Chapter 4: More Inductive Reasoning

Let’s review. You’ve learned about the structure of arguments (premises and a conclusion), how to recognize arguments, and about deductive and inductive arguments. Now, we will focus on inductive arguments in more depth. Remember (it never hurts to remind you again), while deductive arguments aim to provide certainty to their conclusions, inductive arguments aim to provide conclusions that are likely or probable based on the best possible evidence or support.

Barrel of Apples Example

A nice way to illustrate inductive reasoning is through something called “the barrel of apples example” (this is much easier to get than the prisoners and the hats). Although this example leads us into inductive reasoning as a whole, the example itself follows the form of a sample argument, as discussed below. Consider the following. There is a barrel full of 100 apples. Without looking inside the barrel itself, we start picking apples out of it. Let’s say we find that the first few apples we pick out are rotten. It would, clearly, be bad reasoning to conclude that the other 99 apples will be rotten, if we'd only picked out one that was rotten. But as you pick apples out of the barrel, at what point would you bet that every other apple in the barrel will be rotten?

Of course, we might have other reasons to believe that all the rest of the apples will be rotten after only seeing one that is rotten. Maybe we know that the apples have been sitting out in the sun for days. Or maybe we know that there are bacteria in the barrel that have likely infected the apples. Scientists, after all, often come to conclusions based on a small amount of experimental data, usually in conjunction with other beliefs about the way the world works.

The more apples we pull, the better conclusion we will be able to reach. If we have 5 rotten apples (as opposed to one) out of 95, our conclusion about the rest of the apples will be stronger, especially if we have pulled the 5 apples from different sections of the barrel. Consider

elections, where a representative sample of voter belief is taken from different parts of the country, rather than from one area (more on this later).

What if we pulled 50 rotten apples from the barrel? Would you be sure yet that the other 50 apples are rotten? Would you bet money on it?

What if we pulled 99 rotten apples from the barrel? Would you be sure that the last apple is rotten? Would you bet your life savings on it? Would you bet your life on it?

No, I'm guessing, you wouldn't bet your life on it. Because you know, as I do, that it's *possible* that that last apple is not rotten, despite the fact that the last 99 were rotten. Maybe it's a lucky apple, maybe someone threw a perfectly good apple into a barrel of rotten ones. The fact that we can't be 100% certain (as the blind man from The Prisoners and the Hats scenario in the last chapter is certain), is why we wouldn't bet our lives on it. And if you would, you've got a death wish, and not a very good ability to reason.

The fact that we can't be 100% certain is also why this argument is still inductive, even when we have 99 rotten apples. Our conclusion about the last apple, whatever it is, can only be about the probability that the last apple will be one way or another. It can never prove what the last apple is like (as it could if it were deductive).

For clarity, let's see what a *deductive* argument would look like regarding the barrel of apples:

1. There are 100 apples in a barrel.
2. I know that half the apples are rotten, and half are not rotten.
3. I have pulled 50 apples out of the barrel and they were all rotten.

So the next apple I pull out of the barrel will not be rotten.

Let me stress that the previous argument is deductive, meaning that, if the premises are assumed to be true, then the conclusion follows with certainty, in contrast to the inductive arguments below.

Another way to get the point across is this: premises of inductive arguments do not *prove* their conclusions, but rather *support* them. If someone ate your last microwave burrito, and your roommate Joey loves microwave burritos, then it's likely that Joey ate your last burrito. But it's not certain. The fact that Joey likes microwave burritos supports the conclusion that he is the one who ate yours. But it's really not very strong support for that conclusion. Now, if you later find the burrito wrapper on Joey's desk, this is even stronger support for the conclusion. But it's still not as strong as a deductive argument where the conclusion is proven with certainty. A good deductive argument is *sound*, as you should know, but a good inductive argument is *strong*.

> *Strong inductive argument*: more support is given for the conclusion.
>
> *Weak inductive argument*: less support is given for the conclusion.

Here is a strong inductive argument using the barrel of apples example:

> 1. There are 100 apples in a barrel.
> 2. I know 75 are not rotten because I pulled them out of the barrel.
> Thus, the last 25 will not be rotten either.

Here is a weak inductive argument using the barrel of apples example:

> 1. There are 100 apples in a barrel.
> 2. I know 3 are not rotten because I pulled them out of the barrel.
> Thus, the last 97 will not be rotten either.

The latter argument is weak because there is less support for the conclusion—only 3 apples were pulled out of the barrel. The former argument is strong because there is more support given—75 apples is a larger sample size.

With each piece of evidence we find, the conclusion of an inductive argument becomes more likely to be true, just like taking more and more apples out of the barrel. You might also see inductive arguments as building more and more evidence for a conclusion, though this evidence does not always have to be physical/scientific. Evidence must be defined broadly here since, in an inductive argument, it might come in the form of a non-empirical reason/principle. For example, consider this argument against slavery:

1. Most people who kill other humans are morally bad.
2. In court, Terrance claimed that he killed his victim in self-defense.

Thus, Terrance is morally bad.

There are two forms of evidence given: in premise 1, we find evidence in the form of a reason/principle. It is not physical/scientific evidence. But premise 2 does count as physical evidence since it refers to an observable event. The argument is inductive because the conclusion is not proven. Even if the premises are assumed to be true, Terrance may not be morally bad since it is arguable whether killing in self-defense makes one a morally bad person. The key point about evidence or support, though, is that in an inductive argument it should be conceived broadly.

Unstated Premises (again)

Sometimes it's not immediately apparent whether an argument is aiming to be inductive or deductive. This ambiguity often results from the fact that arguments can have unstated premises, as discussed in the first chapter. Context and content sometimes make it unclear whether an unstated premise makes an argument deductive or inductive. When it’s unclear whether the argument is deductive or inductive, it’s best to attribute a believable unstated premise to the speaker. In other words, it's best to ask yourself what the speaker is really trying to say.

Let's say, for example, that my friends and I are planning to go to the beach. One of my friends says she doesn't want to go to East Beach, because every time we've gone in the past, the beach has been crowded. She is probably making an inductive, rather than deductive argument, that looks like this:

> 1. The times we've gone to East Beach in the past, it's been crowded.
> 2. If we go now, it will also be crowded.
> So let's not go to East Beach.

It's debatable, but given the context, my friend was probably not intending the first premise to be: "*Every* time we've gone to East Beach in the past, it's been crowded" which would make it deductive; instead, as noted, she seems to be intending an inductive argument. So if the speaker does not make it clear whether her argument is deductive or inductive, then the nature of the argument depends on the context and your judgment of the speaker's intentions.

There are three common forms of inductive arguments: statistical syllogisms, sample arguments (sometimes called inductive generalizations), and arguments by analogy.

Statistical Syllogism

Many deductive arguments go from more general premises, to a more specific conclusion. For example, if all men are mortal and Socrates is a man, then Socrates is mortal—notice that the conclusion of this deductive argument is quite specific. Inductive arguments are often the opposite of deductive in that they go from specific premises to a more general conclusion. However, this is not always the case, as we can see with the statistical syllogism:

> 1. Most Xs are also Ys.
> 2. This is an X.
> Therefore, this X is also a Y.

Inductive arguments, as you can see from the above, can occasionally go from general premises to a specific conclusion. So what really makes an argument inductive is simply that the conclusion is based on probability.

Don't get confused by the variables like “X” and “Y.” These are just, well, variables, and they stand for terms, just like with symbolic logic from the last chapter. We could replace “X” in the above argument with “male philosophy professors” and we could replace “Y” with “wear glasses” (we might have to rephrase the argument a bit to fit, but the point remains the same).

With this inductive argument form, the statistical syllogism, the likelihood of the conclusion depends on how many Xs are actually Ys. Most people living in the US are US citizens, but we know there are some who are not. As you should know from the discussion above, an inductive argument is not weakened by the fact that there are some exceptions (in fact, exceptions are what *make* the argument inductive rather than deductive). However, the argument *is* weakened if the generalization in the premise includes, say, a stereotype, as does my example: it's a stereotype, an inappropriate generalization, to say that all male philosophy professors wear glasses. So that argument would be a weak inductive syllogism. Later we will see that it is considered a fallacy, a mistake in reasoning, to generalize from a sample that is too small or not representative.

Sample Arguments

The sample argument does, in fact, go from specific claims in the premises to a general claim in the conclusion. We generalize from specific samples to establish general statements about a population when we haven’t observed all its members. Generalizing is not as simple a thing as it seems. It requires probabilities and statistics, but we're not going to get into all the complications. It's enough for this class that you get the gist behind each argument form. Here is the form of a sample argument:

1. Such and such a percentage of observed Xs are Ys.
Therefore, the same percentage of all Xs are Ys.

First of all, note once again that the strong and weak versions of the barrel of apples argument above are sample arguments.

How can we evaluate these types of arguments? The observed Xs make up what we call the *sample*. All the Xs make up the *population*. The *attribute of interest* is the property of having or being Y. And *n* is equal to the sample size. So the sample in, say, the strong barrel of apples argument consists of the 75 apples pulled out of the barrel, where n=75. The population is the full barrel of 100 apples. The attribute of interest is "lack of rottenness" or "non-rottenness" since that is the feature we are generalizing about from premise to conclusion.

The primary question in evaluating such arguments becomes: how likely is it that the same proportion of the target population has that feature/attribute of interest? In making these evaluations, it helps to know as much as you can about the sample.

Is it a *biased sample*? A sample is biased when it contains a disproportionate number of things with a given attribute of interest. If a sample has more New Englanders than anything else, then the sample is biased with respect to New Englanders. To have a good sample argument, the sample should be as diverse as possible, and large enough to contain that diversity in the first place.

A *random sample* is a sample in which every member of the population has a chance of being included. This can be quite complicated to determine, but for our purposes think of the barrel of apples. Is it more likely that you'll get a random sample if you pull every apple from one side of the barrel or if you pull them out from all different parts of the barrel? It's pretty clear that pulling apples from as many parts of the barrel as possible will give you the best random sample, leading to a stronger argument. If you pull them all from one area, it might be that that is the only area with rotten apples—just as if you take a US opinion

poll on some topic and include only people from Kentucky, your results will be biased.

While it's true that sample sizes in studies and polls are a small representation of the whole, the generalizations we make in everyday life are usually based on an even smaller representation of the whole—like one or two cases. This is something to be aware of when making generalizations (see the discussion of the *hasty generalization* fallacy below).

When samples start getting large (500 or more) it gets increasingly difficult to lower what's called *the error margin*, or the chance you have of being wrong about the nature of the population. For the error margin to go down even a percent, it takes another 500 added to the sample. This is why most samples generalize from 1000 or 1500 to the whole, as Gallup polls do.[1] Gallup polls are real-life examples of sample arguments. Gallup takes random samples from state to state of American opinions, asking many different types of questions. They have a good track record of predicting wins and losses of political candidates, among other things (though they do occasionally get it wrong, with the help of overzealous news anchors, as happened with the 2016 US presidential election). Research in social science and the conclusions drawn from it, including those in psychology and sociology, are often expressed with sample arguments.

The inductive sampling method works! But interestingly, some have questioned the samples that Western scholars use to draw conclusions about human nature. For example, many samples in psychology and other related fields are drawn from Western, educated, industrialized, rich, and democratic societies (the acronym, appropriately, is WEIRD). Can we draw convincing conclusions about *all* human nature from such biased samples?[2]

---

[1] See the Gallup website here: http://www.gallup.com/home.aspx.

[2] See the original journal article discussing this topic here: Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences, 33*(2-

Argument from Analogy

1. X and Y both have attribute a.
2. X has attribute f.

So Y has attribute f.

The fact that two things are similar in one way increases the probability that they are similar in other ways—this is the basic idea behind arguments by analogy. But the strength of the argument depends on how alike the attributes are being compared. If two people are famous, then this increases the likelihood that they are talented musicians (since a good amount of famous people are talented musicians), though not by much. But if two people are famous, then this increases the likelihood that they are both rich, we might say significantly. After all, fame is ordinarily more likely to be correlated with wealth than with talent. Moreover, being a talented musician is even less probable, given the other premise, since such a musician is part of a subset of people who are famous and talented generally (one can be talented in things other than music).

Not every similarity between two things shows that other similarities might be shared. If Paris Hilton and Neil Young are both Christians, this would *not* increase the probability that they are both talented musicians (this isn't to say that there aren't some talented Christian musicians out there).

A good test of your understanding of arguments by analogy is your ability to determine the difference between an analogy by itself and an argument by analogy. The homework asks you to make this distinction. An analogy just compares things, but an argument from analogy draws a conclusion from that comparison. For example, here is an analogy:

---

3), 61-83. And see a less technical version of the ideas within that article here: Brookshire, B. (2013, May 8). Psychology is WEIRD. In *Slate*. Retrieved from: http://www.slate.com/articles/health_and_science/science/2013/05/weird_psychology_social_science_researchers_rely_too_much_on_western_college.html.

“Taxing people is like theft.” This claim just compares taxation to stealing, it doesn't draw a conclusion about it. But this is an argument: “Taxing people is like theft. Theft is wrong, so taxing people is wrong too.” Now, an additional premise and conclusion were added, making it an argument by analogy.

The Argument by Analogy for Other Minds

If you're still not seeing the significance of arguments by analogy, consider some famous ones. Let's first recall Descartes' and Avicenna's point from the chapter 2 that we can know only our own internal thoughts with certainty (I think, therefore I am). This implies that we can't even know that other people are thinking thoughts at all; for all we know all other people are robots that've been programmed to imitate what people would look and act like if they actually had internal thoughts. How can we get around this? Ladies and gentlemen, I present to you the argument for other minds:

> 1. I am a human being with a brain that, presumably, is in some way the cause of my internal thoughts.
> 2. Other human beings have brains (presumably within their heads).
> So other human beings, like me, have internal thoughts.

Notice that the argument for other minds doesn't *prove* that other people think, but it gives decent reasons to believe that they do.[3] The conclusion of an argument does not have to be certain (deductive) to be reasonable. Also, if you believe that what causes thoughts is not the brain, but a soul, or some combination of the two, just replace "brain" with "brain/soul" in the above argument—the general logic remains the same.

[3] The philosopher Bertrand Russell made this argument. See the first page of this link for his reasoning: http://hermitmusic.tripod.com/notes_russell_ryle_dennett.pdf.

The Argument by Analogy for Design

Another famous argument by analogy is the argument from design, famously articulated in William Paley's watchmaker analogy.[4] Imagine that you are walking through the woods, and you stumble over something. You look down, and there is a rock. You might curse, and keep walking, but think nothing more about it. However, what if you look down and there is a watch? You might pick up the watch and, whether it still works or not, one question that might strike you is: I wonder who *designed* this watch? What's the difference between the watch and the rock? Well, the watch has working parts that function together, and the rock apparently doesn't. Similarly, you might look around and observe that the world apparently has "working parts" like the watch (the laws of nature, some observable order, beauty, etc.). From this, you might argue that, as the watch requires a designer, the world does too. Thus, you might conclude, there *is* a God, the being that designed the universe.

If the above logic has struck you before, you are the latest of a long line of thinkers who've been struck by the same logic over the years including Socrates, St. Thomas Aquinas, and, as noted, William Paley. As should be clear, the watchmaker scenario is an argument by analogy (for practice, try putting it in premise/conclusion format on your own).The argument is sometimes referred to as the teleological argument (deriving from the Greek words telos meaning "purpose" and logos meaning in this context "the study of").

Philosopher David Hume famously evaluated the argument in depth.[5] For one, he pointed out that no one clearly defines what they mean by "God" when making the argument. Who is the creator? Allah? Zeus? Yahweh? Hume also pointed out that the logic does not exclude a polytheistic (many Gods) view. After all, the argument compares the

---

[4] Paley, W. (2010). *Natural Theology*. M. D. Eddy & D. Knight (Eds.). Oxford, UK: Oxford University Press.

[5] Hume, D. (2012). *Dialogue Concerning Natural Religion*. CreateSpace.

way humans design things to the way a higher being(s) may have designed the universe. When humans create stuff, there are often many people involved (like on an assembly line). Humans also mess up when they create stuff. So what if God messed up on our universe? Are we living in a crappy universe?

Finally, Hume argued that comparing humans to God is a bad idea, and a bit presumptuous. How can we presume to compare the way we create to the way God might create? The most we can say from the argument, Hume claimed, is "I don't know if there is some kind of higher being."

Whatever view you take on the strength of the argument by design, it should be clear that evaluating arguments by analogy is not an exact science. General information must be known to evaluate arguments by analogy. In the example above about musicians, for example, we need to know that many famous people are musicians, and many of those musicians are talented. The argument is made stronger or weaker depending on how well we understand these general claims.

To attack an argument from analogy you must attack the analogy to show that, given the weakness of the analogy, the conclusion isn't likely to follow from the premises. Hume does just this with the argument by design. However, notice that Hume's critique applies primarily to a more specific creator, not as much to a more general creative force or power.

Some Last Words on Inductive Reasoning

We'll discuss fallacies (or mistakes in reasoning) in more depth later, but it's worth pointing out now two common but related inductive fallacies: the *hasty generalization* and the *generalization from an exceptional case*. If our sample is biased, or too small, and we try to draw a strong conclusion from it, we have committed a hasty generalization or a generalization from an exceptional case, since we are not statistically justified in drawing a strong conclusion from such samples. If I use only the example of my cousin to conclude that "All Irish people are drunks"

then my sample size is too small, and I've committed a fallacy. Hopefully you can see the way bad reasoning underlies much prejudice and racism (though clearly this is intermixed with purely emotional prejudices as well). A common variation on the hasty generalization is anecdotal evidence, a fallacy in which we assume that our personal story (or anecdote) successfully generalizes to other people or events.

Something else about inductive claims is that they can be shown to be good or bad based on their relation to alternatives. If we are trying to figure out how to fix our computer, will we listen to our mother who can't use the DVD player, or someone with a degree in computer science? Why do I have a headache? In the absence of a better explanation, I might say my headache is due to being out in the sun too long. (Recall that inferences to best explanations, IBEs, are usually inductive).

Thankfully, from the barrel example we can see that there are some objective standards that help us determine when one inductive inference is better than another: for example, a greater representative sample lends itself to a stronger conclusion. We can have a sliding scale of reasonableness where some inductive arguments are objectively more convincing.

Finally, it is important to notice here that very weak ideas can appear to work (that is, they can appear to be reasonable). If we only take a small sample, we can conclude that what we are trying to argue for might be true. Take, for instance, UFO (spacecrafts from another planet, for the purposes of this example) sightings. Many believers in UFOs will tell you that, given the evidence they have, there is no other conclusion except that UFOs exist. But what sort of evidence are they working with? Usually, they have a few pictures, a personal story, and maybe some quotes from government officials. Is this "undeniable" evidence? The UFO proponent thinks he has created a deductive argument, when in fact he has only created a very weak inductive argument (some would say a fallacy). Of course he might still be right—given that inductive

arguments are based on probability and not certainty, *anyone* making such an argument *might* be right—but we should still be honest about the overall strength of the argument he makes.

Major Ideas for More Inductive Reasoning

Although anything from the readings or homework might appear on the assessments, the following **major ideas** should be clearly understood.

- Barrel of apples example
- Statistical syllogisms
- Sample arguments
- Arguments by analogy