of argument based on an if-then statement. Essentially, it says that whenever an if-then statement is true and the if-clause of the statement is true, then the then-clause of the statement must also be true.

It has the form:

P1- If P, then Q.
P2- P.
C- So, Q.

Here is an example:

P1- If our galaxy has millions of habitable planets, then it seems likely that life has evolved on some planet in our galaxy other than our own.
P2- Our galaxy has millions of habitable planets.
C- So, it seems likely that life has evolved on some planet in our galaxy other than our own.

Modus Tollens

This argument should not be confused with Modus Ponens. It is, in fact, the exact opposite of a Modus Ponens argument. Essentially, it says that whenever an if-then statement is true and the then-clause of the statement is false, the if-clause of the statement must also be false.

It has the form:

P1- If P, then Q.
P2- Not Q.
C- So, not P.

Here is an example:

P1- If it rained last night, the streets would be wet.
P2- The streets are not wet.
C- So, it did not rain last night.

Hypothetical Syllogism

This is another way of arguing from if-then statements. But it does not lead to a simple statement of fact. Instead, it only leads to another if-then statement. In order to establish a statement of fact, we would need to an additional premise, establishing the if-clause of the conclusion. Nonetheless, even hypothetical or if-then statements can be informative.

This argument has the form:

P1- If P, then Q.
P2- If Q, then R.
C- So, if P, then R.

Here is an example:

P1- If you make a budget, then you will learn how you can save money.
P2- If you learn how you can save money, then you can set aside money to spend on the things you want.
C- If you make a budget, then you can set aside money to spend on the things you want.

Disjunctive Syllogism

This argument is commonly known as "process of elimination." It is a perfectly valid form of reasoning. However, you should be cautious: the conclusion is only true if all the premises are true, AND one of the premises lists ALL of the options. The fallacy of a "false dichotomy" results from asserting a false disjunct in a disjunctive syllogism.

This argument has the form:

P1- P or Q.
P2- Not P.
C- So, Q.

Here is an example:

P1- Either John is a liar or the project is due next week.
P2- John isn't a liar.
C- So, the project is due next week.

Dilemma

A dilemma is a genuine, valid argument. It essentially asserts that there are two different sorts of paths you can take. Each of these paths leads to a result. So, you will either wind up with either one result or the other. This is a very useful argument even though it doesn't establish a state of affairs. Usually, philosophers talk about dilemmas as having "horns" (like a bull). When you encounter a dilemma, you must either reject that it is a real dilemma or you have to tangle with one of its horns.

This argument has the form:

P1- P or R.
P2- If P, then Q.
P3- If R, then S.
C- So, Q or S.

Here is an example:

P1- Tonight, I can either go to the movies or the party.
P2- If I go to the movies, I'll probably see an action movie.
P3- If I go to the party, I'll probably see my ex-girlfriend Jane.
C- So, tonight I'll probably either see an action movie or my ex-girlfriend Jane.

Reductio ad Absurdum

This may not look like a valid argument since it appears to involve reasoning to an impossible state of affairs (an absurdity). However, it is in fact a valid argument and a very powerful one. Sometimes it is not possible to establish the truth of a conclusion directly. So, what we need is an indirect method of establishing the truth of the conclusion. Reductio ad absurdum arguments offer an indirect method of argumentation. Effectively, you assume the opposite of what you want to prove. Then you show that this assumption leads to a contradiction (which is impossible or absurd). And so, you reason, my assumption must be false. Hence, the opposite of my assumption must be true.

Here is an example:

> I will show that every human being has a mother:
> P1- Assume that there is some human being, Carl, who has no mother.
> P2- There is at present no other process of human generation than biological reproduction.
> P3- If a human being is not generated, then it cannot exist.
> P4- Biological reproduction requires a mother.
> P5- Carl was not generated through biological reproduction.
> C1/P6- So, Carl does not exist.
> C2/P7- So, Carl both exists and does not exist.
> C3- So, there is no human being who does not have a mother.

Fallacies

Fallacies are arguments that may appear to be valid, but in fact are invalid. So, they cannot reliably lead to true conclusions. I will provide some examples of common fallacies below. It would be a useful exercise to try to identify a counterexample that would show each of these arguments to be invalid.

Fallacy of the Red Herring

A "red herring" is an argument that diverts attention from the matter at hand. In other words, when a person argues using a red herring, she offers an argument that appears to support the conclusion, but is actually irrelevant to the conclusion.

The following three arguments can be considered as examples of red herring arguments.

Ad Hominem

Consider the following conversation:

> Joe: I'm what you would call a "constitutionalist." I believe that we should restrict federal power from anything beyond what the founding fathers intended in the Constitution.
>
> Fred: Are you kidding? Michele Bachmann, Sarah Palin, and Glen Beck believe in that nonsense. And those people are crazy!

Here, Fred is attacking the people who hold this view rather than the view itself. This is an 'ad hominem' argument. In principle, the character of an individual should not be relevant to assessing the correctness of the view espoused by that individual.

Ad Ignorantium

The appeal to the absence of evidence as proof of the non-existence of evidence.

> P1- If Obama was not born in the United States, he is not eligible for the presidency.
> P2- Barack Obama has not produced a copy of the original, long-form birth certificate from the hospital where he was born.
> C- So Barack Obama is not eligible for the presidency of the United States.

Here, the argument appeals to the absence of evidence of the fact that Barack Obama was born in the United States in order to support the claim that Barack Obama was not born in the United States. Again, the fact that we lack some specific sort of evidence for a state of affairs is not necessarily reason to reject the belief in the existence of that state of affairs.

Ad Populum

The appeal to the opinion of the majority as proof.

P1- Everybody hates Ke$ha.
C- So, Ke$ha sucks.

Whether or not Ke$ha sucks should be established on the merits (or not) of Ke$ha's musical performances, not on popular opinion. So, the appeal to popular opinion is not really relevant to the conclusion that is being established.

Begging the Question

A different sort of fallacy that is also quite common. It is also known as reasoning in a circle or circular reasoning.

P1- All human life is sacred.
P2- Sacred life should be protected at all cost.
C1/P3- So, all human life should be protected at all cost.
C2- So, abortion and euthanasia are always wrong.

It may not be obvious how this argument begs the question, but a little bit of reflection will show that it does. Consider what is under dispute in the issues of abortion and euthanasia. Isn't the dispute really about whether or not certain kinds of life ought to be protected at all costs? But the argument simply asserts that every human life should be protected at all costs, since it asserts (without evidence) that all human life is sacred. So, the premises effectively assume a position on the very issue that is under dispute, namely, whether or not there are certain sorts of human life that should not be protected at all costs.

NOTE: "Begging the question" is a frequently misused phrase. Properly used, it refers to circular reasoning, NOT some "natural" question raised by known facts.

False Dichotomy

False Dichotomy

A false dichotomy results from an improper use of a disjunctive syllogism. In this case, the argument asserts a disjunct (an 'or' statement) that is false. Effectively, the or-statement assumes that there is no other alternative, no third possibility. If there is some genuine alternative possibility, then the argument is not valid.

Here is an example:

> P1- Either you are with us or you are against us.
> P2- You are not with us.
> C- So, you must be against us.

Clearly, a third possibility has been ruled out without warrant: I am not with you, but I am also not against you. For instance, what if I agree with your goals, but disagree with your methods? In this case, I might not be on your side ('with you'), but I might also not be opposed to your side ('against you').

Induction

So far we have been considering what are called "deductive" arguments. The logic of deductive arguments is fairly simple and so it is easier to explain basic concepts of logic with deductive arguments. However, when you consider ordinary thought and reasoning, deductive arguments are actually fairly rare. In fact, it seems that other sorts of reasoning are far more typical. More common methods of inference are inductive and abductive inference. In this section we will discuss inductive arguments.

Inductive arguments typically appeal to experience or observation. They include generalizations as well as forecasts about the future. As such, even if the premises of an inductive argument are perfectly true, their truth does not necessitate the truth of the conclusion. Good inductive arguments provide conclusions with only some degree of certainty. The stronger (or more cogent) an inductive argument, the higher the degree of certainty that the conclusion is true; the weaker (or less cogent) the argument, the lower the degree of certainty that the conclusion is true.

Generalizations

A generalization reasons from particular cases to some general truths about all cases of that kind. The fact that generalizations are the result of an inductive inference is clear from the common wisdom that all (or nearly all) generalizations have exceptions. But this is just what you would expect from an inductive argument. So exceptions to a generalization are not like counterexamples; even a good generalization can tolerate some exceptions. As a result, the counterexample method discussed in the section on deductive arguments will not work for inductive arguments.

Nevertheless, we want generalizations to be strong (or result from cogent reasoning) and so we have to ensure that the particular cases we use to support are generalization are accurate, representative, and that the generalization is properly applied to the cases. Weston, in the *Rulebook for Arguments* provides three helpful rules to follow when making generalizations:

1. Use an adequate number of examples
2. Use examples that are representative.
3. Provide context.

Next we will look at some examples of weak generalizations and then some examples of strong generalizations.

Weak Generalizations

Weak generalizations:

> I have had three different cats. Each of them was fat and lazy. So, all cats are fat and lazy.

While I can't be sure, from the example above, whether or not my three cats are representative of all cats (i.e., I don't know whether they are typical kinds or not), I can be sure that three is not an adequate sample size to make a generalization about all cats. Compare this with the example below.

> A recent poll of over 5,000 people in the United States found that 85% of them were NRA members and 98% of them were either strongly or very strongly in favor of extensive gun rights for American citizens. This shows that support for gun rights is still very high in the United States.

5,000 people sounds like a pretty large sample size. Typical political or opinion polls (like Gallup) survey much smaller numbers of people. But the results of this poll look counterintuitive. What is going on? What if I told you that this poll was conducted at a gun show? If that was the case, then no matter how large the sample size, it's clear that it is not a representative sample of people in the United States. So, the generalization to all people in the United States is weak, when based on this sample.

Strong Generalizations

Strong generalizations control for these factors such that the generalization uses a sample that is both representative and large enough to support the generalization. Here are two examples:

> Based on a study of at least one member of every known cat species, scientists have determined that the common house cat does not possess taste buds capable of tasting "sweetness." So, your cat does not have any special desire for sweet or sugary things.

> According to Gallup polling from 1991-2011, support for stricter gun laws has decreased, while support for the status quo has increased. So, from 1991-2011, Americans grew less favorable of gun control legislation.

Even though these generalizations appear to be strong (supported by scientific study of cat species and Gallup polling), they are still inductive generalizations and, as such, not 100% reliable. Moreover, it is very important to interpret the conclusions of generalizations appropriately. For instance, it may not be accurate to infer from the properties of house cats to properties of tigers; and it may not be accurate to infer from the Gallup polling above to current preferences.

The Problem of Induction

Generalizations are not the only form of inductive arguments. Some inductive arguments lead to conclusions about particulars. Consider the following examples:

> Paul has a gene that most people with hair loss have. So, Paul will probably lose his hair as he gets older.
>
> Most Republicans favor lower taxes. John's Dad is a Republican. So, John's Dad probably favors lower taxes.
>
> The sun will rise tomorrow. It always has.

These arguments may appear to be deductive, since they use general truths to reason to particular conclusions. However, if you look closely, you will see that in each case the evidence supporting the argument comes from experience or observation (of the correlation between a gene and hair loss, the opinions of Republicans, or the rising of the sun).

Any argument whose reasoning relies on the appeal to experience runs into the "problem of induction," as it has been known since Hume. We will discuss the problem of induction in more detail later in the course. For now, we simply need to recognize that whenever we reason from some experiences that we have had in the past to experiences we may have in the future, we have to assume that certain basic features about these two sets of experiences will remain unchanged. This is what is called the "inductive principle." The inductive principle can be stated in the following way: all arguments from induction assume that some basic features of the natural world will remain unchanged from the past to the future. All inductive arguments assume that the future will be (roughly) similar to the past, such that experiences in the future will be similar to experiences in the past.

The inductive principle leads us to some surprising conclusions. For example, we should realize that the belief that the sun will rise tomorrow is a belief that only holds true as long as some basic features of the world remain tomorrow the way they have been in the past (that the earth continues to rotate, for instance). The problem is finding justification for this necessary continuity. Hume's recognition of the problem of induction led him to a form of skepticism. For our purposes, we only need to realize that since inductive arguments rely on the inductive principle, they do not provide the same kind of certainty as deductive arguments do.

Arguments about Causes

In talking about inductive arguments, we have already made appeal to arguments about strong correlations. Given that correlations are the kinds of things that we have evidence for on the basis of experiment and observation and given that strong correlations indicate the possibility of causal relations, we might want to start talking about the logic of causality.

The issue of causality is a philosophically complex one and we can't hope to get very deep into the issue here. Nonetheless, it will be helpful to distinguish some bad forms of causal argument from some better forms. After all, the relation of cause and effect is a basic feature of the natural world and clearly something that we want to be able to talk about. So, we should outline some better and worse ways of talking about causality.

The first thing to recognize about causality is that even though our recognition of a causal relation begins with the observation of some correlation between two events or facts, correlation is not the same thing as causation.

Correlation is not Causation

Consider the following scenarios:

> After careful genetic study, scientists have isolated a gene that many different species of mammal (including humans) possess and they have determined that there is a strong correlation between being a carrier of this gene and having hair loss.
>
> Climate scientists, through a careful study of ice core samples and the historical record, have determined that periods of high CO_2 concentration in the earths atmosphere are strongly correlated with periods of rising global surface temperature averages.

In these cases the existence of a strong correlation appears to be reason to believe that there is a causal connection between these events. Indeed, we think that there is a causal connection between these events. But correlation is not enough to give us reason to hold that there is a causal relation, we need more.

Consider the following examples of real correlations:

> Increased use of tissue paper is strongly correlated with having a cold virus.
>
> Increased ice cream sales are strongly correlated with increased drowning deaths.

Human population growth is strongly correlated with the expansion of the universe.

While these examples are genuine examples of strong correlation, it is probably obvious that they are not causal relations. Just because some event is strongly correlated with another does not mean that this event is the cause of the other. For one thing, correlations do not tell us anything about the direction of causation. Effects are just as strongly correlated with causes as causes are with effects. For another thing, if two events are mutual effects of a common cause, then these events will be strongly correlated but not necessarily causal related to each other. Finally, events that may have nothing to do with one another may still happen to show instances of strong correlation. So, it is clear that correlation, while closely connected to causation, is not the same thing as causation.

Advice in Making Causal Arguments

Then, how should we approach arguments about causality? Short answer: very carefully. Seriously, the following three rules gives a pretty good guide of to considering arguments about causality.

1. Recognize instances in which correlation may not imply causation. These include: a) correlations that do not distinguish between cause and effect (there is no causal direction), b) correlations that exist between effects of a common cause, and c) correlations that may be coincidental.
2. Try to isolate a causal mechanism. Consider plausible candidates and imagine what could be the reason for a causal relationship between two events.
3. Work toward the most likely explanation: given multiple plausible causes for a given effect, select the one that is most likely. Principles of likelihood might be things like: simplicity (Ockham's Razor), the existence of some causal mechanism, or the existence of some theoretical framework for explaining the causal relation.
4. Expect complexity. Remember that, philosophically speaking, there are likely many causes for any given effect.

Arguments from Analogy

Finally, sometimes an argument can be made by using just one example. This is not an inductive argument, but a deductive one. It depends entirely on the strength of the premises and holds only insofar as the premises lead to the conclusion.

Analogies can be useful to explain a complex or abstract phenomenon by comparing it to a simpler or more concrete phenomenon. An analogy works by comparing two distinct phenomena based on certain relevant features. Since they are two distinct phenomena, there must be some features that are distinct. In order for the analogy to work, the differences have to be irrelevant to the analogy. In other words, the two phenomena have to be similar in the relevant respects.

Consider the following analogy that compares the regular service required for cars with the regular checkups that may be required by a doctor.

> People take their car for servicing and checkups every few months without complaint. Why shouldn't they take similar care of their bodies?

In order to evaluate this analogy, we need to ask two questions: What is the feature of these two phenomena that is relevant to the argument? Are the two phenomena similar in the relevant way, i.e., are their differences irrelevant to the analogy?

What do you think?