Analyzing and Assessing Arguments

A Primer for Students in Introduction to Philosophy

G. Salmieri

§1. Overview

Over your first few years of life, you learned how to do a lot of things—how to speak, how to think, how to walk, how to grasp and manipulate objects, etc. But you learned most of these things *implicitly*, without having an articulate understanding of what you were doing and how. For example, you learned many words before you learned the word "word" and you learned how to put words together into sentences before you were taught about the parts of speech and the rules of grammar. Many of you have probably never thought much about *how to* walk, but if you've ever had physical therapy to recover from an injury or taken a movement class as part of learning how to dance or to act, you may have learned to analyze the act of walking into component parts and this will have given you more control over the way you walk. Likewise, when athletes train, they often learn to analyze movements they learned as a child into simpler movements, and this process gives them finer grained control over their movements. It is possible to gain greater control of your thinking in the same way, by learning to analyzing the complex activity of thinking into the various mental acts that make it up.

In this primer we're going to focus on a part of our mental activity that looms large in our daily lives, in the sciences, and in philosophy: the activity of arguing—of producing and evaluating arguments. We will learn what an argument is, how to analyze an argument into its parts, and how to use this analysis to methodically assess an argument.

In §1, I provide an overview of the whole process, and §2–5 deal in greater detail with different aspects of the process.

§1.1 Arguments and Their Parts

Before we discuss arguments in earnest, it will be helpful to say a bit about **propositions**, which are the smaller units of thought from which arguments are composed. A proposition is the sort of thought that is capable of being true or false, believed or disbelieved, and asserted or denied. Such thoughts are asserted by declarative sentences. Here are some examples:

1. Healthy grass is green.
2. O. J. Simpson killed Nicole Brown.
3. Twice two is four.
4. Twice two is five.
5. Many diseases are caused by bacteria.
6. Stalin was evil.
7. Donald Trump is the 45th President of the United States.
8. Either Trump or a Democrat will win the 2020 election.
9. Donald Trump should be impeached.
10. Hillary Clinton would have been an awful president.

The proposition is not the same thing as the sentence asserting it because the same thought can be asserted by different sentences. For example, here are several different ways of asserting Proposition 3 from the list above:

3a. Twice two is four.

3b. Two times two is four.

3c. 2 x 2 = 4

3d. Deux fois deux c'est quatre.

3e. 兩次兩次是四次。

Sentences 3a and 3b are alternative ways of asserting the same proposition in English, 3c is a way of asserting it in mathematical notation, 3d is a way of asserting it in French, and 3e is a way of asserting it in Chinese.

Another reason why a proposition is not the same thing as a sentence asserting it, is that propositions can be asserted as parts of more complex sentences that include multiple propositions. For example, here's a sentence that asserts both Propositions 9 and 10 from the list above: "Hillary Clinton would have been an awful president, but Donald Trump should be impeached."

Some of the propositions from the list above are uncontroversially true and others are uncontroversially false. I expect that everyone in the class will all agree that Propositions 1, 3, 5 and 7 are true and that Proposition 4 is false. We will probably disagree over some of the others—with some students thinking they're true and others thinking they're false.

Some of you may think that some of the propositions we may disagree over aren't *really* the sort of things that can be true or false at all. In particular, some of you may think this about Propositions 6, 9, and 10, because these propositions express evaluations and you might think that evaluations aren't the sorts of things that can be true or false. But even if you think this, you'll have noticed that many people *believe* or *disbelieve* each of these propositions as though they were true or false, and they make arguments to try to convince other people of them. So, to understand their role in arguments you'll have to treat them as the sort of thing that can be true or false.

As the term is used in philosophy, an **argument** is a set of related propositions (called **premises**) that is given as a reason for believing a further proposition (called the **conclusion**). Consider, for example, the following, simple argument:

> Whoever murdered Carl had to have access to his rose garden at midnight, and the only person who did was Natalie, therefore Natalie must be the murderer.

Here the conclusion is that Natalie murdered Carl, and there are two premises: (1) that whoever murdered Carl had access to Carl's rose garden at midnight, and (2) that only Natalie had such access.

There are a few ways in which we can represent an argument that make its structure clearer. One popular way is called standard form. Here's what the argument we have been discussing looks like in standard form:

Argument written in standard form:

1. Whoever murdered Carl had access to his rose garden at midnight.
2. Only Natalie had access to Carl's rose garden at midnight.

---

3. Natalie murdered Carl.

Each proposition has been written on its own line and labeled with a number. The premises are written first, and a line is drawn to separate them from the conclusion. This line represents the **inference**—the mental act of moving in thought from the premises to the conclusion.

You will occasionally find arguments written out in this way in some of the readings for this class. However, we will be making more use of a newer way of representing the structure of arguments, which is called "argument mapping." Here's a map of the argument we've been discussing.

Map 1:

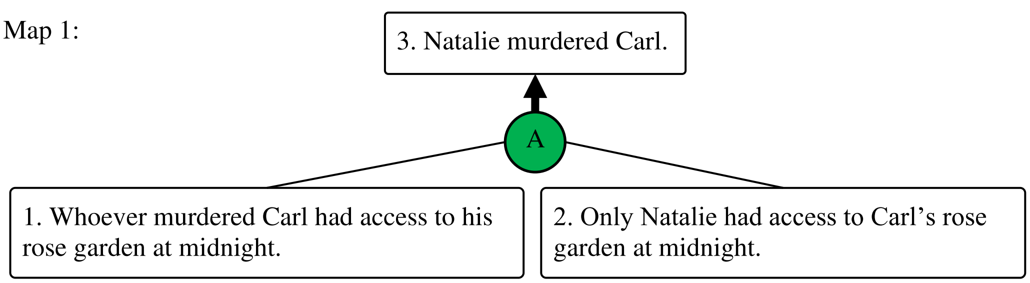

In the map, each proposition is written in its own box and labeled with a number. The inference is represented by a green circle and labeled with a letter. Thin lines connect the circle to the boxes containing the premises, and a thicker line with an arrow at the end connects the circle to box containing the conclusion.[1]

Both of these ways of representing the argument are meant to make it easy to see that Propositions 1 and 2 combine to provide a reason to believe Proposition 3. Another way we can express this point is to say that Propositions 1 and 2, taken together, are meant to help us to tell that Proposition 3 is true.

[1] There are different competing conventions for argument mapping, so you may find maps drawn a bit differently elsewhere, with different symbols used. But for the purposes of this course we will keep to the conventions in this Primer.

Notice that neither of these propositions on its own would give us any reason at all to think that Natalie murdered Carl. If we had no idea that Natalie had access to Carl's rose garden at midnight, then knowing that the murderer had access to it at this time, wouldn't give us any reason to think that Natalie was the murderer. Likewise, if we had no idea that the murderer had access to the rose garden, then knowing that Natalie was the only one with access wouldn't give us a reason to think she committed the murder. It's only when we put the two propositions together that they give us a reason to believe the conclusion. This argument map represents this relationship by having the lines from the two premises meet at the circle representing the inference, and then having a single arrow go from the circle to the conclusion.

§1.2 The Uses of Arguments

If you were trying to convince someone, perhaps a jury, of Natalie's guilt, you might present them with the argument we discussed in the last section. Or, you might formulate this argument silently to yourself in the course of trying to figure out who killed Carl. In either case, we would say that you *inferred* that Natalie was Carl's murderer from the *premises* that the murderer had access to the rose garden at midnight and that only Natalie had such access.

We tend to think of making arguments in situations where people disagree and are trying to convince the other (or to convince some third party). This is the case when people argue in court or around a dinner table; but this is not the only use of arguments, nor is it their fundamental use. Much of what we know is based on other knowledge and inference is the process by which we reach new knowledge from old. It is only because of arguments that we know many of the things we do, and it is only by further arguments that we can come to know many of the things we would like to learn.

Let's survey some of the areas in which much of our knowledge is due to inference. Arguments like the one concerning Carl and Natalie enable us to know who committed crimes for which there were no eyewitnesses. In the natural sciences, almost everything we know beyond the basic data (consisting in observations and measurements) is reached by inference.[2] So is much of our knowledge in higher mathematics. Our knowledge of the future is inferred from what we know about the past and present. And our knowledge of the distant past (beyond the scope of our memories) is also based on argument. Here, in some cases, we have firsthand accounts from people who claim to have observed or taken part in various historical events, but there are questions about the honesty and accuracy of these reports (especially when some contradict others), so we need arguments to determine which accounts we can rely on.

Inferring is a crucial part of how you as an individual come to know the world. It's a process that's fallible, and sometimes it results in errors—beliefs that aren't really knowledge. When two people argue with one another—if they're arguing honestly—they are sharing with some of the

[2] We will discuss later in the course whether any of the content of science are known by means other than inference or observation.

reasons they each have for believing the things they do. Each is trying to teach the other some of what he thinks he knows, and each can help the other identify errors in his own thinking.

Of course, people also often argue with one another dishonestly. Think of a conman who sells a bogus medicine and gives arguments to trick sick people into believing that his medicine will cure their illnesses. The conman knows that the conclusion he's trying to convince them of is false—or, at least, he doesn't really care if it is true. His goal is not to help them to tell what's true, but to get them to believe a conclusion that he wants them to believe, regardless of whether it is true. It is possible to be dishonest with oneself in this same way. You may not be able to work to convince yourself of something you self-consciously know is a lie, but people often make arguments to convince themselves of things that they *want* to believe—or that they think they are *supposed to* believe—regardless of whether the things are really true. For example, someone who suspects her husband of having cheated may not want to believe it because she finds the idea too painful, or she may feel guilty for believing it because she thinks she is supposed to trust him. In either case, she may fish around for arguments to convince her that he is faithful.

This brings me to an important point of clarification. Earlier I said that an argument is meant to give us a "reason to believe" its conclusion. When we speak of the reasons we have for performing an *action*, we usually have in mind either some *goal* that we hope to achieve by performing the action, or some *obligation* that we think we have for taking it. These are called pragmatic or practical reasons. By contrast, we can call the reasons that arguments give epistemic reasons. They are the sort of reasons that help one to *tell that a conclusion is true* and thereby to put one in a position to *know* it.

The primary use of arguments is to *tell what's true*. A secondary use is to convince oneself or others that something is true (whether or not it really is). In this course we will be focused only on the primary use.

§1.3 Mapping the Relations Between Arguments

Our thinking isn't usually made up of single, isolated arguments, but of complex sets of interrelated arguments. And argument mapping helps us to visualize and understand these relations. In this section we'll consider how the argument we looked at earlier concerning Natalie might be related to other arguments. Before we do that, let's review the simple map of that argument and get clear on how we are going to refer to each of its parts.

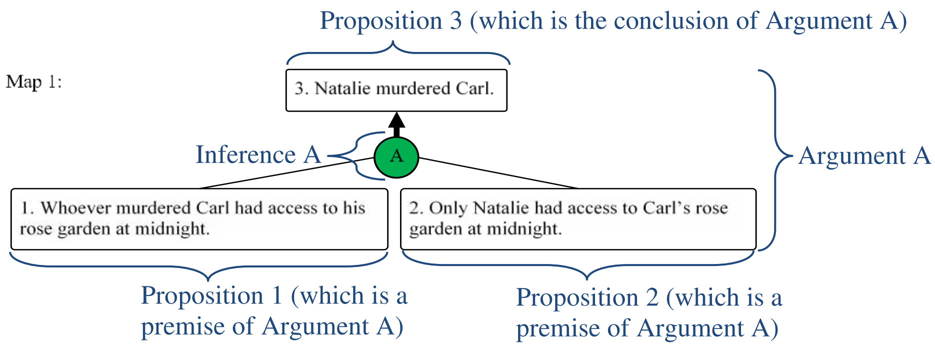

Each box contains a numbered proposition. We'll refer to them as "Proposition 1," "Proposition 2," and "Proposition 3." The green circle represents the act of inference, and we'll refer to it as "Inference A." We'll use the phrase "Argument A" to refer to the whole argument, which includes the inference, along with its premises and the conclusion.

With this terminology under our belts we can look at and discuss more complex maps that show relations between arguments.

One way in which arguments can be related is by sharing the same conclusion. Here's an example:

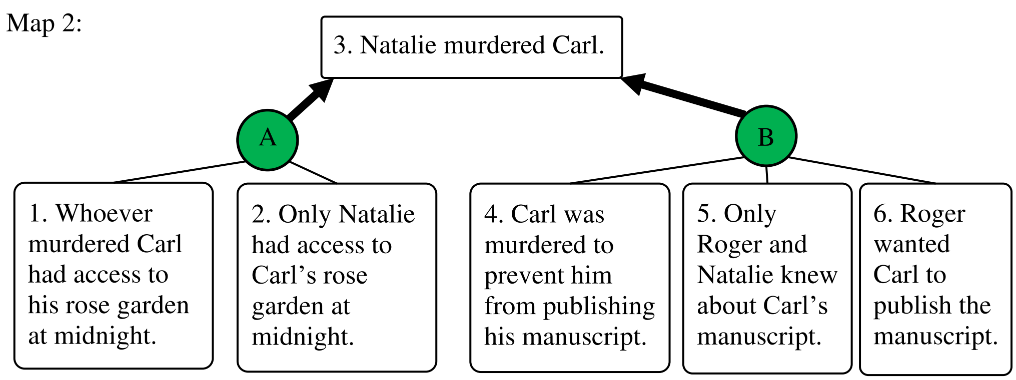

Argument B has the same conclusion as Argument A. They are two distinct arguments because they give separate reasons for believing the conclusion. To see that the reasons are separate, pick one of the premises from either argument, and think about what other premises you would need to combine it with before it would give you a reason to think that Natalie murdered Carl. You will see that in order for Proposition 1 to convince you that Natalie murdered Carl, you would need to also know Proposition 2, but that you wouldn't need to know Propositions 4, 5, or 6. Likewise, in order for Proposition 5 to convince you, you'd need to also know Propositions 4 and 6, but not Propositions 1 or 2.

A second way arguments can be related is that a premise of one argument can be the conclusion of another. Here's an example of that:

Map 3:

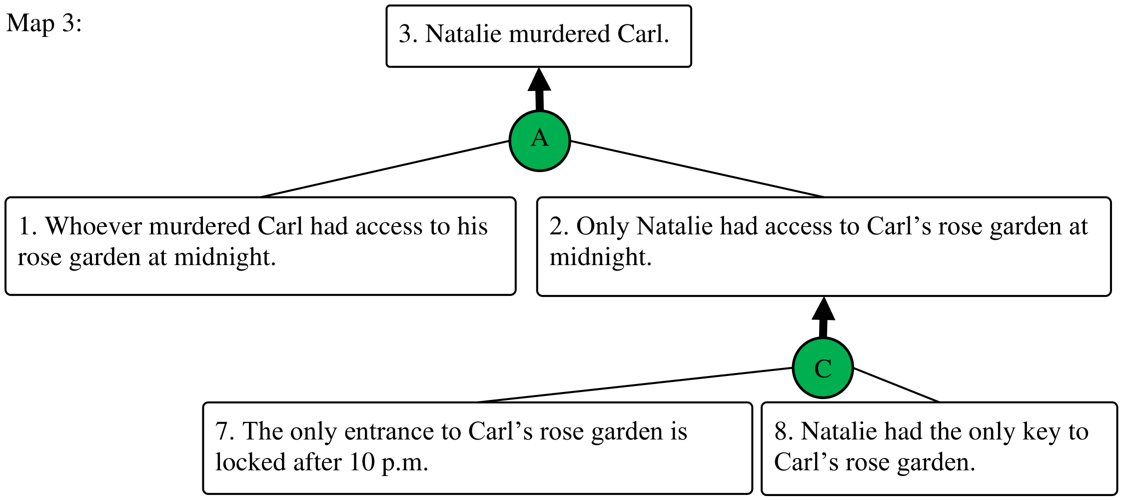

Notice that Argument A (from Map 1) is part of this larger map. One of its premises is now the conclusion of another argument, labeled C. So, we can think of C and A taken together as a larger, two-step argument for Proposition 3.

Now let's consider a larger map that includes all the arguments we've looked at plus two others.

Map 4:

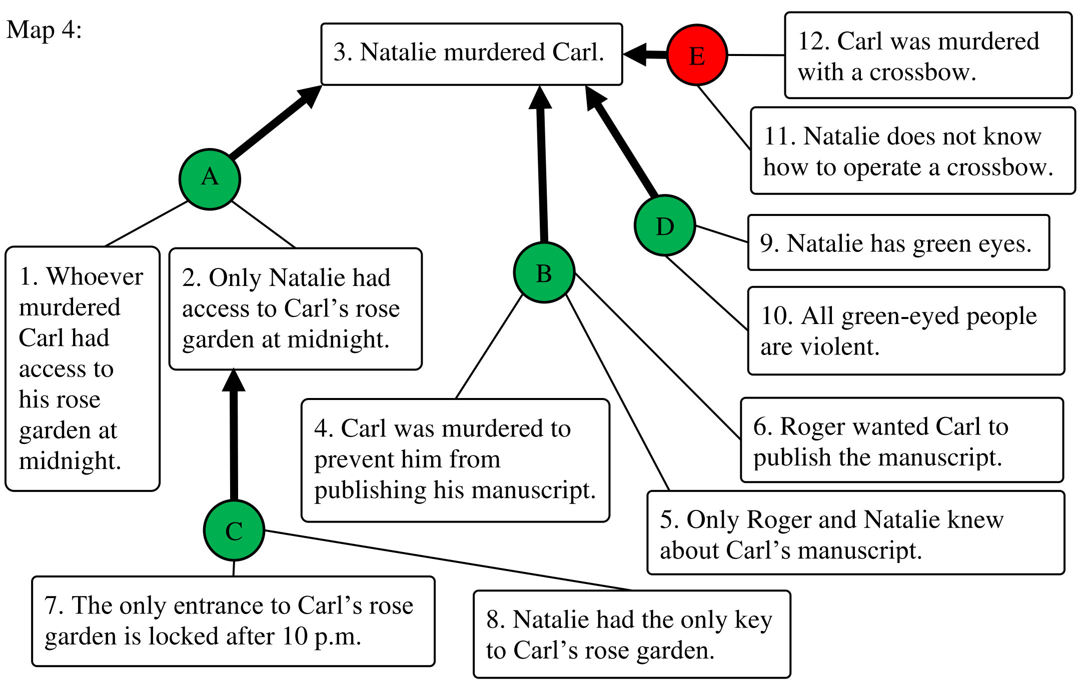

Argument D is a third argument for Proposition 3. I expect you'll agree that it a bad argument. We will discuss what makes some arguments better than others in the next section.

Inference E is represented by a red circle instead of the usual green one because Argument E is an argument *against* proposition 3, rather than for it. The conclusion of the Argument E is

"Natalie did *not* murder Carl." We could write this out as its own proposition in a separate box, but since this conclusion is equivalent to saying that Proposition 3 is false, it's nice to have a way to relate Argument E to Proposition 3 on the map. That's what we use the red circle for. It helps us to see at a glance that Argument E is arguing against the very same conclusion that Arguments A, B, and D are arguing for.

The point of an argument is to help us tell what's true. So, if we end up in a situation where we have arguments with opposite conclusions as we do in Map 4, then there must be something wrong with at least one of the arguments. Let's turn now to the things that can go wrong with arguments, and why some arguments are better than others.

§1.4 Why Some Arguments are Stronger than Others

Mapping an argument is a way of analyzing it—breaking it up into its parts and showing how those parts fit together into a whole. The point of analyzing an argument is that it makes it easier for us to then *assess* the argument. When we assess an argument, we are asking how strongly it supports the conclusion—how good a reason it gives us to think that the conclusion is true. The strongest arguments enable us to tell that their conclusions are true. In doing so, they establish the conclusions as *knowledge*.

If an argument supports its conclusion strongly enough to establish it as knowledge, the argument is called **conclusive** and is said to be a **proof** or to **prove** the conclusion.[3] These arguments are especially valuable. On the opposite extreme would be an argument that gives us no reason to believe the conclusion and so leaves us no closer to knowing it than we were before. There is no special name for such arguments, but let's call them **worthless** because they don't do *any* of what an argument should do. Many arguments fall in between these two extremes. That's because knowledge isn't an all-or-nothing affair. There is a whole spectrum between really knowing a proposition and being entirely ignorant of it. We can call this spectrum the proposition's **epistemic status**.

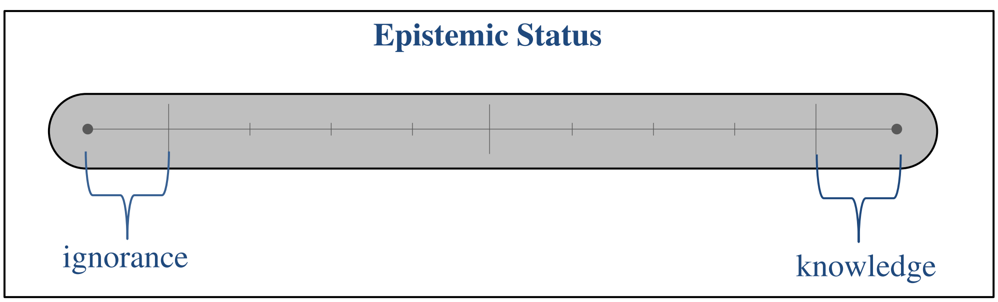

[3] People often call arguments that they come up with proofs if they think the arguments prove their conclusions, but that doesn't mean that the arguments really do prove the conclusions. We have to assess them for ourselves to see.

Consider the proposition "Natalie murdered Carl," which was the conclusion of many of the arguments in the maps we looked at earlier. Suppose that this proposition is true and that you're a detective investigating Carl's murder. When you begin the case, you probably don't know the proposition at all. At this point you have no idea who killed Carl, and no reason to suspect Natalie; you might not even know who Natalie is. So, the proposition would be at the extreme left end of the scale picture above. But the investigation proceeds and you learn more about Carl's life and death, at some point you formulate the theory that Natalie murdered Carl because there is some evidence pointing to her. Perhaps at this point, it's not much evidence. We certainly wouldn't say that you *know* she killed him or even that you *believe* it (or have reason to believe it), but you now suspect her, so we wouldn't say that you're totally ignorant of her having murdered him either. The proposition is now somewhere on the scale to the right of ignorance but to the left of the half-way mark.

Eventually, as you accumulate more evidence, she becomes the lead suspect. Now if you had to bet, you'd say that she did it, and you'd have good reasons to support your bet. We might say you believed she did it, but you don't *know* it yet. Finally, at a certain point you might get enough evidence to really be certain that she did it. Now the proposition has progressed to the right-most region on the scale—you know it.

All the evidence that you accumulated along the way could be spelled out as arguments. And we can think of what arguments do as helping us advance along the scale from ignorance to knowledge. An argument that's strong enough to take us all the way to knowledge is a conclusive argument or proof. An argument that doesn't take us any of the way is worthless. But many arguments take us part of the way—they give us *some* reason to believe the conclusion without giving us *conclusive* reason.

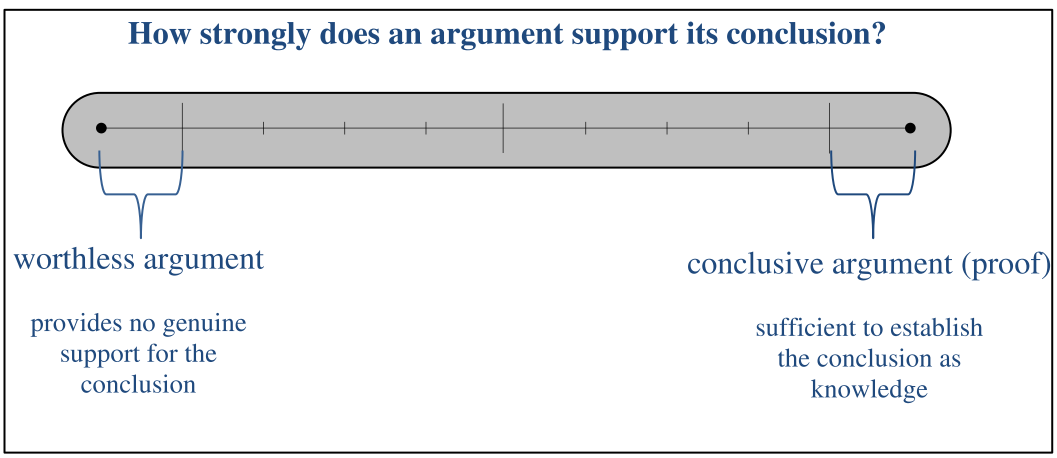

Notice that in the scale for epistemic status above knowledge is represented by a range and not by a point. This is because, even among the things we know, we think of ourselves as knowing somethings better than others. For example, you probably think you know both that Trump is the 45^{th} President of the United States and that twice two equals 4. But you might think that you know the second of these propositions better than the first, since you can probably imagine some bizarre scenario in which Trump isn't really the president and you're the victim of an elaborate hoax, but it's hard to imagine any scenario in which you can be mistaken that twice two is four. Perhaps some of you think that you don't really *know* that Trump is the 45th President because you think you can't totally rule out this hoax scenario. We'll discuss these sorts of skeptical worries later in the course. For now, my point is just that to saying that you know something is not to rule out the possibility that there are other things that you know even *better*. That's why I'm representing knowledge as a range rather than as a point. Similarly, to say that an argument is conclusive is just to say that it's *enough* to establish its conclusion as knowledge. It is not to say that there cannot be some other argument that is even stronger.

The two factors that contribute to the strength of an argument are its premises and its inferences. So, to assess an argument we need to assess each premise and each inference.

The strongest premises are ones that we *know* to be true independent of knowing the conclusion. In order for an argument to prove its conclusion all of its premises must be like this. On the other extreme if we have no reason at all to think that a premise is true (or if we know that it is false), then it will make any argument it is part of worthless. If, on the other hand, the premise has an intermediate epistemic status, an argument containing it could still support the conclusion to some extent, without proving it. This is illustrated in the map below.

Map 1a:

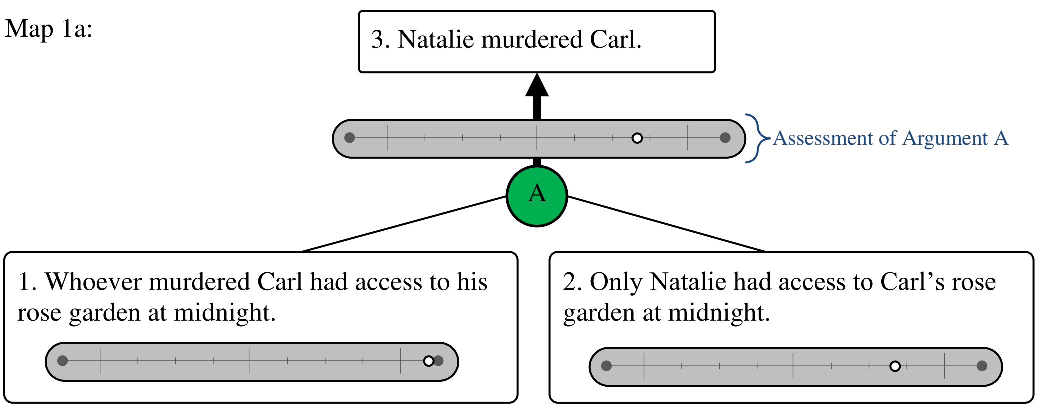

The scales placed in the boxed for Propositions 1 and 2 indicate the epistemic status of those propositions. The scale in the box for Proposition 1 has a mark in the rightmost region indicating that that proposition is known to be true. Proposition 2's scale shows that the proposition isn't quite known, though there is some reason to believe it. The scale drawn above the green circle indicates the strength of Argument A as a whole. We see that it is no stronger than the weakest premise, which is Proposition 2.

Since the case of Natalie and Carl is fictitious, there's no way for us to actually assess the premises. The epistemic statuses in the map above are made up (as is everything else in the example). In §3, below, we'll discuss how to assess actual premises of actual arguments.

In addition to assessing premises, we need to assess inferences. In order for an argument to support its conclusion, the premises and conclusion need to be related in such a way that it is unlikely for the conclusion to be false if the premises are true. The more unlikely it is for the conclusion to be false if the premises are true, the stronger the inference. Sometimes the premises and conclusion are related in such a manner that one would be caught in a contradiction if one held that the premises were true, but the conclusion was false. These inferences are called **deductions** and are said to necessitate their conclusions. *Deductions are as strong as it is possible for an inference to be*, so if we have a scale assessing the strength of an inference, we should represent deduction not as a range, but as a point at the end of the scale. (We will discuss how deductions work in §4.1, below.)

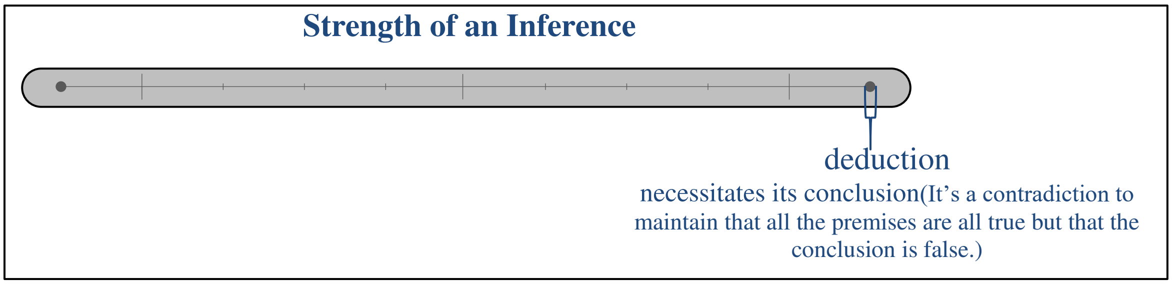

Inference A is a deduction. If whoever murdered Carl had access to his rose garden at midnight, and Natalie was the only person who had access then, then Natalie *has to be* the murderer. If we said she wasn't we would be saying that the murderer was someone *other than Natalie* who according to Proposition 1 had access to Carl's rose garden at midnight, but Proposition 2 tells us that *only Natalie* had access to the rose garden at midnight. So, to hold both premises and deny the conclusion would be to say that Natalie *was* and *wasn't* the only person with access to the rose garden at midnight, and that's a contradiction. If the premises are true, the conclusion has to be. The only way to consistently deny the conclusion is to deny one of the premises. So, Inference A is as strong as an inference can be. We can add this assessment into our map of Argument A as follows.

Map 1b:

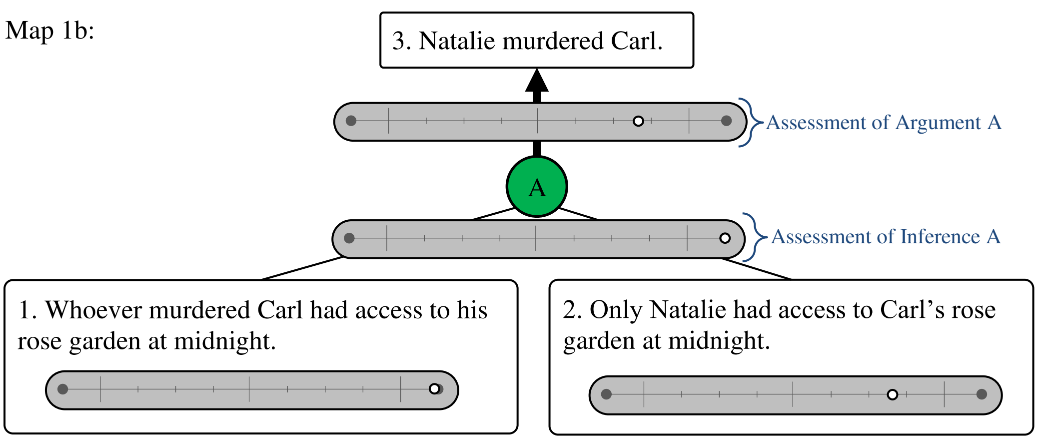

The scale below the green circle represents our assessment of Inference A, and it is marked at the rightmost point to show that the inference is a deduction. The scale above the circle represents our assessment of Argument A as a whole. Notice that, even though the inference is a deduction, the argument taken as a whole isn't conclusive, because we don't know that Premise 2 is true. The strength of the argument as a whole depends on the strength of *all* its elements, whereas the strength of the inference is independent of the strength of the premises. That's illustrated by the following map:

Map 5:

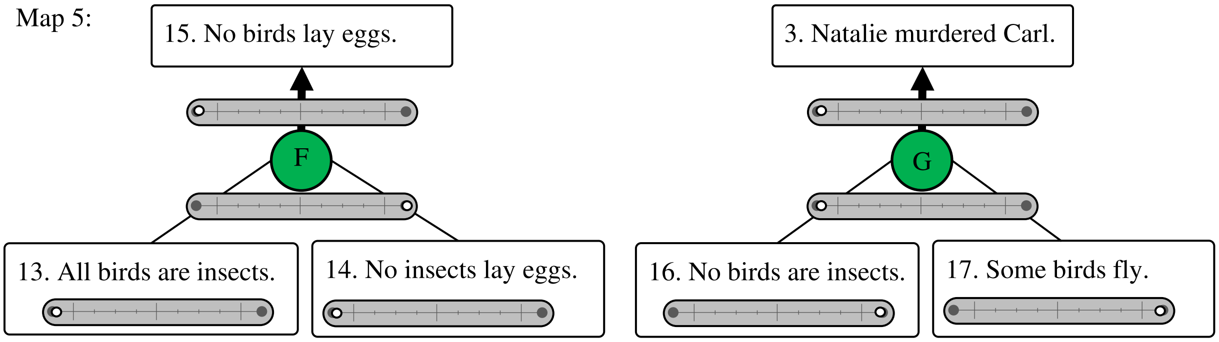

This map contains two separate arguments (Arguments F and G). Argument F has two awful premises—premises that no one has any reason to believe and that we all know to be false. But Inference F is as strong as can be; it's a deduction. Propositions 13 and 14 are obviously false, but if they *were* true, then Proposition 15 would *have to be* true also. Nevertheless, Argument F is worthless, because its premises are so bad. Argument G is also worthless, but for an opposite reason. We know that both of its premises are true, but the premises aren't related to one another and to the conclusion such that their being true gives us any reason to think that the conclusion is true as well. The problem here is with the inference. You can't infer anything about Carl and Natalie from premises about birds and insects. Inferences this bad are called **non-sequiturs**.

Some inferences are extremely strong without being deductions. Consider the following argument:

Map 6:

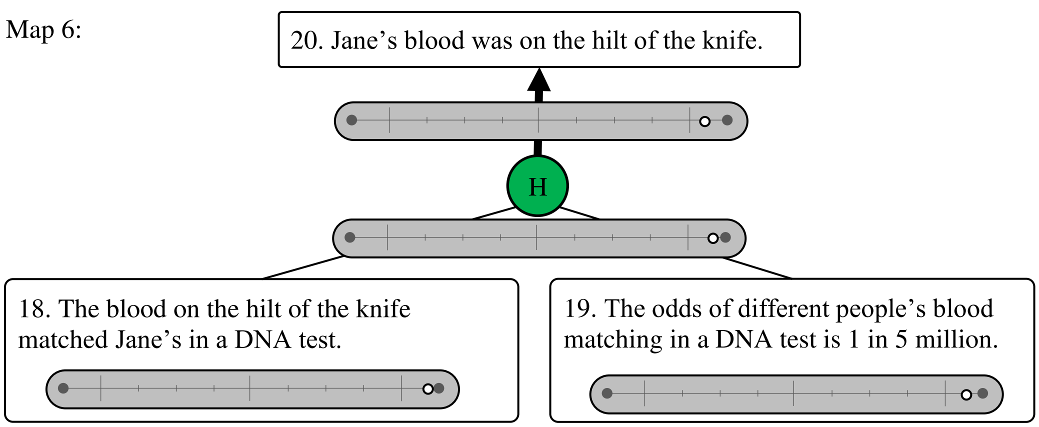

Propositions 18 and 19 do not necessitate Proposition 20. But knowing them would give us an extremely strong reason to believe Proposition 20. The reason is so strong that in most context we would say that Argument H would establish Proposition 20 as knowledge. Let's use the word **compelling** for inferences that are strong enough to establish their conclusions as knowledge, if their premises are true. If so, here's what our scale of inference strength looks like:

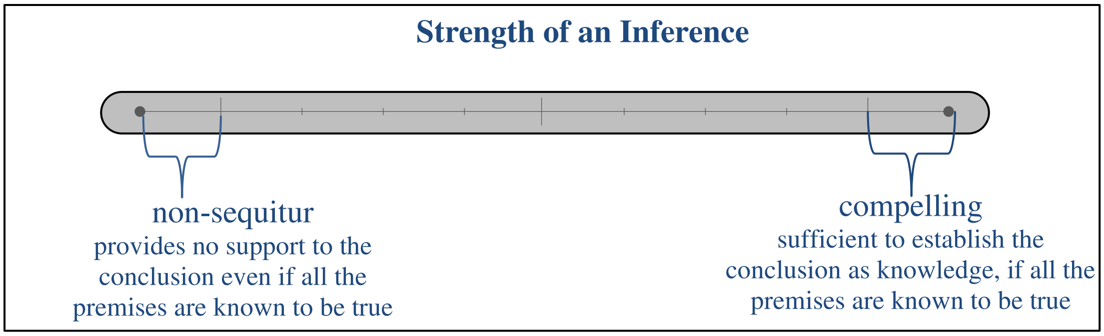

Even inferences that aren’t compelling can be useful. Consider the following argument:

Map 7:

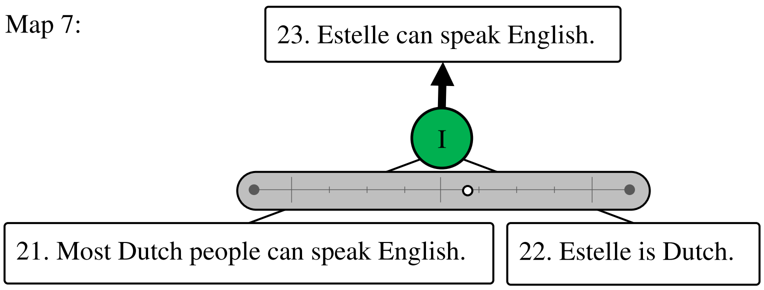

Inference I isn’t compelling, but it’s not bad, either. If you knew the premises to be true, this argument wouldn’t put you in a position to *know* the conclusion, but (unless you had other relevant knowledge of Estelle), it ought to lead you to expect that she can speak English.

In §3 below, we’ll discuss how to identify and assess different sorts of inference. The examples in this section are intended just to give you a sense that some are stronger than others.

To review then, arguments range in strength from worthless to conclusive, with conclusive arguments being the ones that are strong enough to establish their conclusions as knowledge. The strength of an argument is determined by the strength of its premises and of its inference. The strength of a premise is its epistemic status. The highest epistemic status is knowledge, and the lowest is that of a proposition one is either wholly ignorant of or knows to be false. The strength of an inference is determined by the relationship between the premises and the conclusion, and this is separate from whether the premises are true. The weakest inferences are called non-sequiturs and the strongest are called compelling.

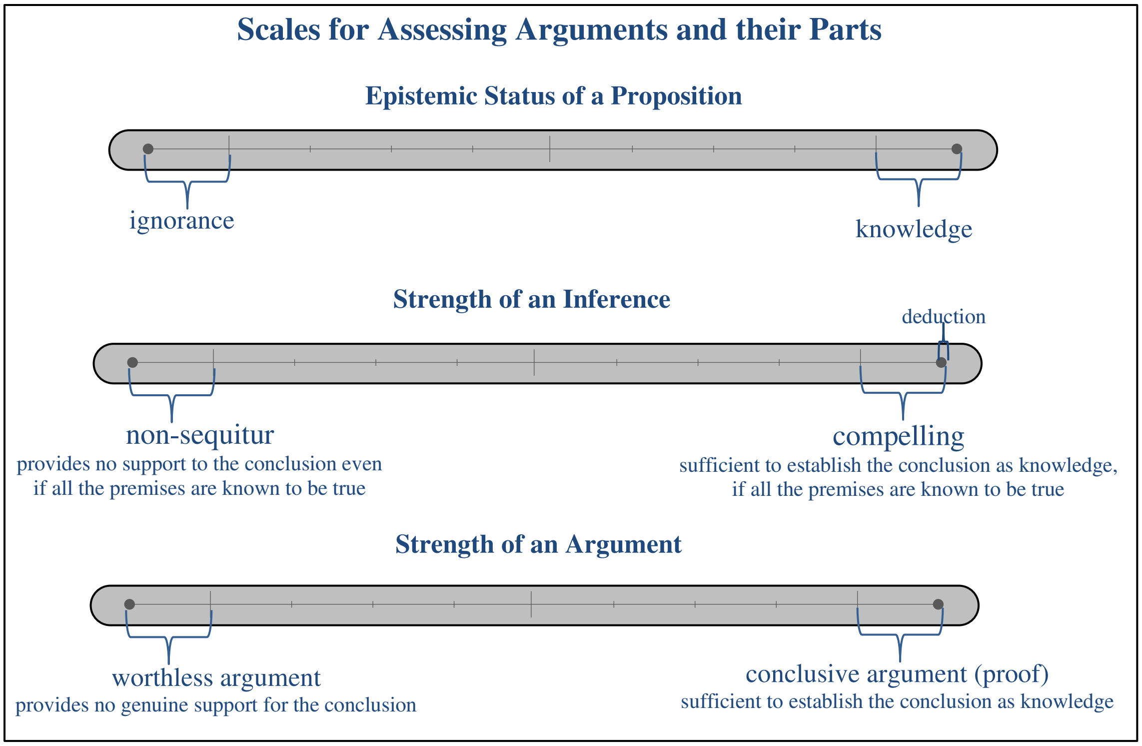

An argument as a whole can be no stronger than its weakest element (premise or inference). And in an argument the weaknesses compound, so if multiple elements have weaknesses, the whole will be weaker than the weakest part.

§1.5 How to Assess Argument

Once you have mapped an argument in order to assess it you have to first identify all of the unsupported premises—the propositions that serve as premises in arguments, without themselves being conclusions of other arguments. On the map, these will be all the boxes that do not have arrows pointing to them. You then must assess each unsupported premise and each inference. To assess a premise is to determine its epistemic status. Since the examples in this primer are fictitious, the assessments of the premises are fictitious as well. See section §3, below, on how to assess actual premises.

To assess an inference, assume that you knew all of its premises are true, and then ask yourself how strong a reason they would give you to believe that the conclusion is also true. How to assess different sort of inferences is discussed in §4, below, but you should be able to get an intuitive sense of how strong an inference is just by asking yourself if the premises were true how strong a reason would they give you to believe the conclusion.

Once you have assessed all the premises and inferences, you can then assess each argument as a whole. The argument can be no stronger than its weakest element (premise or inference).

Weaknesses within an argument compound, so if there are weaknesses in more than one element, the argument will be weaker than its weakest element.

For some sorts of premises and arguments, there are precise mathematical ways to evaluate their strength, but that sort of precision is not always possible. It is enough for our purposes to place premises, inferences, and arguments in rough regions of the scales that we are using to evaluate them.

Let's try this process, with Map 3 from above. Here the map is again:

Map 3:

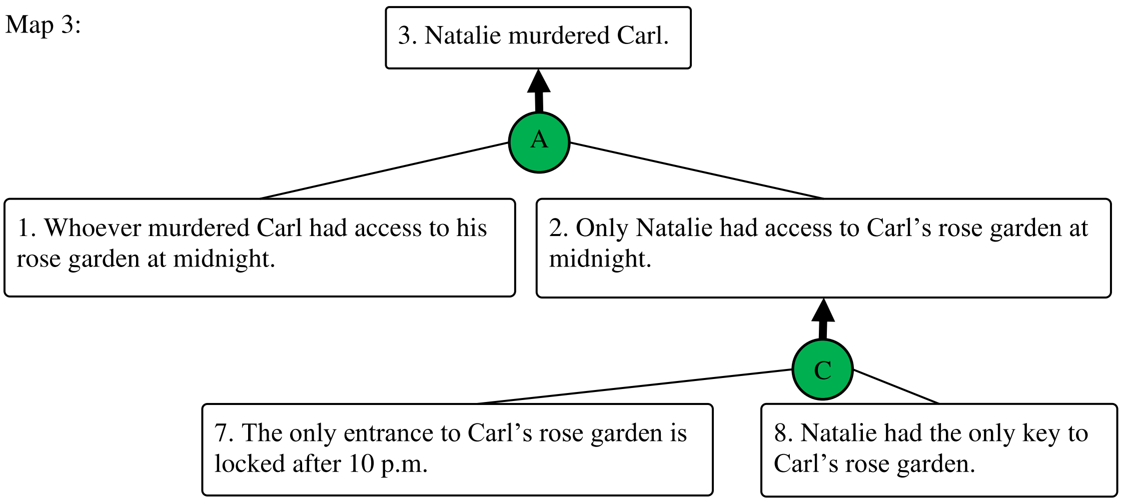

The first step is to identify all the inferences and unsupported premises. There are two inferences (A and C) and three unsupported premises (Propositions 1, 7, and 8). Proposition 3 isn't a premise at all, and Proposition 2 is a premise for Argument A, but it isn't unsupported, because it is the conclusion of Argument C. So, we will need to assess the premises and the inferences. In the map below I've added blank scales for the elements we will need to assess.

Map 3b:

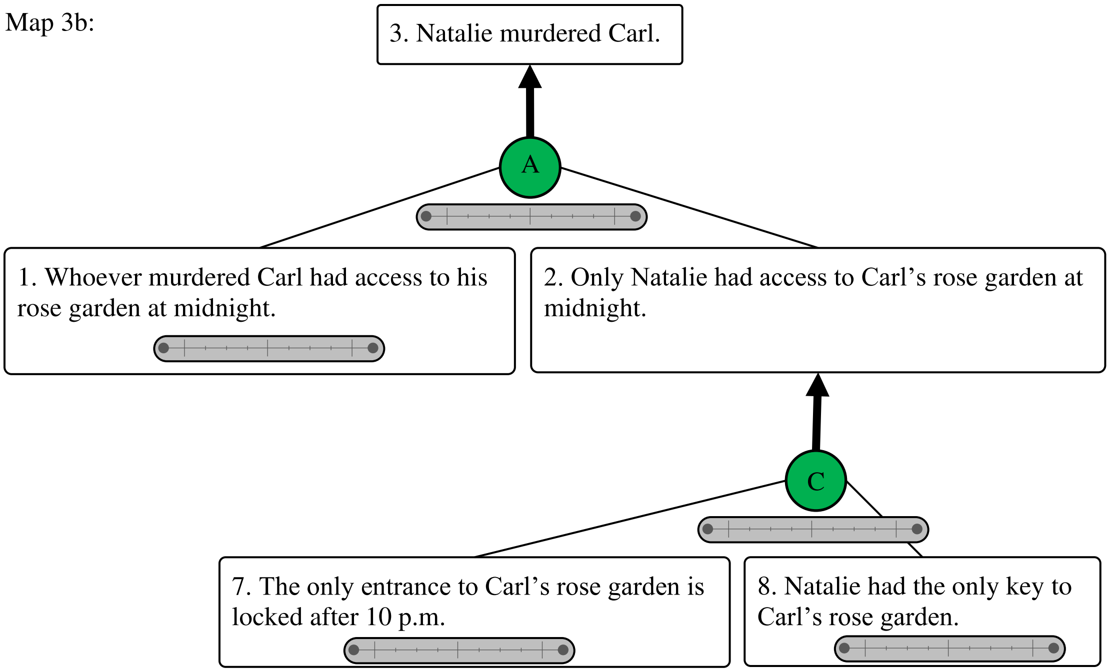

Let's assess the premises first. If this were a real-life argument, we would have to reflect on how strong a reason we have for believing Propositions 1, 7 and 8, but since the example is fictitious, we'll have to make up their epistemic statuses as well. I've done that in the map below:

Map 3a:

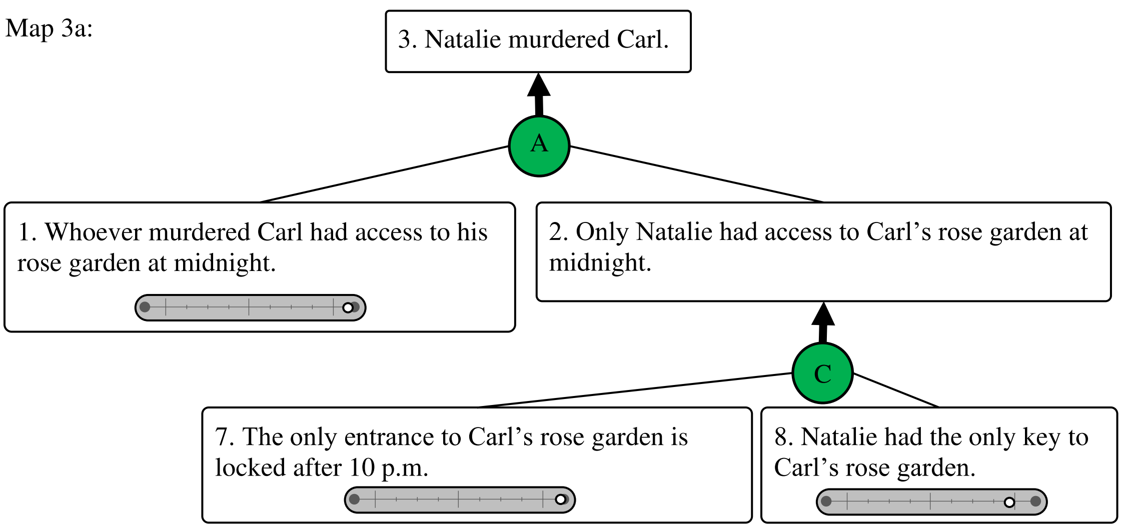

As the premises are assessed on this map, we know Propositions 1 and 7, but we don't quite know 8. Perhaps we think there's some possibility that a second key was made or that Natalie's key was stolen from her.

Now that the premises have been assessed we'll turn to assessing the two inferences.

Map 3b:

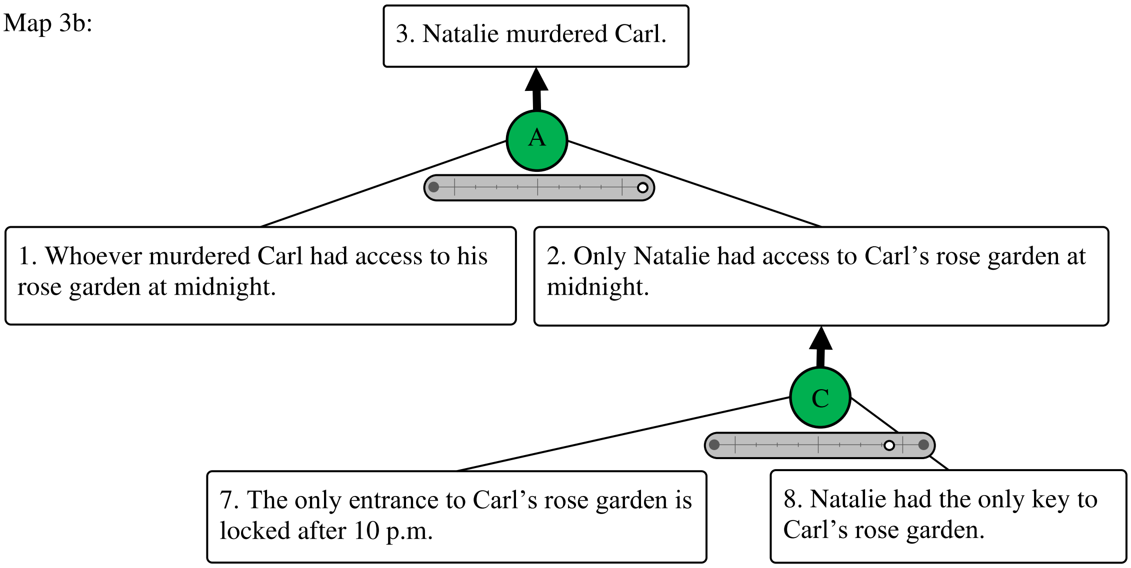

We've already seen that Inference A is a deduction, so it is as strong as can be. What about Inference C? It is definitely not a deduction, there is no *contradiction* involved in holding that someone other than Natalie had access to the garden at midnight, even though it is locked after 10pm, and Natalie had the only key. It's just *unlikely* that someone else had access, since keys are the normal way of accessing locked places, and locks are designed to keep people without keys out. Still it's possible for people to enter locked places without keys—locks can be picked, and presumably the rose garden has walls that can be climbed. To determine how strong Inference C is we'd need to think about how plausible these alternative routes of access are, and that would require some background knowledge. Assessing non-deductive inferences is more difficult than assessing deductions because it requires making use of such knowledge. In this case, since the example is fictitious, there is no background knowledge to rely on, so we'll have to make up more about the example. If we took it for granted that the lock is of a kind that is almost impossible to pick without leaving marks (that weren't found), that picking it would have taken time in which someone doing it would likely have been observed, and that the walls of the garden couldn't be scaled without sounding an alarm (that didn't sound), then I think this inference would be compelling. But in that case, it would be a lot clearer if the person making the argument had made these assumptions explicit by including them as premises. In any case, for the sake of the example, let's assume that we aren't in a position to quite rule out lock-picking and that the inference is strong but not compelling. That's how I marked it on the map above.

Once we have assessed all the unsupported premises and inferences, we can go on to assess the arguments as wholes. The map below incorporates all the assessments we've discussed so far, plus blank scales for assessing the two arguments.

Map 3c:

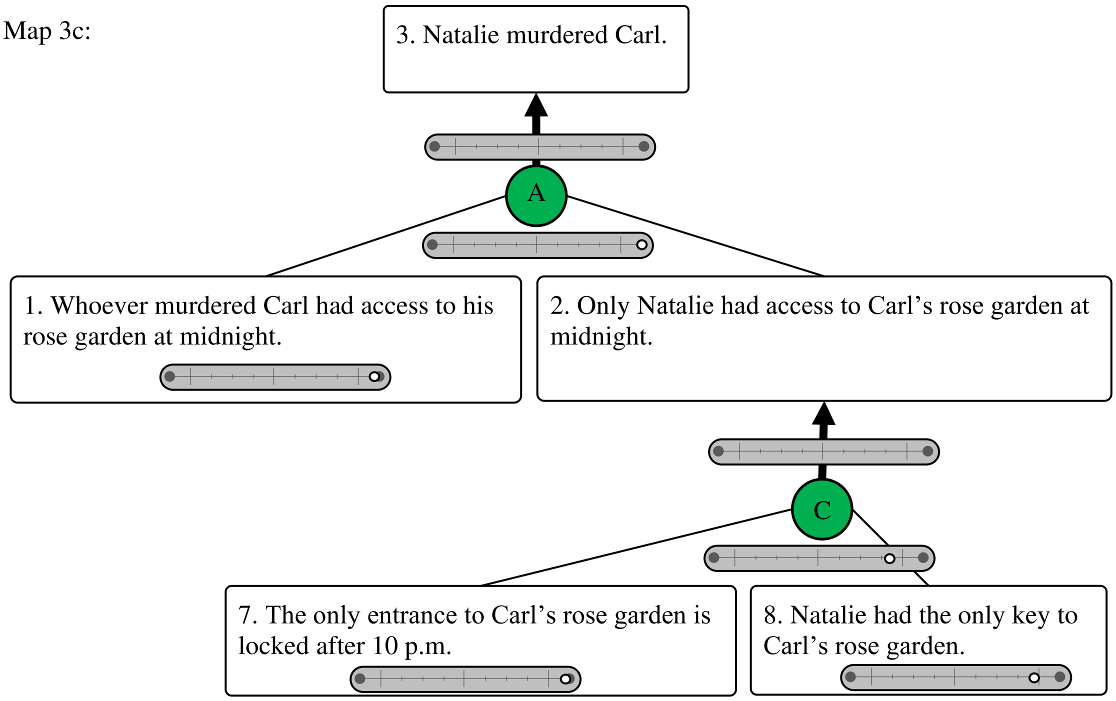

If any of an argument's premises are themselves supported by other arguments, we need to assess those other arguments as wholes before assessing the initial argument. So, in this case, we'll need to assess Argument C before we assess Argument A.

Argument C has one premise (Proposition 7) that is known, but it has two weaknesses. Proposition 8 is not (quite) known to be true, and Inference C is not compelling. Either of these weaknesses taken on its own is sufficient to prevent the argument from being conclusive. The argument as a whole can be no stronger than its weakest element, which (as we've filled out the scales above) is Inference C. But since the weaknesses in an argument compound, and inference C is not the only weak point, in this argument, the argument as a whole is weaker than Inference C. We can represent this on the map below by putting a mark on the scale for Argument C a bit to the left of the mark on the scale for Inference C.

Map 3d:

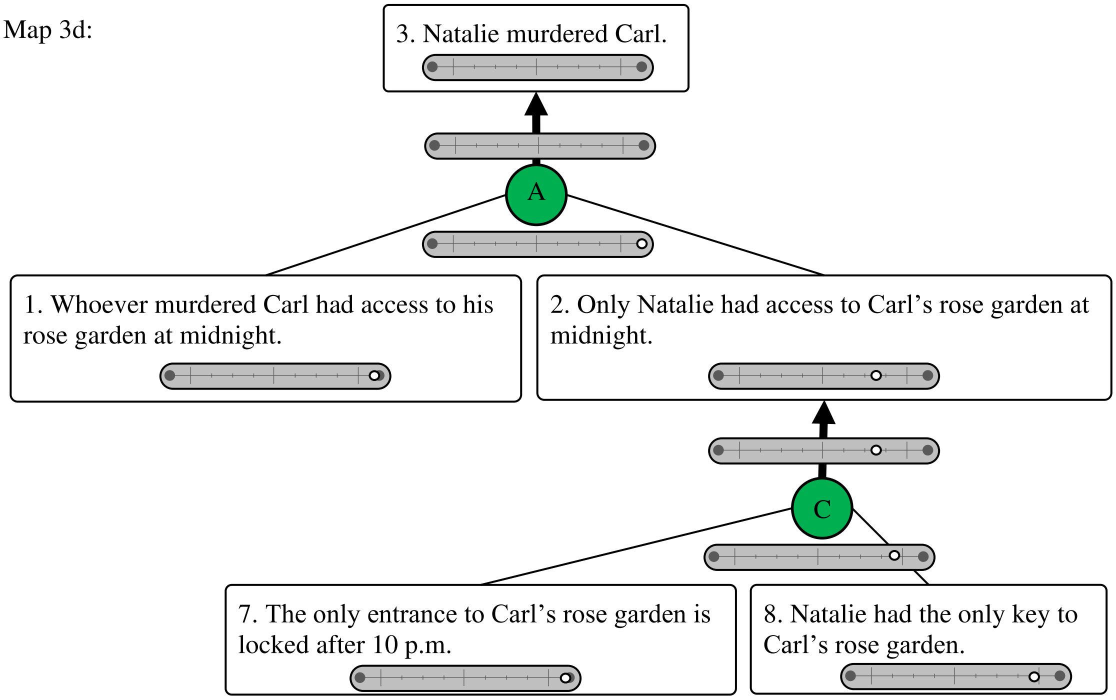

In addition to adding the assessment of Argument C to Map 3d, you'll see that I added a scale for Proposition 2 and marked it with the same assessment. This is because this map shows Argument C as our reason for believing Proposition 2. The reason for having separate scales for Argument C and Proposition 2 is that sometimes we will have multiple arguments for the same proposition. We'll discuss a case like this in a moment. Before we do, let's finish assessing the arguments in this map.

Map 3e:

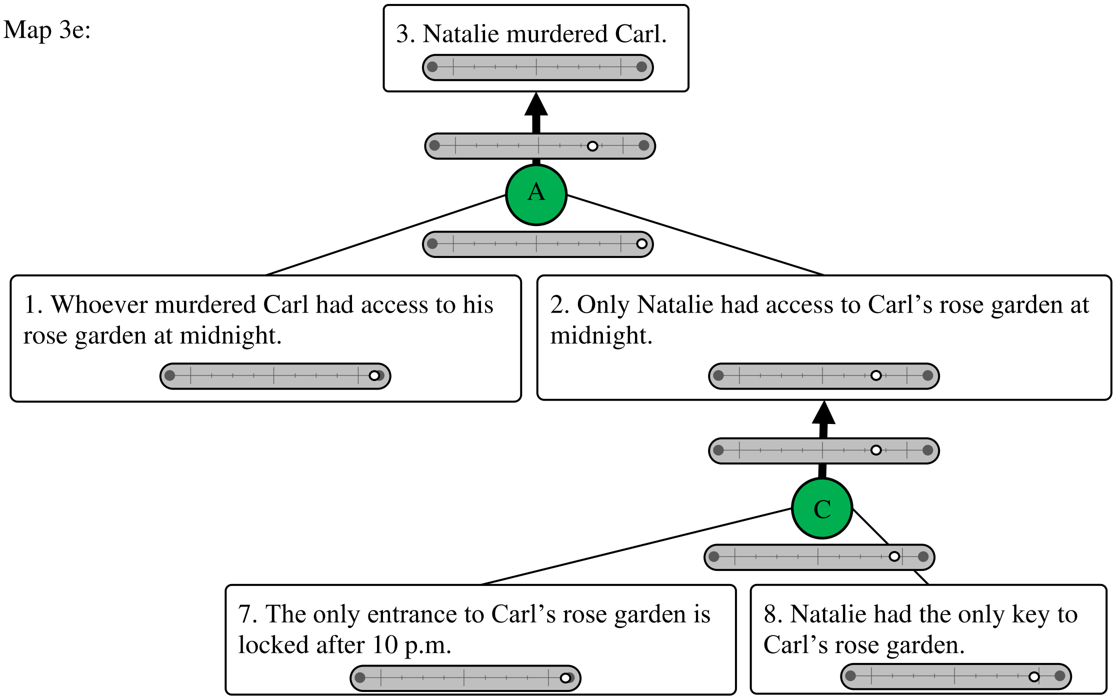

What remains is to assess Argument A. Here there is only one weak element. One of its premises is known the be true and the inference is a deduction, so the argument as a whole will be as strong as the remaining element, Proposition 2. The dot indicating our evaluation of Argument A is therefore placed in the same spot on its scale as we placed the dot on Proposition 2's scale.

If Argument A were the only reason given to accept Proposition 3, then we would give Proposition 3 the same epistemic status we've given Argument A. However, as we've already mentioned, conclusions are often supported by many, separate lines of reasoning. If that were the case, we would have to assess Proposition 3 in light of both the strength of Argument A *and* the strength of the additional arguments supporting it. This is shown on the map below (2a) which includes Argument A along with another argument (Argument B) for Proposition 3. (I'm omitting Argument C from this map to save space.)

Map 2a:

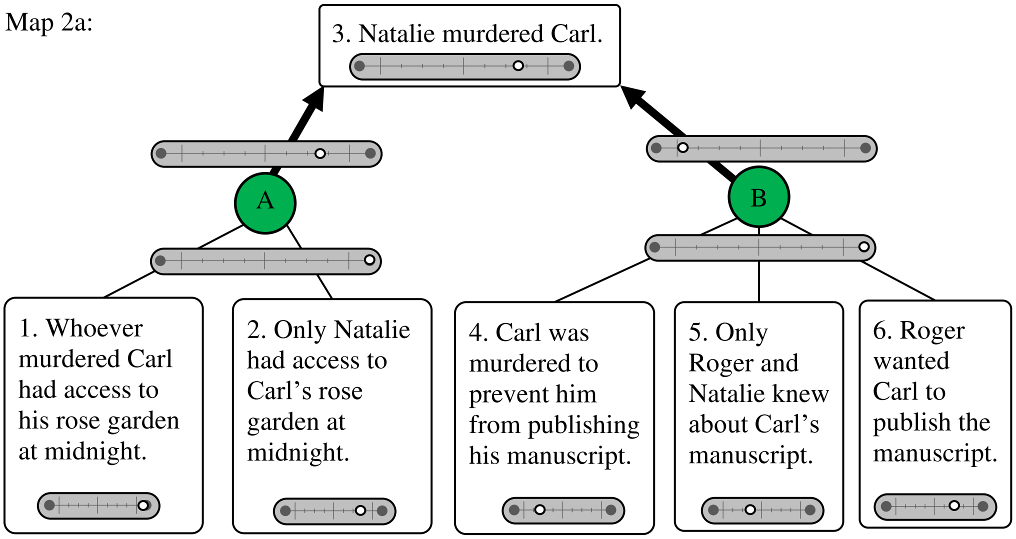

Since Proposition 3 is supported (in this map) by two different arguments, to establish its epistemic status, we need to take both into account. Argument A is fairly strong, though not compelling.

But as I've represented it here, Argument B is very weak. Inference B is a deduction, but the epistemic statuses I've given to its premises are much weaker than those of the premises of Argument A. (Recall that since these are fictitious arguments about fictitious people, the epistemic statuses are also fictitious.) Proposition 6 is approximately as strong as Proposition 2, but Propositions 4 and 5 are considerably weaker, and in an argument the weaknesses compound, so the argument as a whole is considerably weaker than its weakest premise. It's not quite worthless, but it's not worth very much. At best it could give one reason to *suspect* Natalie of the murder.

So, what epistemic status does Proposition 3 have based on these two arguments? Argument A gives us a pretty strong (though not conclusive) reason. Argument B doesn't add much to it, but a weak argument doesn't take away from the reasons given by a strong one, so over all we have about as much reason to believe the conclusion as Argument A gives us. It's not knowledge, but we should regard it as something that's probably true.

An argument is at least as weak as its weakest element, and if an argument has multiple weaknesses the weaknesses compound. But a conclusion is as at least as strong as its strongest argument, and if there are multiple arguments, their strength can compound.

There is one important caveat to this claim that a conclusion is as strong as its strongest argument. A conclusion's strength can be diminished if you have an argument *against* it. (An argument against a proposition is often called an objection or counterargument.) Consider the map below:

Map 8:

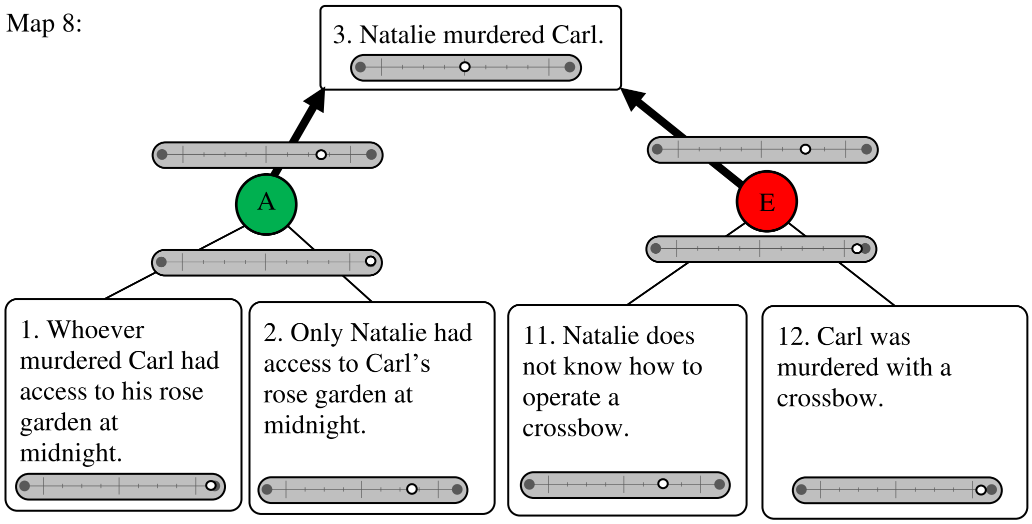

Argument E is an argument against Proposition 3, so whatever strength it has is going to counteract the strength of Argument A. Like Argument A, it is not conclusive. Its inference isn't quite a deduction, but it is compelling. Proposition 12 is known to be true, so the argument would be conclusive, if proposition 11 was also known. But it isn't, so the argument is correspondingly weaker. However, it is still pretty strong. Taken on its own it would lead us to think that Natalie probably didn't murder Carl (since to murder him she would have had to have used a crossbow, and she probably doesn't know how to use one). Taking it in the context of the whole map, it considerably reduces the epistemic status that Proposition 3 would otherwise have due to Argument A. That's why I marked it as just about in the middle in the map above.

This raises an interesting question. What would happen, if both Arguments A and E had been conclusive? That would mean that Argument A would establish that Proposition A is true, while Argument E would establish that it is false. But it can't be both true and false. So, if we find we're in that situation we know we've made a mistake somewhere in our assessment. As Ayn Rand puts it: "Contradictions do not exist. Whenever you think that you are facing one, check your premises. You will find that one of them is wrong."[4] I'll add that you should also check your inferences, since you may have misevaluated one of them. We will discuss how to check your premises in §3 below, as part of our wider discussion of assessing premises. We'll discuss how to assess different types of inferences in §4, and in §5 we'll revisit the issue of how to assess a conclusion in light of multiple arguments for and against it.

Before any of this, though, it will be helpful to say a bit about how to identify arguments in things that you read, and how to analyze them and construct maps. That's the subject of §2.

---

[4] *Atlas Shrugged*, 199. (Penguin Publishing Group. Kindle Edition.)

Glossary of Key Terms from §1

argument – a set of related propositions that is given as a reason for believing a further proposition

compelling – said of any inference that is strong enough that it would establish its conclusion as knowledge, if its premises were known to be true.

conclusion – a proposition that supported by an argument

conclusive – said of an argument that is strong enough to establish its conclusion as knowledge.

deduction – a kind of inference in which the premises necessitate the conclusion because it is contradictory to hold that the premises are true but the conclusion is false. Such inferences are as strong as it is possible for an inference to be.

epistemic status – the scale, running from complete ignorance to knowledge, along which propositions can be assessed.

inference – the mental act of moving from a premise to a conclusion

premise – a proposition given in an argument as part of a reason to believe the conclusion

proof – an argument that is strong enough to establish its conclusion as knowledge

proposition – a unit of thought of the sort which can be true or false, believed or disbelieved, and asserted or denied. Such units of thought can be expressed by declarative sentences.

non-sequitur – a kind of inference in which the premises (even if true) offer no support at all to the conclusion

worthless– said of an argument that offers no support for its conclusion

§2. How to Analyze an Argument

To analyze something is to break it down into its constituent parts, and to analyze an argument is to break it down into its premises and conclusion. I will focus on analyzing arguments presented by other people in written form. Of course, one hears arguments in conversations all the time, and the process by which one analyses them is essentially the same as with written arguments, but it is more difficult because one does not have a “fixed target” which one can take the time to study at one’s own pace. One also can analyze one’s own arguments as well as those offered by other people, and, indeed we will have occasion to do this over the course of the term, but in doing so it is important to achieve a certain critical distance from the argument, and this is best achieved by writing it out, and then treating it as though it were written by someone else.

§2.1 Finding the Arguments

When trying to analyze the arguments in a given text the first step is to identify which passages contain arguments. You need to single out those stretches of text in which one or more propositions are cited as a reason to believe another. There are many ways in English to indicate that one proposition is being offered in support of another. For example, we might say "I should respect her, because she's my mother," or "She's my mother, so I should respect her," "I should respect her, for she's my mother, or "She's my mother; therefore, I should respect her." In all of these cases "She's my mother" is being offered as a premise in support of the conclusion "I should respect her."

Map 8:

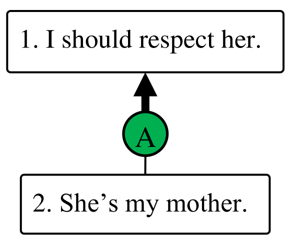

Some of ways this argument might be expressed:

- I should respect her, because she's my mother.
- She's my mother, so I should respect her.
- I should respect her, for she's my mother.
- She's my mother. Therefore, I should respect her.
- I should respect her; she is my mother after all.

Words like "so," "therefore," "thus," "hence," and "consequently" are often used to introduce conclusions; and words like "because," "for," and "since" often introduce premises. "Surely," "certainly," "no doubt," and other words that signal confidence in what one's about to say are also often used to introduce premises. Words of the sorts we've been discussing are sometimes called **particles**.[5] Looking out for particles can help you to identify arguments and adding particles to your own writing is a good way to convey the structure of your own arguments to readers. However, all of these particles also have other uses in English, and people sometimes argue without using particles at all.

Premises and conclusions can be indicated in other ways. For example, in some contexts, one can indicate that a proposition is a conclusion by saying that it "must" or "has to be" the case, but like particles, these words have other uses as well. Or someone could be very explicit and say, "I conclude that I have to respect her, on the basis of the premise that she's my mother." Or, swinging from one extreme to the other, he might express the same argument by saying simply: "She's my mother. I should respect her," or "I should respect her. She's my mother." And, in most contexts, if someone said this, you would recognize that he probably meant one proposition

[5] The word "particle" is used primarily in Classics (the study of the Greek and Latin). The Classicist J. D. Denniston defines it as "a word expressing a mode of thought, considered either in isolation or in relation to another thought, or a mood of emotion" (*The Greek Particles*, xxxvii).

to support the other, and you would be able to tell which was which, because you understand enough about the relations between the propositions to figure out what the author probably intends. Again, sometimes premises or conclusions can be expressed in the form of rhetorical questions: “Shouldn’t I respect her? After all, isn’t she my mother?” There is a wide variety of ways in which premises and conclusions can be expressed, and in which we are able to recognize that this is what is being done.

Some Particles often used in Argument

Used to introduce conclusions:

- so
- therefore
- thus
- hence
- consequently
- in conclusion
- it must be that
- then

Used to introduce premises:

- because
- since
- for
- surely
- certainly
- no doubt
- assuming that
- given that

§2.2 Identifying the Conclusion and All the Premises

Once you are confident that you have found an argument, you need to identify its premises and conclusion. In order to recognize that a passage contains an argument in the first place, you must have already noticed that at least one proposition is intended either to support or to be supported by another. Thus, you will have already identified either a conclusion or a premise. Now you need to identify any remaining premises or conclusions that there may be. In doing this keep in mind that they may be introduced with inferential particles, but that they needn’t be.

The argument may be presented in any order. For example, each of the sentences expresses the same argument as the map on the right.

(i) Alcohol should be illegal, because it’s a drug and all drugs should be illegal.

(ii) Alcohol is a drug, and all drugs should be illegal, so alcohol should be.

(iii) All drugs should be illegal, so alcohol should be, since it’s a drug.

Map 9:

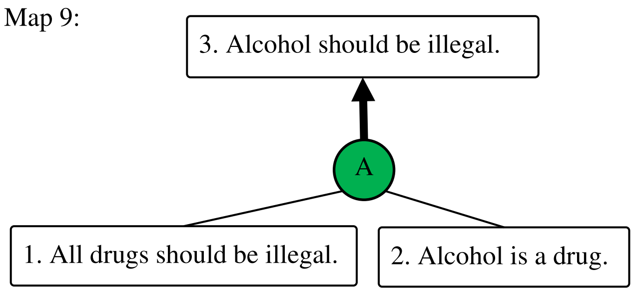

In (i), the conclusion is written first, followed by the two premises; in (ii) the conclusion is written after the premises; and in (iii), it is placed in between them.

To ensure that you have found all of the premises and the conclusion, read through the passage carefully, focusing separately on each proposition—each claim that *could be expressed as a separate sentence* (however it is actually formulated in the passage as written). Then ask yourself *why* the proposition is there. Is it intended as a part of the argument or as some sort of aside? If it is part of the argument, then *what role is it playing*: is it meant to be supporting some conclusion, or to be supported by some other proposition?

If you are having trouble figuring out whether one proposition is intended to support another or to be supported by it, it can help to ask yourself which proposition is more obviously true? In arguments we try to *establish* propositions that we are less sure of by inferring them from ones that we are surer of.

§2.3 Implicit Premises

You may have noticed that there's something a bit unnatural about the three sentences we looked at above expressing the argument that alcohol should be illegal. It is unlikely that anyone making this argument would state it so longwindedly. More likely he'd simply say: (iv) "Alcohol should be illegal because it's a drug" or perhaps (v) "Alcohol should be illegal because all drugs should be." Both of these ways of stating the argument omit one of the premises. People often do this when they think it is obvious what premise would be needed to complete their argument and when they think the person they're speaking with will agree to that premise. If one maps these arguments as written, here's what one would get:

Map 9a of (iv):

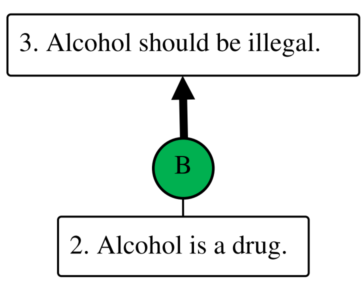

Map 9b of (v):

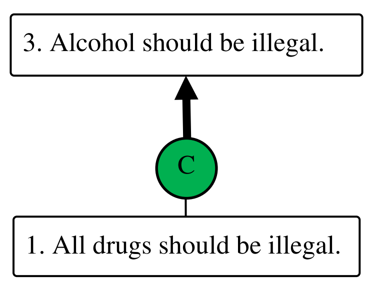

Inferences B and C are non-sequiturs, whereas Inference A (in Map 9, above) is a deduction. Moreover, it is obvious what premise you could need to add to Argument B (or Argument C) to make a very strong argument (namely, Argument A).

We encountered another example of this phenomenon in map 8, above. Here is that map again:

Map 8:

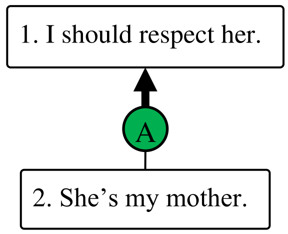

Some of ways this argument might be expressed:

- I should respect her, because she's my mother.
- She's my mother, so I should respect her.
- I should respect her, for she's my mother.
- She's my mother. Therefore, I should respect her.
- I should respect her; she is my mother after all.

This map is a map of the argument expressed in different ways by each of the sentences on the right. But the argument is clearly incomplete as written. Inference A is a non-sequitur, which makes Argument A worthless. But if someone said any of the sentences on the right, you would recognize that he was giving you some reason to believe Proposition 1. This is because there's another premise, which is plausible that when combined with Proposition 2 would make for a stronger argument as follows.

Map 8a:

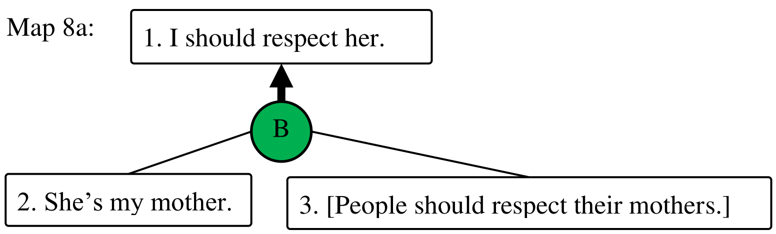

In this map, Proposition 3 is enclosed in brackets to indicate that it isn't stated in the passage we are analyzing and that we have added it ourselves, because we think that the author of the passage intended us to assume it as a premise of the argument. Such unstated premises are called **implicit**.

When analyzing an argument, it is important to make any implicit premises *explicit*—that is, to state them. This is necessary because, when you go on to assess the argument, you will need to assess *all* of the premises in order to determine how strong the argument is. Some arguments appear to be stronger than they are, because their weakest premises are left implicit.

Not every unstated belief held by a person making an argument is an implicit premise of that argument, nor even is every unstated belief that is relevant to the subject of the argument. We can probably imagine all sorts of reasons that the person making this argument has for believing that people should respect their mothers, but none of these reasons count as implicit premises of the argument mapped above. Something is an implicit premise *only if it needs to be added to an argument to prevent one of its inferences from being a non-sequitur*, and if is likely that the person making the argument intended you to assume it.

Thus, the process of finding implicit premises is closely related to the process of assessing the inference. Once you have identified the stated premises and the conclusion, you may notice that the conclusion *does not follow* from the premises. At this point, there are two possibilities: either the inference is a non-sequitur; or there is an implicit premise, which does make the conclusion follow from the premises. You need to use your judgment as to which is the case. Is it more likely that the author of the argument made non-sequitur or that he left one of his premises unstated? People rarely make arguments that include obvious non-sequiturs, so in such cases, it is likely that there is an implicit premise that the author intended you to assume. There are some subtle situations where it is difficult to determine whether an argument is bad or whether there is some implicit premise, and there are cases where it is hard to tell which of several different premises might be implicit; but more often than not it is very clear when someone is relying on an implicit premise and what that premise is.

§2.4 Multi-Step Arguments and Multiple Arguments to the Same Conclusion

In §1.3 above we saw that there are multi-step arguments, where a premise of one argument is supported by a further argument. You need to be on the lookout for this sort of structure when mapping. Here's an example:

> Jane's visit must have been over a weekend, since she spent two full days here, and she wouldn't have been able to do so during the week. But Rob wasn't in town, so the visit had to be on the first weekend in July, since that's the only one when he wasn't here.

Map 10:

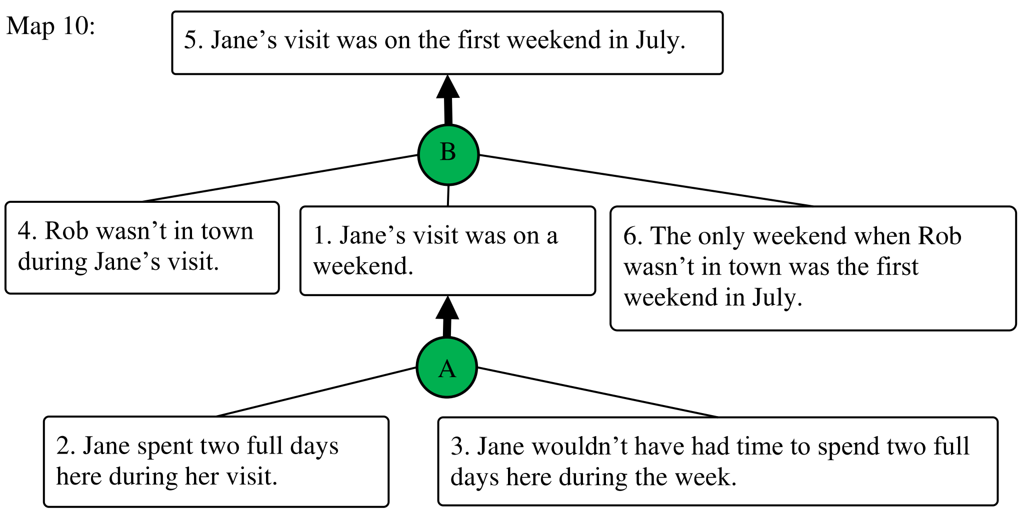

The first sentence of the passage gives us Argument A, with Proposition 1 as its conclusion, and the second sentence then gives us the remaining propositions in Argument B.

It is not uncommon in such multi-step arguments for some of the propositions to be left implicit. For example, here's another way in which someone might express the same argument mapped above.

> Jane wouldn't have been able to spend two whole days here during the week. But the only weekend when Rob was out of town was the first one in July, so her visit must have been over that weekend.

Here Propositions 1, 2 and 4 (from Map 10) are left implicit, but it is reasonably clear that the author of the passage intended the argument expressed by that map.

§2.5 Distinguishing Arguments from Explanations

Consider the following passage and the map to its right:

The sun is hot, because it is a ball of gases undergoing nuclear fusion, and nuclear fusion releases a great deal of heat.

Map 11:

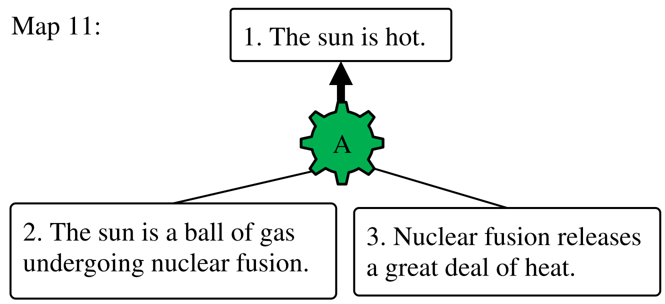

The use of the particle "because" may lead us to interpret this as argument, along the lines illustrated in the map. And if we knew Propositions 2 and 3, they would in fact give us a reason to believe Proposition 1. However, it is hard to imagine a situation in which someone would know Propositions 2 and 3 without already knowing Proposition 1, so it is unlikely that anyone would ever make this argument. The more natural way to interpret this passage is as giving us an **explanation** of Proposition 1.

An argument gives one a reason to believe that its conclusion is true, whereas an explanation cites the *causes* of a phenomenon. That Map 7 contains an explanation, rather than an argument, is illustrated by the use of a gear shape rather than an ordinary circle.

Often when we're trying to reach conclusions about things in the future, we use premises that are also causes. For example, we might conclude that it's about the rain by noticing that there are dark clouds and that such clouds cause rain. But when we're not reasoning about the future, we usually need to know that a proposition is true, before we try to discover its causes. That's certainly the case in the passage above. We first know that the sun is hot, and then we try to discover the causes that explain why it is.

You can usually tell from context (and sometimes from the nuances of how inferential particles are used) whether a passage is meant to explain a proposition or to argue for it. If you're unsure, it can help to ask yourself what question the passage is answering about the relevant proposition. If it's an argument, it will be answering the question "How do you know it?" (or "What reason do you have for believing it?"). If it is an explanation, it will be answering the question "What caused it?"

It is very easy to confuse an explanation for an argument when what is being explained is a person's beliefs or actions. It is possible to think of a person's actions or beliefs as effects and to try to explain them, often by citing biographical facts. Suppose that Charlie spanks his children, and we ask ourselves *why* he does this. Here are two answers we might come up with:

(i) Charlie spanks his children because he was spanked by his father.

(ii) Charlie spanks his children because he thinks it is the most effective way to discipline them.

Notice that (i) gives us an explanation of Charlie's behavior, by citing things in Charlie's past that might cause him to behave as he does, but it doesn't give us Charlie's reasons for acting in this way. By contrast, (ii) indicates what Charlie's reasons might be.

Now consider another example in that concerns a belief rather than an action. Suppose that Dana believes that it is wrong to eat meat, and we ask *why* she believes this. Here are two answers we might get.

(i) Dana's parents believed that it is wrong to eat meat.

(ii) Eating meat causes suffering.

We can map these two answers as follows:

Map 12 of (i):

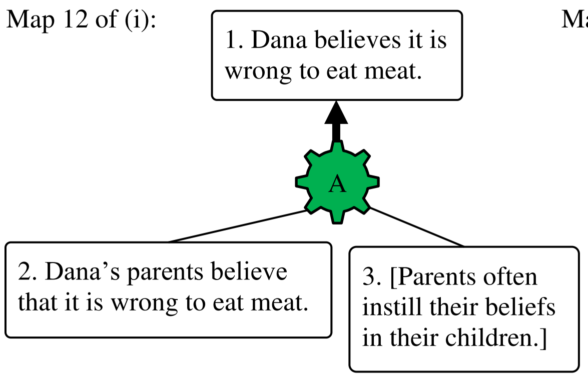

Map 13 of (ii):

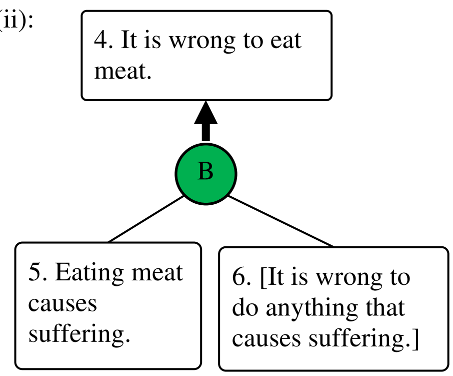

Notice that (i) explains Dana's belief by citing a factor in her biography that caused her to come to this belief, but it doesn't give Dana any reason for believing as she does. It doesn't help Dana or us to tell whether her belief is true. By contrast, (ii) gives something that might be Dana's reason for believing as she does. It gives an *argument* that the belief is true.

Explanations of our actions or beliefs treat these behaviors and beliefs as things that just *happen* to us. But our beliefs and actions don't just happen to us. You are *responsible* for the things you do and for the things you believe. This is why you need to think about the reasons you have for your beliefs and actions, and why you need to think about and evaluate other people's reasons as well, when judging them. Confusing explanations of beliefs (or behaviors) with arguments for them can obscure these reasons.

§2.6 An example of a complex map

An example of a more complex map than we've looked at so far. Here's a brief passage, followed by a map of all the arguments it contains.

> "The laws prohibiting cannabis use should be repealed. They're illegitimate in the first place, because the only proper basis for outlawing an activity is that it violates someone else's rights, and you're not violating anyone's rights if you smoke a joint. Anyway, cannabis is way less dangerous than substances that it's legal to buy and use. Tobacco causes cancer, whereas cannabis is being researched as a potential cancer cure! No one's heard of a "cannabis overdose," but it's easy to kill yourself by overdosing on ibuprofen, which you can buy over the counter, and people die every year of alcohol poisoning. Some studies show that there are risks to driving under the influence of cannabis, but a stoned driver is way safer than a drunk driver. Yet people are allowed to go into any super market and buy a bottle of wine without being harassed by the cops, and our government is spending untold sums arresting people who buy or sell pot. Even people opposed to cannabis use should be able to see that this is money is wasted, since it's not stopping anyone from smoking up. And, by the way, white people smoke up every but as much as anyone else, but somehow the majority of people arrested for cannabis-related offenses are black or Latino, which shows how racist the law is in practice. Instead of throwing away money on half-assed, racist enforcement of these illegitimate laws, the government could be making money on weed, by legalizing it and taxing it. That's what Colorado did, and their tax revenues are way up, so we know it works. Crime and poverty rates also went down in Colorado, so if we want to make the rest of the country safe and rich, we might try following their example."

Map 14

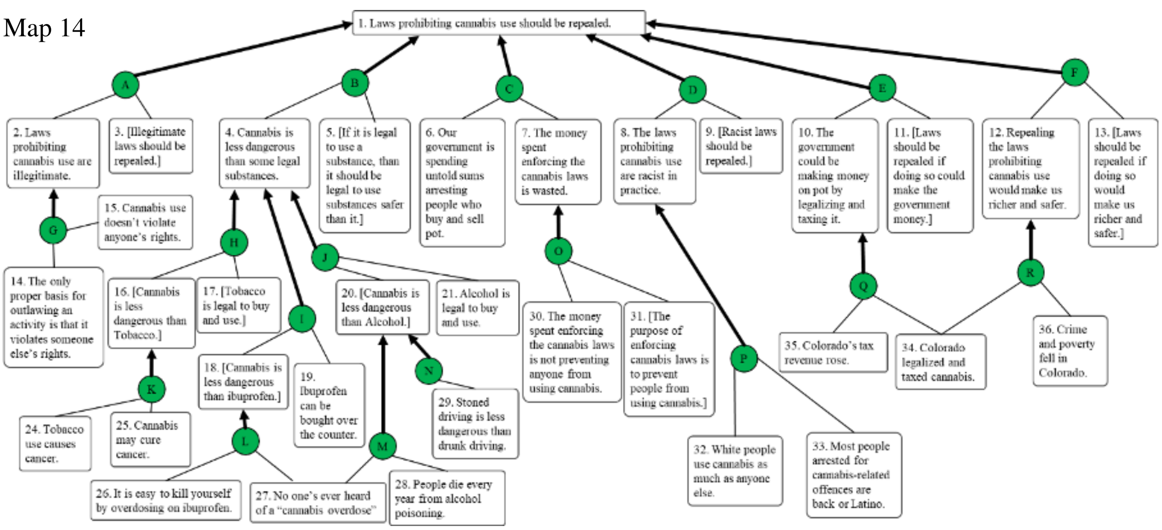

§3. Assessing Premises

An argument is only as good as its premises. Therefore, if the conclusion is to be *known* on the basis of the argument, then premises must themselves be known to be true. However, arguments from premises that are not quite known, can still give us reason to believe their conclusions. This is because there is a continuum of states between fully knowing something and being wholly ignorant of it. We can call each state along this continuum an *epistemic status*. An argument cannot confer on its conclusion an epistemic status stronger than that of its weakest premise. Thus, the primary task in assessing the premises of an argument is to determine the epistemic status of each. We will discuss the different epistemic statuses shortly, after dealing with a few preliminary points.

§3.1 Epistemic Status as Relative to An Audience but Objective

To be known is to be known *to someone*, and different people know different things. In an argument, premises are being offered to *someone* as a reason to accept some conclusion. Let's call the person or people to whom an argument is being given the audience. Clearly, then, in assessing the premises of an argument, what we need to figure out is whether *the audience knows them*—or, more generally, what their epistemic status is for the audience. When we speak of epistemic status, we always mean the epistemic status of a particular proposition for a particular audience. And the same proposition may have a different epistemic status for different audiences.[6]

We must take care, however, not to confuse a proposition's epistemic status for an audience with how firmly the audience believes it. By the firmness of a belief, I mean how steadfast the audience is in holding it. We can think of this as how easy it would be to talk someone out of the belief. The stronger a proposition's epistemic status is, the more steadfast it is reasonable to be in believing it. It would be unreasonable of someone to get talked out of something he knows to be true. (Indeed, if he was talked out of it, we would probably conclude that he didn't really *know* it in the first place.) However, people can sometimes be quite firm in beliefs that have a very low epistemic status. One example would be a fervent racist who, despite all the evidence to the contrary, persists in the belief that the members of other races are inferior to members of his own. Another example would be a self-deceived husband who clings to the belief that his wife is faithful to him, despite strong evidence that she is having an affair.

One way to formulate the distinction between firmness of belief and epistemic status is to say that the former is *subjective* and the latter *objective*. Something is subjective if it is determined by someone's feelings or opinions. It is objective if it is determined by the facts. Though different people know different things, whether a particular person knows something or not is not determined simply by whether he *thinks* or *feels like* he knows it, since many people think they know things that they don't. Nor is something's epistemic status for someone a matter of his

[6] To illustrate this, imagine a large party at which Al, Barbara, Carla, and Dan are all in attendance. Al and Barbara hardly ever socialize without one another, and both Carla and Dan know this. While Barbara is waiting in line to get refreshments, Al heads to the bathroom, running into Carla on the way. When Carla sees Al, she infers that Barbara is also at the party, though she can't be quite certain of this, because it could be one of those rare occasions when one of the two socializes without the other. In the meantime, Dan sees Barbara waiting in line and infers that Al is at the party. Like Carla he is not quite certain of his conclusion. So, Carla and Dan both believe that Al and Barbara are both at the party. Carla knows that Al is at the party and the proposition that Barbara is there has a weaker epistemic status, whereas the situation is reversed for Dan.

feelings or opinions about it. The epistemic status of a proposition for a given person is determined by facts about such things as the observations he has made and the arguments he has.

When evaluating arguments for certain purposes, is important to consider the (subjective) firmness of an audience's beliefs. The more firmly the audience believes the premises of an argument, the more persuasive it will find the argument. So, if we were interested in arguments primarily for the purposes of persuading other people, we would have to evaluate the premises with an eye to how firmly the audience believes them. This is how we might proceed in a rhetoric class or if we were considering whether to use an argument in an advertisement or a political campaign. However, this is not why we are interested in arguments in this class. We are interested in arguments because inference is a crucial *means of knowledge*. From this point of view, what makes an argument good is not that is persuades anyone of anything, but that it puts the audience in a position to know that the conclusion is true (or else brings the audience close to such a position). Therefore, what is relevant to us when assessing premises is not how strongly the audience believes them, but what their epistemic status is for the audience.

And the audience we are primarily interested in is ourselves. After all, we are studying arguments because we want to assess the arguments on which our own beliefs are based and to reach new knowledge. Therefore, when assessing the premises of the arguments we consider in this class, you should be focused on the epistemic status of these premises *for you*. Do you know the premises to be true? Are you entirely ignorant as to whether they're true? Or does your state with respect to them fall somewhere along the continuum between knowledge and ignorance? If so, where along that continuum? In asking and answering these questions, you need to keep in mind that whether you know something is not simply a matter of how sure you *feel* about it, but a matter of how objectively strong your *reasons* are.

§3.2 The Continuum Between Knowledge and Ignorance

To get a sense of the continuum between knowing something and being wholly ignorant of it, it helps to consider some examples. Let's start with the following scenario: While walking across the quad one day at 8:00am, you see someone at a distance whom you almost recognize as your acquaintance Bob, but you lose sight of him before you can be sure. If you had gotten a better look at the man and recognized him as Bob, then you would know that Bob was on the quad at 8:00am.[7] As things stand, however, you do not know it. Still, you are not in the same position with respect to the proposition "Bob was on the quad at 8:00am" as you would be if you hadn't seen what you did. Your experience on the quad gives you *some reason* to believe the proposition, without making you *certain* of it.

---

[7] At least, most people ordinarily take people in this sort of situation to have knowledge. There are some philosophers who call into question whether we really do know, in such cases, but we can put their views aside for the time being. We will have occasion to discuss them later in the term.

Notice that you might go on to use this proposition in arguments. Suppose, for example, that a bank robbery was committed in Manhattan at 8:03 and Bob was suspected of being behind the wheel of the getaway car. If you knew that Bob was on the quad at 8:00, you could be certain, via the following argument, that these allegations are false.

Map 15:

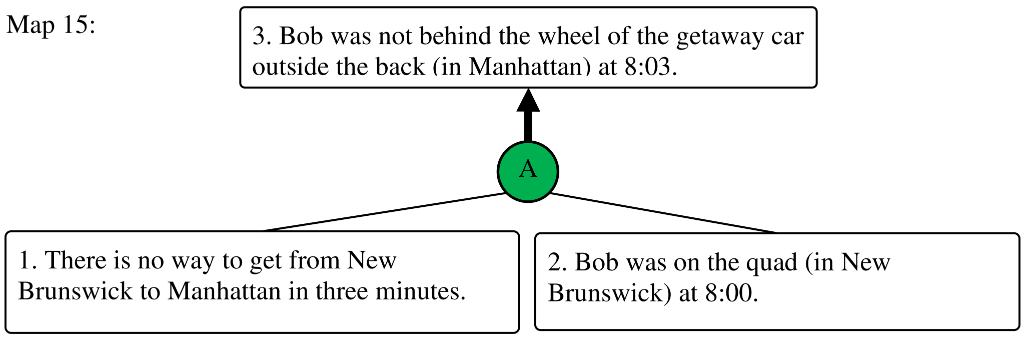

As things stand, however, you do not know Proposition 2, and so you do not know the conclusion, Proposition 3. But, because you have some reason to believe Proposition 2, you have some reason to believe the conclusion as well. Perhaps, then, you should go to the police and make a statement. If you do, how much weight should the police give to your testimony? Let's consider some of the questions that they might ask you (and that you could ask of yourself):

> For how long did you see him? At what distance? How good was the visibility? Did you focus on him, or was he in the periphery of your vision? How familiar are you with Bob's appearance in the first place? Is it based on vision alone that you think you recognized him, or did you hear his voice as well, or smell that unusual cologne that he's always wearing?

In addition to these questions, you might also consider how distinctive Bob's appearance is: is he someone who is easily recognized in a crowd, or someone who can easily be mistaken for someone else? Notice that for each of these questions there is a range of possible answers, and that how sure you could be that Bob was on the quad depends on where along that range your answer falls.

Another way in which you might have some reason to believe something without being certain of it is by remembering it vaguely. This would be the case with someone who did see Bob on the quad recently and recognized him but couldn't recall with certainty *when* the encounter occurred. To get a sense of the considerations that are relevant to determining the epistemic status to this person of the proposition that Bob was on the quad at 8:00am we could make up a list of questions similar to those that we considered above for the case of seeing Bob on the quad. A third way that one might come to have an intermediate epistemic status with respect to a proposition is by inferring it from earlier propositions. In the example we've just been considering, your conclusion that Bob couldn't have been behind the wheel of the getaway car, has such a status because it is inferred from uncertain premises. But, as was mentioned briefly earlier, there are inferences in which, even if the premises are certain, the conclusion cannot be inferred with certainty. Later, when we consider some of these types of arguments in greater detail, we will get a sense of the factors that would lead to the conclusions being nearer to or further from certainty.

§3.3 The Epistemic Statuses: Certainty, Probability, Possibility, and Unfoundedness

In evaluating arguments, it is necessary to determine the epistemic status of each premise—where it falls along the scale running from ignorance to knowledge.

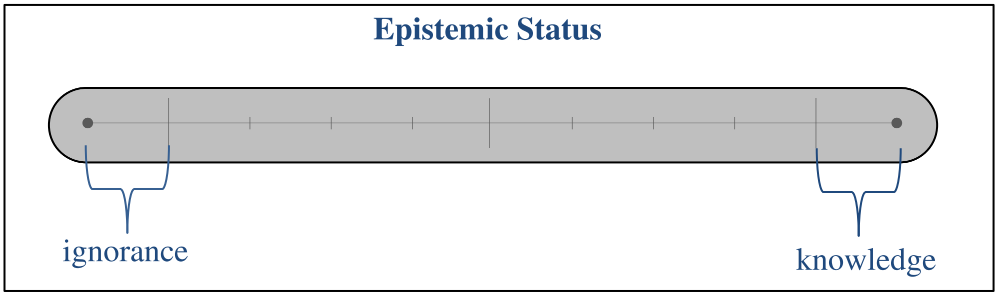

In some unusual cases we might be able to specify this to such a fine degree of resolution that we can assign a number to it, and different scales have been devised for assigning such numbers. In most cases, however, this is not possible, and we distinguish epistemic statuses more coarsely using terms like "unfounded", "possible", "probable", and "certain". In order to get clearer on epistemic status, we will need to discuss each of these terms.

Let's start with **certainty**. We've been using this word a lot in the last few paragraphs, in a way that makes it seem to be nearly synonymous with "knowledge". Understanding the relation between these two concepts will help us to understand epistemic status better. To be *certain* of a proposition is to *regard it as knowledge*, as opposed to regarding it as doubtful. When you are certain of something you act on it and draw inferences from it confidently, whereas when you are uncertain you are more tentative, always keeping in mind the possibility that the proposition is false and planning for that contingency. In this sense, people are sometimes certain of things unreasonably. The fervent racist discussed above, for example, might not hesitate to act on his belief. However, when we speak of certainty as an epistemic status, we are referring not to the way a person actually does regard the proposition but to the way it is *reasonable* for him to regard it. We might call this *rational certainty* or *objective certainty* (as opposed to *subjective* certainty). In any case, this is how I am going to use the word "certain" going forward.

There are some disputes among epistemologists about the relationship between certainty and knowledge. However, the following is comparatively uncontroversial: *the beliefs that a person is entitled to classify as knowledge and those that he is entitled to classify as certain are the same*. The difference between classifying them as certain and classifying them as knowledge is that, in classifying them as certain, he is contrasting them specifically with beliefs that have a lower epistemic status and is focusing on the fact that he is in a position to act on these beliefs and to draw inferences on these beliefs without hesitation.[8]

In thinking about the different epistemic statuses, it is often helpful to consider how they come up in the criminal justice process, when police officers, judges, and juries often need to weigh evidence, and to specify how sure they are of various propositions. The concept of "certainty" is central to criminal trials, where the jury is instructed to return a guilty verdict only if they are certain that the defendant is guilty—only, to use a familiar phrase, if his guilt has been established "beyond a reasonable doubt." In general, we can think of certainty as the state we are in with respect to a proposition when there is no reasonable doubt about it. Thus, the proposition that Bob was on the quad is not certain because there is a reasonable doubt about it: you didn't get a good enough look at the person in question to eliminate the possibility that it was someone else who vaguely resembled Bob. Whereas, if you had seen Bob clearly and heard and recognized his voice, these doubts would have been assuaged, and you could be certain that Bob was on the quad.

Of course, there are still doubts that someone might raise here, based on such farfetched scenarios as Bob impersonators, holograms, or hallucinogens; but most of us, in most contexts, would consider such doubts *unreasonable*, and certainly they would be ruled out as unreasonable in court (unless there was some specific evidence for them in a certain situation). As you might imagine, there is room for argument about what sorts of doubts are reasonable in what contexts and about when we have certainty. Indeed, some philosophers think that there is very little about which we can be genuinely certain—that there is very little that we really *know*—while others think that we know a great deal. We will have occasion to discuss this debate later. For now, let's put this issue aside and proceed on the assumption that we can be certain of such things as that the people we see in front of us are really there. This is an assumption that we all do make in our daily lives.

If a proposition fails to be certain, it may still be **probable**. A proposition is probable if the evidence is strong enough that it is more likely to be true than not, so that it is reasonable to assume it provisionally, while still making allowances for the possibility that it is false. For example, suppose that on your way to class you and one of your classmates, who you don't know well, both stopped at an ATM to make a withdrawal and you happened to see his receipt and notice that the balance was $602.47. Half an hour later, during the class, it would be probable that this was still his balance. Since you have been with him in the intervening time, you would know that he hasn't made any further withdrawals or deposits or used a debit card. But you could not be certain of the balance, since you do not know whether anyone else has access to the account, or whether any previous transactions posted to the account during this period; and you know that these sorts of events regularly happen with bank accounts. Still, because thirty minutes

[8] Among the disputed issues are whether it is possible to be certain of something that one doesn't actually know and whether it is possible to know something without being certain of it.

is so short a time, the odds are against any of these things having happened in this period, so it is more likely than not that the balance has remained the same.

I discussed earlier how juries are instructed to convict in criminal cases only when they are certain that the defendant is guilty. In civil cases (that is, lawsuits) there is a less rigorous standard and only probability is required. In legal terms, a jury should find for the plaintiff and order the defendant to pay damages if "the preponderance of the evidence" favors the plaintiff's case. This means: if the plaintiff's case is more likely to be true than false, given all the evidence presented.[9]

A proposition that is not probable may still qualify as **possible**. To call a proposition "possible" in the relevant sense is to say that there is *some reason* to believe it, and that it is therefore reasonable to regard it as something that "might be" true and to take it into consideration in our thinking. Thus, in the example from above concerning Bob, it is at least possible that he was on the quad at 8:00am (even if it turns out that it is not probable). He *might* have been there.

Perhaps you're thinking that "anything is possible" and that you would be in a position to say that Bob "might have been on the quad" even if you hadn't had the experience of seeing someone who you thought you recognized as him there. There is a sense of the word "possible" in which you could say that it was possible that Bob was on the quad, even if you hadn't had the experience, but there is another sense of the word in which this would be false. And it is this sense—let's call it **epistemic possibility** that is relevant for our present purposes.

To see this, imagine calling the police to tell them of the possibility that Bob was on the quad at 8:00am, three minutes before the robbery in Manhattan. They would ask you *why* you thought it was possible, and they would not react favorably if you responded that "anything's possible". They would have quite a different response, however, if you told them about your experience of thinking you recognized him. (Indeed, this report could prove quite useful to them. If Bob has been claiming that he was on the quad at that time, your report could make this alibi considerably more credible.)

There are other contexts in which we can see the idea of epistemic possibility at work in the law. We are all familiar with the idea of a "suspect"—that is, of a person whom the police think *might* have committed a certain crime, and whom they set to work investigating. There may be a number of suspects for a given crime but notice that the police don't regard *everyone* as a suspect, nor do they consider everyone who had the opportunity and ability to commit it a

---

[9] It is worth reflecting on the reason for this difference between criminal and civil law. In a criminal case, what is being decided is whether a person deserves to be punished, and it would be a grave injustice to punish him for a crime he didn't commit, so it is important to be certain that he is guilty. In a civil case, however, what is being disputed is which of two parties should have to bear a certain cost. For example, I might sue you for $1,000 to repair my car, claiming that you caused the damages. In this case, the court needs to decide between forcing you to pay the money and leaving me to pay it. And, if the preponderance of the evidence favors the conclusion that you caused the damage, then even if the jurors had some reasonable doubts, it would be unjust of them to leave me to pay for the damage that you probably caused.

suspect, unless there are only a few such people. (Recall the case discussed earlier, in which the fact that Natalie was one of the several million people who had access to the garden in which Carl was killed didn't give us any reason to suspect that Natalie was the murderer.) In general, the police need some specific reason to suspect someone of a crime, and likewise we need some specific reason to suspect a proposition of being true—that is, to classify it as "possible".

Epistemically possible propositions are to be contrasted with **unfounded** ones. A proposition is unfounded (or "arbitrary" or "baseless") when there is *no reason* to believe it. This is the status of things that you make up out of thin air—for example, that there are monsters under your bed, that your roommate committed a murder five years ago, or that his second cousin once lived in Manhattan. Some of these propositions are more far-fetched than others—there are no such things as monsters, and very few people commit murders, but many people live in Manhattan. However, all these propositions have in common that you have *no reason* to think they are true, or even to consider them, and the proper course of action is to dismiss them out of hand.

Because we have no reason to believe them, unfounded premises can offer *no support whatsoever* to any conclusions that we might infer from them. Thus, unfoundedness is the lowest epistemic status.

In some cases, we not only have no reason to believe that a proposition is true, we actually have reason to believe that it is false. This happens when the proposition contradicts something that we know (or have reason to believe) to be true. For example, if you knew that that your roommate only had one second cousin and that she spent her whole life in Wyoming, you would know that the proposition that she lives in Manhattan was false—or, to put it more simply, you would know that she did not live in Manhattan. (By the same token, if it were *probable* that she had spent her whole life in Wyoming, it would be *probable* that she didn't live in Manhattan.) In certain respects, propositions that are certainly or probably false can be thought of as having an even lower status than that of unfounded propositions. However, as far as their value as premises is concerned, an unfounded premise is no better than one that is certainly false: neither gives us any reason at all to believe any conclusion.

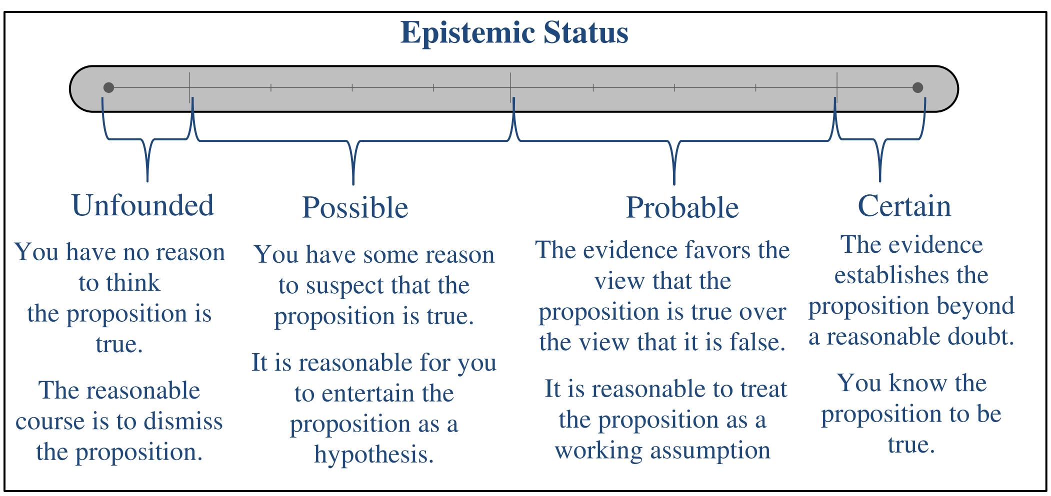

Thus, we can think of the continuum of epistemic statuses as running from unfoundedness to certainty, with the "possible" and "probable" denoting ranges in between. Of course, among the propositions which are possible, some will be nearer to being probable than others, and, within the probable propositions, some will be nearer than others to being certain. To capture this, we often use the words "probable" and "certain" in a comparative (rather than absolute) sense and say that one proposition is "more probable" or "more certain" than another. This should not be taken to imply that either proposition is probable or certain. For example, when two propositions are merely possible, one may nevertheless be more probable than the other. A proposition is probable in the absolute (rather than comparative) sense when it is more probable that it is true than that it is false—or, as we put it earlier when the preponderance of the evidence favors it. And a proposition is certain (in the absolute rather than the comparative sense), when it is no longer epistemically possible that it is false (that is, when it has been established beyond a reasonable doubt.)

We often assert propositions tentatively, indicating that they are probable rather than certain, by qualifying them with adverb "probably" (for example, "Bob was probably on the quad"), and we usually assert possibilities by saying that they "may" or "might be" the case ("Bob may have been on the quad").

Map 13:

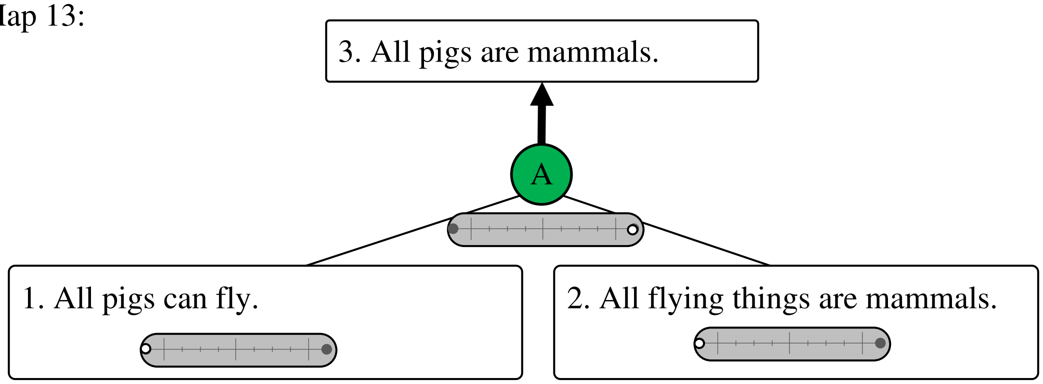

Additionally, the fact that a conclusion can be inferred from false premises does not give us any reason to believe that the conclusion is also false. In fact, you can take any true proposition and make up arguments that infer it from false premises. Here's an example:
Since the premises are false this argument doesn't give us any reason to believe that pigs are mammals, but it certainly doesn't show that they're *not* mammals!

§3.4 The Structure of Knowledge and the Danger of Circular Reasoning

Because we can acquire knowledge by inference, our knowledge of some things depends on our knowledge of others. If, to return to an earlier example, we know that Natalie murdered Carl, because we know that the murderer had access to his rose garden and that only she had such access, then our knowledge that she is the murderer depends on our knowledge of these two premises. The same pattern holds for lower epistemic statuses. If it were merely probable that the killer had access to the garden, then it would be merely probable that Natalie was guilty, and the probability of her guilt would depend on that of the premises. For the sake of simplicity however, let's ignore these other epistemic statuses for the time being and restrict our discussion to knowledge.

Our ability to use the conclusion of one argument as a premise for another leads to knowledge having a sometimes-complicated structure, in which one piece of putative knowledge can depend on an earlier piece, which depends on a still earlier one. It is important when making and evaluating arguments, to keep this structure in mind and to consider which earlier beliefs the premises depend on.

To see why this is important, imagine asking someone who had inferred Natalie's guilt by the above argument how he knew that the murder had access to the rose garden at midnight. Suppose he answered as follows: "Natalie was the murderer, and her bedroom opens directly onto the rose garden." He's using Natalie's guilt as a premise from which to infer that the murderer had access to the garden, and then going on to use this conclusion as a premise from which to infer Natalie's guilt. The structure of his reasoning is as follows:

Map 14:

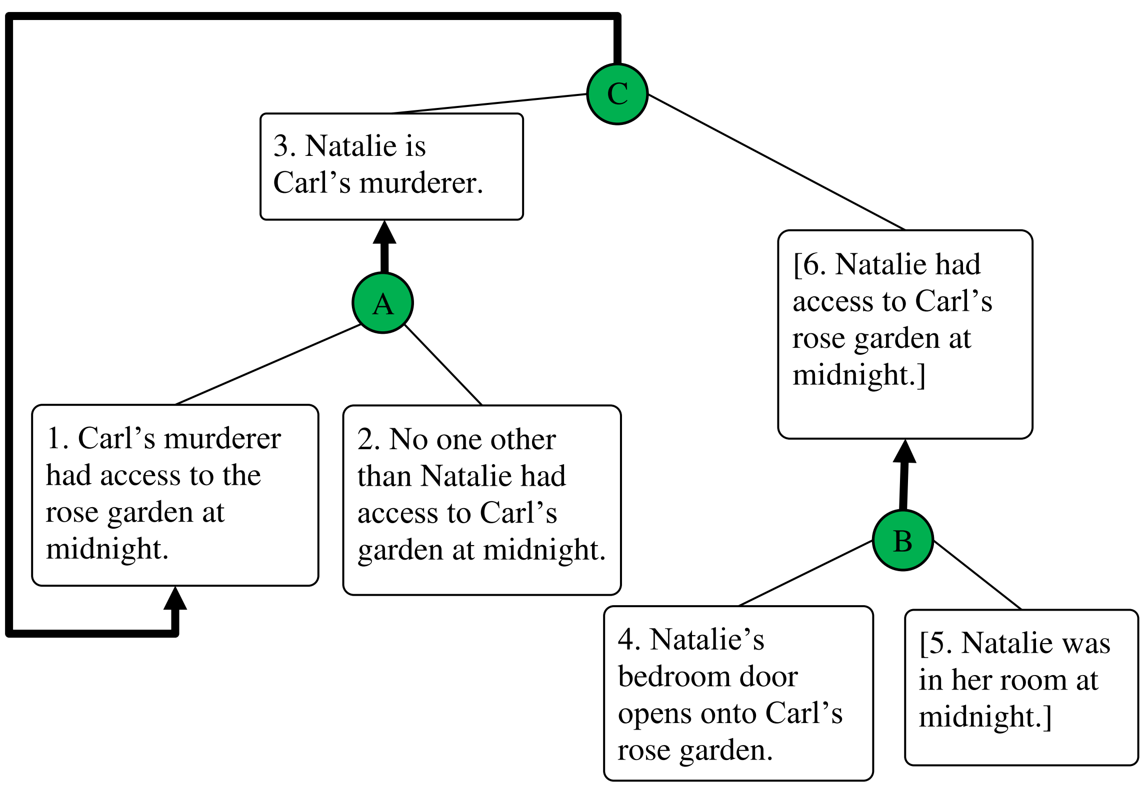

A fallacy is a mistake in reasoning, and one name for the fallacy here is "circular reasoning." The reason for this name may be obvious from the map above, were we can see the line circling around from Inference C to Proposition 1. The argument proceeds in a circle, because Proposition 3 is being used to support itself. It is the conclusion of Argument A, but it is a premise for Argument C, which has one of Argument A's premises as its conclusion. If we knew that Natalie was the murderer independently of Argument A, then of course we could use it (in Argument C) to establish Proposition 1, but then we cannot go on to use Proposition 1 to establish that Natalie's guilt. This fallacy is also called "begging the question", because the person making the argument asks (or "begs") us to accept as a premise the very thing that's in question—namely, Natalie's guilt.

The lesson we can draw from this is that when evaluating a proposition as a premise in an argument we need to consider the epistemic status it would have *if we did not know the argument's conclusion*. To illustrate this let's return to another example we looked at earlier:

Map 15:

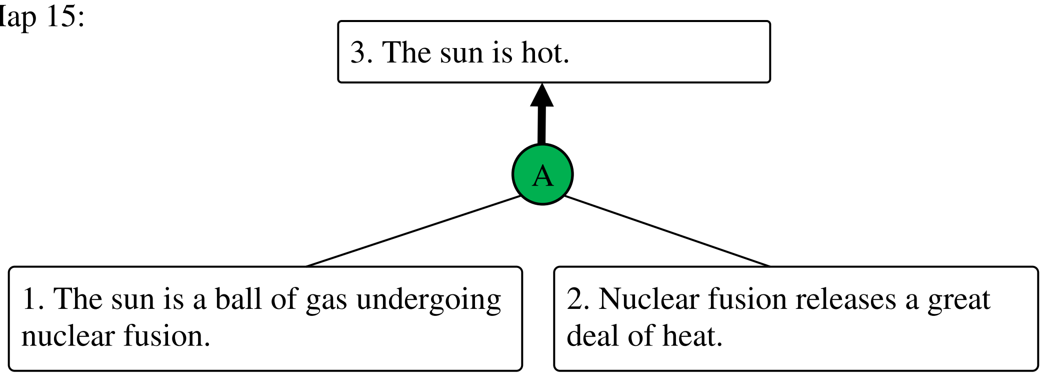

When we considered this example earlier, we said that it was an explanation rather than an argument, but let's suppose that someone tried to use it as an argument to prove the conclusion and assess the premises. Scientists are certain of both of the premises, but part of their basis for believing Proposition 1 is that they know that the sun is hot. So, this argument would beg the question.

In both of these example cases, the fallacy is comparatively obvious, and real-life cases of circular reasoning can be more difficult to detect. They happen because we so often infer unselfconsciously, without even realizing that we have done it, and so we fail to take notice of what other beliefs our conclusions rest on.

Glossary of Key Terms from §3

unfounded – the epistemic status of a proposition for which there is no reason to believe

possible – the epistemic status of a proposition for which there is reason to believe

probable – the epistemic status of a proposition for which there is enough evidence to establish that it is more likely to be true than false

certain – the epistemic status of a proposition for which the evidence supports beyond a reasonable doubt

begging the question – to assume as a premise in an argument the conclusion one is attempting to establish; also known as "circular reasoning"

§4. Assessing Inferences

Recall from earlier that an inference is good when the premises are related to one another and to the conclusion in such a manner that it is impossible or unlikely for the conclusion to be false if the premises are true. The more unlikely it is for the conclusion to be false when the premises are true, the stronger the inference is. We will now discuss several different types of inference, considering how they work and what makes some inferences of that type stronger than others.

§4.1 Deduction

Deduction is the strongest type of inference—the type in which it is impossible for the conclusion to be false when the premises are true. This is sometimes expressed by saying that, in a deduction, the premises *necessitate* the conclusion or that the conclusion *follows necessarily* from them. It is this feature of necessitation that Aristotle, the first logician, focused on when defining deduction.[10]

We will discuss shortly what it is about the premises of deductive arguments that makes the conclusions follow necessarily, but first let's make a few observations and introduce a few terms. Unlike the sorts of inferences that will be discussed in the remaining sections, deductive inferences do not differ from one another in strength. Either an inference is deductive, or it is not. However, there are some arguments that may seem like deductions when they are not. To

[10] Aristotle, the first logician, defined deduction as follows: "A deduction is an argument in which, certain things being laid down, something other than these necessarily comes about through them." (*Topics* I.1 100a25-27)

differentiate between the genuine deductions and the imposters, logicians call the genuine ones *valid* deductions and the imposters *invalid*. (We will turn soon to examples of each.) All valid deductive arguments have equally strong inferences, but this doesn't mean that the arguments as a whole are equally strong, since their premises may still differ in epistemic status. Thus, the two things to do in evaluating a deductive argument are to determine whether it is valid and to assess the premises.

If the deduction is valid and the premises are certain then the argument will make the conclusion certain. If it is valid and all but one of the premises are certain, then it will elevate the conclusion to the epistemic status of the remaining premise. If more than one of the premises is uncertain, then the argument will give us less reason to believe the conclusion than we have to believe the least certain of its premises.

Now let's consider how the premises of deductive arguments necessitate their conclusions, by looking at some examples. (For the time being ignore the symbolic representation on the right.)

Map 16:

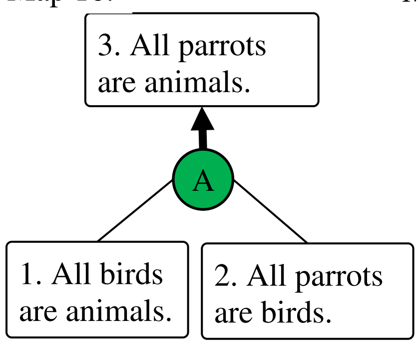

Map 17:

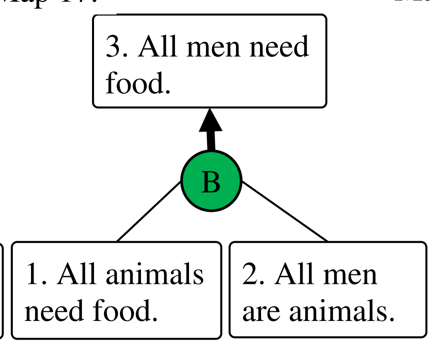

Map 18:

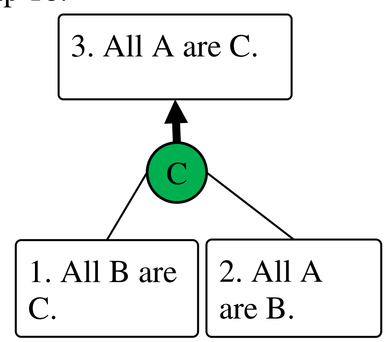

All birds are animals.	All animals need food.	All B are C
All parrots are birds.	All men are animals.	All A are B
All parrots are animals.	All men need food.	All A are C

Let's focus on the first of these examples. Notice that it is not because of anything special about the subject matter that the premises necessitate the conclusion. If we changed the argument to be about Volkswagens, cars, and Jettas, rather than birds, animals, and parrots, the premises would still necessitate the conclusion. The same holds if we changed the argument to make it about musicians, matadors, and marriage counselors. In this last case, both premises would be false, but if they *were* true (that is, if all musicians were matadors and all matadors, marriage counselors), then the conclusion would have to be true: all musicians would have to be marriage counselors. What makes the deduction work doesn't have to do with the subject matter of the propositions involved, but with their structure and interrelation. This is called the *form* of the argument.

Notice now that the argument immediately to the right of the one we have been considering has the same logical form. (If you don't see this immediately, try replacing the phrase "need food" with "are things that need food" or "are food-needers".) In order to focus on the forms of deductive arguments rather than their content, Aristotle introduced the practice of using letters to stand for the subjects and predicates of the propositions, leaving behind only words like "some", "all", "no", "is" and "not" (and their variants). Thus, we can arrive at the symbolic representation, presented in the right above, of the form of the arguments we've been discussing.

Here are some other forms of deductive arguments. (Again, the arguments next two each other share the same form which is represented symbolically on the right.)

Map 19:

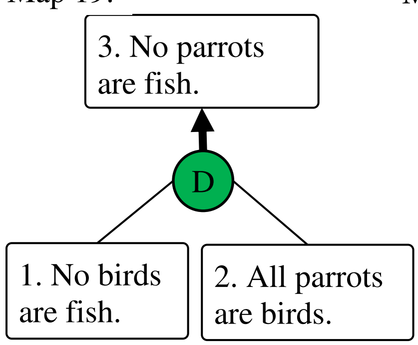

Map 20:

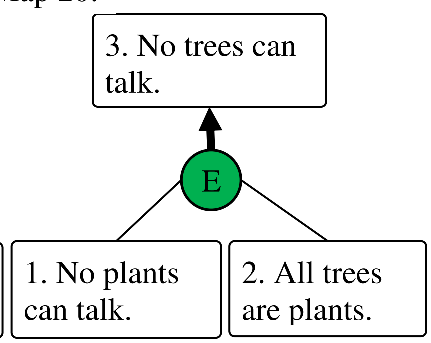

Map 21:

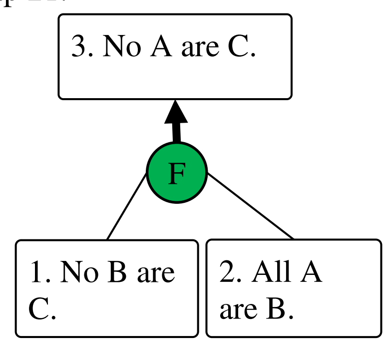

Map 22:

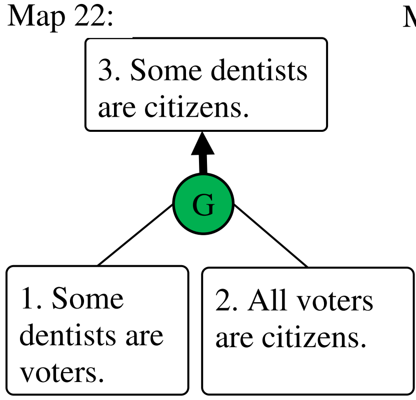

Map 23:

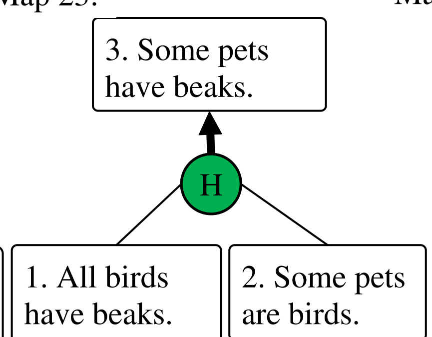

Map 24:

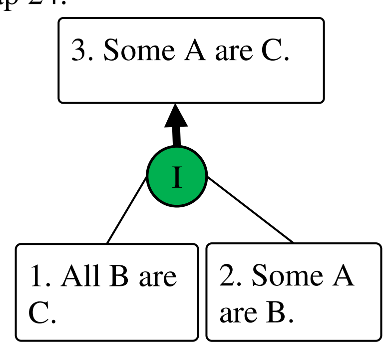

Map 25:

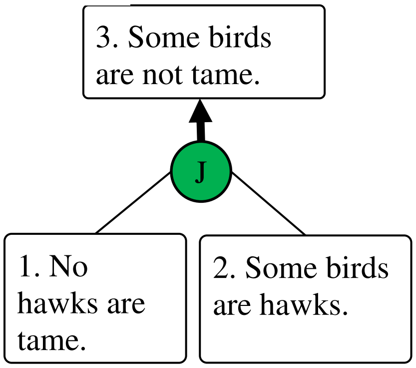

Map 26:

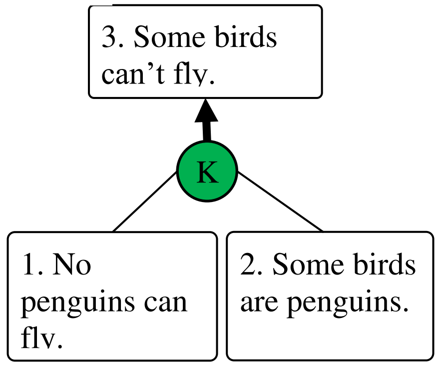

Map 27:

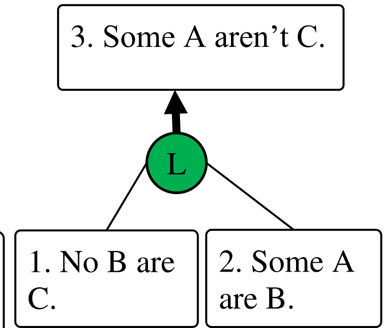

No birds are fish.
All parrots are birds.
____
No parrots are fish.

No plants can talk.
All trees are plants.
____
No trees can talk.

No B are C
All A are B
____
No A are C

All voters are citizens.
Some dentists are voters.
____
Some dentists are citizens.

All birds have beaks.
Some pets are birds.
____
Some pets have beaks.

All B are C
Some A are B
____
Some A are C

No hawks are tame.
Some birds are hawks.
____
Some birds are not tame.

No penguins can fly.
Some birds are penguins.
____
Some birds can't fly.

No B are C
Some A are B
____
Some A aren't C

Aristotle went through all possible combinations of two premises and a conclusion that could be made with propositions of the form “All/Some A are/aren’t B” and figured out which were valid deductions. It turned out that all of the valid forms could be restated in terms of the four that we’ve already looked at.

Why, in each of these cases, do the premises necessitate the conclusion? It is because one premise is a universal proposition saying something about everything of a certain sort and the other premise tells us that something else is that sort of thing. The conclusion then applies the universal premise about all things of the sort to the thing that the other premise tells us is a member of that sort. If the universal premise is really true, and the thing in question really belongs to the relevant sort, then the universal premise will, of course, have to apply to that thing.

Though most deductive arguments involve the application of universal propositions to particular cases, not all do. Some, which were focused on by later Greek and Medieval logicians, involve the application of hypothetical statements to actual cases or the application of statements about alternative possibilities to cases in which some of the alternatives have been ruled out. Here are some examples:

If you get eat too much, you’ll get sick.	If p, then q.
You ate too much.	p.
You’ll get sick.	q.
If he entered the yard, the dog would have barked.	If p, then q.
The dog didn’t bark.	Not q.
He didn’t enter the yard.	Not p.
Joe is either a soldier or a civilian.	p or q.
Joe isn’t a civilian.	Not p.
Joe is a soldier.	q.

Map 28:

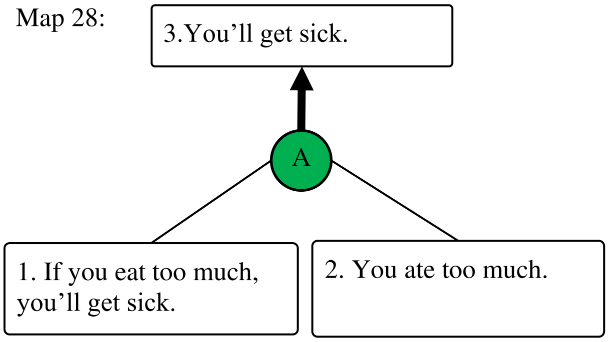

Map 29:

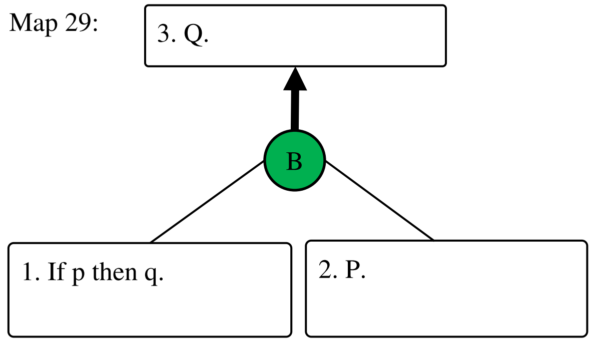

Map 30:

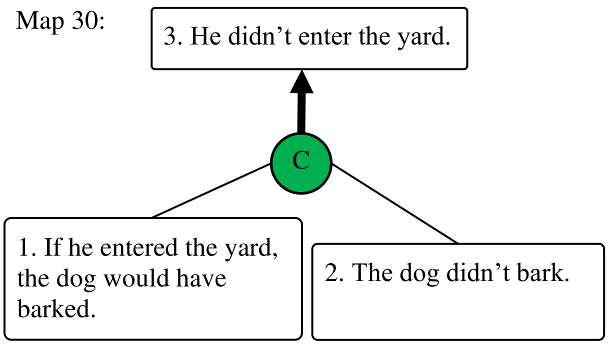

Map 31:

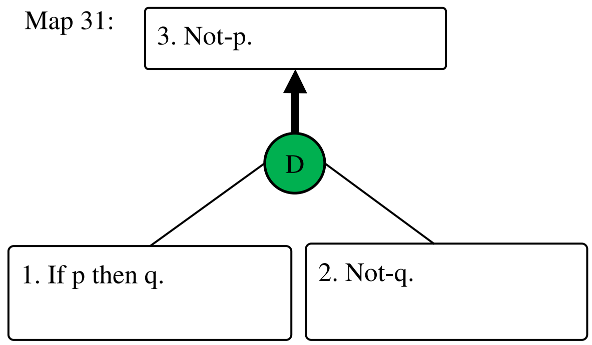

Map 32:

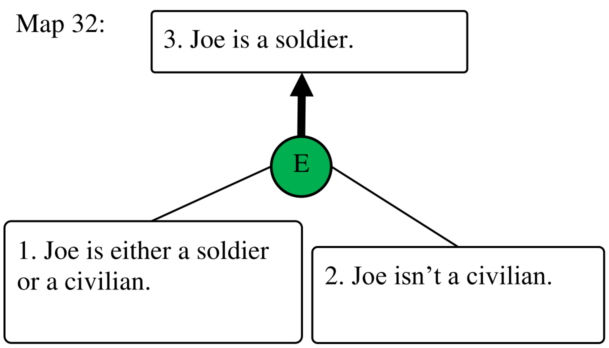

Map 33:

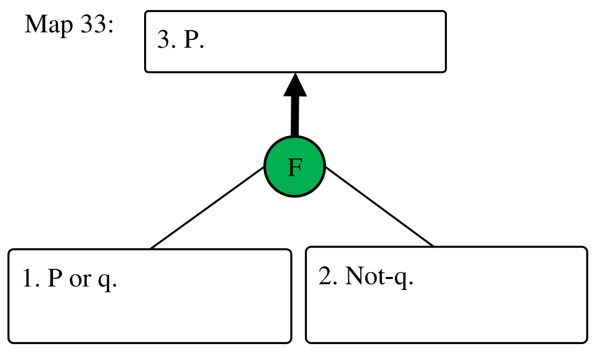

These argument forms involve complex propositions that are built up from simpler ones. When representing the arguments symbolically, we use lowercase letters to represent the simpler propositions.

Thanks to Aristotle and later logicians, the way deductions of different sorts work is well understood, and the various valid and invalid forms have been catalogued. But even among people who have not studied their works it is relatively rare to find sustained disagreement about whether an argument is valid. People are generally quite good at recognizing whether or not the premises follow necessarily from the conclusion.

That said, we do sometimes argue invalidly, particularly when we are not paying attention. Here are the three most common invalid argument forms (or “deductive fallacies”):

Map 34:

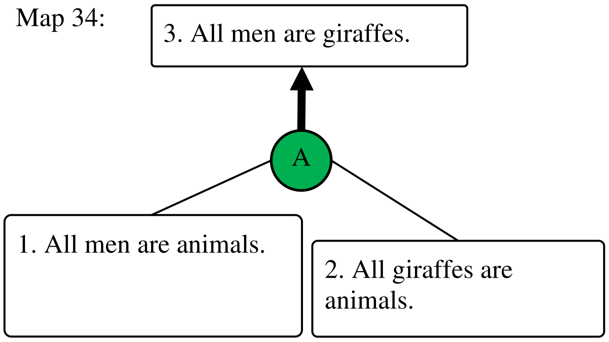

Map 35:

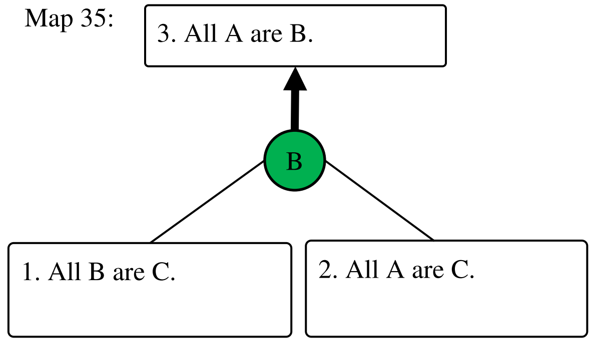

Map 36:

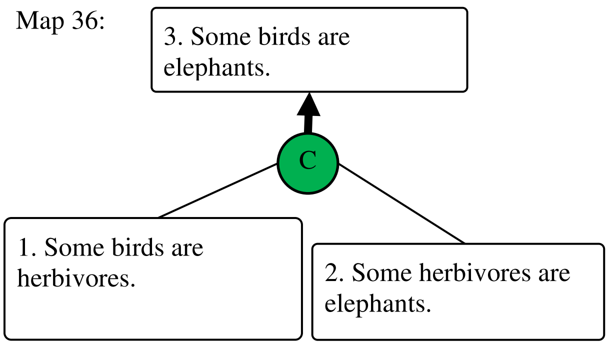

Map 37:

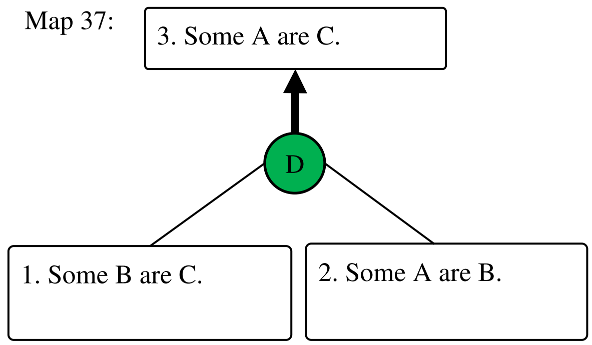

Map 38:

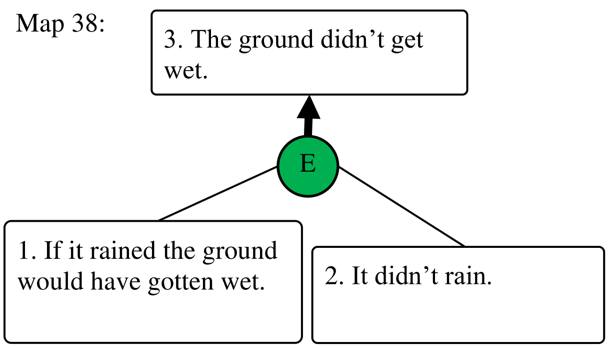

Map 39:

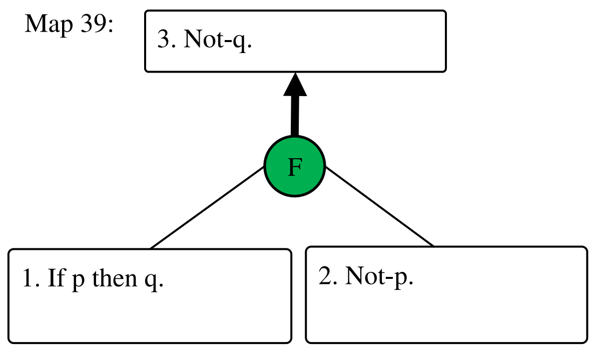

All men are animals.

All giraffes are animals.

All men are giraffes.

All B are C

All A are C

All A are B

Some birds are herbivores.

Some herbivores are elephants.

Some birds are elephants.

Some B are C

Some A are B

Some A are C

If it rained the ground would have gotten wet.	If p, then q.
It didn't rain.	Not p.
The ground didn't get wet.	Not q.

For the purposes of this course, it is not necessary to memorize the samples of valid and invalid forms of deductive inference discussed in this section. They do not exhaust all the possible deductive forms, nor even all the forms that we will encounter this semester. When dealing with a deductive argument, instead of trying to classify it under one of the forms we've discussed, simply ask yourself whether the premises and conclusion are so related that it is impossible for the conclusion to be false, if the premises are true. If the answer isn't immediately obvious, try to make up an argument of the same form with true premises and a false conclusion.

You might have gotten the impression from this section that deductive arguments are better than other kinds of arguments. There's a respect in which this is true: their inferences are stronger. However, an argument is only as good as its premises, and the premises of deductions are (for the most part, at least) generalizations, all (or nearly all) of which are established by other forms of argument.[11] Given this, is it more accurate to view deduction as the easier and more straightforward part of a process whose more difficult part is the arguments that establish the general propositions used as premises in the deductions. The primary form of argument by which these general propositions are established is *induction*, which we will go on to discuss. Before turning to it, it will be useful to briefly address another topic.

§4.2 Arguments Applying Statistics

Consider the following argument, and let's assume that the premises are certain:

[11] Later in the term we will discuss whether there are any general propositions that are known independent of argument.

Map 40a:

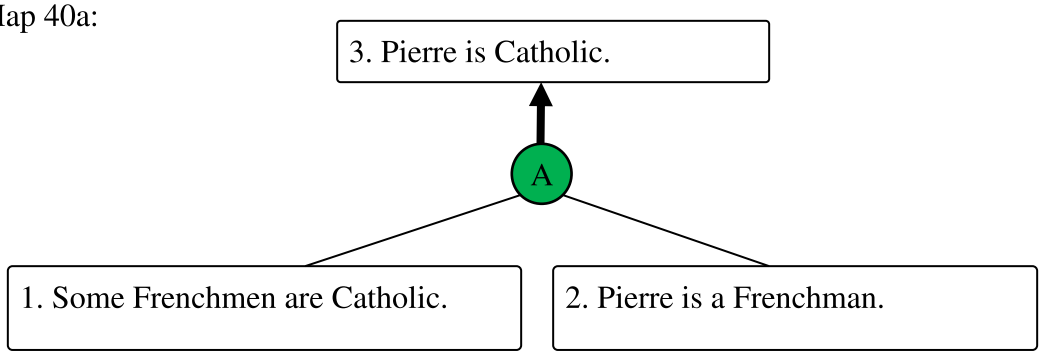

If P1 said that all Frenchmen were Catholic the argument would be a valid deduction, but, as it stands, the argument is invalid. Does it, though, give us any reason to believe its conclusion? Does it make the conclusion even epistemically *possible*? For all that P1 says, there may be only two Catholics among the tens of millions of Frenchmen. So, if these premises were the *only* support we had for the conclusion, the conclusion would be unfounded. But suppose we changed the argument as follows:

Map 40b:

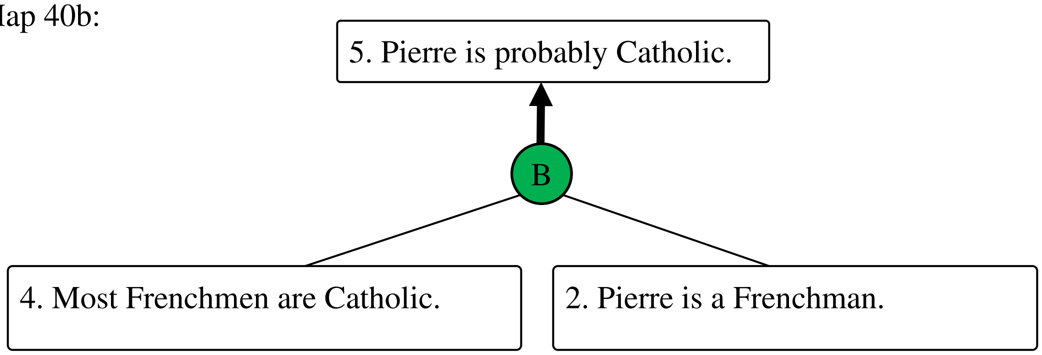

This is clearly an improvement. If *most* Frenchmen are Catholic, then isn't the conclusion that Pierre is *probable*? After all, given that he's French, isn't he more likely to be Catholic than not? And if we change P1 again to be more specific as to how many Frenchmen are Catholic, can't we then say more precisely how probable it is that Pierre is Catholic?

Map 40c:

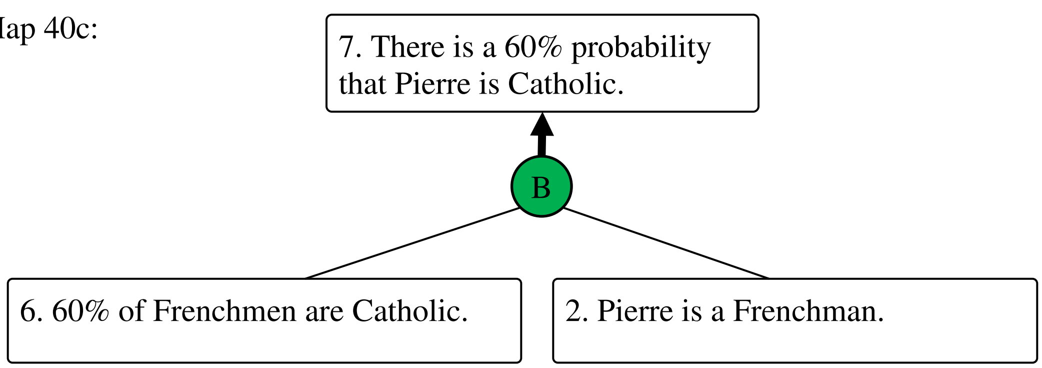

We have here an extremely simple example of *statistical reasoning*. Statistics is the science that deals with the interpretation and analysis of quantitative data about the prevalence of traits within groups, and it is sometimes used to draw conclusions about how certain we can be that an individual has a certain trait.

Notice that we haven't yet assessed either of the last two arguments. Since what we're interested in here is the inference, let's assume for the sake of argument that we know the premises to be true. If you know that Pierre is French and that most Frenchmen are Catholic, does this mean that the proposition "Pierre is Catholic" is has an epistemic status of probable for you? If you know that 60% of Frenchmen are Catholic, does it have a probability of 60%? Not necessarily.

It depends on what *else* you know about Pierre. Suppose that Pierre is a Muslim and that you know this about him in addition to knowing the two premises above. Then, far from being probable, the proposition that Pierre is Catholic would be certainly false. Even if you didn't already know whether Pierre was Catholic, there are lots of other things you might know about him that would make it unreasonable for you to infer that he was probably Catholic, from the knowledge that most Frenchmen are. Suppose you knew, for example, that Pierre was a Communist and that most communists are atheists, or that he comes from a particular part of France that is predominantly protestant or that he's extremely intelligent and that intelligent people often hold views outside of the mainstream of their society.

From the premise that 60% of Frenchmen are Catholic, you can infer with 60% probability about a Frenchman selected absolutely at random, that he is Catholic. But, as any decent statistician knows, if Pierre is not a random Frenchman, you cannot infer it about him without first consulting the rest of your knowledge about him to see whether any of it is relevant to determining his religion. It is important to keep this in mind, since we are so often bombarded with statistics. For example, if you are a woman, and you read online that unmarried women are fifty percent more likely to die in any given year than married women, then even if this statistic is true, it would be unreasonable to infer that you will probably increase your longevity by accepting a marriage proposal.

There is an important point to be make here that applies far beyond statistical arguments. It is a special feature of deductive inferences that one can analyze them in isolation from the rest of one's knowledge. This is because deductive arguments are just those in which, if the premises are true, then, no matter what else may be the case, the conclusions must also be true. But this is not the case with any other kind of argument. Because of this, when assessing any non-deductive inference, you need to make a point to consider whether you have any knowledge other than the premises that impacts the degree to which the premises support the conclusion. This makes assessing non-deductive inferences much more difficult than assessing deductive ones.

§4.3 Induction

An induction is an argument in which a universal conclusion pertaining to all objects of a certain sort is inferred from premises about particular objects of that sort. For example, you might infer from premises about particular men being mortal, to the conclusion that all men are mortal.

It is either primarily or exclusively from induction that we acquire the universal knowledge that then gets applied in deductions. However, whereas deduction is well understood, the nature of induction and the standards by which inductive inferences can be assessed are subject to a great deal of confusion and controversy. For example, it is debated whether inductive arguments can ever establish their conclusions with certainty. It is hard to see how they could, since it is not clear what it is about knowing that some (or even many) members of a group have a certain feature that can assure us that the other members of the group must have it as well. On the other hand, since all or most of our deductive arguments depend on induced propositions, and arguments are only as good as their premises, if induction cannot establish conclusions with certainty, then we cannot be certain of the conclusions of (almost any) deductions either. As a result, few if any arguments would be able to establish their conclusions with certainty, and almost all of the things that we think we know by inference would be cast into doubt.

Most contemporary philosophers embrace this skeptical conclusion. A few go further, arguing that induction not only fails to give us certainty but fail to give us any reason whatsoever to believe their conclusions. If they are right, then much of what we ordinary take ourselves to know is not only uncertain, but entirely unfounded. Most philosophers reject this position. Many hold that, though induced conclusions are never certain, they can have such a high degree of probability that we can think of them as certain for most purposes. Some hold that induction can establish certainty, but only when the premises are considered in the context of a great deal of background knowledge that is too vast to be enumerated into premises.

These differing views of the status of induced conclusions are based on differing views of what (if anything) it is about knowing facts about particulars that gives us reasons to believe universal conclusions. For example, the philosophers who have the most positive view of induction tend to focus on the processes by which we form concepts (such as “man” and “moral”) in the first place and classify particular things under them. Some of them stress the role played in this process by an awareness of cause and effect and argue that it is an understanding of causes and of the reasons we had for forming the relevant concepts that justify us in inferring from facts about some instances of those concepts to universal conclusions applying to all the instances. Philosophers who have a less positive attitude towards induction usually think we know less about cause and effect in the first place and tend to regard the classification of particulars into kinds as unfounded. The different theories of the nature of induction, lead to differences of opinion about some of the standards that should be employed in evaluating inductive arguments. For all of these reasons, induction is a difficult subject to cover briefly, and I will not say much about it here. However, I will mention a few factors which all but the most skeptical philosophers agree are crucial to the strength of an inductive inference.

First, as with the case of statistical arguments (discussed above) inductive arguments cannot be assessed in isolation from the rest of one's knowledge. In particular, no matter how many members of a certain kind one might know to have a certain trait, one cannot rationally infer that *all* members of the kind have it, if one knows of any member that doesn't (or if one has good reason to suspect that any member doesn't). So, to induce that all the members of a kind have a certain trait, it is not sufficient that you know that *some* members have it, it must be that *all* the members you know about have it.

Second, you must be on guard against the possibility that it is a *coincidence* that all the members you know about have the trait. This can easily occur when you don't know many members, so (all other things being equal) knowing about more members the more of the members of a kind one is aware of, the stronger one's inductions about it will be.

However, even when one knows something to be true of a great many members of a kind, it is still easy for this to be a coincidence. For example, consider the case of a teacher who has a great deal of experience with autistic children but who has never taught other children. Suppose that this teacher found that a certain pedagogical technique was extremely effective with all of the students on which he tried it. Does this put him in a position to infer that this technique works well for *all children*? No, because all of the students on which he has tried the technique have something in common other than being children—their autism. Moreover, this common trait is one that there is reason to think might make a difference to which techniques are effective. The teacher would not be in a position to draw a conclusion about children *in general* until he has tried the technique on children who do not have the disability. Indeed, he will not have a very strong inference about *all* children, until the group of children on which he has tried the technique is *diverse* enough that all of its members do not share any trait (other than being children) that there is reason to think might be relevant to which techniques they respond well to. In general, when assessing inductions, you should consider whether the examples used have anything in common that might be relevant to the conclusion other than that they are members of the kind being induced about.

§4.4 Statistical Generalization

A statistical generalization infers a conclusion about the proportion of members of a subject kind have a certain predicate, from knowledge about the proportion of subject-members in a sample group that have the relevant predicate. For example:

Map 41:

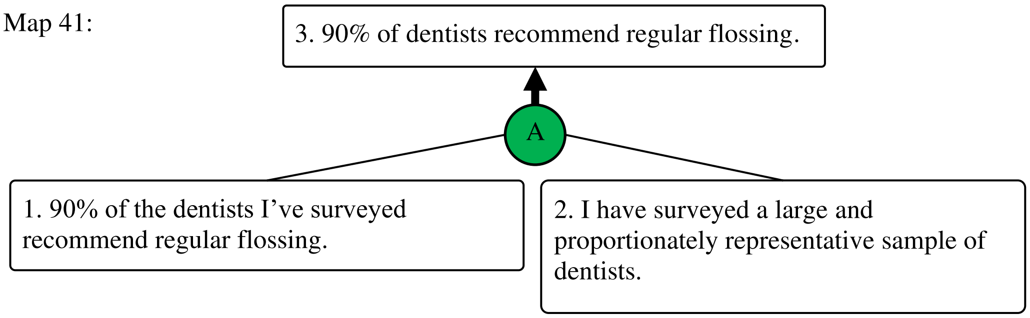

To assess a statistical generalization, you must consider whether the sample is large and proportionally representative. All things being equal, a larger sample makes for a stronger generalization, but whether the sample is representative matters more than how large it is. To be representative, the sample has to range over all the differences there are among members of the subject kind which there is reason to think might make a difference to whether the predicate applies to them.

In a statistical generalization, the sample must not just be representative in the sense of including some subject-members which each of the characteristics that may be relevant to the predicate, it must include the same proportion of members with each of these characteristics as exists in the whole population of subject-members. For example, if there is reason to think that whether a dentist practices in an urban or rural area is relevant to whether he would recommend flossing, then it is not enough that one has both urban and rural dentists in one's sample, the proportion of urban to rural dentists in one's sample has to match the proportion in the general population of dentists.

§4.5 Inference to the Best Explanation

When you look at a patch of ground and see regularly spaced shoe-shaped impressions in the pattern that we call "foot-prints", you immediately infer that someone walked there. What argument are you using? Perhaps this one:

P1. There are regularly spaced shoe-shaped impressions in the ground.

P2. The impressions had to be caused by something.

P3. The only thing that could cause such impressions is a person walking by.

C. Therefore, a person walked by.

This is a valid deductive argument, but P3 is false. There are other things that could cause such a pattern of impressions in the ground. For example, a machine could easily be made that could do the task. Or they could be produced by a chimpanzee walking upright with shoes on. There may be other ways as well. But even though there are other things that *could* cause impressions of the relevant sort, it is clear that, in most circumstances at least, the *best* explanation of the footprints would be that someone walked by. Because of this, it is eminently reasonable to conclude that a person walked by; indeed, unless one had some special evidence to the contrary, it would be irrational not to draw this conclusion. However, the inference taking place is not a deduction; it is an "inference to the best explanation". If we wanted to lay out the argument it would be as follows:

P1. There are regularly spaced foot-shaped impressions in the ground.

P2. The impressions had to be caused by something.

P3. A person walking by would cause such impressions.

P4. This is the best available explanation of the impressions.

C. A person walked by.

Inference to the best explanation is constantly used in the sciences, in solving crimes, and in other contexts. In such inferences, one concludes that a certain proposition is true because it would explain a known effect better than any alternative explanation. Recall that an explanation explains an effect by citing causes. An inference to the best explanation concludes that a certain putative cause exists, because it would explain a known effect.

Such an argument depends on knowing (or having excellent reason to believe) several things: (1) that the effect in question exists, (2) that it is an *effect* (that is, something which is caused by something else), (3) what sorts of things could cause the effect, (4) which of these causes explains the effect best. Notice how these four correspond to P1-P4 in the argument above. We can restate the form of argument involved as follows:

P1. p.

P2. p is an effect, requiring an explanation.

P3. q would cause p.

P4. q is best available explanation of p.

C. q.

If all the premises are certain, then the degree of support the argument provides for the conclusion is proportionate to *how much better* the explanation in question is than the sum of all the alternative explanations available. In the example of the footprints, the alternative explanations are quite poor, so the conclusion is either certain or nearly so. But in a case where there were several decent explanations, the conclusion would only be probable, or merely possible, depending on how good the other explanations were. For example, think of a murder which any of three people could have committed. The murder is the effect, and there are three explanations corresponding to the three suspects—let's call them Ed, Fran, and George. Suppose that the best of these three explanations is that Fran did it (perhaps she had a stronger motive than either of the others), but that this explanation was only slightly better. If so, then argument would only make the proposition that Fran committed the murder *possible*, because though it is more likely that she did it that than Ed did or that George did, it is still more likely that one of the two men did it than it is that she did. Indeed, Fran is probably innocent. So, to assess an inference to the best explanation we need to know not only that the explanation in question is the best one, we need to know how much better it is than competing alternatives.

What makes some explanations better than others in the first place? There are at least three factors: (i) the degree of detail in which the effect is explained; (ii) how much independent reason there is to believe that the cause exists and is operative in the relevant context; and (iii) how well the statement of the cause is circumscribed.

Let's begin with the first of these factors and consider it in connection with a variant of our footprint example. While on a hike, you come across what we would normally describe as animal tracks. These are the effect that you want to explain. Notice that in describing them as animal tracks, we're already explaining them as effects of an animal, so for now don't think of them as animal tracks but as a certain pattern of impressions in the ground. Notice that there are different levels of detail at which this pattern can be described. At the one extreme, they could be described simply as impressions in the ground. A more detailed description would include the approximate size of the impressions and their foot-like shape and it would indicate pattern in which the impressions occur—for example, it might say that they occur at regular intervals along two roughly parallel lines, and that the impressions are staggered somewhat, so that the impressions in the left line are slightly ahead of those in the right. A still more detailed description would specify the shape, size and pattern more precisely, including such details as whether there are toe marks and how many, the precise shape of each part of each the impression, how deep the impressions are, just how far apart, how exactly each impression is oriented relative to the others, etc. The description could be more or less detailed depending on how many of these aspects of the impression it described and the degree of specificity with which it described each—for example, whether numerical measurements are given, and, if so, with what degree of precision. (There are entire books on documenting animal tracks and some people make this their life's work.) The upshot of the preceding is that we can describe the effect at different levels of detail. The relevance of this to assessing explanations is that, all other factors being equal, one explanation of an effect is better than another if it can explain the effect in greater detail. So, consider several different explanations that someone might give for the animal tracks:

(A) The impressions were made by one or more entities pressing into the ground.
(B) The impressions were made by a walking animal.
(C) The impressions were made by a charging elephant.
(D) The impressions were made by a relatively small animal with four toes.
(E) The impressions were made by a canine.
(F) The impressions were made by a fox.
(G) The impressions were made by an adult female kit fox moving at top speed.

Explanation (A) does explain why there are impressions in the ground, but it does so only at the most generalized level. It explains why there are impressions without explaining any of the details of these impressions. This may be the best explanation someone could give if *all* he knew about the effect was that there were impressions in the ground, but in the example, we know a lot more about the effect than that. Explanation (B) is a better explanation because it explains a lot more about the impressions. Our knowledge of animals, feet, and walking allows us to figure out what *sort* of impressions would be caused by an animal walking, and since the impressions we see on the ground are of this sort, this explanation explains the effect in more detail than the previous one. Explanation (C) gives further details of the cause, from which we could infer further details that would have to be true of the effect. The tracks made by a charging elephant would quite large and deep. Let's suppose that this is not so of the tracks we are looking at. If so, (C) will be ruled out entirely as an explanation, because the effect to be explained couldn't have been produced by the specified cause. Like (C), Explanation (D) gives us details about what kind of animal caused the tracks and, again, we know more or less what sort of tracks an animal of the sort specified would make. Let's suppose that in this case, the tracks we observe are of the right sort to have been caused by a relatively small four-toed animal. If so, then this is the best of the explanations so far, because it explains the effect accurately and in greater detail than any of the others. In fact, this is probably the best explanation that a layman would be in a position to give. Someone who knows a little more about the feet of different animals and how they walk would be able to give an explanation like (E) or (F) which would explain further details of the tracks, and an experienced woodsman could explain subtler details of the tracks with an explanation like (G).

In this example, all of the explanations other than (C) are consistent with one another. Animal feet are things that press into the ground; small, four-toed animals are animals; canines are small and four-toed, foxes are canines, female kit foxes are foxes, and running at top speed is one of the ways in which female kit foxes move. Thus, though one explanation is better than the others in that it is more detailed, all the explanations can be simultaneously true. There are cases, however, in which two competing explanations that are both consistent with what is known about an effect, but cannot both be true. In some such cases, one explanation explains the effect in greater detail. For example, suppose that you already knew that the animal tracks were caused either by a jack rabbit or a fox (perhaps because you know that these are the only two sorts of animals in the area of the right approximate size) and that you know next to nothing about jackrabbits' feet but enough about foxes' to know that they would make tracks of roughly the shape observed. In this situation, the explanation that the tracks were made by a fox would explain the tracks in greater detail than the explanation that they were made by a jackrabbit.

Let's move on now to the second factor that makes some explanations better than others. The more independent reason we have to believe that the cause specified by an explanation exists and is operative in the relevant context, the better the explanation is. For example, suppose that you're looking at a photograph taken in Alaska of a set of large tracks through the snow. Two explanations for the tracks occur to you: (A) "They were caused by a polar bear", (B) "It's the abominable snow man!" Clearly (A) is a far better explanation than (B), because you know that polar bears exist and live in Alaska, whereas the idea that there's an abominable snow man is unfounded (or, at any rate, it has a much lower status than the idea that there are polar bears).

Now consider a case in which you know that both of the causes you're considering as explains of an effect really do exist: You're standing on a dude ranch in Texas and you hear hoof beats behind you. Here are two explanations for the sound: (A) "A horse is approaching", (B) "A zebra is approaching". You know that both horses and zebras exist, but (A) is still the better explanation, because in addition to knowing that horses exist, you know that horses are comparatively common in North America, especially on dude ranches, whereas zebras are rare. This is what I mean by saying you have independent reason to believe that "the cause is operative in the relevant context"—you not only know that horses exist and cause effects like hoof beats, you know that you're in the sort of situation in which there are likely to be horses causing these effects.

I chose this particular example because there's a saying in medicine: "When you hear hoof beats, think horses, not zebras." Diagnosing a patient is an example of inference to the best explanation: the patient comes to the doctor with symptoms, and the doctor needs to infer their cause. Young doctors fresh out of medical school often make the mistake of inferring that the patient has some exotic disease, even though the symptoms can be explained by a much more common condition. The exotic diseases, which are jokingly called "zebras," are bad explanations because, other than the fact that they *could* cause the patient's symptoms, there is no reason to expect to encounter them in (for example) a clinic in an American suburb, and the symptoms can be explained by other conditions ("horses") that there is independent reason to expect to encounter when working in such a clinic.

It is worth noting, however, that sometimes the best explanation of something we observe is unusual or even unprecedented. There are animals whose very existence was first inferred from their tracks (or, in some cases, fossilized remains of tracks), and the existence of certain microbes were inferred because they explained many the details of how certain diseases (especially typhus) spread much better than any competing theories.

The third factor that makes some explanations better than others is how well the statement of the cause is circumscribed. To get a sense of what this means, suppose that, after inferring from a set of footprints that a man walked by, we went on to infer from the size, shape, and arrangement of the footprints something about the man's weight, shoe size, and gate. So far, so good, but then suppose we went on further to describe his taste in literature, his hat-size and his mother's maiden name. Now our explanation of the footprints would be as follows: “They were made by a 200-pound man, walking briskly in size 11½ Bruno Magli Moc-Toe Oxfords, who adores Dostoevsky, has a hat-size of 7½, and whose mother’s maiden name was Schwartz.” The extra details given in the last three clauses make the explanation worse than it would otherwise be, because *they don’t explain anything about the footprints*.[12] At best such superfluous details in an explanation are distracting irrelevancies; but, in the context of an inference to the best explanation, they are worse than this. In this kind of argument, the reason we have for believing in the existence of the cause is that it would explain the effect. Therefore, the argument only gives us reasons to believe in those features of the proposed cause that play a role in explaining the effect. Thus, if having a hat-size of 7 ½ explains nothing about the footprints, then the argument can give us *no reason* to believe that *a man with this hat-size* walked past, though it does give us a reason to believe that *a man* walked past.

Let’s sum up by reviewing some of the things we need to consider when evaluating an inference to the best explanation: (1) Do we know that the effect being explained exists at all? (Or, more generally, what is the epistemic status of the proposition that it exists?) (2) Do we know that it is an *effect* (something that was caused by something else)? (3) Do we know enough about this kind of effect to speculate about what might cause it and to evaluate alternative explanations? (4) Would the cause proposed in the explanation explain the effect? (5) In how much detail does it explain it? (6) Do we have any independent reason to believe that this cause exists and is operative in this context? (7) Is the explanation properly circumscribed, or does it include features that don’t contribute to explaining the effect? (8) What other causes could explain the effect? (9) Is the proposed explanation really better than all of these explanations? (10) How much better is it?

§4.6 Analogical Argument

An analogy is a likeness between things, and an analogical argument is one that exploits an analogy to infer something about one of the like items from premises about the other. For example, if you know that Socrates was sentenced to death, and that Bertrand Russell is like Socrates, you might infer either that Bertrand Russell was sentenced to death or (more plausibly) that something *like* being sentenced to death happened to Russell.

As stated, neither of the two arguments just given concerning Russell and Socrates is very good. Part of the difficulty is that it is not clear *in what respect* the premise is claiming that Russell is like Socrates, and this makes it difficult to know what features of Socrates it is reasonable to attribute to Russell. After all, the two were alike in some respects: they were both human beings, males, philosophers, and people with a reputation for holding controversial positions that people worried would have a negative influence on the morals of the young. On the other hand, they are

[12] Perhaps you can imagine a situation in which these details would explain something—for example, if the footprints were leading from the site of a Schwartz family reunion to the site of a seminar on *Crime and Punishment*, and a 7½ hat was found next to them. But let’s assume that we are not dealing with this sort of situation.

different in many respects: one was Greek and the other British, one lived in the 5th Century B.C. and the other (mostly) in the 20th Century A.D., and their philosophical views were different on many points. Since they are like in some ways and not in others, we would expect some, but not all, of the things that are true of Socrates either to be true of Russell or to have close parallels in Russell's case.

Now in the case of Socrates' being sentenced to death, this was a result of his being thought to have controversial views that might corrupt the youth. Since this was also true of Russell, we might expect something somewhat similar to have happened to him. In fact, he was never sentenced to death, but he did suffer professionally in various ways because of parents who were concerned that he might corrupt their children. He was fired from a position at the City College of New York after the parents of some students at the college expressed outrage over his liberal views on sex and marriage. Is this like what happened to Socrates? It is in some ways, but not in others. What happened to the two philosophers is similar in that both suffered negative consequences which were due to concern over worries that they were immoral and would corrupt the morals of young people. But they are different in other respects—for example, Socrates was put to death and Russell merely dismissed from a job, and the action against Socrates was taken by the government of his city, whereas the action against Russell was taken by his employer.

When we know of a similarity between two items, it is reasonable to expect to find other, related similarities, but it can be difficult to specify the respects in which the items are similar, and to know what sorts of further similarities to expect to find, or how strongly to expect them. All other things being equal, the more one specifies the respect in which the like items are similar, what other similarities one should expect to find between the items, and why one should expect to find them, the stronger the argument from analogy.

Notice, however, that if one fully specifies these things, what one is doing is first inducing from the case of the one item (or set of items) to a universal conclusion, and then using this conclusion as a premise in a deduction about the remaining item. To return to our previous example, we might identify the likeness between Socrates and Russell as that they were both well-known figures with controversial views on moral subjects who were in a position to influence the young. And we might conclude that Socrates' death sentence was a case of his being treated unjustly at the hands of concerned parents with more traditional moral views. Further, we might conclude from Socrates' case (perhaps with the addition of a few others) that people with radical moral views in a position to influence young people are treated unjustly by concerned parents with traditional values. This would be an induction. And then we might deduce from this general conclusion and the premise that Russell was a figure with radical moral views in a position to influence the young that he too was treated unjustly by concerned parents with traditional values. In this case, rather than an argument from analogy, we have an induction followed by a deduction. But we arrived at this sequence of arguments by specifying as fully as we could something that began as an analogical argument. So, we might think of analogical arguments as crude forms of induction and deduction, in which one never states explicitly the universal proposition that is the conclusion of the induction and a premise in the deduction. Or, by the same token, we might think of induction and deduction as a refined form of analogical argument.

Analogical reasoning is especially useful in contexts in which one does not know of anything that could have caused an effect that one is trying to explain. Thus, it is often used as part of inferences to the best explanation. In the last section I mentioned that someone might infer the existence of a previously unknown type of animal on the basis of knowledge of its tracks. Consider how such an argument might work. We would know that the tracks could not have been made by any of the animals that we know of, but we would try to specify the ways in which the tracks were like those made by various known animals, and we would infer that the animal that made these tracks was like those animals in ways that would lead to the similarities in the tracks.

§4.7 Reductio ad Absurdum

"*Reductio ad absurdum*" ("reductio" for short) is Latin for "reduction to absurdity". It refers to the strategy of proving a proposition false by deducing an absurd conclusion (often a contradiction) from it in conjunction with other known premises. Such arguments are common in math. They were used to prove that the square root of two is not a rational number (and thus that there are such things as irrational numbers) and that there is no largest prime number. Here's a non-mathematical example of a reductio.

> There cannot be a chess-playing computer program good enough to win every game against any opponent, regardless of whether it plays as white or black. If there were such a program, it could be pitted against itself, and both sides would have to win. But it's impossible for both sides to win in a chess game, therefore there cannot be a chess-playing program such as the one described.[13]

And here is an example from mathematics, first in paragraph form and then laid out:

> Suppose that the set of prime numbers is finite. If so, it must contain a largest member—call it L. We can then multiply all the prime numbers together. Call the result P. P will be divisible by each of the prime numbers, so P+1, won't be divisible by any prime number. Therefore, it won't be divisible by any number, which means that it will be prime. But P+1 is larger than L. So, L isn't the largest prime number. But this is a contradiction, so the assumption that lead to it must be false, and the set of prime numbers must be infinite.

[13] This example was originated by Jim Pryor.

Map 42:

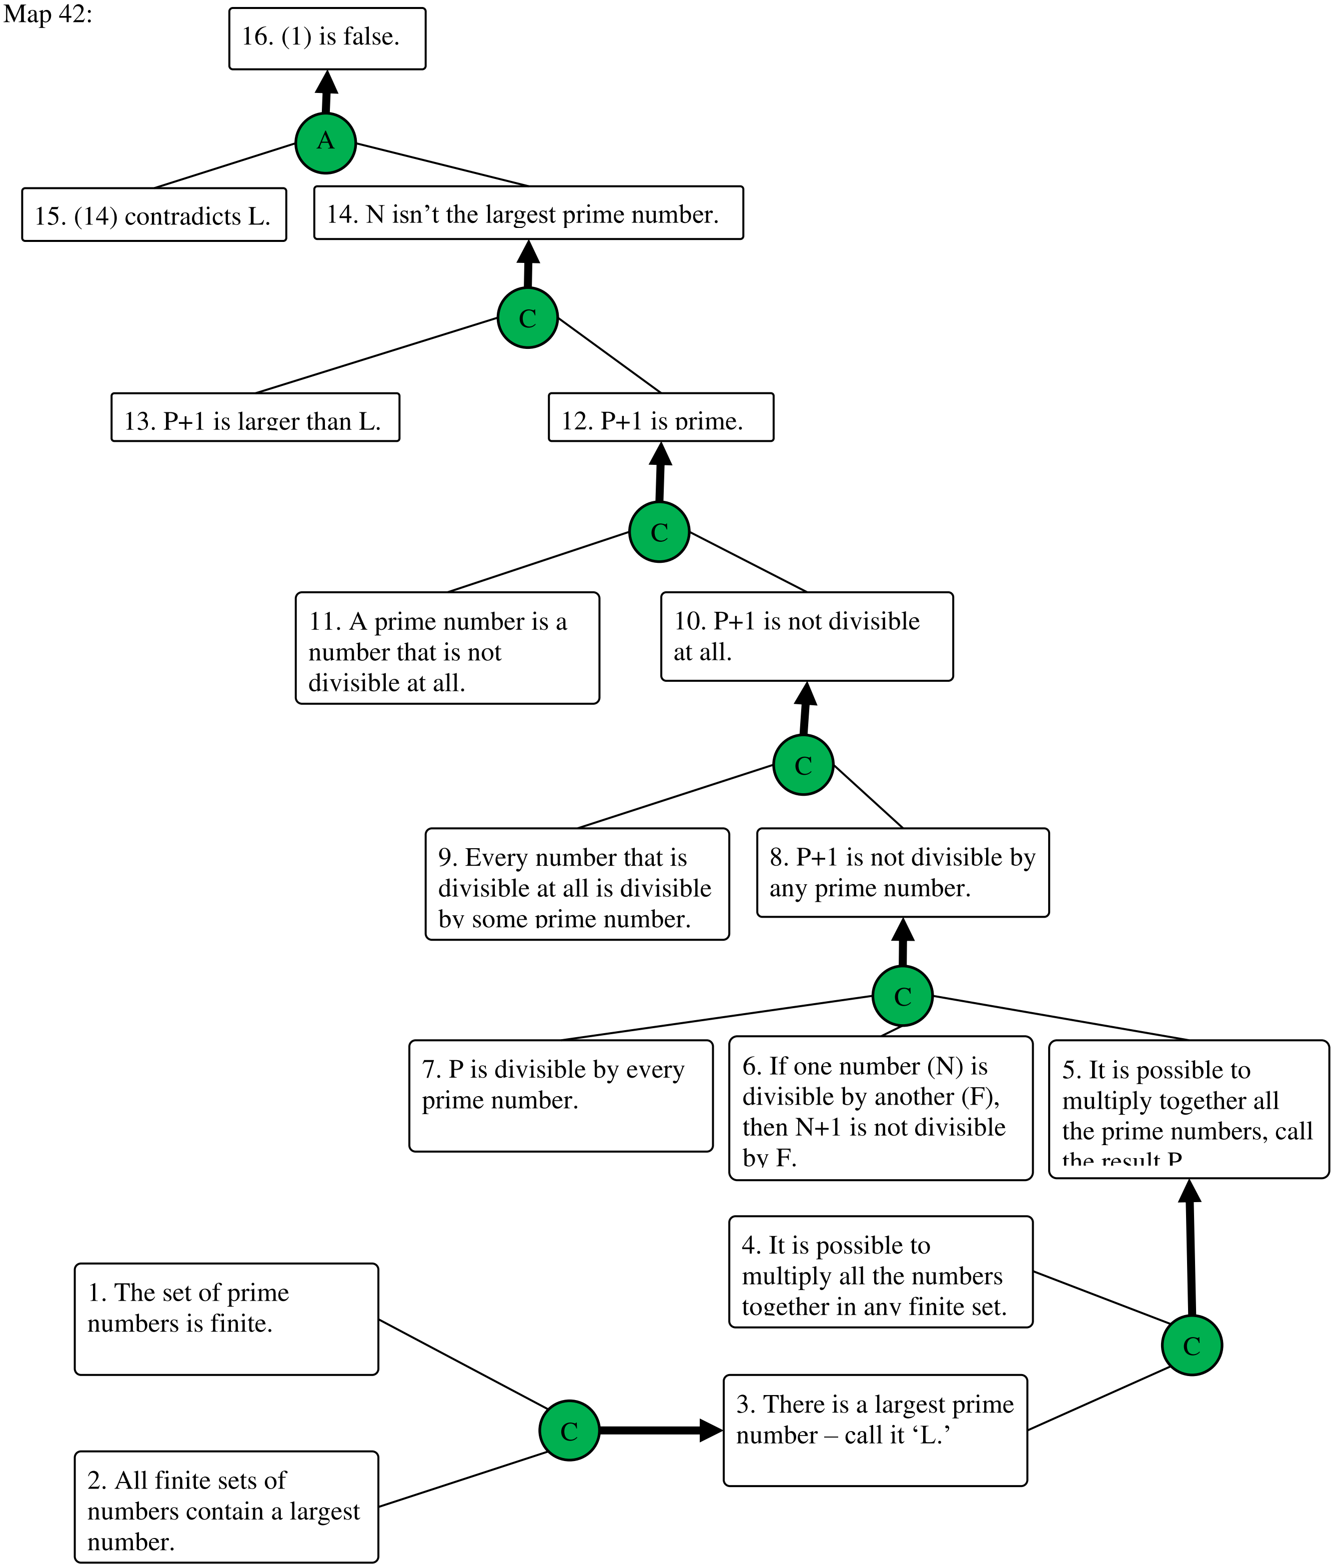
16. (1) is false.
A
15. (14) contradicts L.
14. N isn't the largest prime number.
C
13. P+1 is larger than L.
12. P+1 is prime.
C
11. A prime number is a number that is not divisible at all.
10. P+1 is not divisible at all.
C
9. Every number that is divisible at all is divisible by some prime number.
8. P+1 is not divisible by any prime number.
C
7. P is divisible by every prime number.
6. If one number (N) is divisible by another (F), then N+1 is not divisible by F.
5. It is possible to multiply together all the prime numbers, call the result P.
4. It is possible to multiply all the numbers together in any finite set.
C
1. The set of prime numbers is finite.
C
3. There is a largest prime number – call it 'L.'
2. All finite sets of numbers contain a largest number.

In assessing a *reductio ad absurdum* you need to determine two things: (1) whether the supposedly absurd conclusion is really false, and (2) whether it really follows necessarily from the proposition that the argument sets out to disprove. With regard to the first of these issues, the strongest reductios deduce contradictions, which it is easy to see cannot be true, but in other cases it is less certain that the "absurd" conclusion is false. With regard to the second issue, you assess a reductio in the same way that you would access any other deductive argument: you determine whether the inference is valid and how certain each of the premises is other than the one that has been assumed for the sake of refuting it. In order for the argument to prove that the assumed premise is false, all the other premises must be certain, and it must be certain that the "absurd" conclusion is false. However, even if this is not the case, the reductio can still give you strong reason to disbelieve the assumed premise if all of the other premises and the falsehood of the absurd conclusion all have a much higher epistemic status than the assumption that the argument is trying to show to be false.

§5. Assessing Conclusions in Light of Multiple Arguments

A proposition is unfounded unless we have some reason to believe it, and unless it is the sort of proposition that we can know to be true by direct observation or in some other way, we only have reason to believe it if we have an argument for it. Let us put aside the propositions that we can know independent of argument and consider what epistemic status propositions that require arguments have in different circumstances.

We have already said that such propositions are unfounded if we have no arguments for them. On the other end of the spectrum, if we have a proof of a proposition, it is certain, and we should regard it as knowledge. In between these extremes are propositions for which we have inconclusive arguments. They qualify as possible or probable, depending on the strength of the arguments. Finally, there are propositions that we have arguments against, and these arguments may prove the falsehood of the proposition or establish it with some lesser degree of probability.

So far, we have been thinking of these arguments one by one, but when a proposition is controversial, there are usually multiple arguments concerning it, some of them trying to establish it and others trying to refute it. In order to assess these propositions—to determine their epistemic status—we need to consider all of these arguments and that raises some questions that we have not yet addressed and will need to take up now.

§5.1 Refutations of Arguments (and the Burden of Proof Principle)

It is important to distinguish between refutations of arguments and refutations of propositions. A refutation *of a proposition* is a proof (or attempted proof) that the proposition is false. A refutation *of an argument* is a proof that the argument does not actually support its conclusion. The word "refutation" is usually reserved for cases in which the falsehood of the proposition or the badness of the argument is established *conclusively*. The word "counterargument" is a broader term for any argument that casts doubt on a proposition or argument.

Counterarguments to propositions are simply arguments that the propositions are false, and one assesses them just as one would any other argument. We will discuss later how to deal with situations in which there are arguments for and against the same proposition. I want to focus now refutations of arguments.

Most of this primer has been about evaluating arguments to determine how strongly, if at all, they support their conclusions, so you already know how to refute an argument or give a counterargument to it. An argument can be refuted by showing either that one of its premises is false or unfounded, or by showing that the conclusion does not follow from the premises. Doing either of these things would show that the argument provides *no support* for its conclusion. Weaker counterarguments would show that the argument supports its conclusion less strongly than might have been supposed by either showing that the premises have a lower epistemic status than supposed or that the inference is weaker than supposed.

Just as one must assess arguments, one must assess counterarguments, and there can be counterarguments in defense of the original arguments, and then further counterarguments to these, forming a sometimes elaborate back and forth. No doubt you're familiar with this process from debates or the sorts of conversations that we ordinarily call "arguments".

If an argument has been successfully refuted, then the epistemic status of the argument's conclusion returns to what it would have been if the argument had never been given. Thus, if all the arguments in favor of a proposition have been refuted (assuming that all of our reasons to believe the proposition come from arguments) the proposition is *unfounded*.

In some cases, it may be possible to go on to prove that the proposition is false, but it is important to recognize that one does not need to do this in order to be justified in rejecting the proposition. This is an important principle of logic called "The Burden of Proof Principle" (or sometimes the "Onus of Proof Principle"). The principle states that the *burden of proof* is on the person who *asserts* a proposition. The point is usually made in the context of a debate between two people. If one of them makes a claim, then it is his responsibility to give the other person reasons to accept it, and the other person is not required to give reasons to reject it. The second person's responsibility begins only after the first person has given such reasons. Once these reasons have been given, the second person must either then accept the proposition, or give some counterargument either to the first person's argument or to the proposition itself. But though often formulated in the language of a debate between two parties the point applies equally well to a solitary person considering the epistemic status of a proposition: he needs a reason to regard it as certain, probable, or even possible; and, if unless he has some such reason, he needs no further argument in order to reject it as unfounded.

To appreciate the importance of the burden of proof, consider its application in the law, where it is familiar in the form of the slogan that a defendant is "presumed innocent until proven guilty". When a prosecutor charges you with a crime, it is his responsibility to provide arguments that you are guilty. Once he has done so, your job (or your lawyer's) is to refute those arguments. In some cases, you may be able to prove that you did not commit the crime—for example, you might produce witnesses who testify that you were elsewhere when the crime occurred—but you have no obligation to do so. If the prosecutor makes his charges arbitrarily, without supporting them with arguments, then it is a sufficient defense to simply point out that he has done so. And, if he does make arguments that you are guilty, it is sufficient to refute them.

Of course, in a criminal court, any counterargument good enough to show that there is a reasonable doubt about your guilt is sufficient for an acquittal. If this were the best you could do, however, people would rightly continue to think that you probably committed the crime. But, suppose that you entirely refuted the prosecutions' arguments, showing that they provided no reason whatsoever to believe that you were guilty. Then, even if you provided no positive argument that you couldn't have committed the crime, wouldn't it be unreasonable and unjust for anyone to continue even to *suspect* you? Wouldn't you have cleared your name?

You may be thinking that refuting the arguments for a proposition is good enough for most purposes, but that the matter remains unsettled until the proposition has been positively disproven. If so, consider the following point: if one was not entitled to dismiss unfounded propositions *without argument*, it would be impossible to disprove anything. Why not? Suppose that after being accused arbitrarily of committing a murder, you tried to prove your innocence producing witnesses who could testify that you were elsewhere when the murder was committed. Your accuser could respond by alleging, without any evidence, that the witnesses are accomplices. Even if you had a whole football stadium full of witnesses and time-stamped video footage (for example, if you were on the Jumbotron during a game), he could allege that they are all part of some vast conspiracy and that the footage is forged. Or he might say (again without any evidence) that you have an identical twin who was in the stadium, or that you have magic powers and are able to be in two places at once, or that you committed the murder remotely by telekinesis. If he is allowed to assert whatever he likes without giving reasons, and it cannot be dismissed until it is disproven, then no matter what arguments you give to disprove each assertion, he can always dream up further unfounded assertions that will allow him to maintain it, and then the matter won't be settled until you disprove these new assertions. On the other hand, if you are entitled to dismiss his unfounded comebacks simply because they are unfounded, then you are entitled to dismiss his initial accusation on the same grounds.

Because unfounded propositions are to be dismissed, once all the arguments in favor of a proposition have been refuted, the proposition should be dismissed, even if there are no further arguments against it.

However, it must be emphasized that it is not legitimate to dismiss a proposition as unfounded until all of the arguments for it have been refuted. If there are multiple arguments, it is not sufficient to refute only some of them, and it is not sufficient to merely cast some doubt on the arguments with weak counterarguments. In situations in which many arguments have been offered for a proposition, the task of evaluating (and possibly refuting) all of them may seem daunting, but usually these arguments fall into families with similar structures, and the families can often be assessed as wholes. For example, as we will discuss in class, there have probably been thousands of distinct arguments offered over the centuries for the existence of God, but they are all variations of a handful of archetypical arguments. Thus, one needn't assess each argument individually.

§5.2 Assessing Arguments for Incompatible Conclusions

We have been discussing why refutations of propositions are not needed (though they are sometimes possible) once the arguments in favor of a proposition have been refuted. But how should we assess a conclusion when we have two arguments one for it and one against, neither of which have been refuted. There are two ways in which this situation can come about. First, someone who denies the conclusion of an argument might present an independent argument against it, instead of trying to refute the original argument. Second, different people might independently present you with arguments for logically incompatible conclusions. For example, there may be one argument that pleasure is the only thing of value and another argument that wisdom is the only thing of value. Since two different things cannot each be the only thing of value, the conclusion of each argument entails that the conclusion of the other is false. Thus, for each of the two conclusions, you have an argument that it is true and another that it is false.

How then should we proceed when we have separate arguments for incompatible conclusions? The first step is to make sure that the conclusions are in fact in compatible—they may merely appear to be so. But assuming that you have done this you should proceed to the following steps.

First, evaluate each argument on its own. In doing so you might refute one or both of them, in which case there is no further question of how to proceed. You assess the remaining argument (if there is one) as though it was the only argument that was given.

If neither argument is refuted, the next question to ask is whether either argument, when taken in isolation, seems to establish its conclusion as certain. If both arguments seem to do so, then you know that there must be an error that you have not found in one of them. Until you identify the error, both conclusions remain possible, but neither is probable or certain.

Now let's consider a situation in which one of two arguments—we'll call it Argument A—seems to establish its conclusion with certainty, while the other—Argument B—seems to establish the contrary proposition as merely possible or probable. In this situation, you should reexamine Argument A in light of Argument B, to see if there is anything in B that casts doubt on any of A's premises or in its inference. (This last part is especially important if A is not a deduction.) If there is nothing in B to cast doubt on A, then A's conclusion remains certain and B's is disproven.

In a situation where each argument, taken independently, seems to establish its conclusion as probable, you need to consider each in the light of the other. Either one will undermine the premises or inference of the other in which case that argument's conclusion emerges as probable and the others as either possible or unfounded, or the two arguments will partially undermine one another, in which case both conclusions will probably emerge as possible, with neither as probable. For example, if both arguments are inferences to the best explanation, then it cannot be true that two incompatible explanations are each the best: either one is better than the other, in which case only that one is the best explanation, and only the conclusion of that argument is probable; or else the explanations are equally good, in which case neither is best, and the each of two conclusions is (at most) possible.

§5.3 Combining Multiple Arguments for the Same Proposition

When you have multiple arguments for the same proposition, first assess each individually. If any of them taken by itself establishes the proposition as certain, then the proposition is certain, regardless of the quality of the others.

In some cases, when multiple arguments each independently establish a conclusion as possible or probable, their force can add up, making the conclusion more and more probable, and ultimately certain. This is pattern of reasoning is common in induction and especially in inference to the best explanation. However, in order for the force of arguments to add up in this manner, they cannot remain wholly separate arguments, each offered in isolation from the others. They need to be brought together in what we might call a “master argument,” to show how each of the arguments reinforces the others. For example, when a certain hypothesis turns out to be the best explanation not just of one effect, but of many very different ones, and when it can be shown that the other possible explanations of each effect cannot explain the others, this strengthens our reason for believing that the hypothesis is the correct explanation for each of the effects, and so several inferences to the best explanation are integrated together into one overarching and increasingly powerful argument for the hypothesis. This is how the theory that matter is made of atoms was proven over the course of the 19th Century, as it became apparent that it could explain a wider and wider range of phenomena.