Clear and Present Thinking

A
handbook
in logic
and
rationality
.

Version 1.1

Northwest Passage Books

Project Director:
Brendan Myers (CEGEP Heritage College)

Authors:
Brendan Myers
Charlene Elsby (Queen's University)
Kimberly Baltzer-Jaray (University of Waterloo)
Nola Semczyszyn (Franklin & Marshall College)

Editor / Proofreader:
Natalie Evans (University of Guelph - Humber)

Layout and Design:
Nathaniel Winter-Hébert, Lana Winter-Hébert
www. winterhebert.com

Version 1.1 (21st May 2013)

Northwest Passage Books

For all other enquiries, please visit *brendanmyers.net*

Chapter 3: Basics of Argumentation

Let's define argumentation as the process of seriously debating the worth and the merits of some proposition. The word 'argument' here does not refer to an angry shouting match. Rather, it refers to **any two (or more) statements in which one is the reason for the other**; one is supported by the other(s), or one follows from the other(s). We 'build' arguments by assembling together basic statements into particular structures, and having assembled them together that way, we can more easily test to see whether the ideas being discussed are worth your time.

3.1. Propositions

Arguments have various parts. And the part that's easiest to identify is called the **proposition**: also sometimes called the **statement**, or the **claim**. (For the purpose of understanding argumentation, these terms mean the same thing, and are often used interchangeably.) A proposition is **a simple sentence that has just one meaning**, for it expresses one thought according to the rules of grammar in one's language. Also, a proposition **asserts that something is the case, or is not the case**. When a proposition asserts that something *is* the case, it is also called an **affirma- tion**; when a proposition asserts that something *is not* the case, it is also called a **negation** or **denial**.

Not all sentences are propositions. Some sentences are questions, some are commands, some are emotional exclamations, and some are poetic devices like metaphors. One way to recognize a statement is to look for sentences that could be given as a direct answer to a straightforward question. Another is to look for sentences that could be either true or false; a sentence that one could agree with, or disagree with.

With that in mind, which of the following sentences are propositions?

- The lamp on my table is switched on.
- Good morning everyone!
- My sweater is green.
- How many cars are parked outside right now?
- Smoking is bad for your health.
- Smoking is good for your health.
- Stop driving on the wrong side of the road.
- The revolution will not be televised.
- My love is like a red, red rose.
- WTF?

Also, it is possible for a single sentence to contain within it more than one proposition.

- It's raining today, and I'm feeling blue. (Two propositions)
- The book on my table is well-read, but boring. (Two propositions.)
- This new kitchen gadget can slice any vegetable, as well as any fruit, but it can't handle meat. (Three propositions.)

And, it's also possible to have a paragraph of dialogue in which only one or two sentences are propositions, and the rest of the paragraph is made of expressions that, while they might help communicate the speaker's feelings, are not expressions that can be used to build an argument. Consider this example:

> "The other day, I was really pissed off. I ordered this new computer from the Internet. And it took three weeks to get here, which was bad enough. Then when it arrived I got so mad again! Because the one I ordered was silver, but the one they sent me was black! Somebody in that company is asleep at the wheel."

Clearly, the speaker here is angry about this situation. But if the speaker wanted to draw any logical conclusions from this discussion, for instance about what to do, or about whether to trust the company again, the only relevant sentences here are the ones which stick to the facts. Here's the same discussion again, with the irrelevant expressions crossed out:

> "~~The other day, I was really pissed off~~. I ordered this new computer from the Internet. And it took three weeks to get here, ~~which was bad enough. Then when it arrived I got so mad again! Because~~ the one I ordered was silver, but the one they sent me was black! ~~Somebody in that company is really asleep at the wheel~~."

As you can see (I hope!), it's really easy to tell the difference between a sentence that is a decent and useful proposition, and another that isn't. Logic starts to look complicated when there are lots of propositions with lots of relations to each other. But even the argument with thousands of lines is still made of simple, straightforward true-or-false sentences like these. The other parts of the argument have to do with the way that propositions are used, or the way they are positioned in relation to other propositions in the general structure of the argument. If you can figure out this part of the textbook, you can figure out everything else!

Once we have sorted out which sentences are propositions and which sentences are not, we are almost ready to put them together into arguments. It's possible to have a sentence which is a proposition, but which you can't use in an argument because of vagueness or an ambiguity in its words or grammar. For example:

> "Women are stronger than men."

This looks like a perfectly ordinary proposition: it could be either true or false. We could stage arm-wrestling or weight lifting competitions to test it. But is that what the word 'stronger' means here? Or, does it mean that women have more willpower than men? Does it mean that women have thicker and tougher bones than men? Does this statement generalize about the 'average man' or the 'average woman'? If we do not have the context or the meaning of the word 'stronger' here, then this proposition is probably too vague to be used in an argument. The various uses of the word "stronger" are homonyms, and the sentence is vague because we don't know which sense it is that the speaker or writer means. That is an issue separate from the issue of whether the proposition, once properly understood, is true or false.

> "People who get good marks in school are very intelligent."

Again, this looks like a decent proposition, but one might want to clarify the meaning of the word 'intelligent' before using it in an argument. The problem here isn't just that someone could counter-propose that some intelligent people get bad marks in school, or that some stupid people get good marks. Those kinds of issues can come up when the argumentation is underway. But before we get that far, we have to know what the speaker means by the word 'intelligent'. Is it just a matter of the ability to perform well on school tests? Is it the ability to speak clearly and sound like you know what you're talking about? Is it the ability to solve problems quickly? Is it something else?

> "Beer is better than wine."

A judgment of value can act as a decent proposition. But in an example like this one, we would need to know what measure of value is being used here. Is beer considered better because it is cheaper? Or because it has less alcohol content? Or because it's easier for people to make their own beer at home? Or, is this person merely expressing a personal taste preference? Also, given that there are thousands of recipes for beer, and thousands of recipes for wine, it might not be clear what kind of beer and what kind of wine is being compared.

It is often the case that propositions like these are clarified by introducing the argument with a few handy definitions. The definitions might not form part of the argument, but they can provide the context or the background information that will allow debaters to understand each other and then decide whether they agree or disagree.

Propositions can also be clarified by their position in the argument, and their relationship to other propositions.

3.2 Parts of Arguments

Once we have figured out what a proposition is we can build arguments by arranging propositions into particular relationships with other propositions. Remember, an argument needs at least two propositions, not just one.

The first type of proposition that an argument needs is a **premise**. This is a statement given in support of another statement; it is the *reason* why another statement should be accepted as true. Propositions can come from your world view, or your personal experience, or some other trustworthy source. Most arguments have more than one premise and most arguments state the premises first.

The other type of proposition that an argument needs is a **conclusion**. This is the 'point' of an argument; it is that which is supported by the premises; it is that which the speaker is trying to persuade another person to believe is the case. Rather than coming from your experience or your world view or some other source, the conclusion follows from the premises of the argument.

The difference between the premises of an argument and its conclusion are not differences in the statements themselves. Rather, to identify which are the premises and which is the conclusion, you have to rely on context. What is being used as a reason, and what is supposed to follow from those reasons? Sometimes a conclusion that follows from a number of premises is then used as a premise for another conclusion. Consider the following argument:

> "I don't believe he's telling the truth. You see how his eyebrow twitches, and he's sweating a little more than normal. If he is lying, you shouldn't give him your money."

In this example there are two arguments. The speaker intends to support the conclusion that "he is not telling the truth/he is lying" with the premises that "his eyebrow twitches" and "he's sweating more than normal". The conclusion that "he is lying" is used again as a premise, to support the conclusion that "you shouldn't give him your money", which is the *overall conclusion* of the argument.

Stories, poems, explanations, speeches, and so on, can sometimes look like arguments. They might even be made up of statements. But if they do not have premises giving you reasons for accepting conclusions, then they are not arguments. This, in case I haven't mentioned it yet, is why thinking logically about something is often called 'reasoning' about it.

The other parts of arguments have to do with the way premises and conclusions are put together.

An **inference** is the name for the relationship between statements in an argument. It is a line of logic between propositions that lead you from the premises to the conclusion. Inferences are often embodied in certain **indicator words**, which show you which way the direction of the argument is flowing. Here are a few examples of indicator words:

- Because
- Since
- Given that...
- Which means that...
- We can conclude that...
- Hence
- It follows that...

- Therefore
- Consequently...
- This implies...

...and so on. I've mentioned that an argument needs at least two propositions. But two propositions placed side by side do not make an argument. There must be a relationship between them, showing that one leads you to the other, one supports the other, and one follows from the other. That relationship is called an inference; and between its propositions an argu- ment must have inferences too, or else it is not an argu- ment. The indicator words "Because", "Since", "Given that" (and many others) indicate that what follows the indicator word is being used as a premise or reason to support a conclusion. Indicator words that indicate the conclusion are "Which means that", "We can conclude that", "Hence", "Therefore", "Consequentl y", etc.

3.2 Truth and Validity

Truth, in this way of understanding logic, is a property of propositions. As we've already seen, arguments must be made of sentences that could be either true or false, and not from other kinds of sentences. And there are various ways we could find out whether a given proposition is true. For example:

- The proposition corresponds to the facts, as you are able to observe them or somehow prove them (this is called the **Correspondence** theory of truth).
- The proposition is acceptably consistent, or 'coheres well', with other statements that form part of your world view (the **Coherence** theory).
- When put to some kind of test, the proposition turns out to be a very useful and practical thing to believe (the **Pragmatic** theory).

As truth is a property of sentences, so **validity** is a property of inferences. We say that an argument is valid if its inferences lead you properly from premises to conclusions. Validity is determined by looking at the form, or the structure of the argument, and *not* the content – those are two separate issues.

And finally, **soundness** is a property of argu- ments as a whole. An argument is sound if it has true premises and valid inferences. Both of these conditions must be met

Arguments themselves also come in two main types: **deduction** and induction. A deduction, or a deductive argument, is a type of argument that, if it begins with true premises, logically guarantees that the conclusion is also true. Deduction works because in a deductive argument, nothing appears in the conclu- sion that was not already present in at least one of the premises. You can think of a deductive argument as a kind of 'unpacking' or 'synthesizing' of the premises.

An **induction**, or an inductive argument, is a type of argument that asserts the likelihood of the conclusion. In an inductive argument, if the premises are true, then the conclusion is probably true. Unlike a deduction, an induction can go beyond what is asserted in the premises. Its conclusion can say more than what the premises say. For example, you can use an induction to make a prediction about the future. But an induction cannot guarantee the truth of a conclusion, as a deduction can do.

[…]

3.5 Some Common Deductive Argument Forms

Earlier we stated that the definition of an argument is "**any two (or more) statements in which one is the reason for the other**". This section will introduce some valid deductive argument forms. In deductive argumentation, we take some number of premises as given, and from these we are able to make other claims according to certain logical rules of inference. If the conclusion that results comes out of the given premises as a result of applying the accepted rules of inference, then we say that the conclusion follows necessarily from the premises, or that the argument is "valid".

The validity of an argument is determined not by what it says, but by its *form*. That means that when we assess the validity of an argument, we assume that the premises are true. If, on the other hand, we want to question the truth of the premises, we would be evaluating not its validity, but its *soundness*. Consider the following argument:

All Pigs can fly.
Babe is a Pig.
Therefore, Babe can fly.

This argument is valid. That is, assuming that the premises are true, the conclusion necessarily follows. Of course, we can question the *soundness* of the argu- ment. If we can disprove the premise that "All pigs can fly", then the argument would be unsound. We might also question whether we want to consider Babe a pig, rather than a fictional character resembling a pig. In either case, if either one of the premises is not true, then the argument is not sound. But that does not mean it is not valid. An argument can be valid without being sound. Let's look at an example of the same *form*:

All humans are mortal.
Brendan is a human.
Therefore, Brendan is mortal.

This argument is both valid and sound. […]

3.6.1 Modus Ponens or Affirming the Antecedent

Modus Ponens is a valid argument form taking a conditional statement as one premise, and the affirma- tion of its antecedent as another premise. So, if I claim "If something, then another thing" and then affirm "something", I can logically deduce that "another thing". If the conditional statement and the affirmation of its antecedent are both true, the truth of the conclusion is guaranteed.

Let's take an example.

(P1) If the dog is barking, then there's an intruder in the house.
(P2) The dog is barking!
(C) Therefore, there's an intruder in the house!

Of course, there might be other reasons why the dog might bark. But according to Premise 1, the fact that the dog is barking implies that there is definitely an intruder in the house. And we are assuming that P1 is true.

This argument takes the general form:

(P1) If P, then Q.
(P2) P.
(C) Therefore, Q.

Let's look at an example:

(P1) If it is raining, then I will need my umbrella. (P2) It is raining.
(C) Therefore, I will need my umbrella.

There might be other reasons why you might need your umbrella. Perhaps it's to be used as a prop in a theatrical performance. But nothing in this argument tells you that. And besides, whether or not that's the case, the first premise still tells you that you need it when it rains.

Affirming the Consequent: Modus Ponens' Invalid Half Brother
There's a sneaky invalid argument out there that looks a lot like Modus Ponens. What would happen if instead we affirmed the consequent, instead of the antecedent? We would have an argument like this:

> (P1) If it is raining, then I will need my umbrella. (P2) I will need my umbrella.
> (P3) Therefore, it is raining.

We tend to make this logical leap and equate the fact that we need our umbrella with the fact that it's raining. But though it is not equally likely that we might need the umbrella for a theatrical performance, it is still a possibility. That is, the fact that I need my umbrella does not *absolutely guarantee* that it's raining. This argument form is therefore invalid.

Practical Uses of Modus Ponens:
Every circuit in your computer uses this pattern of ar- gument to make calculations. In effect, the diodes and transistors in your computer CPU are like 'switches', which operate as if they are reasoning like this:

> If a signal comes in from direction X, then send it out again in direction Y.
> A signal just came in from direction X.
> Therefore, the thing to do is send it out in direction Y.

3.6.2 Modus Tollens or Denying the Consequent

Modus Tollens is a valid argument form taking a conditional statement as one premise, and the denial of its consequent as another premise. So, if I claim "If something, then another thing" and then deny "another thing", I can logically deduce that "not something". Here I'm recognizing that if the relation between "something" and "another thing" holds,
and if "another thing" failed to happen, or is false (depending on what that thing is), then "something" must not have happened, or must not be true.

Let's take an example.

> (P1) If you gave me a diamond tiara, I'd be the happiest girl in the world!
> (P2) I am not the happiest girl in the world.
> (C) Therefore, you did not give me a diamond tiara.

This argument takes the general form:

(P1) If P, then Q.
(P2) Not Q.
(C) Therefore, not P.

Like Modus Ponens's evil half brother, there's another bad argument out there attempting at every turn to pass itself off as valid.

Denying the Antecedent: Fallacy!
Again, when we see a conditional statement and a negation, we're immediately tempted to think 'Modus Tollens'. But what happens if we deny the antecedent instead of the consequent? We get an argument like this:

> (P1) If you gave me a diamond tiara, I'd be the happiest girl in the world!
> (P2) You did not give me a diamond tiara.
> (C) Therefore, I am not the happiest girl in the world.

Again, the truth of these premises does not absolutely guarantee the truth of the conclusion. Even if you did not give me a diamond tiara, I might still be the happiest girl in the world for some other reason. I might have been the happiest girl in the world all along, and there's quite possibly nothing you could do to change that. This argument form is invalid.

3.6.3 Categorical Syllogisms

The four standard statements in categorical logic can be combined into 24 possible valid logical argument forms. But we can just look at a few of them; once you get the idea behind how categorical syllogisms are judged as valid or invalid, it's easy to discern the difference.

One valid categorical syllogism was already given in the introduction to this section. That was:

All humans are mortal.
Brendan is a human.
Therefore, Brendan is mortal.

This argument is valid. We can, in general, con- clude that if an entire class of things has some quality, and if something is a member of that class, it has that quality.

But we can also generalize further. If an entire class of things has some quality, and all of the things that have that quality have some other quality, then we can make a valid inference that the entire class also has that other quality.

For example:

All farm animals are cannibalistic.
All cows are farm animals.
Therefore all cows are cannibalistic.

If you accept the validity of the first argument, then you must also accept the validity of this argument. This makes sense, because if every individual cow is a farm animal and therefore cannibalistic, then the whole cow species is cannibalistic.

Now let's try some negative statements.

No human is immortal.
Brendan is a human.
Therefore Brendan is not immortal.

What this argument says is that if none of the members of the class of humans is immortal, then neither is a specific individual of that class. Again, we can generalize. If no specific member of the class is immortal, then the whole class is excluded from immortality.

No human is immortal.
All philosophy professors are humans.
Therefore no philosophy professor is immortal.

These are only some of the possible combinations of categorical statements that result in valid syllogisms. If you can keep track of what thing or what kind of thing belongs to what class, then you're in pretty good shape for evaluating the validity of categorical syllogisms.

3.6.1 Enthymemes

An enthymeme is a categorical syllogism in which one of the premises is missing. People use them all the time, often without realizing it, when they want to get a certain point across quickly, or when they can assume the listeners know what the they are talking about. It's really easy to commit a fallacy called 'undistributed middle' when making an enthymeme, because we aren't always keeping close track of where the premises are. So to analyze an enthymeme, one has to lay out all the propositions in the place where they would stand in a categorical syllogism, fill in the missing proposition, and then determine whether the inferences are valid or invalid.

"Many songs by Justin Timberlake are popular. So this new song will be popular too."

P1. Some Justin Timberlake songs are popular.
P2. This new song is composed by Justin Timberlake.
C. Therefore, this new song will be popular.

"He is a leprous man, for he is unclean." (Leviticus 13)

P1. Leprous men are unclean.
P2. He is unclean.
C. Therefore, he is a leprous man.

"Yond Cassius has a lean and hungry look. He thinks too much. Such men are dangerous." (Shakespeare, Julius Caesar, III.2)

P1. Cassius has a lean and hungry look and thinks too much.
P2. Men who have lean and hungry looks and who think too much are dangerous.
C. Therefore, Cassius is dangerous.

By the way: which of these enthymemes are sound, and which are not?

3.6.2 Hypothetical Syllogism

A hypothetical syllogism is a valid argument form that takes as premises two conditional statements and concludes a third, where the consequent of the first premise is identical to the antecedent of the second.

For instance, if I make the claim,

(P1) If it gets below freezing outside, I can make ice out there.

And I also make the claim that,

(P2) If I can make ice, my soft drinks will be deliciously refreshing.

Then I can conclude that,

(C) If it gets below freezing outside, my soft drinks will be deliciously refreshing.

Essentially, we are demonstrating the transitive property of conditional statements. That is, if we have two conditional statements where the consequent of one is identical to the antecedent of another, we can eliminate them and mash the rest of the two premises together to get a conclusion that is definitely true.

This argument takes the general form

(P1) If P, then Q
(P2) If Q, then R
(C) If P, then R

But this could all be made clearer by taking a few examples. We can apply the hypothetical syllogism to categorical thinking:

(P1) If Socrates is a man, Socrates is an animal.
(P2) If Socrates is an animal, Socrates is a substance. (C) If Socrates is a man, Socrates is a substance.

We could also apply the hypothetical syllogism to causal relations:

(P1) If I set the house on fire, it will burn down.
(P2) If the house burns down, I'll collect insurance money.
(C) If I set the house on fire, I'll collect insurance money.

In any case, the transitive property of the implication relation that constitutes a conditional statement guarantees that the hypothetical syllogism is valid. That is, the hypothetical syllogism can be proven valid just by the definition of conditional statements.

3.6.1 Disjunctive Syllogism

This argument establishes the truth of some proposi- tion by ruling out all other possibilities until there's just one left still standing.

60

Form:

Either P is true, or Q is true. P is false.
Therefore, Q is true.

Either P is true, or Q is true. Q is false.
Therefore, P is true.

Examples:

(P1) This tree is either coniferous or it is deciduous. (P2) I see by its flat leaves that it is not coniferous.
(C) Therefore, this tree is deciduous.

(P1) One of us is going to die here, Mister Bond. It's either you or me.
(P2) And it isn't going to be me.
(C) So it will have to be you!

[…]

3.7 Induction

All of the argument forms we have looked at so far have been deductively valid. That meant, we said, that the conclusion follows from necessity if the premises are true. But to what extent can we ever be sure of
the truth of those premises? Inductive argumentation is a less certain, more realistic, more familiar way of reasoning that we all do, all the time. Inductive argu- mentation recognizes, for instance, that a premise like "All horses have four legs" comes from our previous experience of horses. If one day we were to encounter a three-legged horse, deductive logic would tell us that "All horses have four legs" is false, at which point the premise becomes rather useless for a deducer. In fact, deductive logic tells us that if the premise "All horses have four legs" is false, even if we know there are many, many four-legged horses in the world, when we go to the track and see hordes of four-legged horses, all we can really be certain of is that "There is at least *one* four-legged horse."

Inductive logic allows for the more realistic premise, "The vast majority of horses have four legs". And inductive logic can use this premise to infer other useful information, like "If I'm going to get Chestnut booties for Christmas, I should probably get four of them." The trick is to recognize a certain amount of uncertainty in the truth of the conclusion, something for which deductive logic does not allow. In real life, however, inductive logic is used much more frequently and (hopefully) with some success. Let's take a look at some of the uses of inductive reasoning.

Predicting the Future

We constantly use inductive reasoning to predict the future. We do this by compiling evidence based on past observations, and by assuming that the future will resemble the past. For instance, I make the observation that every other time I have gone to sleep at night, I have woken up in the morning. There is actually no certainty that this will happen, but I make the inference because of the fact that this is what has hap- pened every other time. In fact, it is not the case that "All people who go to sleep at night wake up in the morning". But I'm not going to lose any sleep over that. And we do the same thing when our experience has been less consistent. For instance, I might make the as- sumption that, if there's someone at the door, the dog will bark. But it's not outside the realm of possibility that the dog is asleep, has gone out for a walk, or has been persuaded not to bark by a clever intruder with sedative-laced bacon. I make the assumption that if there's someone at the door, the dog will bark, because that is what *usually* happens.

Explaining Common Occurrences

We also use inductive reasoning to explain things that commonly happen. For instance, if I'm about to start an exam and notice that Bill is not here, I might explain this to myself with the reason that Bill is stuck in traffic. I might base this on the reasoning that being stuck in traffic is a common excuse for being late, or because I know that Bill never accounts for traffic when he's estimating how long it will take him to get somewhere. Again, that Bill is actually stuck in traffic is not certain, but I have some good reasons to think it's probable. We use this kind of reasoning to explain past events as well. For instance, if I read somewhere that 1986 was a particularly good year for tomatoes, I assume that 1986 also had some ideal combination of rainfall, sun, and consistently warm temperatures. Although it's possible that a scientific madman circled the globe planting tomatoes wherever he could in 1986, inductive reasoning would tell me that the former, environmental explanation is more likely. (But I could be wrong.)

Generalizing

Often we would like to make general claims, but in fact it would be very difficult to prove any general claim with any certainty. The only way to do so would be to observe every single case of something about which we wanted to make an observation. This would be, in fact, the only way to prove such assertions as, "All swans are white". Without being able to observe every single swan in the universe, I can never make that claim with certainty. Inductive logic, on the other hand, allows us to make the claim, with a certain amount of modesty.

3.7.1 Inductive Generalization

Inductive generalization allows us to make general claims, despite being unable to actually observe every single member of a class in order to make a certainly true general statement. We see this in scientific studies, population surveys, and in our own everyday reasoning. Take for example a drug study. Some doctor or other wants to know how many people will go blind if they take a certain amount of some drug for so many years. If they determine that 5% of people in the study go blind, they then assume that 5% of all people who take the drug for that many years will go blind. Likewise, if I survey a random group of people and ask them what their favourite colour is, and 75% of them say "purple", then I assume that purple is the favourite colour of 75% of people. But we have to be careful when we make an inductive generalization. When you tell me that 75% of people really like purple, I'm going to want to know whether you took that survey outside a Justin Bieber concert.

Let's take an example. Let's say I asked a class of 400 students whether or not they think logic is a valuable course, and 90% of them said yes. I can make an inductive argument like this:

> (P1) 90% of 400 students believe that logic is a valuable course.
> (C) Therefore 90% of students believe that logic is a valuable course.

There are certain things I need to take into account in judging the quality of this argument. For instance, did I ask this in a logic course? Did the respondents have to raise their hands so that the professor could see them, or was the survey taken anonymously? Are there enough students in the course to justify using them as a representative group for students in general?

If I did, in fact, make a class of 400 *logic* students raise their hands in response to the question of whether logic is valuable course, then we can identify a couple of problems with this argument. The first is **bias**. We can assume that anyone enrolled in a logic course is more likely to see it as valuable than any random student. I have therefore skewed the argument in favour of logic courses. I can also question whether the students were answering the question honestly. Perhaps if they are trying to save the professor's feelings, they are more likely to raise their hands and assure her that the logic course is a valuable one.

Now let's say I've avoided those problems. I have assured that the 400 students I have asked are randomly selected, say, by soliciting email responses from randomly selected students from the university's entire student population. Then the argument looks stronger.

Another problem we might have with the argument is whether I have asked *enough* students so that the whole population is well-represented. If the student body as a whole consists of 400 students, my argument is very strong. If the student body numbers in the tens of thousands, I might want to ask a few more before assuming that the opinions of a few mirror those of the many. This would be a problem with my **sample size**.

Let's take another example. Now I'm going to run a scientific study, in which I will pay someone $50 to take a drug with unknown effects and see if it makes them blind. In order to control for other variables, I open the study only to white males between the ages of 18 and 25.

A bad inductive argument would say:

(P1) 40% of 1000 people who took the drug went blind.
(C) Therefore 40% of people who take the drug will go blind.

A better inductive argument would make a more modest claim:

(P1) 40% of the 1000 people who took the drug went blind.
(C) Therefore 40% of white males between the ages of 18 and 25 who take the drug will go blind.

The point behind this example is to show how inductive reasoning imposes an important limitation on the possible conclusions a study or a survey can make. In order to make good generalizations, we need to ensure that our sample is **representative**, **non-biased**, and **sufficiently sized**.

3.7.2 Statistical Syllogism

Where in an inductive generalization we saw statement expressing a statistic applied to a more general group, we can also use statistics to go from the general to the particular. For instance, if I know that most computer science majors are male, and that some random individual with the androgynous name "Cameron" is an computer science major, then we can be reasonably certain that Cameron is a male. We tend to represent the uncertainty by qualifying the conclusion with the word "probably". If, on the other hand, we wanted to say that something is unlikely, like that Cameron were a female, we could use "probably not". It is also possible to temper our conclusion with other similar qualifying words.

Let's take an example.

(P1) Of the 133 people found guilty of homicide last year in Canada, 79% were jailed.
(P2) Socrates was found guilty of homicide last year in Canada.
(C) Therefore, Socrates was probably jailed.

In this case we can be reasonably sure that Socrates is currently rotting in prison. Now the certainty of our conclusion seems to be dependent on the statistics we're dealing with. There are definitely more certain and more uncertain cases.

(P1) In the last election, 50% of voting Americans voted for Obama, while 48% voted for Romney.
(P2) Jim is a voting American.
(C) Therefore, Jim probably voted for Obama.

Clearly, this argument is not as strong as the first. It is only slightly more likely than not that Jim voted for Obama. In this case we might want to revise our conclusion to say:

(C) Therefore, it is slightly more likely than not that Jim voted for Obama.

In other cases, the likelihood that something is or is not the case approaches certainty. For example:

(P1) There is a 0.00000059% chance you will die on any

single flight, assuming you use one of the most poorly rated airlines.
(P2) I'm flying to Paris next week.
(C) There's more than a million to one chance that I will die on my flight.

Note that in all of these examples, nothing is ever stated with absolute certainty. It is possible to improve the chances that our conclusions will be accurate by being more specific, or finding out more information. We would know more about Jim's voting strategy, for instance, if we knew where he lived, his previous voting habits, or if we simply asked him for whom he voted (in which case, we might also want to know how often Jim lies).

3.7.3 Induction by Shared Properties

Induction by shared properties involves noting the similarity between two things with respect to their properties, and inferring from this that they may share other properties.

A familiar example of this is how a company might recommend products to you based on other customers' purchases. Amazon.com tells me, for instance, that customers who bought the complete Sex and the City DVD series also bought Lipstick Jungle and Twilight.

Assuming that people buy things because they like them, we can rephrase this as:

(P1) There are a large number of people who, if they like Sex and the City and Twilight, will also like Lipstick Jungle.

I could also make the following observation:

(P2) I like Sex and the City and Twilight.

And then infer from there two premises that:

(C) I would also like Lipstick Jungle.

And I did. In general, induction by shared properties assumes that if something has properties w, x, y, and z, and if something else has properties w, x, and y, then it's reasonable to assume that that something else also has property z. Note that in the above example all of the properties were actually preferences with regard to entertainment. The kinds of properties involved in the comparison can and will make an argument better or worse. Let's consider a worse induction.

(P1) Lisa is tall, has blonde hair, has blue eyes, and rocks out to Nirvana on weekends.
(P2) Gina is tall, has blonde hair, and has blue eyes.
(C) Therefore Gina probably rocks out to Nirvana on weekends.

In this case the properties don't seem to be related in the same way as in the first example. While the first three are physical characteristics, the last property instead indicates to us that Lisa is stuck in a 90's grunge phase. Gina, though she shares several properties with Lisa, might not share the same undying love for Kurt Cobain. Let's try a stronger argument.

(P1) Bob and Dick both wear plaid shirts all the time, wear large plastic-rimmed glasses, and listen to bands you've never heard of.
(P2) Bob drinks PBR.
(C) Dick probably also drinks PBR.

Here we can identify the qualities that Bob and Dick have in common as symptoms of hipsterism. The fact that Bob drinks PBR is another symptom of this affectation. Given that Dick is exhibiting most of the same symptoms, the idea that Dick would also drink PBR is a reasonable assumption to make.

Practical Uses

A procedure very much like Induction by Shared Properties is performed by nurses and doctors when they diagnose a patient's condition. Their thinking goes like this:

(P1) Patients who have elephantitus display an increased

heart rate, elevated blood pressure, a rash on their skin, and a strong desire to visit the elephant pen at the zoo.
(P2) The patient here in front of me has an increased heart rate, elevated blood pressure, and a strong desire to visit the elephant pen at the zoo.
(C) It is probable, therefore, that the patient here in front of me has elephantitus.

The more that a patient's symptoms match the 'textbook definition' of a given disease, then the more likely it is that the patient has that disease. Caregivers then treat the patient for the disease that they think the patient probably has. If the disease doesn't respond to the treatment, or the patient starts to present different symptoms, then they consider other conditions with similar symptoms that the patient is likely to have.

3.7.4 Induction by Shared Relations

Induction by shared relations is much like induction by shared properties, except insofar that what is shared are not properties, but relations. A simple example is the causal relation, from which we might make an inductive argument like this:

(P1) Percocet, Oxycontin and Morphine reduce pain, cause drowsiness, and may be habit forming.
(P2) Heroin also reduces pain and causes drowsiness.
(C) Heroin is probably also habit forming.

In this case the effects of reducing pain, drowsiness, and addiction are all assumed to be caused by the drugs listed. We can use an induction by shared relation to make the probable conclusion that if heroin, like the other drugs, reduces pain and causes drowsiness, it is probably also habit forming.

Another interesting example are the relations we have with other people. For instance, Facebook knows everything about you. But let's focus on the "friends with" relation. They compare who your friends are with the friends of your friends in order to determine who else you might actually know. The induction goes a little like this:

(P1) Donna is friends with Brandon, Kelly, Steve, and Brenda.
(P2) David is friends with Brandon, Kelly, and Steve.
(C) David probably also knows Brenda.

We could strengthen that argument if we knew that Brandon, Kelly, Steve, and Brenda were all friends with each other as well. We could also make an alternate conclusion based on the same argument above:

(C) David probably also knows Donna.

They do, after all, know at least three of the same people. They've probably run into each other at some point.

3.8 Scientific Method

The procedure that scientists use is also a standard form of argument. Part of it is inductive, and so like other inductions, its conclusions only give you the likelihood or the probability that something is true, and not the certainty that it's true. But when it is done correctly, the conclusions it reaches are very well grounded in experimental evidence. Another part of it is deductive; and like other deductions, it gives you certain knowledge - but it gives you certainty about what's false, not what's true! These two parts have to be put together in a particular way. Here's a rough outline of how the procedure works.

Observation: Something is observed in the world which invokes your curiosity.
Theory: An idea is proposed which could explain why the thing which you observed happened, or why it is what it is. This is the part of the procedure where scientists can get quite creative and imaginative.
Prediction: A test is planned which could prove or disprove the theory. As part of the plan, the scientist will offer a proposition in this form: "If my theory is true, then the experiment will have [whatever] result."
Experiment: The test is performed, and the results are recorded.
5(a). Successful Result: If the prediction you

made at stage 3 came true, then the theory devised at step 2 is strengthened. This part of scientific method is inductive, and not deductive. And then we go back to step 3 to make more predictions and do more and more tests, to see if the theory can get stronger yet.

5(b). Failed Result: If the prediction did not come true, then the theory is falsified. This part of scientific method is deductive: scientists can't always be certain about what's true but they can be absolutely certain about what's false. When our predictions fail, we go back to step 2 and devise a new theory to put to the test,
and a new prediction to go with it.

Actually, a failed experimental result is really a kind of success, because falsification rules out the impossible. And that frees up the scientist to pursue other, more promising theories.

Scientists often test more than one theory at the same time, so that they can eventually arrive at the "last theory standing." In this way, scientists can use a form of disjunctive syllogism (see 3.6.6 above) to arrive at definitive conclusions about what theory is the best explanation for the observation. Here's how that part of the procedure works.

(P1) Either Theory 1 is true, or Theory 2 is true, or Theory 3 is true, or Theory 4 is true. (And so on, for however many theories are being tested.)
(P2) By experimental observation, Theories 1 and 2 and 3 were falsified.
(C) Therefore, Theory 4 is true.

Or, at least, Theory 4 is strengthened to the point where it would be quite absurd to believe anything else. After all, there might be other theories that we haven't thought of, or tested yet. But until we think of them, and test them, we're going to go with the best theory we've got.

There's a bit more to scientific method than this. There are paradigms and paradigm shifts, epistemic values, experimental controls and variables, and the various ways that scientists negotiate with each other as they interpret experimental results. There are also a few differences between the experimental methods used by physical scientists (such as chemists), and social scientists (such as anthropologists). But these things will be discussed in the expanded edition of this textbook.

Scientific method is the most powerful and suc- cessful form of knowing ever devised. Every advance in engineering, medicine, and technology has been made possible by people applying science to their problems. It is adventurous, curious, rigorously logical, and inspirational – it is even possible to be artistic about scientific discoveries. And the best part about science is that anyone can do it. Science can look difficult because there's a lot of jargon involved, and a lot of math. But even the most complicated quantum physics and the most far-reaching astronomy follows the same method, in principle, as that primary school project in which you played with magnets or built a model volcano.

3.7 Exercises for Chapter Three

1. Identify which of the following statements are propositions:

(a) Tea time is at 2pm.
(b) Why don't you love me anymore?
(c) Please keep off the grass.
(d) There's something wrong with kids today.
(e) Thou shalt not kill.
(f) Those 6 swans are looking at me funny.
(g) Some people have trouble with propositions.
(h) Can you pass the salt?
(i) There's a hole in my bucket.
(j) Could you be any more ridiculous?
(k) 67% of statistics are made up on the spot.
(l) Don't you dare kick that puppy.
(m) Puppy kickers are evil.
(n) This cat is my white whale.
(o) My feet hurt.
(p) There will be a sea battle tomorrow.
(q) Parades are stupid.
(r) You should probably not kidnap children.
(s) Kidnapping is illegal.
(t) Don't go into that barn.
(u) Fa la la la la, la la la la.

2. Identify the following statements as a *simple statement*, *negation*, *conjunction*, *disjunction*, *conditional*, or *biconditional*.

(a) Lois is awesome.
(b) If you don't eat your meat, you can't have any pudding.
(c) You can go to the party if and only if your homework is done.
(d) You said you would give me a pony, but you didn't.
(e) Either you're going to the dentist, or I'll rip that tooth out myself.
(f) I'm a wussy little girl.
(g) "Hoser" is not an acceptable Scrabble word.
(h) Your professor is dreamy, and also so smart.
(i) If he kisses the puppy, he'll get the votes; and if he doesn't, he won't.
(j) Having a computer is necessary if you want to Skype with your grandmother.
(k) Happy faces are so 90's.
(l) Either you're going to eat this candy, or I will.
(m) I keyed your car, and I boil bunnies.
(n) You're not special.
(o) He didn't know what he was doing.
(p) If you hear sirens, you're supposed to pull over.
(q) You're going to work today, or you're not getting paid.
(r) I have a test tomorrow, and my paper is due.

3. Identify the form of the following deductive arguments. (Modus Ponens, Modus Tollens, Hypothetical Syllogism, Categorical Syllogism, Disjunctive Syllogism, Adjunction, Constructive Dilemma, or Destructive Dilemma)

(a) If you don't have a pencil, you can't write the exam. You don't have a pencil. So you can't write the exam.
(b) If you buy the farm, you can get kittens. If you buy a boat, you can go sailing. You're either going to buy the farm, or buy a boat. Therefore you can either have kittens or go sailing.
(c) If Lois has a bicycle, she also has a bicycle helmet. If Lois has a bicycle helmet, her hair will be flat. Therefore, if Lois has a bicycle, her hair will be flat.
(d) If you robbed that store, you would be found guilty. You were not found guilty. Therefore, you didn't rob that store.
(e) Kittens are either cute, or kittens are ugly. Kittens are not ugly. Therefore kittens are cute.
(f) I have two buttons missing. I have a tail. Therefore I have two buttons missing and I have a tail.
(g) All good muffins have chocolate chips. This is a good muffin. Therefore this muffin has chocolate chips.

4. Supply the conclusion that results from the following premises:

(a) P1: All monkeys like bananas.
P2: George is a monkey.
(b) P1: If this cupcake is less than a week old, George will eat it.
P2: George will not eat that cupcake.
(c) P1: Either you're lying to me, or I'm stupid.
P2: I'm not stupid.
(d) P1: If there's a monkey in the room, you can smell bananas.
P2: If there's a cake in the room, you can smell cake.
P3: There's either a monkey in the room, or some cake.
(e) P1: If you want to get ahead in life, you have to know your argument forms.
P2: You want to get ahead in life.
(f) P1: If you have a boat, people call you "Captain".
P2: If people call you "Captain", you get a lot of street cred.

5. Identify a problem with the following inductive arguments.

(a) P1: 79% of men who take drugs prefer cocaine.
P2: Princess Peach takes drugs.
C: Therefore Princess Peach prefers cocaine.
(b) P1: 60% of people who shop at Mountain Equip ment Co-Op like mountain climbing.
C: Therefore 60% of people like mountain climbing.
(c) P1: 100% of the people I asked said their name was Joe Brown.

C: Therefore 100% of people are named Joe Brown.

6. Identify these arguments as either: *inductive general- ization*, *statistical syllogism*, *induction by shared properties*, or *induction by shared relations*.

(a) P1: Of the 10% of the population surveyed, most said they support the "kittens for all" movement. C: Therefore most people support the "kittens for all" movement.

(b) P1: Kant's *Critique of Pure Reason* is a heavy book, densely worded, has a boring cover and if you read it in a coffee shop, people think you're cool.
P2: Heidegger's *Being and Time* is a heavy book, densely worded, and has a boring cover.
C: Reading Heidegger's *Being and Time* in a coffee shop will make people think you're cool.

(c) P1: 67% of people who attend university never have the opportunity to commit armed robbery.
P2: Bob went to university.
C: Therefore, Bob has probably never committed an armed robbery.