


Introduction to Logic and Critical Thinking

Version 2.0



Matthew J. Van Cleave

vancleave@mac.com



Introduction to Logic and Critical Thinking by Matthew J. Van Cleave is licensed under a Creative Commons Attribution 4.0 International License. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

Table of contents

Preface

Chapter 1: Reconstructing and analyzing arguments

- 1.1 What is an argument?
- 1.2 Identifying arguments
- 1.3 Arguments vs. explanations
- 1.4 More complex argument structures
- 1.5 Using your own paraphrases of premises and conclusions to reconstruct arguments in standard form
- 1.6 Validity
- 1.7 Soundness
- 1.8 Deductive vs. inductive arguments
- 1.9 Arguments with missing premises
- 1.10 Assuring, guarding, and discounting
- 1.11 Evaluative language
- 1.12 Evaluating a real-life argument

Chapter 2: Formal methods of evaluating arguments

- 2.1 What is a formal method of evaluation and why do we need them?
- 2.2 Propositional logic and the four basic truth functional connectives
- 2.3 Negation and disjunction
- 2.4 Using parentheses to translate complex sentences
- 2.5 "Not both" and "neither nor"
- 2.6 The truth table test of validity
- 2.7 Conditionals
- 2.8 "Unless"
- 2.9 Material equivalence
- 2.10 Tautologies, contradictions, and contingent statements
- 2.11 Proofs and the 8 valid forms of inference
- 2.12 How to construct proofs
- 2.13 Short review of propositional logic
- 2.14 Categorical logic
- 2.15 The Venn test of validity for immediate categorical inferences
- 2.16 Universal statements and existential commitment
- 2.17 Venn validity for categorical syllogisms

Chapter 3: Evaluating inductive arguments and probabilistic and statistical fallacies

- 3.1 Inductive arguments and statistical generalizations
- 3.2 Inference to the best explanation and the seven explanatory virtues

- 3.3 Analogical arguments
 - 3.4 Causal arguments
 - 3.5 Probability
 - 3.6 The conjunction fallacy
 - 3.7 The base rate fallacy
 - 3.8 The small numbers fallacy
 - 3.9 Regression to the mean fallacy
 - 3.10 Gambler's fallacy
- Chapter 4: Informal fallacies
- 4.1 Formal vs. informal fallacies
 - 4.1.1 Composition fallacy
 - 4.1.2 Division fallacy
 - 4.1.3 Begging the question fallacy
 - 4.1.4 False dichotomy
 - 4.1.5 Equivocation
 - 4.2 Slippery slope fallacies
 - 4.2.1 Conceptual slippery slope
 - 4.2.2 Causal slippery slope
 - 4.3 Fallacies of relevance
 - 4.3.1 Ad hominem
 - 4.3.2 Straw man
 - 4.3.3 Tu quoque
 - 4.3.4 Genetic
 - 4.3.5 Appeal to consequences
 - 4.3.6 Appeal to authority

Answers to exercises

Glossary/Index

Preface

This is an introductory textbook in logic and critical thinking. The goal of the textbook is to provide the reader with a set of tools and skills that will enable them to identify and evaluate arguments. The book is intended for an introductory course that covers both formal and informal logic. As such, it is not a formal logic textbook, but is closer to what one would find marketed as a “critical thinking textbook.” The formal logic in chapter 2 is intended to give an elementary introduction to formal logic. Specifically, chapter 2 introduces several different formal methods for determining whether an argument is valid or invalid (truth tables, proofs, Venn diagrams). I contrast these formal methods with the informal method of determining validity introduced in chapter 1. What I take to be the central *theoretical* lesson with respect to the formal logic is simply that of understanding the difference between formal and informal methods of evaluating an argument’s validity. I believe there are also *practical* benefits of learning the formal logic. First and foremost, once one has internalized some of the valid forms of argument, it is easy to impose these structures on arguments one encounters. The ability to do this can be of use in evaluating an argumentative passage, especially when the argument concerns a topic with which one is not very familiar (such as on the GRE or LSAT).

However, what I take to be of far more practical importance is the skill of being able to reconstruct and evaluate arguments. This skill is addressed in chapter 1, where the central ideas are that of using the *principle of charity* to put arguments into *standard form* and of using the *informal test of validity* to evaluate those arguments. Since the ability to reconstruct and evaluate arguments is a skill, one must practice in order to acquire it. The exercises in each section are intended to give students some practice, but in order to really master the skill, one must practice much, much more than simply completing the exercises in the text. It makes about as much sense to say that one could become a critical thinker by reading a critical thinking textbook as that one could become fluent in French by reading a French textbook. Logic and critical thinking, like learning a foreign language, takes practice because it is a *skill*.

While chapters 1 and 2 mainly concern deductive arguments, chapter 3 addresses inductive arguments, including probabilistic and statistical fallacies. In a world in which information is commonly couched within probabilistic and statistical frameworks, understanding these basic concepts, as well as some of the common mistakes is essential for understanding our world. I have tried to

write chapter 3 with an eye towards this understanding. As with all the chapters, I have tried to walk the fine line between being succinct without sacrificing depth.

Chapter 4 picks out what I take to be some of the most common fallacies, both formal and informal. In my experience, many critical thinking textbooks end up making the fallacies sound obvious; one is often left wondering how anyone could commit such a fallacy. In my discussion of the fallacies I have tried to correct this not only in the particular examples I use in the text and exercises, but also by discussing what makes a particular fallacy seductive.

I have used numerous different textbooks over the years that I have been teaching logic and critical thinking courses. Some of them were very good; others were not. Although this textbook is my attempt to improve on what I've encountered, I am indebted to a number of textbooks that have shaped how I teach logic and critical thinking. In particular, Sinnott-Armstrong and Fogelin's *Understanding Arguments: An Introduction to Informal Logic* and Copi and Cohen's *Introduction to Logic* have influenced how I present the material here (although this may not be obvious). My interest in better motivating the seductiveness of the fallacies is influenced by Daniel Kahneman's work in psychology (for which he won the Nobel Prize in economics in 2002).

This textbook is an "open textbook" that is licensed under the Creative Commons Attribution 4.0 license (CC BY 4.0). Anyone can take this work and alter it for their own purposes as long as they give appropriate credit to me, noting whether or not you have altered the text. (If you would like to alter the text but have come across this textbook in PDF format, please do not hesitate to email me at vancleave@mac.com and I will send you the text in a file format that is easier to manipulate.) Many colleges and universities have undertaken initiatives to reduce the cost of textbooks. I see this as an issue of access to education and even an issue of justice. Some studies have shown that without access to the textbook, a student's performance in the class will suffer. Many students lack access to a textbook simply because they do not buy it in the first place since there are more pressing things to pay for (rent, food, child care, etc.) and because the cost of some textbooks is prohibitive. Moreover, both professors and students are beholden to publishers who profit from selling textbooks (professors, because the content of the course is set by the author of the textbook, and perhaps market forces, rather than by the professor herself; students, because they must buy the newest editions of increasingly expensive

Preface

textbooks). If education is necessary for securing certain basic human rights (as philosophers like Martha Nussbaum have argued), then lack of access to education is itself an issue of *justice*. Providing high quality, low-cost textbooks is one, small part of making higher education more affordable and thus more equitable and just. This open textbook is a contribution towards that end.

Matthew J. Van Cleave

January 4, 2016

1.1 What is an argument?

This is an introductory textbook in logic and critical thinking. Both logic and critical thinking centrally involve the analysis and assessment of arguments. “Argument” is a word that has multiple distinct meanings, so it is important to be clear from the start about the sense of the word that is relevant to the study of logic. In one sense of the word, an argument is a heated exchange of differing views as in the following:

Sally: Abortion is morally wrong and those who think otherwise are seeking to justify murder!

Bob: Abortion is not morally wrong and those who think so are right-wing bigots who are seeking to impose their narrow-minded views on all the rest of us!

Sally and Bob are *having an argument* in this exchange. That is, they are each expressing conflicting views in a heated manner. However, that is not the sense of “argument” with which logic is concerned. Logic concerns a different sense of the word “argument.” An argument, in this sense, is *a reason for thinking that a statement, claim or idea is true*. For example:

Sally: Abortion is morally wrong because it is wrong to take the life of an innocent human being, and a fetus is an innocent human being.

In this example Sally has given an argument against the moral permissibility of abortion. That is, she has given us a reason for thinking that abortion is morally wrong. The **conclusion** of the argument is the first four words, “abortion is morally wrong.” But whereas in the first example Sally was simply asserting *that* abortion is wrong (and then trying to put down those who support it), in this example she is offering a reason for *why* abortion is wrong.

We can (and should) be more precise about our definition of an argument. But before we can do that, we need to introduce some further terminology that we will use in our definition. As I’ve already noted, the conclusion of Sally’s argument is that abortion is morally wrong. But the reason for thinking the conclusion is true is what we call the **premise**. So we have two parts of an argument: the premise and the conclusion. Typically, a conclusion will be supported by two or more premises. Both premises and conclusions are statements. A **statement** is a type of sentence that can be true or false and

corresponds to the grammatical category of a “declarative sentence.” For example, the sentence,

The Nile is a river in northeastern Africa

is a statement. Why? Because it makes sense to inquire whether it is true or false. (In this case, it happens to be true.) But a sentence is still a statement even if it is false. For example, the sentence,

The Yangtze is a river in Japan

is still a statement; it is just a false statement (the Yangtze River is in China). In contrast, none of the following sentences are statements:

Please help yourself to more casserole

Don't tell your mother about the surprise

Do you like Vietnamese pho?

The reason that none of these sentences are statements is that it doesn't make sense to ask whether those sentences are true or false (rather, they are requests or commands, and questions, respectively).

So, to reiterate: all arguments are composed of premises and conclusions, which are both types of statements. The premises of the argument provide a reason for thinking that the conclusion is true. And arguments typically involve more than one premise. A standard way of capturing the structure of an argument is by numbering the premises and conclusion. For example, recall Sally's argument against abortion:

Abortion is morally wrong because it is wrong to take the life of an innocent human being, and a fetus is an innocent human being.

We could capture the structure of that argument like this:

1. It is morally wrong to take the life of an innocent human being
2. A fetus is an innocent human being
3. Therefore, abortion is morally wrong

By convention, the last numbered statement (also denoted by the “therefore”) is the conclusion and the earlier numbered statements are the premises. This is what we will call **standard argument form**. We can now give a more precise definition of an argument. An **argument** is a set of statements, some of which (the premises) attempt to provide a reason for thinking that some other statement (the conclusion) is true. Although arguments are typically given in order to convince or persuade someone of the conclusion, the argument itself is independent of one’s attempt to use it to convince or persuade. For example, I have just given you this argument not in an attempt to convince you that abortion is morally wrong, but as an illustration of what an argument is. Later on in this chapter and in this book we will learn some techniques of *evaluating* arguments, but for now the goal is to learn to *identify* an argument, including its premises and conclusion(s). It is important to be able to identify arguments and understand their structure, whether or not you agree with conclusion of the argument. In the next section I will provide some techniques for being able to identify arguments.

Exercise 1: Which of the following sentences are statements and which are not?

1. No one understands me but you.
2. Alligators are on average larger than crocodiles.
3. Is an alligator a reptile or a mammal?
4. An alligator is either a reptile or a mammal.
5. Don’t let any reptiles into the house.
6. You may kill any reptile you see in the house.
7. East Africans are not the best distance runners.
8. Obama is not a Democrat.
9. Some humans have wings.
10. Some things with wings cannot fly.
11. Was Obama born in Kenya or Hawaii?
12. Oh no! A grizzly bear!
13. Meet me in St. Louis.
14. We met in St. Louis yesterday.
15. I do not want to meet a grizzly bear in the wild.

1.2 Identifying arguments

The best way to identify whether an argument is present is to ask whether there is a statement that someone is trying to establish as true by basing it on some other statement. If so, then there is an argument present. If not, then there isn't. Another thing that can help in identifying arguments is knowing certain key words or phrases that are premise indicators or conclusion indicators. For example, recall Sally's abortion argument:

Abortion is morally wrong *because* it is wrong to take the life of an innocent human being, and a fetus is an innocent human being.

The word "because" here is a premise indicator. That is, "because" indicates that what follows is a reason for thinking that abortion is morally wrong. Here is another example:

I know that the student plagiarized *since* I found the exact same sentences on a website and the website was published more than a year before the student wrote the paper.

In this example, the word "since" is a **premise indicator** because what follows it is a statement that is clearly intended to be a reason for thinking that the student plagiarized (i.e., a premise). Notice that in these two cases, the premise indicators "because" and "since" are interchangeable: I could have used "because" in place of "since" or "since" in the place of "because" and the meaning of the sentences would have been the same. In addition to premise indicators, there are also conclusion indicators. Conclusion indicators mark that what follows is the conclusion of an argument. For example,

Bob-the-arsonist has been dead for a year, *so* Bob-the-arsonist didn't set the fire at the East Lansing Starbucks last week.

In this example, the word "so" is a **conclusion indicator** because what follows it is a statement that someone is trying to establish as true (i.e., a conclusion). Here is another example of a conclusion indicator:

A poll administered by Gallup (a respected polling company) showed candidate x to be substantially behind candidate y with only a week left before the vote, *therefore* candidate y will probably not win the election.

In this example, the word “therefore” is a conclusion indicator because what follows it is a statement that someone is trying to establish as true (i.e., a conclusion). As before, in both of these cases the conclusion indicators “so” and “therefore” are interchangeable: I could have used “so” in place of “therefore” or “therefore” in the place of “so” and the meaning of the sentences would have been the same.

Table 1 contains a list of some common premise and conclusion indicators:

Premise indicators	Conclusion indicators
since	therefore
because	so
for	hence
as	thus
given that	implies that
seeing that	consequently
for the reason that	it follows that
is shown by the fact that	we may conclude that

Although these words and phrases can be used to identify the premises and conclusions of arguments, they are not failsafe methods of doing so. Just because a sentence contains them does not mean that you are dealing with an argument. This can easily be shown by examples like these:

I have been running competitively *since* 1999.

I am so happy to have finally finished that class.

Although “since” can function as a premise indicator and although “so” can function as a conclusion indicator, neither one is doing so here. This shows that you can’t simply mindlessly use occurrences of these words in sentences to show that there is an argument being made. Rather, we have to rely on our understanding of the English sentence in order to determine whether an argument is being made or not. Thus, the best way to determine whether an argument is present is by asking the question: Is there a statement that someone is trying to establish as true or explain why it is true by basing it on some other statement? If so, then there is an argument present. If not, then there isn’t. Notice that if we apply this method to the above examples, we will

see that there is no argument present because there is no statement that someone is trying to establish as true *by basing it on some other statement*. For example, the sentence “I have been running competitively since 1999” just contains one statement, not two. But arguments always require at least two separate statements—one premise and one conclusion, so it cannot possibly be an argument.

Another way of explaining why these occurrences of “so” and “since” do not indicate that an argument is present is by noting that both premise indicators and conclusion indicators are, grammatically, conjunctions. A grammatical conjunction is a word that connects two separate statements. So, if a word or term is truly being used as a premise or conclusion indicator, it must connect two separate statements. Thus, if “since” were really functioning as a premise indicator in the above example then what followed it would be a statement. But “1999” is not a statement at all. Likewise, in the second example “so” is not being used as a conclusion indicator because it is not conjoining two separate statements. Rather, it is being used to modify the extent of “happy.” In contrast, if I were to say “Tom was sleeping, so he couldn’t have answered the phone,” then “so” is being used as a conclusion indicator. In this case, there are clearly two separate statements (“Tom was sleeping” and “Tom couldn’t have answered the phone”) and one is being used as the basis for thinking that the other is true.

If there is any doubt about whether a word is truly a premise/conclusion indicator or not, you can use the **substitution test**. Simply substitute another word or phrase from the list of premise indicators or conclusion indicators and see if the resulting sentence still makes sense. If it does, then you are probably dealing with an argument. If it doesn’t, then you probably aren’t. For example, we can substitute “it follows that” for “so” in the Bob-the-arsonist example:

Bob-the-arsonist has been dead for a year, *it follows that* Bob-the-arsonist didn’t set the fire at the East Lansing Starbucks last week.

However, we cannot substitute “because” for “so” in the so-happy-I-finished-that-class example:

I am *because* happy to have finally finished that class.

Obviously, in the latter case the substitution of one conclusion indicator for another makes the sentence meaningless, which means that the “so” that occurred originally wasn’t functioning as a conclusion indicator.

Exercise 2: Which of the following are arguments? If it is an argument, identify the conclusion of the argument.

1. The woman in the hat is not a witch since witches have long noses and she doesn’t have a long nose.
2. I have been wrangling cattle since before you were old enough to tie your own shoes.
3. Albert is angry with me so he probably won’t be willing to help me wash the dishes.
4. First I washed the dishes and then I dried them.
5. If the road wasn’t icy, the car wouldn’t have slid off the turn.
6. Albert isn’t a fireman and he isn’t a fisherman either.
7. Are you seeing that rhinoceros over there? It is huge!
8. The fact that obesity has become a problem in the U.S. is shown by the fact that obesity rates have risen significantly over the past four decades.
9. Bob showed me a graph with the rising obesity rates and I was very surprised to see how much they’ve risen.
10. Albert isn’t a fireman because Albert is a Greyhound, which is a kind of dog, and dogs can’t be firemen.
11. Charlie and Violet are dogs and since dogs don’t sweat, it is obvious that Charlie and Violet don’t sweat.
12. The reason I forgot to lock the door is that I was distracted by the clown riding a unicycle down our street while singing Lynyrd Skynyrd’s “Simple Man.”
13. What Bob told you is not the real reason that he missed his plane to Denver.
14. Samsung stole some of Apple’s patents for their smartphones, so Apple stole some of Samsung’s patents back in retaliation.
15. No one who has ever gotten frostbite while climbing K2 has survived to tell about it, therefore no one ever will.

1.3 Arguments vs. explanations

So far I have defined arguments in terms of premises and conclusions, where the premises are supposed to provide a reason (support, evidence) for accepting the conclusion. Many times the goal of giving an argument is simply to establish *that* the conclusion is true. For example, when I am trying to convince someone that obesity rates are rising in the U.S. I may cite evidence such as studies from the Center for Disease Control (CDC) and the National Institute of Health (NIH). The studies I cite would function as premises for the conclusion that obesity rates are rising. For example:

We know that obesity is on the rise in the U.S. because multiple studies carried out by the CDC and NIH have consistently shown a rise in obesity over the last four decades.

We could put this simple argument into standard form like this:

1. Multiple studies by the CDC and NIH have consistently shown a rise in obesity over the last four decades.
2. Therefore, obesity is on the rise in the U.S.

The standard form argument clearly distinguishes the premise from the conclusion and shows how the conclusion is supposed to be supported by the evidence offered in the premise. Again, the goal of this simple argument would be to convince someone that the conclusion is *true*. However, sometimes we already know that a statement or claim is true and we are trying to establish *why* it is true rather than *that* it is true. An argument that attempts to show *why* its conclusion is true is an **explanation**. Contrast the previous example with the following:

The reason that the rate of obesity is on the rise in the U.S. is that the foods we most often consume over the past four decades have increasingly contained high levels of sugar and low levels of dietary fiber. Since eating foods high in sugar and low in fiber triggers the insulin system to start storing those calories as fat, it follows that people who consume foods high in sugar and low in fiber will tend to store more of the calories consumed as fat.

This passage gives an explanation for *why* obesity is on the rise in the U.S. Unlike the earlier example, here it is taken for granted that obesity *is* on the rise in the U.S. That is the claim whose truth we are trying to explain. We can put the obesity explanation into standard form just like any other argument. In order to do this, I will make some paraphrases of the premises and conclusion of the argument (for more on how to do this, see section 1.5 below).

1. Over the past four decades, Americans have increasingly consumed foods high in sugar and low in fiber.
2. Consuming foods high in sugar and low in fat triggers the insulin system to start storing those calories as fat.
3. When people store more calories as fat, they tend to become obese.
4. Therefore, the rate of obesity is on the rise in the U.S.

Notice that in this explanation the premises (1-3) attempt to give a reason for *why* the conclusion is true, rather than a reason for thinking *that* the conclusion is true. That is, in an explanation we assume that what we are trying to explain (i.e., the conclusion) is true. In this case, the premises are supposed to show *why* we should *expect* or *predict* that the conclusion is true. Explanations often give us an *understanding* of why the conclusion is true. We can think of explanations as a type of argument, we just have to distinguish two different types of argument: those that attempt to establish *that* their conclusion is true (arguments), and those that attempt to establish *why* their conclusion is true (explanations).

Exercise 3: Which of the following is an explanation and which is an argument? Identify the main conclusion of each argument or explanation. (Remember if the premise(s) seems to be establishing *that* the conclusion is true, it is an argument, but if the premise(s) seems to be establishing *why* the conclusion is true, it is an explanation.)

1. Wanda rode the bus today because her car was in the shop.
2. Since Wanda doesn't have enough money in her bank account, she has not yet picked up her car from the shop.
3. Either Bob or Henry rode the bus to work today. But it wasn't Henry because I saw him riding his bike to work. Therefore, it was Bob.
4. It can't be snowing right now since it only snows when it is 32 degrees or below and right now it is 40 degrees.

5. The reason some people with schizophrenia hear voices in their head is that the cognitive mechanism that monitors their own self-talk is malfunctioning and they attribute their own self-talk to some external source.
6. Fracking should be allowed because, although it does involve some environmental risk, it reduces our dependence on foreign oil and there is much greater harm to the environment due to foreign oil drilling than there is due to fracking.
7. Wanda could not have ridden the bus today because today is a city-wide holiday and the bus service is not operating.
8. The Tigers lost their star pitcher due to injury over the weekend, therefore the Tigers will not win their game against the Pirates.
9. No one living in Pompeii could have escaped before the lava from Mt. Vesuvius hit. The reason is simple: the lava was flowing too fast and there was nowhere to go to escape it in time.
10. The reason people's allergies worsen when they move to Cincinnati is that the pollen count in Cincinnati is higher than almost anywhere else in the surrounding area.

1.4 More complex argument structures

So far we have seen that an argument consists of a premise (typically more than one) and a conclusion. However, very often arguments and explanations have a more complex structure than just a few premises that directly support the conclusion. For example, consider the following argument:

No one living in Pompeii could have survived the eruption of Mt. Vesuvius. The reason is simple: the lava was flowing too fast and there was nowhere to go to escape it in time. Therefore, this account of the eruption, which claims to have been written by an eyewitness living in Pompeii, was not actually written by an eyewitness.

The **main conclusion** of this argument—the statement that depends on other statements as evidence but doesn't itself provide any evidence for any other statement—is:

- A. This account of the eruption of Mt. Vesuvius was not actually written by an eyewitness.

However, the argument's structure is more complex than simply having a couple of premises that provide evidence directly for the conclusion. Rather, some statement provides evidence directly for the main conclusion, but that statement itself is supported by another statement. To determine the structure of an argument, we must determine which statements support which. We can use our premise and conclusion indicators to help with this. For example, the passage contains the phrase, "the reason is..." which is a premise indicator, and it also contains the conclusion indicator, "therefore." That conclusion indicator helps us to identify the main conclusion, but the more important thing to see is that statement A does not itself provide evidence or support for any of the other statements in the argument, which is the clearest reason why statement A is the main conclusion of the argument. The next question we must answer is: which statement most directly supports A? What most directly supports A is:

- B. No one living in Pompeii could have survived the eruption of Mt. Vesuvius.

However, there is also a reason offered in support of B. That reason is that:

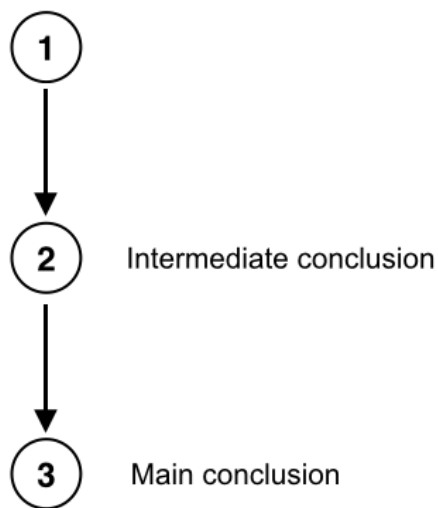
- C. The lava from Mt. Vesuvius was flowing too fast and there was nowhere for someone living in Pompeii to go in order to escape it in time.

So the main conclusion (A) is directly supported by B, and B is supported by C. Since B acts as a premise for the main conclusion but is also itself the conclusion of further premises, we refer to B as an **intermediate conclusion**. The important thing to recognize here is that one and the same statement can act as both a premise and a conclusion. Statement B is a premise that supports the main conclusion (A), but it is also itself a conclusion that follows from C. Here is how we would put this complex argument into standard form (using numbers this time, as we always do when putting an argument into standard form):

1. The lava from Mt. Vesuvius was flowing too fast and there was nowhere for someone living in Pompeii to go in order to escape it in time.
2. Therefore, no one living in Pompeii could have survived the eruption of Mt. Vesuvius. (from 1)

3. Therefore, this account of the eruption of Mt. Vesuvius was not actually written by an eyewitness. (from 2)

Notice that at the end of statement 2 I have written in parentheses “from 1” (and likewise at the end of statement 3 I have written “from 2”). This is a shorthand way of saying: “this statement follows from statement 1.” We will use this convention as a way of keeping track of the structure of the argument. It may also help to think about the structure of an argument spatially, as [figure 1](#) shows:



The main argument here (from 2 to 3) contains a subargument, in this case the argument from 1 to 2. In general, the **main argument** is simply the argument whose premises *directly* support the main conclusion, whereas a **subargument** is an argument that provides *indirect* support for the main conclusion by supporting one of the premises of the main argument. You can always add further subarguments to the overall structure of an argument by providing evidence that supports one of the unsupported premises.

Another type of structure that arguments can have is when two or more premises provide direct but independent support for the conclusion. Here is an example of an argument with that structure:

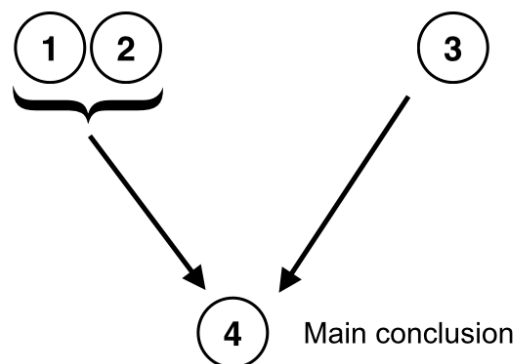
I know that Wanda rode her bike to work today because when she arrived at work she had her right pant leg rolled up (which cyclists do in order to

keep their pants legs from getting caught in the chain). Moreover, our coworker, Bob, who works in accounting, saw her riding towards work at 7:45 am.

The conclusion of this argument is “Wanda rode her bike to work today” and there are two premises that provide independent support for it: the fact that Wanda had her pant leg cuffed and the fact that Bob saw her riding her bike. Here is the argument in standard form:

1. Wanda arrived at work with her right pant leg rolled up.
2. Cyclists often roll up their right pant leg.
3. Bob saw Wanda riding her bike towards work at 7:45.
4. Therefore, Wanda rode her bike to work today. (from 1-2, 3 independently)

Again, notice that next to statement 4 of the argument I have written the premises from which that conclusion follows. In this case, in order to avoid any ambiguity, I have noted that the support for the conclusion comes independently from statements 1 and 2, on the one hand, and from statement 3, on the other hand. It is important to point out that an argument or subargument can be supported by one or more premises. We see this in the present argument since the conclusion (4) is supported jointly by 1 and 2, and singly by 3. As before, we can represent the structure of this argument spatially, as figure 2 shows:



There are endless different argument structures that can be generated from these few simple patterns. At this point, it is important to understand that arguments can have these different structures and that some arguments will be longer and more complex than others. Determining the structure of very

complex arguments is a skill that takes some time to master. Even so, it may help to remember that any argument structure ultimately traces back to some combination of these.

Exercise 4: Write the following arguments in standard form and show how the argument is structured using a diagram like the ones I have used in this section.

1. There is nothing wrong with prostitution because there is nothing wrong with consensual sexual and economic interactions between adults. Moreover, since there's no difference between a man who goes on a blind date with a woman, buys her dinner and then has sex with her and a man who simply pays a woman for sex, that is another reason for why there is nothing wrong with prostitution.
2. Prostitution is wrong because it involves women who have typically been sexually abused as children. We know that most of these women have been abused from multiple surveys done with women who have worked in prostitution and that show a high percentage of self-reported sexual abuse as children.
3. There was someone in this cabin recently because there was warm water in the tea kettle and because there was wood still smoldering in the fireplace. But the person couldn't have been Tim because Tim has been with me the whole time. Therefore, there must be someone else in these woods.
4. It is possible to be blind and yet run in the Olympic Games since Marla Runyan did it at the 2000 Sydney Olympics.
5. The train was late because it had to take a longer, alternate route since the bridge was out.
6. Israel is not safe if Iran gets nuclear missiles since Iran has threatened multiple times to destroy Israel and if Iran had nuclear missiles it would be able to carry out this threat. Moreover, since Iran has been developing enriched uranium, they have the key component needed for nuclear weapons—every other part of the process of building a nuclear weapon is simple compared to that. Therefore, Israel is not safe.
7. Since all professional hockey players are missing front teeth and Martin is a professional hockey player, it follows that Martin is missing front teeth. And since almost all professional athletes who are missing their front teeth have false teeth, it follows that Martin probably has false teeth.

8. Anyone who eats the crab rangoon at China Food restaurant will probably have stomach troubles afterward. It has happened to me every time, which is why it will probably happen to other people as well. Since Bob ate the crab rangoon at China Food restaurant, he will probably have stomach troubles afterward.
9. Albert and Caroline like to go for runs in the afternoon in Hyde Park. Since Albert never runs alone, we know that any time Albert is running, Caroline is running too. But since Albert looks like he has just run (since he is panting hard), it follows that Caroline must have ran too.
10. Just because Jeremy's prints were on the gun that killed Tim and the gun was registered to Jeremy, it doesn't follow that Jeremy killed Tim since Jeremy's prints would certainly be on his own gun and someone else could have stolen Jeremy's gun and used it to kill Tim.

1.5 Using your own paraphrases of premises and conclusions to reconstruct arguments in standard form

Although sometimes we can just lift the premises and conclusion verbatim from the argument, we cannot always do this. Paraphrases of premises or conclusions are sometimes needed in order to make the standard form argument as clear as possible. A **paraphrase** is the use of different words to capture the same idea in a clearer way. There will always be multiple ways of paraphrasing premises and conclusions and this means that there will never be just one way of putting an argument into standard form. In order to paraphrase well, you will have to rely on your understanding of English to come up with what you think is the best way of capturing the essence of the argument. Again, typically there is no *single* right way to do this, although there are certainly better and worse ways of doing it. For example, consider the following argument:

Just because Jeremy's prints were on the gun that killed Tim and the gun was registered to Jeremy, it doesn't follow that Jeremy killed Tim since Jeremy's prints would certainly be on his own gun and someone else could have stolen Jeremy's gun and used it to kill Tim.

What is the conclusion of this argument? (Think about it before reading on.) Here is one way of paraphrasing the conclusion:

The fact that Jeremy's prints were on the gun that killed Tim and the gun was registered to Jeremy doesn't mean that Jeremy killed Tim.

This statement seems to capture the essence of the main conclusion in the above argument. The premises of the argument would be:

1. Jeremy's prints would be expected to be on a gun that was registered to him
2. Someone could have stolen Jeremy's gun and then used it to kill Tim

Notice that while I have paraphrased the first premise, I have left the second premise almost exactly as it appeared in the original paragraph. As I've said, paraphrases are needed in order to try to make the standard form argument as clear as possible and this is what I've tried to do in capturing premise 1 as well as the conclusion of this argument. So here is the reconstructed argument in standard form:

1. Jeremy's prints would be expected to be on a gun that was registered to him
2. Someone could have stolen Jeremy's gun and then used it to kill Tim
3. Therefore, the fact that Jeremy's prints were on the gun that killed Tim and the gun was registered to Jeremy doesn't mean that Jeremy killed Tim. (from 1-2)

However, as I have just noted, there is more than one way of paraphrasing the premises and conclusion of the argument. To illustrate this, I will give a second way that one could accurately capture this argument in standard form. Here is another way of expressing the conclusion:

We do not know that Jeremy killed Tim.

That is clearly what the above argument is trying to ultimately establish and it is a much simpler (in some ways) conclusion than my first way of paraphrasing the conclusion. However, it also takes more liberties in interpreting the argument than my original paraphrase. For example, in the original argument there is no occurrence of the word "know." That is something that I am introducing in my own paraphrase. That is a totally legitimate thing to do, as long as introducing new terminology helps us to clearly express the essence of the premise or

conclusion that we're trying to paraphrase.¹ Since my second paraphrase of the conclusion differs from my first paraphrase, you can expect that my premises will differ also. So how shall I paraphrase the premises that support this conclusion? Here is another way of paraphrasing the premises and putting the argument into standard form:

1. Tim was killed by a gun that was registered to Jeremy and had Jeremy's prints on it.
2. It is possible that Jeremy's gun was stolen from him.
3. If Jeremy's gun was stolen from him, then Jeremy could not have killed Tim.
4. Therefore, we do not know that Jeremy killed Tim. (from 1-3)

Notice that this standard form argument has more premises than my first reconstruction of the standard form argument (which consisted of only three statements). I have taken quite a few liberties in interpreting and paraphrasing this argument, but what I have tried to do is to get down to the most essential logic of the original argument. The paraphrases of the premises I have used are quite different from the wording that occurs in the original paragraph. I have introduced phrases such as "it is possible that" as well as **conditional statements** (if...then statements), such as premise 3. Nonetheless, this reconstruction seems to get at the essence of the logic of the original argument. As long as your paraphrases help you to do that, they are good paraphrases. Being able to reconstruct arguments like this takes many years of practice in order to do it well, and much of the material that we will learn later in the text will help you to better understand how to capture an argument in standard form, but for now it is important to recognize that there is never only one way of correctly capturing the standard form of an argument. And the reason for this is that there are multiple, equally good, ways of paraphrasing the premises and conclusion of an argument.

¹ How do we know that a paraphrase is accurate? Unfortunately, there is no simple way to answer this question. The only answer is that you must rely on your mastery and understanding of English in order to determine for yourself whether the paraphrase is a good one or not. This is one of those kinds of skills that is difficult to teach, apart from just improving one's mastery of the English language.

1.6. Validity

So far we have discussed what arguments are and how to determine their structure, including how to reconstruct arguments in standard form. But we have not yet discussed what makes an argument good or bad. The central concept that you will learn in logic is the concept of validity. Validity relates to how well the premises support the conclusion, and it is the golden standard that every argument should aim for. A **valid argument** is an argument whose conclusion cannot possibly be false, assuming that the premises are true. Another way of putting this is as a conditional statement: A valid argument is an argument in which *if* the premises are true, the conclusion *must* be true. Here is an example of a valid argument:

1. Violet is a dog
2. Therefore, Violet is a mammal (from 1)

You might wonder whether it *is* true that Violet is a dog (maybe she's a lizard or a buffalo—we have no way of knowing from the information given). But, for the purposes of validity, it doesn't matter whether premise 1 is *actually* true or false. All that matters for validity is whether the conclusion follows from the premise. And we can see that the conclusion, Violet is a mammal, does seem to follow from the premise, Violet is a dog. That is, given the truth of the premise, the conclusion has to be true. This argument is clearly valid since if we assume that "Violet is a dog" is true, then, since all dogs are mammals, it follows that "Violet is a mammal" must also be true. As we've just seen, whether or not an argument is valid has nothing to do with whether the premises of the argument are actually true or not. We can illustrate this with another example, where the premises are clearly false:

1. Everyone born in France can speak French
2. Barack Obama was born in France
3. Therefore, Barack Obama can speak French (from 1-2)

This is a valid argument. Why? Because when we *assume* the truth of the premises (everyone born in France can speak French, Barack Obama was born in France) the conclusion (Barack Obama can speak French) *must* be true. Notice that this is so even though none of these statements is *actually* true. Not everyone born in France can speak French (think about people who were born there but then moved somewhere else where they didn't speak French and

never learned it) and Obama was not born in France, but it is also false that Obama can speak French. So we have a valid argument even though neither the premises nor the conclusion is actually true. That may sound strange, but if you understand the concept of validity, it is not strange at all. Remember: validity describes the *relationship* between the premises and conclusion, and it means that the premises imply the conclusion, whether or not that conclusion is true. In order to better understand the concept of validity, let's look at an example of an *invalid* argument:

1. George was President of the United States
2. Therefore, George was elected President of the United States (from 1)

This argument is **invalid** because it is possible for the premise to be true and yet the conclusion false. Here is a counterexample to the argument. Gerald Ford was President of the United States but he was never elected president, since Ford Replaced Richard Nixon when Nixon resigned in the wake of the Watergate scandal.² So it doesn't follow that just because someone is President of the United States that they were *elected* President of the United States. In other words, it is possible for the premise of the argument to be true and yet the conclusion false. And this means that the argument is invalid. If an argument is invalid it will always be possible to construct a counterexample to show that it is invalid (as I have done with the Gerald Ford scenario). A **counterexample** is simply a description of a scenario in which the premises of the argument are all true while the conclusion of the argument is false. If you can construct a counterexample to an argument, the argument is invalid.

In order to determine whether an argument is valid or invalid we can use what I'll call the **informal test of validity**. To apply the informal test of validity ask yourself whether you can *imagine* a world in which all the premises are true and yet the conclusion is false. If you *can* imagine such a world, then the argument is invalid. If you *cannot* imagine such a world, then the argument is valid. Notice: it is possible to imagine a world where the premises are true even if the premises aren't, as a matter of actual fact, true. This is why it doesn't matter for validity whether the premises (or conclusion) of the argument are actually true. It will help to better understand the concept of validity by applying the informal test of validity to some sample arguments.

² As it happens, Ford wasn't elected Vice President either since he was confirmed by the Senate, under the twenty fifth amendment, after Spiro Agnew resigned. So Ford wasn't ever elected by the Electoral College—as either Vice President or President.

1. Joan jumped out of an airplane without a parachute
2. Therefore, Joan fell to her death (from 1)

To apply the informal test of validity we have to ask whether it is possible to imagine a scenario in which the premise is true and yet the conclusion is false (if so, the argument is invalid). So, can we imagine a world in which someone jumped out of an airplane without a parachute and yet did not fall to her death? (Think about it carefully before reading on.) As we will see, applying the informal test of validity takes some creativity, but it seems clearly possible that Joan could jump out of an airplane without a parachute and not die—she could be perfectly fine, in fact. All we have to imagine is that the airplane was not operating and in fact was on the ground when Joan jumped out of it. If that were the case, it would be a) true that Joan jumped out of an airplane without a parachute and yet b) false that Joan fell to her death. Thus, since it is possible to imagine a scenario in which the premise is true and yet the conclusion is false, the argument is invalid. Let's slightly change the argument, this time making it clear that the plane is flying:

1. Joan jumped out of an airplane traveling 300 mph at a height of 10,000 ft without a parachute
2. Joan fell to her death (from 1)

Is this argument valid? You might think so since you might think that anyone who did such a thing would surely die. But is it possible to not die in the scenario described by the premise? If you think about it, you'll realize that there are lots of ways someone could survive. For example, maybe someone *else* who was wearing a parachute jumped out of the plane after them, caught them and attached the parachute-less person to them, and then pulled the ripcord and they both landed on the ground safe and sound. Or maybe Joan was performing a stunt and landed in a giant net that had been set up for that purpose. Or maybe she was just one of those people who, although they did fall to the ground, happened to survive (it has happened before). All of these scenarios are consistent with the information in the first premise being true and also consistent with the conclusion being false. Thus, again, any of these counterexamples show that this argument is invalid. Notice that it is also possible that the scenario described in the premises ends with Joan falling to her death. But that doesn't matter because all we want to know is whether it is possible that she doesn't. And if it is possible, what we have shown is that the

conclusion does not logically follow from the premise alone. That is, the conclusion doesn't have to be true, even if we grant that the premise is. And that means that the argument is not valid (i.e., it is invalid).

Let's switch examples and consider a different argument.

1. A person can be President of the United States only if they were born in the United States.
2. Obama is President of the United States.
3. Kenya is not in the United States.
4. Therefore, Obama was not born in Kenya (from 1-3)

In order to apply the informal test of validity, we have to ask whether we can imagine a scenario in which the premises are both true and yet the conclusion is false. So, we have to imagine a scenario in which premises 1, 2, and 3 are true and yet the conclusion ("Obama was not born in Kenya") is false. Can you imagine such a scenario? You cannot. The reason is that if you are imagining that it is a) true that a person can be President of the United States only if they were born in the United States, b) true that Obama is president and c) true that Kenya is not in the U.S., then it *must* be true that Obama was not born in Kenya. Thus we know that on the assumption of the truth of the premises, the conclusion *must* be true. And that means the argument is valid. In this example, however, premises 1, 2, and 3 are not only assumed to be true but are actually true. However, as we have already seen, the validity of an argument does not depend on its premises *actually* being true. Here is another example of a valid argument to illustrate that point.

1. A person can be President of the United States only if they were born in Kenya
2. Obama is President of the United States
3. Therefore, Obama was born in Kenya (from 1-2)

Clearly, the first premise of this argument is false. But if we were to imagine a scenario in which it is true and in which premise 2 is also true, then the conclusion ("Obama was born in Kenya") must be true. And this means that the argument is valid. We cannot imagine a scenario in which the premises of the argument are true and yet the conclusion is false. The important point to recognize here—a point I've been trying to reiterate throughout this section—is that the validity of the argument does not depend on whether or not the

premises (or conclusion) are *actually* true. Rather, validity depends only on the logical relationship between the premises and the conclusion. The actual truth of the premises is, of course, important to the quality of the argument, since if the premises of the argument are false, then the argument doesn't provide any reason for accepting the conclusion. In the next section we will address this topic.

Exercise 5: Determine whether or not the following arguments are valid by using the informal test of validity. If the argument is invalid, provide a counterexample.

1. Katie is a human being. Therefore, Katie is smarter than a chimpanzee.
2. Bob is a fireman. Therefore, Bob has put out fires.
3. Gerald is a mathematics professor. Therefore, Gerald knows how to teach mathematics.
4. Monica is a French teacher. Therefore, Monica knows how to teach French.
5. Bob is taller than Susan. Susan is taller than Frankie. Therefore, Bob is taller than Frankie.
6. Craig loves Linda. Linda loves Monique. Therefore, Craig loves Monique.
7. Orel Hershizer is a Christian. Therefore, Orel Hershizer communicates with God.
8. All Muslims pray to Allah. Muhammad is a Muslim. Therefore, Muhammad prays to Allah.
9. Some protozoa are predators. No protozoa are animals. Therefore, some predators are not animals.
10. Charlie only barks when he hears a burglar outside. Charlie is barking. Therefore, there must be a burglar outside.

1.7 Soundness

A good argument is not only valid, but also sound. Soundness is defined in terms of validity, so since we have already defined validity, we can now rely on it to define soundness. A **sound argument** is a valid argument that has all true premises. That means that the conclusion of a sound argument will always be true. Why? Because if an argument is valid, the premises transmit truth to the

conclusion on the assumption of the truth of the premises. But if the premises are actually true, as they are in a sound argument, then since all sound arguments are valid, we know that the conclusion of a sound argument is true. Compare the last two Obama examples from the previous section. While the first argument was sound, the second argument was not sound, although it was valid. The relationship between soundness and validity is easy to specify: all sound arguments are valid arguments, but not all valid arguments are sound arguments.

Although soundness is what any argument should aim for, we will not be talking much about soundness in this book. The reason for this is that the only difference between a valid argument and a sound argument is that a sound argument has all true premises. But how do we determine whether the premises of an argument are actually true? Well, there are lots of ways to do that, including using Google to look up an answer, studying the relevant subjects in school, consulting experts on the relevant topics, and so on. But none of these activities have anything to do with logic, per se. The relevant disciplines to consult if you want to know whether a particular statement is true is almost never logic! For example, logic has nothing to say regarding whether or not protozoa are animals or whether there are predators that aren't in the animal kingdom. In order to learn whether those statements are true, we'd have to consult biology, not logic. Since this is a logic textbook, however, it is best to leave the question of what is empirically true or false to the relevant disciplines that study those topics. And that is why the issue of soundness, while crucial for any good argument, is outside the purview of logic.

1.8 Deductive vs. Inductive arguments

The concepts of validity and soundness that we have introduced apply only to the class of what are called "deductive arguments". A **deductive argument** is an argument whose conclusion is supposed to follow from its premises with absolute certainty, thus leaving no possibility that the conclusion doesn't follow from the premises. For a deductive argument to fail to do this is for it to fail as a deductive argument. In contrast, an **inductive argument** is an argument whose conclusion is supposed to follow from its premises with a high level of probability, which means that although it is possible that the conclusion doesn't follow from its premises, it is unlikely that this is the case. Here is an example of an inductive argument:

Tweets is a healthy, normally functioning bird and since most healthy, normally functioning birds fly, Tweets probably flies.

Notice that the conclusion, Tweets probably flies, contains the word “probably.” This is a clear indicator that the argument is supposed to be inductive, not deductive. Here is the argument in standard form:

1. Tweets is a healthy, normally functioning bird
2. Most healthy, normally functioning birds fly
3. Therefore, Tweets probably flies

Given the information provided by the premises, the conclusion does seem to be well supported. That is, the premises do give us a strong reason for accepting the conclusion. This is true even though we can imagine a scenario in which the premises are true and yet the conclusion is false. For example, suppose that we added the following premise:

Tweets is 6 ft tall and can run 30 mph.

Were we to add that premise, the conclusion would no longer be supported by the premises, since any bird that is 6 ft tall and can run 30 mph, is not a kind of bird that can fly. That information leads us to believe that Tweets is an ostrich or emu, which are not kinds of birds that can fly. As this example shows, inductive arguments are **defeasible arguments** since by adding further information or premises to the argument, we can overturn (defeat) the verdict that the conclusion is well-supported by the premises. Inductive arguments whose premises give us a strong, even if defeasible, reason for accepting the conclusion are called, unsurprisingly, **strong inductive arguments**. In contrast, an inductive argument that does not provide a strong reason for accepting the conclusion are called **weak inductive arguments**.

Whereas strong inductive arguments are defeasible, valid deductive arguments aren't. Suppose that instead of saying that *most* birds fly, premise 2 said that *all* birds fly.

1. Tweets is a healthy, normally function bird.
2. All healthy, normally functioning birds can fly.
3. Therefore, Tweets can fly.

This is a valid argument and since it is a valid argument, there are no further premises that we could add that could overturn the argument's validity. (True, premise 2 is false, but as we've seen that is irrelevant to determining whether an argument is valid.) Even if we were to add the premise that Tweets is 6 ft tall and can run 30 mph, it doesn't overturn the validity of the argument. As soon as we use the **universal generalization**, "all healthy, normally function birds can fly," then when we assume that premise is true and add that Tweets is a healthy, normally functioning bird, it *has* to follow from those premises that Tweets can fly. This is true even if we add that Tweets is 6 ft tall because then what we have to imagine (in applying our informal test of validity) is a world in which all birds, including those that are 6 ft tall and can run 30 mph, can fly.

Although inductive arguments are an important class of argument that are commonly used every day in many contexts, logic texts tend not to spend as much time with them since we have no agreed upon standard of evaluating them. In contrast, there is an agreed upon standard of evaluation of deductive arguments. We have already seen what that is; it is the concept of validity. In chapter 2 we will learn some precise, formal methods of evaluating deductive arguments. There are no such agreed upon formal methods of evaluation for inductive arguments. This is an area of ongoing research in philosophy. In chapter 3 we will revisit inductive arguments and consider some ways to evaluate inductive arguments.

1.9 Arguments with missing premises

Quite often, an argument will not explicitly state a premise that we can see is needed in order for the argument to be valid. In such a case, we can supply the premise(s) needed in order so make the argument valid. Making missing premises explicit is a central part of reconstructing arguments in standard form. We have already dealt in part with this in the section on paraphrasing, but now that we have introduced the concept of validity, we have a useful tool for knowing when to supply missing premises in our reconstruction of an argument. In some cases, the missing premise will be fairly obvious, as in the following:

Gary is a convicted sex-offender, so Gary is not allowed to work with children.

The premise and conclusion of this argument are straightforward:

1. Gary is a convicted sex-offender
2. Therefore, Gary is not allowed to work with children (from 1)

However, as stated, the argument is invalid. (Before reading on, see if you can provide a counterexample for this argument. That is, come up with an imaginary scenario in which the premise is true and yet the conclusion is false.) Here is just one counterexample (there could be many): Gary is a convicted sex-offender but the country in which he lives does not restrict convicted sex-offenders from working with children. I don't know whether there are any such countries, although I suspect there are (and it doesn't matter for the purpose of validity whether there are or aren't). In any case, it seems clear that this argument is relying upon a premise that isn't explicitly stated. We can and should state that premise explicitly in our reconstruction of the standard form argument. But what is the argument's missing premise? The obvious one is that no sex-offenders are allowed to work with children, but we could also use a weaker statement like this one:

Where Gary lives, no convicted sex-offenders are allowed to work with children.

It should be obvious why this is a "weaker" statement. It is weaker because it is not so universal in scope, which means that it is easier for the statement to be made true. By relativizing the statement that sex-offenders are not allowed to work with children to the *place* where Gary lives, we leave open the possibility that other places in the world don't have this same restriction. So even if there are other places in the world where convicted sex-offenders are allowed to work with children, our statements could still be true since in *this* place (the place where Gary lives) they aren't. (For more on strong and weak statements, see section 1.10). So here is the argument in standard form:

1. Gary is a convicted sex-offender.
2. Where Gary lives, no convicted sex-offenders are allowed to work with children.
3. Therefore, Gary is not allowed to work with children. (from 1-2)

This argument is now valid: there is no way for the conclusion to be false, assuming the truth of the premises. This was a fairly simple example where the

missing premise needed to make the argument valid was relatively easy to see. As we can see from this example, a **missing premise** is a premise that the argument needs in order to be as strong as possible. Typically, this means supplying the statement(s) that are needed to make the argument valid. But in addition to making the argument valid, we want to make the argument plausible. This is called “the principle of charity.” The **principle of charity** states that when reconstructing an argument, you should try to make that argument (whether inductive or deductive) as strong as possible. When it comes to supplying missing premises, this means supplying the most plausible premises needed in order to make the argument either valid (for deductive arguments) or inductively strong (for inductive arguments).

Although in the last example figuring out the missing premise was relatively easy to do, it is not always so easy. Here is an argument whose missing premises are not as easy to determine:

Since children who are raised by gay couples often have psychological and emotional problems, the state should discourage gay couples from raising children.

The conclusion of this argument, that the state should not allow gay marriage, is apparently supported by a single premise, which should be recognizable from the occurrence of the premise indicator, “since.” Thus, our initial reconstruction of the standard form argument looks like this:

1. Children who are raised by gay couples often have psychological and emotional problems.
2. Therefore, the state should discourage gay couples from raising children.

However, as it stands, this argument is invalid because it depends on certain missing premises. The conclusion of this argument is a **normative statement**—a statement about whether something *ought to be true*, relative to some standard of evaluation. Normative statements can be contrasted with **descriptive statements**, which are simply factual claims about what *is true*. For example, “Russia *does not* allow gay couples to raise children” is a descriptive statement. That is, it is simply a claim about what is *in fact the case* in Russia today. In contrast, “Russia *should not* allow gay couples to raise children” is a normative statement since it is not a claim about what is true, but what ought to

be true, relative to some standard of evaluation (for example, a moral or legal standard). An important idea within philosophy, which is often traced back to the Scottish philosopher David Hume (1711-1776), is that statements about what ought to be the case (i.e., normative statements) can never be derived from statements about what is the case (i.e., descriptive statements). This is known within philosophy as **the is-ought gap**. The problem with the above argument is that it attempts to infer a normative statement from a purely descriptive statement, violating the is-ought gap. We can see the problem by constructing a counterexample. Suppose that in society x it is true that children raised by gay couples have psychological problems. However, suppose that in that society people do not accept that the state should do what it can to decrease harm to children. In this case, the conclusion, that the state should discourage gay couples from raising children, does not follow. Thus, we can see that the argument depends on a missing or assumed premise that is not explicitly stated. That missing premise must be a normative statement, in order that we can infer the conclusion, which is also a normative statement. There is an important general lesson here: *Many times an argument with a normative conclusion will depend on a normative premise which is not explicitly stated.* The missing normative premise of this particular argument seems to be something like this:

The state should always do what it can to decrease harm to children.

Notice that this is a normative statement, which is indicated by the use of the word "should." There are many other words that can be used to capture normative statements such as: good, bad, and ought. Thus, we can reconstruct the argument, filling in the missing normative premise like this:

1. Children who are raised by gay couples often have psychological and emotional problems.
2. The state should always do what it can to decrease harm to children.
3. Therefore, the state should discourage gay couples from raising children. (from 1-2)

However, although the argument is now in better shape, it is *still* invalid because it is still possible for the premises to be true and yet the conclusion false. In order to show this, we just have to imagine a scenario in which both the premises are true and yet the conclusion is false. Here is one counterexample to the argument (there are many). Suppose that while it is true that children of gay couples often have psychological and emotional problems, the rate of

psychological problems in children raised by gay couples is actually lower than in children raised by heterosexual couples. In this case, even if it were true that the state should always do what it can to decrease harm to children, it does not follow that the state should discourage gay couples from raising children. In fact, in the scenario I've described, just the opposite would seem to follow: the state should discourage heterosexual couples from raising children.

But even if we suppose that the rate of psychological problems in children of gay couples is higher than in children of heterosexual couples, the conclusion still doesn't seem to follow. For example, it could be that the reason that children of gay couples have higher rates of psychological problems is that in a society that is not yet accepting of gay couples, children of gay couples will face more teasing, bullying and general lack of acceptance than children of heterosexual couples. If this were true, then the harm to these children isn't so much due to the fact that their parents are gay as it is to the fact that their community does not accept them. In that case, the state should not necessarily discourage gay couples from raising children. Here is an analogy: At one point in our country's history (if not still today) it is plausible that the children of black Americans suffered more psychologically and emotionally than the children of white Americans. But for the government to discourage black Americans from raising children would have been unjust, since it is likely that if there was a higher incidence of psychological and emotional problems in black Americans, then it was due to unjust and unequal conditions, not to the black parents, *per se*. So, to return to our example, the state should only discourage gay couples from raising children if they know that the higher incidence of psychological problems in children of gay couples *isn't* the result of any kind of injustice, but is due to the simple fact that the parents are gay.

Thus, one way of making the argument (at least closer to) valid would be to add the following two missing premises:

- A. The rate of psychological problems in children of gay couples is higher than in children of heterosexual couples.
- B. The higher incidence of psychological problems in children of gay couples is not due to any kind of injustice in society, but to the fact that the parents are gay.

So the reconstructed standard form argument would look like this:

1. Children who are raised by gay couples often have psychological and emotional problems.
2. The rate of psychological problems in children of gay couples is higher than in children of heterosexual couples.
3. The higher incidence of psychological problems in children of gay couples is not due to any kind of injustice in society, but to the fact that the parents are gay.
4. The state should always do what it can to decrease harm to children.
5. Therefore, the state should discourage gay couples from raising children. (from 1-4)

In this argument, premises 2-4 are the missing or assumed premises. Their addition makes the argument much stronger, but making them explicit enables us to clearly see what assumptions the argument relies on in order for the argument to be valid. This is useful since we can now clearly see which premises of the argument we may challenge as false. Arguably, premise 4 is false, since the state shouldn't always do what it can to decrease harm to children. Rather, it should only do so as long as such an action didn't violate other rights that the state has to protect or create larger harms elsewhere.

The important lesson from this example is that supplying the missing premises of an argument is not always a simple matter. In the example above, I have used the principle of charity to supply missing premises. Mastering this skill is truly an art (rather than a science) since there is never just one correct way of doing it (cf. section 1.5) and because it requires a lot of skilled practice.

Exercise 6: Supply the missing premise or premises needed in order to make the following arguments valid. Try to make the premises as plausible as possible while making the argument valid (which is to apply the principle of charity).

1. Ed rides horses. Therefore, Ed is a cowboy.
2. Tom was driving over the speed limit. Therefore, Tom was doing something wrong.
3. If it is raining then the ground is wet. Therefore, the ground must be wet.
4. All elves drink Guinness, which is why Olaf drinks Guinness.
5. Mark didn't invite me to homecoming. Instead, he invited his friend Alexia. So he must like Alexia more than me.

6. The watch must be broken because every time I have looked at it, the hands have been in the same place.
7. Olaf drank too much Guinness and fell out of his second story apartment window. Therefore, drinking too much Guinness caused Olaf to injure himself.
8. Mark jumped into the air. Therefore, Mark landed back on the ground.
9. In 2009 in the United States, the net worth of the median white household was \$113,149 a year, whereas the net worth of the median black household was \$5,677. Therefore, as of 2009, the United States was still a racist nation.
10. The temperature of the water is 212 degrees Fahrenheit. Therefore, the water is boiling.
11. Capital punishment sometimes takes innocent lives, such as the lives of individuals who were later found to be not guilty. Therefore, we should not allow capital punishment.
12. Allowing immigrants to migrate to the U.S. will take working class jobs away from working class folks. Therefore, we should not allow immigrants to migrate to the U.S.
13. Prostitution is a fair economic exchange between two consenting adults. Therefore, prostitution should be allowed.
14. Colleges are more interested in making money off of their football athletes than in educating them. Therefore, college football ought to be banned.
15. Edward received an F in college Algebra. Therefore, Edward should have studied more.

1.10 Assuring, guarding and discounting

As we have seen, arguments often have complex structures including subarguments (recall that a subargument is an argument for one of the premises of the main argument). But in practice people do not always give further reasons or argument in support of every statement they make. Sometimes they use certain rhetorical devices to cut the argument short, or to hint at a further argument without actually stating it. There are three common strategies for doing this:

Assuring: informing someone that there are further reasons although one is not giving them now

Guarding: weakening one's claims so that it is harder to show that the claims are false

Discounting: anticipating objections that might be raised to one's claim or argument as a way of dismissing those objections.³

We will discuss these in order, starting with assuring. Why would we want to assure our audience? Presumably when we make a claim that isn't obvious and that the audience may not be inclined to believe. For example, if I am trying to convince you that the United States is one of the leading producers of CO₂ emissions, then I might cite certain authorities such as the Intergovernmental Panel on Climate Change (IPCC) as saying so. This is one way of assuring our audience: by citing authorities. There are many ways to cite authorities, some examples of which are these:

Dentists agree that...

Recent studies have shown...

It has been established that...

Another way of assuring is to comment on the strength of one's own convictions. The rhetorical effect is that by commenting on how sure you are that something is true, you imply, without saying, that there must be very strong reasons for what you believe—assuming that the audience believes you are a reasonable person, of course. Here are some ways of commenting on the strength of one's beliefs:

I'm certain that...

I'm sure that...

I can assure you that...

³ This characterization and discussion draws heavily on chapter 3, pp. 48-53 of Sinnott-Armstrong and Fogelin's *Understanding Arguments*, 9th edition (Cengage Learning).

Over the years, I have become convinced that...

I would bet a million dollars that...

Yet another way of assuring one's audience is to make an audience member feel that it would be stupid, odd, or strange to deny the claim one is making. One common way to do this is by implying that *every sensible person* would agree with the claim. Here are some examples:

Everyone with any sense agrees that...

Of course, no one will deny that...

There is no question that...

No one with any sense would deny that...

Another common way of doing this is by implying that *no sensible person* would agree with a claim that we are trying to establish as false:

It is no longer held that...

No intelligent person would ever maintain that...

You would have to live under a rock to think that...

Assurances are not necessarily illegitimate, since the person may be right and may in fact have good arguments to back up the claims, but the assurances are not themselves arguments and a critical thinker will always regard them as somewhat suspect. This is especially so when the claim isn't obviously true.

Next, we will turn to guarding. Guarding involves weakening a claim so that it is easier to make that claim true. Here is a simple contrast that will make the point. Consider the following claims:

- A. *All* U.S. Presidents were monogamous
- B. *Almost all* U.S. Presidents were monogamous
- C. *Most* U.S. Presidents were monogamous

- D. *Many* U.S. Presidents were monogamous
- E. *Some* U.S. Presidents were monogamous

The weakest of these claims is E, whereas the strongest is A and each claims descending from A-E is increasingly weaker. It doesn't take very much for E to be true: there just has to be at least one U.S. President who was monogamous. In contrast, A is much less likely than E to be true because it requires every U.S. President to have been monogamous. One way of thinking about this is that any time A is true, it is also true that B-E is true, but B-E could be true without A being true. That is what it means for a claim to be stronger or weaker. A **weak claim** is more likely to be true whereas a **strong claim** is less likely to be true. E is much more likely to be true than A. Likewise, D is somewhat more likely to be true than C, and so on.

So, guarding involves taking a stronger claim and making it weaker so there is less room to object to the claim. We can also guard a claim by introducing a probability clause such as, "it is possible that..." and "it is arguable that..." or by reducing our level of commitment to the claim, such as moving from "I know that x" to "I believe that x." One common use of guarding is in reconstructing arguments with missing premises using the principle of charity (section 1.9). For example, if an argument is that "Tom works for Merrill Lynch, so Tom has a college degree," the most charitable reconstruction of this argument would fill in the missing premise with "most people who work for Merrill Lynch have college degrees" rather than "everyone who works for Merrill Lynch has a college degree." Here we have created a more charitable (plausible) premise by weakening the claim from "all" to "most," which as we have seen is a kind of guarding.

Finally, we will consider discounting. Discounting involves acknowledging an objection to the claim or argument that one is making, while dismissing that same objection. The rhetorical force of discounting is to make it seem as though the argument has taken account of the objections—especially the ones that might be salient in a person's mind. The simplest and most common way of discounting is by using the "A but B" locution. Contrast the following two claims:

- A. The worker was inefficient, but honest.
- B. The worker was honest, but inefficient.

Although each statement asserts the same facts, A seems to be recommending the worker, whereas B doesn't. We can imagine A continuing: "And so the manager decided to keep her on the team." We can imagine B continuing: "Which is why the manager decided to let her go." This is what we can call the "A but B" locution. The **"A but B" locution** is a form of discounting that introduces what will be dismissed or overridden first and then follows it by what is supposed to be the more important consideration. By introducing the claim to be dismissed, we are discounting that claim. There are many other words that can be used as discounting words instead of using "but." Table 2 below gives a partial list of words and phrases that commonly function as discounting terms.

although	even if	but	nevertheless
though	while	however	nonetheless
even though	whereas	yet	still

Exercise 7: Which rhetorical techniques (assuring, guarding, discounting) are being using in the following passages?

1. Although drilling for oil in Alaska will disrupt some wildlife, it is better than having to depend on foreign oil, which has the tendency to draw us into foreign conflicts that we would otherwise not be involved in.
2. Let there be no doubt: the entity that carried out this attack is a known terrorist organization, whose attacks have a characteristic style—a style that is seen in this attack today.
3. Privatizing the water utilities in Detroit was an unprecedented move that has garnered a lot of criticism. Nonetheless, it is helping Detroit to recover from bankruptcy.
4. Most pediatricians agree that the single most important factor in childhood obesity is eating sugary, processed foods, which have become all too common in our day and age.
5. Although not every case of AIDS is caused by HIV, it is arguable that most are.
6. Abraham Lincoln was probably our greatest president since he helped keep together a nation on the brink of splintering into two.
7. No one with any sense would support Obamacare.
8. Even if universal healthcare is expensive, it is still the just thing to do.

9. While our country has made significant strides in overcoming explicit racist policies, the wide disparity of wealth, prestige and influence that characterize white and black Americans shows that we are still implicitly a racist country.
10. Recent studies have shown that there is no direct link between vaccines and autism.

1.11 Evaluative language

Yet another rhetorical technique that is commonly encountered in argumentation is the use of evaluative language to influence one's audience to accept the conclusion one is arguing for. Evaluative language can be contrasted with descriptive language. Whereas **descriptive language** simply describes a state of affairs, without passing judgment (positive or negative) on that state of affairs, **evaluative language** is used to pass some sort of judgment, positive or negative, on something. Contrast the following two statements:

Bob is *tall*.

Bob is *good*.

"Tall" is a descriptive term since being tall is, in itself, neither a good nor bad thing. Rather, it is a **purely descriptive term** that does not pass any sort of judgment, positive or negative, on the fact that Bob is tall. In contrast, "good" is a **purely evaluative term**, which means that the only thing the word does is make an evaluation (in this case, a positive evaluation) and doesn't carry any descriptive content. "Good," "bad," "right," and "wrong" are examples of purely evaluative terms. The interesting kinds of terms are those that are both descriptive and evaluative. For example:

Bob is *nosy*.

"Nosy" is a negatively evaluative term since to call someone nosy is to make a negative evaluation of them—or at least of that aspect of them. But it also implies a descriptive content, such as that Bob is curious about other people's affairs. We could re-describe Bob's nosiness using purely descriptive language:

Bob is *very curious about other people's affairs*.

Notice that while the phrase “very curious about other people’s affairs” does capture the descriptive sense of “nosy,” it doesn’t capture the evaluative sense of nosy, since it doesn’t carry with it the negative connotation that “nosy” does.

Evaluative language is rife in our society, perhaps especially so in political discourse. This isn’t surprising since by using evaluative language to describe certain persons, actions, or events we can influence how people understand and interpret the world. If you can get a person to think of someone or some state of affairs in terms of a positively or negatively evaluative term, chances are you will be able to influence their evaluation of that person or state of affairs. That is one of the rhetorical uses of evaluative language. Compare, for example,

Bob is a *rebel*.

Bob is a *freedom fighter*.

Whereas “rebel” tends to be a negatively evaluative term, “freedom fighter,” at least for many Americans, tends to be a positively evaluative term. Both words, however, have the same descriptive content, namely, that Bob is someone who has risen in armed resistance to an existing government. The difference is that whereas “rebel” makes a negative evaluation, “freedom fighter” makes a positive evaluation. [Table 3](#) below gives a small sampling of some evaluative terms.

beautiful	dangerous	wasteful	sneaky	cute
murder	prudent	courageous	timid	nosy
sloppy	smart	capable	insane	curt

English contains an interesting mechanism for turning positively evaluative terms into negative evaluative ones. All you have to do is put the word “too” before a positively evaluative terms and it will all of a sudden take on a negative connotation. Compare the following:

John is *honest*.

John is *too honest*.

Whereas “honest” is a positively evaluative term, “too honest” is a negatively evaluative term. When someone describes John as “too honest,” we can easily imagine that person going on to describe how John’s honesty is actually a liability or negative trait. Not so when he is simply described as honest. Since the word “too” indicates an excess, and to say that something is an excess is to make a criticism, we can see why the word “too” changes the valence of an evaluation from positive to negative.

Evaluative language provides a good illustration of the difference between logic, which is concerned with the analysis and evaluation of arguments, and rhetoric, which is concerned with persuasion more generally. There are many ways that humans can be caused to believe things besides through rational argumentation. In fact, sometimes these other persuasive techniques are much more effective. (Consider advertising techniques in the 1950s, which more often tried to use argument and evidence to convince consumers to buy products, compared to advertising today, which rarely uses argument and evidence.) In any case, evaluative language—especially the use of hybrid terms that have both descriptive and evaluative aspects—can lead people to subtly accept a claim without ever arguing for it. As an analogy for how this could work in conversation, consider the concept of what philosophers⁴ have called “presupposition.” If I say something like

Even Jane could pass

I have asserted that Jane could pass the course. But I have also presupposed that Jane is not a very good student (or not very smart) by using the word “even.” If I were to say “nuh-uh,” this would naturally be taken as rejecting the claim that Jane could pass (i.e., I would be saying that she *couldn’t* pass). And if I were to agree, I would naturally be taken as agreeing that she could pass. But notice that there isn’t any simply yes/no way to disagree with the presupposition that Jane isn’t a smart/good student. Since presuppositions are more difficult to challenge, they can end up influencing what people in the conversation are taking for granted and in this way presupposition can influence what people accept as true *without any argument or evidence*. Of course, a person could explicitly challenge the implicit presupposition that Jane isn’t smart or a good student, but that takes extra effort and many times people don’t realize that a presupposition has just slipped into a conversation.

⁴ For example, see David Lewis’s “Scorekeeping in a Language Game” (1979).

I suggest that hybrid evaluative/descriptive terms can work as a kind of presupposition. If I describe someone as an “insurgent,” for example, I am saying something both descriptive—person who has risen in armed resistance against an existing government—and negatively evaluative since the connotation of the term “insurgent” (as compared to “freedom fighter”) has come to be that of someone doing something bad or negative. In using the term “insurgent” no one has *explicitly* claimed that the individual/group in question is bad, but because the term has (for us) a negative connotation it can lead us to be more receptive to accepting (implicitly) claims such as that the person/group is bad or is doing something bad/harmful.

Thus, like assuring and discounting (section 1.10), evaluative language is a rhetorical technique. As such, it is more concerned with *non-rational persuasion* than it is with giving *reasons*. Non-rational persuasion is ubiquitous in our society today, not the least of which because advertising is ubiquitous and advertising today almost always uses non-rational persuasion. Think of the last time you saw some commercial present *evidence* for why you should buy their product (i.e., never) and you will realize how pervasive this kind of rhetoric is. Philosophy has a complicated relationship with rhetoric—a relationship that stretches back to Ancient Greece. Socrates disliked those, such as the Sophists, who promised to teach people how to effectively persuade someone of something, regardless of whether that thing was true. Although some people might claim that there is no essential difference between *giving reasons* for accepting a conclusion and trying to *persuade by any means*, most philosophers, including the author of this text, think otherwise. If we define rhetoric as the art of persuasion, then although argumentation is a kind of rhetoric (since it is a way of persuading), not all rhetoric is argumentation. The essential difference, as already hinted at, is that argumentation attempts to persuade by giving reasons whereas rhetoric attempts to persuade by any means, including non-rational means. If I tell you over and over again (in creative and subliminal ways) to drink Beer x because Beer x is the best beer, then I may very well make you think that Beer x is the best beer, but I have not thereby given you a *reason to accept that* Beer x is the best beer. Thinking of it rationally, the mere fact that I’ve told you lots of times that Beer x is the best beer gives you no good reason for believing that Beer x is in fact the best beer.

The rhetorical devices surveyed in the last two sections may be effective ways of persuading people, but they are not the same thing as offering an argument.

And if we attempt to see them as arguments, they turn out to be pretty poor arguments. One of the many things that psychologists study is how we are persuaded to believe or do things. As an **empirical science**, psychology attempts to describe and explain the way things are, in this case, the processes that lead us to believe or act as we do. Logic, in contrast, is not an empirical science. Logic is not trying to tell us how we do think, but what *good* thinking is and, thus, how we *ought* think. The study of logic is the study of the nature of arguments and, importantly, of what distinguishes a *good* argument from a *bad* one. “Good” and “bad” are what philosophers call **normative concepts** because they involve standards of evaluation.⁵ Since logic concerns what makes something a *good argument*, logic is sometimes referred to as a **normative science**. The key standard of evaluation of arguments that we have seen so far is that of validity. In chapter 2 we will consider some more precise, formal methods of understanding validity. Other “normative sciences” include ethics (the study of what a good life is and how we ought to live) and epistemology (the study of what we have good reason to believe).

1.12 Analyzing a real-life argument

In this section I will analyze a real-life argument—an excerpt from President Obama’s September 10, 2013 speech on Syria. I will use the concepts and techniques that have been introduced in this chapter to analyze and evaluate Obama’s argument. It is important to realize that regardless of one’s views—whether one agrees with Obama or not—one can still analyze the structure of the argument and even evaluate it by applying the informal test of validity to the reconstructed argument in standard form. I will present the excerpt of Obama’s speech and then set to work analyzing the argument it contains. In addition to creating the excerpt, the only addition I have made to the speech is numbering each paragraph with Roman numerals for ease of referring to specific places in my analysis of the argument.

I. My fellow Americans, tonight I want to talk to you about Syria, why it matters and where we go from here. Over the past two years, what began as a series of peaceful protests against the repressive regime of Bashar al-Assad has turned into a brutal civil war. Over a hundred thousand people have been killed. Millions have fled the country. In that time, America has worked with allies to provide humanitarian support, to help the moderate opposition and to shape a political settlement.

⁵ We encountered normative concepts when discussing normative statements in section 1.9.

II. But I have resisted calls for military action because we cannot resolve someone else's civil war through force, particularly after a decade of war in Iraq and Afghanistan.

III. The situation profoundly changed, though, on Aug. 21st, when Assad's government gassed to death over a thousand people, including hundreds of children. The images from this massacre are sickening, men, women, children lying in rows, killed by poison gas, others foaming at the mouth, gasping for breath, a father clutching his dead children, imploring them to get up and walk. On that terrible night, the world saw in gruesome detail the terrible nature of chemical weapons and why the overwhelming majority of humanity has declared them off limits, a crime against humanity and a violation of the laws of war.

IV. This was not always the case. In World War I, American GIs were among the many thousands killed by deadly gas in the trenches of Europe. In World War II, the Nazis used gas to inflict the horror of the Holocaust. Because these weapons can kill on a mass scale, with no distinction between soldier and infant, the civilized world has spent a century working to ban them. And in 1997, the United States Senate overwhelmingly approved an international agreement prohibiting the use of chemical weapons, now joined by 189 governments that represent 98 percent of humanity.

V. On Aug. 21st, these basic rules were violated, along with our sense of common humanity.

VI. No one disputes that chemical weapons were used in Syria. The world saw thousands of videos, cellphone pictures and social media accounts from the attack. And humanitarian organizations told stories of hospitals packed with people who had symptoms of poison gas.

VII. Moreover, we know the Assad regime was responsible. In the days leading up to Aug. 21st, we know that Assad's chemical weapons personnel prepared for an attack near an area where they mix sarin gas. They distributed gas masks to their troops. Then they fired rockets from a regime-controlled area into 11 neighborhoods that the regime has been trying to wipe clear of opposition forces.

VIII. Shortly after those rockets landed, the gas spread, and hospitals filled with the dying and the wounded. We know senior figures in Assad's military machine reviewed the results of the attack. And the regime increased their shelling of the same neighborhoods in the days that followed. We've also studied samples of blood and hair from people at the site that tested positive for sarin.

IX. When dictators commit atrocities, they depend upon the world to look the other way until those horrifying pictures fade from memory. But these things happened. The facts cannot be denied.

X. The question now is what the United States of America and the international community is prepared to do about it, because what happened to those people, to

those children, is not only a violation of international law, it's also a danger to our security.

XI. Let me explain why. If we fail to act, the Assad regime will see no reason to stop using chemical weapons.

XII. As the ban against these weapons erodes, other tyrants will have no reason to think twice about acquiring poison gas and using them. Over time our troops would again face the prospect of chemical warfare on the battlefield, and it could be easier for terrorist organizations to obtain these weapons and to use them to attack civilians.

XIII. If fighting spills beyond Syria's borders, these weapons could threaten allies like Turkey, Jordan and Israel.

XIV. And a failure to stand against the use of chemical weapons would weaken prohibitions against other weapons of mass destruction and embolden Assad's ally, Iran, which must decide whether to ignore international law by building a nuclear weapon or to take a more peaceful path.

XV. This is not a world we should accept. This is what's at stake. And that is why, after careful deliberation, I determined that it is in the national security interests of the United States to respond to the Assad regime's use of chemical weapons through a targeted military strike. The purpose of this strike would be to deter Assad from using chemical weapons, to degrade his regime's ability to use them and to make clear to the world that we will not tolerate their use. That's my judgment as commander in chief.

The first question to ask yourself is: What is the main point or conclusion of this speech? What conclusion is Obama trying to argue for? This is no simple question and in fact requires a good level of reading comprehension in order to answer it correctly. One of the things to look for is conclusion or premise indicators (section 1.2). There are numerous conclusion indicators in the speech, which is why you cannot simply mindlessly look for them and then assume the first one you find is the conclusion. Rather, you must rely on your comprehension of the speech to truly find the main conclusion. If you carefully read the speech, it is clear that Obama is trying to convince the American public of the necessity of taking military action against the Assad regime in Syria. So the conclusion is going to have to have something to do with that. One clear statement of what looks like a main conclusion comes in paragraph 15 where Obama says:

And that is why, after careful deliberation, I determined that it is in the national security interests of the United States to respond to the Assad regime's use of chemical weapons through a targeted military strike.

The phrase, "that is why," is a conclusion indicator which introduces the main conclusion. Here is my paraphrase of that conclusion:

Main conclusion: It is in the national security interests of the United States to respond to Assad's use of chemical weapons with military force.

Before Obama argues for this main conclusion, however, he gives an argument for the claim that Assad did use chemical weapons on his own civilians. This is what is happening in paragraphs 1-9 of the speech. The reasons he gives for how we know that Assad used chemical weapons include:

- images of the destruction of women and children (paragraph VI)
- humanitarian organizations' stories of hospitals full of civilians suffering from symptoms of exposure to chemical weapons (paragraph VI)
- knowledge that Assad's chemical weapons experts were at a site where sarin gas is mixed just a few days before the attack (paragraph VII)
- the fact that Assad distributed gas masks to his troops (paragraph VII)
- the fact that Assad's forces fired rockets into neighborhoods where there were opposition forces (paragraph VII)
- senior military officers in Assad's regime reviewed results of the attack (paragraph VIII)
- the fact that sarin was found in blood and hair samples from people at the site of the attack (paragraph VIII)

These premises do indeed provide support for the conclusion that Assad used chemical weapons on civilians, but it is probably best to see this argument as a strong inductive argument, rather than a deductive argument. The evidence strongly supports, but does not compel, the conclusion that Assad was responsible. For example, even if all these facts were true, it could be that some other entity was trying to set Assad up. Thus, this first subargument should be taken as a strong inductive argument (assuming the premises are true, of course), since the truth of the premises would increase the probability that the conclusion is true, but not make the conclusion absolutely certain.

Although Obama does give an argument for the claim that Assad carried out chemical weapon attacks on civilians, that is simply an assumption of the main argument. Moreover, although the conclusion of the main argument is the one I have indicated above, I think there is another, intermediate conclusion that

Obama argues for more directly and that is that *if we don't respond to Assad's use of chemical weapons, then our own national security will be put at risk*. We can clearly see this conclusion stated in paragraph 10. Moreover, the very next phrase in paragraph 11 is a premise indicator, "let me explain why." Obama goes on to offer reasons for why failing to respond to Assad's use of chemical weapons would be a danger to our national security. Thus, the conclusion Obama argues more directly for is:

Intermediate conclusion: A failure to respond to Assad's use of chemical weapons is a threat to our national security.

So, if that is the conclusion that Obama argues for most directly, what are the premises that support it? Obama gives several in paragraphs 11-14:

- A. If we don't respond to Assad's use of chemical weapons, then Assad's regime will continue using them with impunity. (paragraph 11)
- B. If Assad's regime uses chemical weapons with impunity, this will effectively erode the ban on them. (implicit in paragraph 12)
- C. If the ban on chemical weapons erodes, then other tyrants will be more likely to attain and use them. (paragraph 12)
- D. If other tyrants attain and use chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield (paragraph 12)
- E. If we don't respond to Assad's use of chemical weapons and if fighting spills beyond Syrian borders, our allies could face these chemical weapons. (paragraph 13)
- F. If Assad's regime uses chemical weapons with impunity, it will weaken prohibitions on other weapons of mass destruction. (paragraph 14)
- G. If prohibitions on other weapons of mass destruction are weakened, this will embolden Assad's ally, Iran, to develop a nuclear program. (paragraph 14)

I have tried to make explicit each step of the reasoning, much of which Obama makes explicit himself (e.g., premises A-D). The main threats to national security that failing to respond to Assad would engender, according to Obama, are that U.S. troops and U.S. allies could be put in danger of facing chemical weapons and that Iran would be emboldened to develop a nuclear program. There is a missing premise that is being relied upon for these premises to validly imply the conclusion. Here is a hint as to what that missing premise is: Are all of these things truly a threat to national security? For example, how is Iran having a

nuclear program a threat to our national security? It seems there must be an implicit premise—not yet stated—that is to the effect that all of these things *are* threats to national security. Here is one way of construing that missing premise:

Missing premise 1: An increased likelihood of U.S. troops or allies facing chemical weapons on the battlefield or Iran becoming emboldened to develop a nuclear program are all threats to U.S. national security interests.

We can also make explicit within the standard form argument other intermediate conclusions that follow from the stated premises. Although we don't have to do this, it can be a helpful thing to do when an argument contains multiple premises. For example, we could explicitly state the conclusion that follows from the four conditional statements that are the first four premises:

1. If we don't respond to Assad's use of chemical weapons, then Assad's regime will continue using them with impunity.
2. If Assad's regime uses chemical weapons with impunity, this will effectively erode the ban on them.
3. If the ban on chemical weapons erodes, then other tyrants will be more likely to attain and use them.
4. If other tyrants attain and use chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield.
5. *Therefore, if we don't respond to Assad's use of chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield. (from 1-4)*

Premise 5 is an intermediate conclusion that makes explicit what follows from premises 1-4 (which I have represented using parentheses after that intermediate conclusion). We can do the same thing with the inference that follows from premises, 1, 7, and 8 (i.e., line 9). If we add in our missing premises then we have a reconstructed argument for what I earlier called the "intermediate conclusion" (i.e., the one that Obama most directly argues for):

1. If we don't respond to Assad's use of chemical weapons, then Assad's regime will continue using them with impunity.
2. If Assad's regime uses chemical weapons with impunity, this will effectively erode the ban on them.

3. If the ban on chemical weapons erodes, then other tyrants will be more likely to attain and use them.
4. If other tyrants attain and use chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield.
5. *Therefore, if we don't respond to Assad's use of chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield.* (from 1-4)
6. If we don't respond to Assad's use of chemical weapons and if fighting spills beyond Syrian borders, our allies could face these chemical weapons.
7. If Assad's regime uses chemical weapons with impunity, it will weaken prohibitions on other weapons of mass destruction.
8. If prohibitions on other weapons of mass destruction are weakened, this will embolden Assad's ally, Iran, to develop a nuclear program.
9. *Therefore, if we don't respond to Assad's use of chemical weapons, this will embolden Assad's ally, Iran, to develop a nuclear program.* (from 1, 7-8)
10. *An increased likelihood of U.S. troops or allies facing chemical weapons on the battlefield or Iran becoming emboldened to develop a nuclear program are threats to U.S. national security interests.*
11. *Therefore, a failure to respond to Assad's use of chemical weapons is a threat to our national security.* (from 5, 6, 9, 10)

As always, in this standard form argument I've listed in parentheses after the relevant statements which statements those statements follow from. The only thing now missing is how we get from this intermediate conclusion to what I earlier called the main conclusion. The main conclusion (i.e., that it is in national security interests to respond to Assad with military force) might be thought to follow directly. But it doesn't. It seems that Obama is relying on yet another unstated assumption. Consider: even if it is true that we should respond to a threat to our national security, it doesn't follow that we should respond with military force. For example, maybe we could respond with certain kinds of economic sanctions that would force the country to submit to our will. Furthermore, maybe there are some security threats such that responding to them with military force would only create further, and worse, security threats. Presumably we wouldn't want our response to a security threat to create even bigger security threats. For these reasons, we can see that Obama's argument, if it is to be valid, also relies on missing premises such as these:

Missing premise 2: The *only* way that the United States can adequately respond to the security threat that Assad poses is by military force.

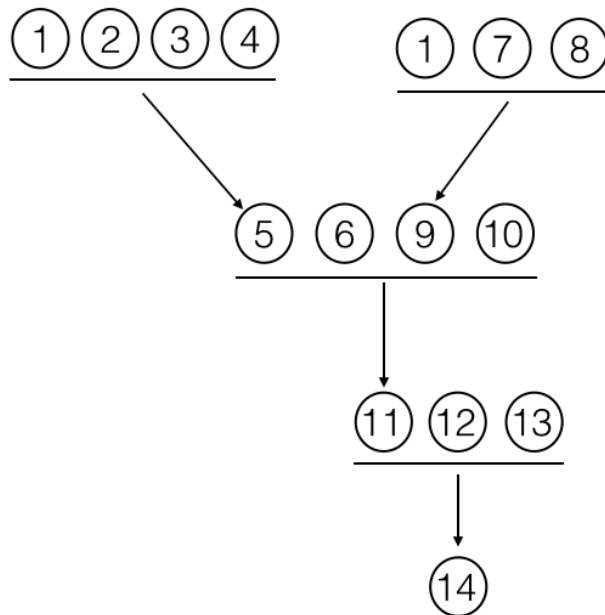
Missing premise 3: It is in the national security interests of the United States to respond adequately to *any* national security threat.

These are big assumptions and they may very well turn out to be mistaken. Nevertheless, it is important to see that the main conclusion Obama argues for depends on these missing premises—premises that he never explicitly states in his argument. So here is the final, reconstructed argument in standard form. I have italicized each missing premise or intermediate conclusion that I have added but that wasn't explicitly stated in Obama's argument.

1. If we don't respond to Assad's use of chemical weapons, then Assad's regime will continue using them with impunity.
2. If Assad's regime uses chemical weapons with impunity, this will effectively erode the ban on them.
3. If the ban on chemical weapons erodes, then other tyrants will be more likely to attain and use them.
4. If other tyrants attain and use chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield.
5. *Therefore, if we don't respond to Assad's use of chemical weapons, U.S. troops will be more likely to face chemical weapons on the battlefield. (from 1-4)*
6. If we don't respond to Assad's use of chemical weapons and if fighting spills beyond Syrian borders, our allies could face these chemical weapons.
7. If Assad's regime uses chemical weapons with impunity, it will weaken prohibitions on other weapons of mass destruction.
8. If prohibitions on other weapons of mass destruction are weakened, this will embolden Assad's ally, Iran, to develop a nuclear program.
9. *Therefore, if we don't respond to Assad's use of chemical weapons, this will embolden Assad's ally, Iran, to develop a nuclear program. (from 1, 7-8)*
10. *An increased likelihood of U.S. troops or allies facing chemical weapons on the battlefield or Iran becoming emboldened to develop a nuclear program are threats to U.S. national security interests.*
11. *Therefore, a failure to respond to Assad's use of chemical weapons is a threat to our national security. (from 5, 6, 9, 10)*

12. *The only way that the United States can adequately respond to the security threat that Assad poses is by military force.*
13. *It is in the national security interests of the United States to respond adequately to any national security threat.*
14. *Therefore, it is in the national security interests of the United States to respond to Assad's use of chemical weapons with military force.* (from 11-13)

In addition to showing the structure of the argument by use of parentheses which show which statements follow from which, we can also diagram the arguments spatially as we did in section 1.4 like this:



This is just another way of representing what I have already represented in the standard form argument, using parentheses to describe the structure. As is perhaps even clearer in the spatial representation of the argument's structure, this argument is complex in that it has numerous subarguments. So while statement 11 is a premise of the main argument for the main conclusion (statement 14), statement 11 is also itself a conclusion of a subargument whose premises are statements 5, 6, 9, and 10. And although statement 9 is a premise in that argument, it itself is a conclusion of yet another subargument whose premises are statements 1, 7 and 8. Almost any interesting argument will be complex in this way, with further subarguments in support of the premises of the main argument.

This chapter has provided you the tools to be able to reconstruct arguments like these. As we have seen, there is much to consider in reconstructing a complex argument. As with any skill, a true mastery of it requires lots of practice. In many ways, this is a skill that is more like an art than a science. The next chapter will introduce you to some basic formal logic, which is perhaps more like a science than an art.

2.1 What are formal methods of evaluation and why do we need them?

In chapter 1 we introduced the concept of validity and the informal test of validity. According to that test, in order to determine whether an argument is valid we ask whether we can imagine a scenario where the premises are true and yet the conclusion is false. If we can, then the argument is invalid; if we can't then the argument is valid. The informal test relies on our ability to imagine certain kinds of scenarios as well as our understanding of the statements involved in the argument. Because not everyone has the same powers of imagination or the same understanding, this informal test of validity is neither precise nor objective. For example, while one person may be able to imagine a scenario in which the premises of an argument are true while the conclusion is false, another person may be unable to imagine such a scenario. As a result, the argument will be classified as invalid by the first individual, but valid by the second individual. That is a problem because we would like our standard of evaluation of arguments (i.e., validity) to be as precise and objective as possible, and it seems that our informal test of validity is neither. It isn't precise because the concept of being able to imagine x is not precise—what counts as imagining x is not something that can be clearly specified. What are the precise success conditions for having imagined a scenario where the premises are true and the conclusion is false? But the informal test of validity also isn't objective since it is possible that two different people who applied the imagination test correctly could come to two different conclusions about whether the argument is valid. As I noted before, this is partly because people's understanding of the statements differ and partly because people have different powers of imagination.

The goal of a formal method of evaluation is to eliminate any imprecision or lack of objectivity in evaluating arguments. As we will see by the end of this chapter, logicians have devised a number of formal techniques that accomplish this goal for certain classes of arguments. What all of these formal techniques have in common is that you can apply them without really having to understand the meanings of the concepts used in the argument. Furthermore, you can apply the formal techniques without having to utilize imagination at all. Thus, the formal techniques we will survey in this chapter help address the lack of precision and objectivity inherent in the informal test of validity. In general, a formal method of evaluation is a method of evaluation of arguments that does not require one to understand the meaning of the statements involved in the argument. Although at this point this may sound like gibberish, after we have

introduced the formal methods, you will understand what it means to evaluate an argument without knowing what the statements of the argument mean. By the end of this chapter, if not before, you will understand what it means to evaluate an argument by its *form*, rather than its *content*.

However, I will give you a sense of what a formal method of evaluation is in a very simple case right now, to give you a foretaste of what we will be doing in this chapter. Suppose I tell you:

It is sunny and warm today.

This statement is a conjunction because it is a complex statement that is asserting two things:

It is sunny today.

It is warm today.

These two statements are conjoined with an “and.” So the conjunction is really two statements that are conjoined by the “and.” Thus, if I have told you that it is both sunny and warm today, it follows logically that it is sunny today. Here is that simple argument in standard form:

1. It is sunny today and it is warm today.
2. Therefore, it is sunny today. (from 1)

This is a valid inference that passes the informal test of validity. But we can also see that the *form* of the inference is perfectly general because it would work equally well for any conjunction, not just this one. This inference has a particular *form* that we could state using placeholders for the statements, “it is sunny today” and “it is warm today”:

1. A and B
2. Therefore, A

We can see that any argument that had this *form* would be a valid argument. For example, consider the statement:

Kant was a deontologist and a Pietist.

That statement is a conjunction of two statements that we can capture explicitly in the first premise of the following argument:

1. Kant was a deontologist and Kant was a Pietist.
2. Therefore, Kant was a deontologist. (from 1)

Regardless of whether you know what the statements in the first premise mean, we can still see that the inference is valid because the inference has the same form that I just pointed out above. Thus, you may not know what “Kant” is (one of the most famous German philosophers of the Enlightenment) or what a “deontologist” or “Pietist” is, but you can still see that since these are statements that form a conjunction, and since the inference made has a particular form that is valid, this particular inference is valid. That is what it means for an argument to be valid in virtue of its *form*. In the next section we will delve into formal logic, which will involve learning a certain kind of language. Don’t worry: it won’t be as hard as your French or Spanish class.

2.2 Propositional logic and the four basic truth functional connectives

Propositional logic (also called “sentential logic”) is the area of formal logic that deals with the logical relationships between propositions. A **proposition** is simply what I called in section 1.1 a statement.¹ Some examples of propositions are:

Snow is white

Snow is cold

Tom is an astronaut

The floor has been mopped

The dishes have been washed

¹ Some philosophers would claim that a proposition is not the same as a statement, but the reasons for doing so are not relevant to what we’ll be doing in this chapter. Thus, for our purposes, we can treat a proposition as the same thing as a statement.

We can also connect propositions together using certain English words, such as “and” like this:

The floor has been mopped *and* the dishes have been washed.

This proposition is called a **complex proposition** because it contains the connective “and” which connects two separate propositions. In contrast, “the floor has been mopped” and “the dishes have been washed” are what are called atomic propositions. **Atomic propositions** are those that do not contain any truth-functional connectives. The word “and” in this complex proposition is a truth-functional connective. A **truth-functional connective** is a way of connecting propositions such that the **truth value** of the resulting complex proposition can be determined by the truth value of the propositions that compose it. Suppose that the floor has not been mopped but the dishes have been washed. In that case, if I assert the conjunction, “the floor has been mopped and the dishes have been washed,” then I have asserted something that is false. The reason is that a conjunction, like the one above, is only true when each **conjunct** (i.e., each statement that is conjoined by the “and”) is true. If either one of the conjuncts is false, then the whole conjunction is false. This should be pretty obvious. If Bob and Sally split chores and Bob’s chore was to both vacuum and dust whereas Sally’s chore was to both mop and do the dishes, then if Sally said she mopped the floor and did the dishes when in reality she only did the dishes (but did not mop the floor), then Bob could rightly complain that it isn’t true that Sally both mopped the floor and did the dishes! What this shows is that *conjunctions are true only if both conjuncts are true*. This is true of all conjunctions. The conjunction above has a certain form—the same form as any conjunction. We can represent that form using placeholders—lowercase letters like p and q to stand for any statement whatsoever. Thus, we represent the form of a **conjunction** like this:

p and q

Any conjunction has this same form. For example, the complex proposition, “it is sunny and hot today,” has this same form which we can see by writing the conjunction this way:

It is sunny today and it is hot today.

Although we could write the conjunction that way, it is more natural in English to conjoin the adjectives “sunny” and “hot” to get “it is sunny and hot today.” Nevertheless, these two sentences mean the same thing (it’s just that one sounds more natural in English than the other). In any case, we can see that “it is sunny today” is the proposition in the “p” place of the form of the conjunction, whereas “it is hot today” is the proposition in the “q” place of the form of the conjunction. As before, this conjunction is true only if both conjuncts are true. For example, suppose that it is a sunny but bitterly cold winter’s day. In that case, while it is true that it is sunny today, it is false that it is hot today—in which case the conjunction is false. If someone were to assert that it is sunny and hot today in those circumstances, you would tell them that isn’t true. Conversely, if it were a cloudy but hot and humid summer’s day, the conjunction would still be false. The only way the statement would be true is if both conjuncts were true.

In the formal language that we are developing in this chapter, we will represent conjunctions using a symbol called the “dot,” which looks like this: “ \cdot ” Using this symbol, here is how we will represent a conjunction in symbolic notation:

$$p \cdot q$$

In the following sections we will introduce four basic truth-functional connectives, each of which have their own symbol and meaning. The four basic truth-functional connectives are: conjunction, disjunction, negation, and conditional. In the remainder of this section, we will discuss only conjunction.

As we’ve seen, a conjunction conjoins two separate propositions to form a complex proposition. The conjunction is true if and only if both conjuncts are true. We can represent this information using what is called a **truth table**. Truth tables represent how the truth value of a complex proposition depends on the truth values of the propositions that compose it. Here is the truth table for conjunction:

p	q	$p \cdot q$
T	T	T
T	F	F
F	T	F
F	F	F

Here is how to understand this truth table. The header row lists the atomic propositions, p and q , that the conjunction is composed of, as well as the conjunction itself, $p \cdot q$. Each of the following four rows represents a possible scenario regarding the truth of each conjunct, and there are only four possible scenarios: either p and q could both be true (as in row 1), p and q could both be false (as in row 4), p could be true while q is false (row 2), or p could be false while q is true (row 3). The final column (the truth values under the conjunction, $p \cdot q$) represents how the truth value of the conjunction depends on the truth value of each conjunct (p and q). As we have seen, a conjunction is true if and only if both conjuncts are true. This is what the truth table represents. Since there is only one row (one possible scenario) in which both p and q are true (i.e., row 1), that is the only circumstance in which the conjunction is true. Since in every other row at least one of the conjuncts is false, the conjunction is false in the remaining three scenarios.

At this point, some students will start to lose a handle on what we are doing with truth tables. Often, this is because one thinks the concept is much more complicated than it actually is. (For some, this may stem, in part, from a math phobia that is triggered by the use of symbolic notation.) But a truth table is actually a very simple idea: it is simply a representation of the meaning of a truth-functional operator. When I say that a conjunction is true only if both conjuncts are true, that is just what the table is representing. There is nothing more to it than that. (Later on in this chapter we will use truth tables to prove whether an argument is valid or invalid. Understanding that will require more subtlety, but what I have so far introduced is not complicated at all.)

There is more than one way to represent conjunctions in English besides the English word “and.” Below are some common English words and phrases that commonly function as truth-functional conjunctions.

but	yet	also	although
however	moreover	nevertheless	still

It is important to point out that many times English conjunctions carry more information than simply that the two propositions are true (which is the only information carried by our symbolic connective, the dot). We can see this with English conjunctions like “but” and “however” which have a contrastive sense. If I were to say, “Bob voted, but Caroline didn’t,” then I am contrasting what

Bob and Caroline did. Nevertheless, I am still asserting two independent propositions. Another kind of information that English conjunctions represent but the dot connective doesn't is temporal information. For example, in the conjunction:

Bob brushed his teeth and got into bed

There is clearly a temporal implication that Bob brushed his teeth first and then got into bed. It might sound strange to say:

Bob got into bed and brushed his teeth

since this would seem to imply that Bob brushed his teeth while in bed. But each of these conjunctions would be represented in the same way by our dot connective, since the dot connective does not care about the temporal aspects of things. If we were to represent "Bob got into bed" with the capital letter A and "Bob brushed his teeth" with the capital letter B, then both of these propositions would be represented exactly the same, namely, like this:

A · B

Sometimes a conjunction can be represented in English with just a comma or semicolon, like this:

While Bob vacuumed the floor, Sally washed the dishes.

Bob vacuumed the floor; Sally washed the dishes.

Both of these are conjunctions that are represented in the same way. You should see that both of them have the form, $p \cdot q$.

Not every conjunction is a truth-function conjunction. We can see this by considering a proposition like the following:

Maya and Alice are married.

If this were a truth-functional proposition, then we should be able to identify the two, independent propositions involved. But we cannot. What would those propositions be? You might think two propositions would be these:

Maya is married

Alice is married

But that cannot be right since the fact that Maya is married and that Alice is married is not the same as saying that Maya and Alice are married to each other, which is clearly the implication of the original sentence. Furthermore, if you tried to add “to each other” to each proposition, it would no longer make sense:

Maya is married to each other

Alice is married to each other

Perhaps we could say that the two conjuncts are “Maya is married to Alice” and “Alice is married to Maya,” but the truth values of those two conjuncts are not independent of each other since if Maya is married to Alice it must also be true that Alice is married to Maya. In contrast, the following is an example of a truth-functional conjunction:

Maya and Alice are women.

Unlike the previous example, in this case we can clearly identify two propositions whose truth values are independent of each other:

Maya is a woman

Alice is a woman

Whether or not Maya is a woman is an issue that is totally independent of whether Alice is a woman (and vice versa). That is, the fact that Maya is a woman tells us nothing about whether Alice is a woman. In contrast, the fact that Maya is married to Alice implies that Alice is married to Maya. *So the way to determine whether or not a conjunction is truth-functional is to ask whether it is formed from two propositions whose truth is independent of each other.* If there are two propositions whose truth is independent of each other, then the conjunction is truth-functional; if there are not two propositions whose truth is independent of each other, the conjunction is not truth-functional.

Exercise 8: Identify which of the following sentences are truth-functional conjunctions. If the sentence is a truth-functional conjunction, identify the two conjuncts (by writing them down).

1. Jack and Jill are coworkers.
2. Tom is a fireman and a father.
3. Ringo Starr and John Lennon were bandmates.
4. Lucy loves steak and onion sandwiches.
5. Cameron Diaz has had several relationships, although she has never married.
6. Bob and Sally kissed.
7. A person who plays both mandolin and guitar is a multi-instrumentalist.
8. No one has ever contracted rabies and lived.
9. Jack and Jill are cowboys.
10. Josiah is Amish; nevertheless, he is also a drug dealer.
11. The Tigers are the best baseball team in the state, but they are not as good as the Yankees.
12. Bob went to the beach to enjoy some rest and relaxation.
13. Lauren isn't the fastest runner on the team; still, she is fast enough to have made it to the national championship.
14. The ring is beautiful, but expensive.
15. It is sad, but true that many Americans do not know where their next meal will come from.

2.3. Negation and disjunction

In this section we will introduce the second and third truth-functional connectives: negation and disjunction. We will start with negation, since it is the easier of the two to grasp. **Negation** is the truth-functional operator that switches the truth value of a proposition from false to true or from true to false. For example, if the statement "dogs are mammals" is true (which it is), then we can make that statement false by adding a negation. In English, the negation is most naturally added just before the noun phrase that follows the linking verb like this:

Dogs are *not* mammals.

But another way of adding the negation is with the phrase, “it is not the case that” like this:

It is not the case that dogs are mammals.

Either of these English sentences expresses the same proposition, which is simply the negation of the atomic proposition, “dogs are mammals.” Of course, that proposition is false since it is true that dogs are mammals. Just as we can make a true statement false by negating it, we can also make a false statement true by adding a negation. For example, the statement, “Cincinnati is the capital of Ohio” is false. But we can make that statement true by adding a negation:

Cincinnati is *not* the capital of Ohio

There are many different ways of expressing negations in English. Here are a few ways of expressing the previous proposition in different ways in English:

Cincinnati *isn't* the capital of Ohio

It's not true that Cincinnati is the capital of Ohio

It is not the case that Cincinnati is the capital of Ohio

Each of these English sentences express the same true proposition, which is simply the negation of the atomic proposition, “Cincinnati is the capital of Ohio.” Since that statement is false, its negation is true.

There is one respect in which negation differs from the other three truth-functional connectives that we will introduce in this chapter. Unlike the other three, negation does not *connect* two different propositions. Nonetheless, we call it a truth-functional connective because although it doesn't actually connect two different propositions, it does change the truth value of propositions in a truth-functional way. That is, if we know the truth value of the proposition we are negating, then we know the truth value of the resulting negated proposition. We can represent this information in the truth table for negation. In the following table, the symbol we will use to represent negation is called the “tilde” (\sim). (You can find the tilde on the upper left-hand side of your keyboard.)

p	$\sim p$
T	F
F	T

This truth table represents the meaning of the truth-functional connective, negation, which is represented by the tilde in our symbolic language. The header row of the table represents some proposition p (which could be any proposition) and the negation of that proposition, $\sim p$. What the table says is simply that if a proposition is true, then the negation of that proposition is false (as in the first row of the table); and if a proposition is false, then the negation of that proposition is true (as in the second row of the table).

As we have seen, it is easy to form sentences in our symbolic language using the tilde. All we have to do is add a tilde to left-hand side of an existing sentence. For example, we could represent the statement "Cincinnati is the capital of Ohio" using the capital letter C , which is called a constant. In propositional logic, a **constant** is a capital letter that represents an atomic proposition. In that case, we could represent the statement "Cincinnati is not the capital of Ohio" like this:

$$\sim C$$

Likewise, we could represent the statement "Toledo is the capital of Ohio" using the constant T . In that case, we could represent the statement "Toledo is not the capital of Ohio" like this:

$$\sim T$$

We could also create a sentence that is a conjunction of these two negations, like this:

$$\sim C \cdot \sim T$$

Can you figure out what this complex proposition says? (Think about it; you should be able to figure it out given your understanding of the truth-functional connectives, negation and conjunction.) The proposition says (literally): "Cincinnati is not the capital of Ohio and Toledo is not the capital of Ohio." In later sections we will learn how to form complex propositions using various

combinations of each of the four truth-functional connectives. Before we can do that, however, we need to introduce our next truth-functional connective, disjunction.

The English word that most commonly functions as disjunction is the word “or.” It is also common that the “or” is preceded by an “either” earlier in the sentence, like this:

Either Charlie or Violet tracked mud through the house.

What this sentence asserts is that one or the other (and possibly both) of these individuals tracked mud through the house. Thus, it is composed out of the following two atomic propositions:

Charlie tracked mud through the house

Violet tracked mud through the house

If the fact is that Charlie tracked mud through the house, the statement is true. If the fact is that Violet tracked mud through the house, the statement is also true. This statement is only false if in fact neither Charlie nor Violet tracked mud through the house. This statement would also be true even if it was both Charlie and Violet who tracked mud through the house. Another example of a disjunction that has this same pattern can be seen in the “click it or ticket” campaign of the National Highway Traffic Safety Administration. Think about what the slogan means. What the campaign slogan is saying is:

Either buckle your seatbelt or get a ticket

This is a kind of warning: buckle your seatbelt or you’ll get a ticket. Think about the conditions under which this statement would be true. There are only four different scenarios:

Your seatbelt is buckled	You do not get a ticket	True
Your seatbelt is not buckled	You get a ticket	True
Your seatbelt is buckled	You get a ticket	True
Your seatbelt is not buckled	You do not get a ticket	False

The first and second scenarios (rows 1 and 2) are pretty straightforwardly true, according to the “click it or ticket” statement. But suppose that your seatbelt is buckled, is it still possible to get a ticket (as in the third scenario—row 3)? Of course it is! That is, the statement *allows* that it could both be true that your seatbelt is buckled and true that you get a ticket. How so? (Think about it for a second and you’ll probably realize the answer.) Suppose that your seatbelt is buckled but your are speeding, or your tail light is out, or you are driving under the influence of alcohol. In any of those cases, you would get a ticket even if you were wearing your seatbelt. So the disjunction, click it or ticket, clearly allows the statement to be true even when both of the **disjuncts** (the statements that form the disjunction) are true. The only way the disjunction would be shown to be false is if (when pulled over) you were not wearing your seatbelt and yet did not get a ticket. Thus, the only way for the disjunction to be false is when both of the disjuncts are false.

These examples reveal a pattern: a **disjunction** is a truth-functional statement that is true in every instance except where both of the disjuncts are false. In our symbolic language, the symbol we will use to represent a disjunction is called a “wedge” (\vee). (You can simply use a lowercase “v” to write the wedge.) Here is the truth table for disjunction:

p	q	$p \vee q$
T	T	T
T	F	T
F	T	T
F	F	F

As before, the header of this truth table represents two propositions (first two columns) and their disjunction (last column). The following four rows represent the conditions under which the disjunction is true. As we have seen, the disjunction is true when at least one of its disjuncts is true, including when they are both true (the first three rows). A disjunction is false only if both disjuncts are false (last row).

As we have defined it, the wedge (\vee) is what is called an “inclusive or.” An **inclusive or** is a disjunction that is true even when both disjuncts are true. However, sometimes a disjunction clearly implies that the statement is true only if either one or the other of the disjuncts is true, *but not both*. For example, suppose that you know that Bob placed either first or second in the race because you remember seeing a picture of him in the paper where he was standing on a podium (and you know that only the top two runners in the race get to stand on the podium). Although you can’t remember which place he was, you know that:

Bob placed either first or second in the race.

This is a disjunction that is built out of two different atomic propositions:

Bob placed first in the race

Bob placed second in the race

Although it sounds awkward to write it this way in English, we could simply connect each atomic statement with an “or”:

Bob placed first in the race or Bob placed second in the race.

That sentence makes explicit the fact that this statement is a disjunction of two separate statements. However, it is also clear that in this case the disjunction would not be true if all the disjuncts were true, because it is not possible for all the disjuncts to be true, since Bob cannot have placed both first and second. Thus, it is clear in a case such as this, that the “or” is meant as what is called an “exclusive or.” An **exclusive or** is a disjunction that is true only if one or the other, but not both, of its disjuncts is true. When you believe the best interpretation of a disjunction is as an exclusive or, there are ways to represent that using a combination of the disjunction, conjunction and negation. The reason we interpret the wedge as an inclusive or rather than an exclusive or is that while we can build an exclusive or out of a combination of an inclusive or and other truth-functional connectives (as I’ve just pointed out), there is no way to build an inclusive or out of the exclusive or and other truth-functional connectives. We will see how to represent an exclusive or in section 2.5.

Exercise 9: Translate the following English sentences into our formal language using conjunction (the dot), negation (the tilde), or disjunction (the wedge). Use the suggested constants to stand for the atomic propositions.

1. Either Bob will mop or Tom will mop. (B = Bob will mop; T = Tom will mop)
2. It is not sunny today. (S = it is sunny today)
3. It is not the case that Bob is a burglar. (B = Bob is a burglar)
4. Harry is arriving either tonight or tomorrow night. (A = Harry is arriving tonight; B = Harry is arriving tomorrow night)
5. Gareth does not like his name. (G = Gareth likes his name)
6. Either it will not rain on Monday or it will not rain on Tuesday. (M = It will rain on Monday; T = It will rain on Tuesday)
7. Tom does not like cheesecake. (T = Tom likes cheesecake)
8. Bob would like to have both a large cat and a small dog as a pet. (C = Bob would like to have a large cat as a pet; D = Bob would like to have a small dog as a pet)
9. Bob Saget is not actually very funny. (B = Bob Saget is very funny)
10. Albert Einstein did not believe in God. (A = Albert Einstein believed in God)

2.4 Using parentheses to translate complex sentences

We have seen how to translate certain simple sentences into our symbolic language using the dot, wedge, and tilde. The process of translation starts with determining what the atomic propositions of the sentence are and then using the truth functional connectives to form the compound proposition. Sometimes this will be fairly straightforward and easy to figure out—especially if there is only one truth-functional operator used in the English sentence. However, many sentences will contain more than one truth-functional operator. Here is an example:

Bob will not go to class but will play video games.

What are the atomic propositions contained in this English sentence? Clearly, the sentence is asserting two things:

Bob will not go to class

Bob will play video games

The first statement is not an atomic proposition, since it contains a negation, "not." But the second statement is atomic since it does not contain any truth-functional connectives. So if the first statement is a negation, what is the non-negated, atomic statement? It is this:

Bob will go to class

I will use the constant C to represent this atomic proposition and G to represent the proposition, "Bob will play video games." Now that we have identified our two atomic propositions, how can we build our complex sentence using only those atomic propositions and the truth-functional connectives? Let's start with the statement "Bob will not go to class." Since we have defined the constant "C" as "Bob will go to class" then we can easily represent the statement "Bob will not go to class" using a negation, like this:

$\sim C$

The original sentence asserts that, but it is also asserts that Bob will play video games. That is, it is asserting both of these statements. That means we will be connecting " $\sim C$ " with "G" with the dot operator. Since we have already assigned "G" to the statement "Bob will play video games," the resulting translation should look like this:

$\sim C \cdot G$

Although sometimes we can translate sentences into our symbolic language without the use of parentheses (as we did in the previous example), many times a translation will require the use of parentheses. For example:

Bob will not both go to class and play video games.

Notice that whereas the earlier sentence asserted that Bob will not go to class, this sentence does not. Rather, it asserts that Bob will not do both things (i.e., go to class and play video games), but only one or the other (and possibly neither). That is, this sentence does not tell us for sure that Bob will/won't go to

class or that he will/won't play video games, but only that *he won't do both of these things*. Using the same translations as before, how would we translate this sentence? It should be clear that we cannot use the same translation as before since these two sentences are not saying the same thing. Thus, we cannot use the translation:

$$\sim C \cdot G$$

since that translation says for sure that Bob *will not* go to class and that he *will* play video games. Thus, our translation must be different. Here is how to translate the sentence:

$$\sim(C \cdot G)$$

I have here introduced some new symbols, the parentheses. Parentheses are using in formal logic to show groupings. In this case, the parentheses represent that the conjunction, " $C \cdot G$," is grouped together and the negation ranges over *that whole conjunction* rather than just the first conjunct (as was the case with the previous translation). When using multiple operators, you must learn to distinguish which operator is the main operator. The **main operator** of a sentence is the one that ranges over (influences) the whole sentence. In this case, the main operator is the negation, since it influences the truth value of all the rest of the sentence. In contrast, in the previous example ($\sim C \cdot G$), the main operator was the conjunction rather than the negation since it influences both parts of sentence (i.e., both the " $\sim C$ " and the " G "). We can see the need for parentheses in distinguishing these two different translations. Without the use of parentheses, we would have no way to distinguish these two sentences, which clearly have different meanings.

Here is a different example where we must utilize parentheses:

Noelle will either feed the dogs or clean her room, but she will not do the dishes.

Can you tell how many atomic propositions this sentence contains? It contains three atomic propositions which are:

Noelle will feed the dogs (F)

Noelle will clean her room (C)

Noelle will do the dishes (D)

What I've written in parentheses to the right of the statement is the constant that I'll use to represent these atomic statements in my symbolic translation. Notice that the sentence is definitely not asserting that each of these statements is true. Rather, what we have to do is use these atomic propositions to capture the meaning of the original English sentence using only our truth-functional operators. In this sentence we will actually use all three truth-functional operators (disjunction, conjunction, negation). Let's start with negation, as that one is relatively easy. Given how we have represented the atomic proposition, D, to say that Noelle will not do the dishes is simply the negation of D:

$$\sim D$$

Now consider the first part of the sentence: Noelle will either feed the dogs or clean her room. You should see the "either...or" there and recognize it as a disjunction, which we represent with the wedge, like this:

$$F \vee C$$

Now, how are these two compound propositions, " $\sim D$ " and " $F \vee C$ " themselves connected? There is one word in the sentence that tips you off—the "but." As we saw earlier, "but" is a common way of representing a conjunction in English. Thus, we have to conjoin the disjunction ($F \vee C$) and the negation ($\sim D$). You might think that we could simply conjoin the two propositions like this:

$$F \vee C \cdot \sim D$$

However, that translation would not be correct, because it is not what we call a well-formed formula. A **well-formed formula** is a sentence in our symbolic language that has exactly one interpretation or meaning. However, the translation we have given is ambiguous between two different meanings. It could mean that (Noelle will feed the dogs) or (Noelle will clean her room and not do the dishes). That statement would be *true* if Noelle fed the dogs and also did the dishes. We can represent this possibility symbolically, using parentheses like this:

$$F \vee (C \cdot \sim D)$$

The point of the parentheses is to group the main parts of the sentence together. In this case, we are grouping the “ $C \cdot \sim D$ ” together and leaving the “ F ” by itself. The result is that those groupings are connected by a disjunction, which is the main operator of the sentence. In this case, there are only two groupings: “ F ” on the one hand, and “ $C \cdot \sim D$ ” on the other hand.

But the original sentence could also mean that (Noelle will feed the dogs or clean her room) and (Noelle will not wash the dishes). In contrast with our earlier interpretation, this interpretation would be *false* if Noelle fed the dogs and did the dishes, since this interpretation asserts that Noelle will not do the dishes (as part of a conjunction). Here is how we would represent this interpretation symbolically:

$$(F \vee C) \cdot \sim D$$

Notice that this interpretation, unlike the last one, groups the “ $F \vee C$ ” together and leaves the “ $\sim D$ ” by itself. These two groupings are then connected by a conjunction, which is the main operator of this complex sentence.

The fact that our initial attempt at the translation (without using parentheses) yielded an ambiguous sentence shows the need for parentheses to disambiguate the different possibilities. Since our formal language aims at eliminating all ambiguity, we must choose one of the two groupings as the translation of our original English sentence. So, which grouping accurately captures the original sentence? It is the second translation that accurately captures the meaning of the original English sentence. That sentence clearly asserts that Noelle will not do the dishes and that is what our second translation says. In contrast, the first translation is a sentence that could be true even if Noelle *did* do the dishes. Given our understanding of the original English sentence, it should not be true under those circumstances since it clearly asserts that Noelle will not do the dishes.

Let’s move to a different example. Consider the sentence:

Either both Bob and Karen are washing the dishes or Sally and Tom are.

This sentence contains four atomic propositions:

Bob is washing the dishes (B)

Karen is washing the dishes (K)

Sally is washing the dishes (S)

Tom is washing the dishes (T)

As before, I've written the constants that I'll use to stand for each atomic proposition to the right of each atomic proposition. You can use any letter you'd like when coming up with your own translations, as long as each atomic proposition uses a different capital letter. (I typically try to pick letters that are distinctive of each sentence, such as picking "B" for "Bob".) So how can we use the truth functional operators to connect these atomic propositions together to yield a sentence that captures the meaning of the original English sentence? Clearly B and K are being grouped together with the conjunction "and" and S and T are also being grouped together with the conjunction "and" as well:

$(B \cdot K)$

$(S \cdot T)$

Furthermore, the main operator of the sentence is a disjunction, which you should be tipped off to by the phrase "either...or." Thus, the correct translation of the sentence is:

$(B \cdot K) \vee (S \cdot T)$

The main operator of this sentence is the disjunction (the wedge). Again, it is the main operator because it groups together the two main sentence groupings.

Let's finish this section with one final example. Consider the sentence:

Tom will not wash the dishes and will not help prepare dinner; however, he will vacuum the floor or cut the grass.

This sentence contains four atomic propositions:

Tom will wash the dishes (W)

Tom will help prepare dinner (P)

Tom will vacuum the floor (V)

Tom will cut the grass (C)

It is clear from the English (because of the “not”) that we need to negate both W and P. It is also clear from the English (because of the “and”) that W and P are grouped together. Thus, the first part of the translation should be:

$$(\sim W \cdot \sim P)$$

It is also clear that the last part of the sentence (following the semicolon) is a grouping of V and C and that those two propositions are connected by a disjunction (because of the word “or”):

$$(V \vee C)$$

Finally, these two groupings are connected by a conjunction (because of the “however,” which is a word that often functions as a conjunction). Thus, the correct translation of the sentence is:

$$(\sim W \cdot \sim P) \cdot (V \vee C)$$

As we have seen in this section, translating sentences from English into our symbolic language is a process that can be captured as a series of steps:

Step 1: Determine what the atomic propositions are.

Step 2: Pick a unique constant to stand for each atomic proposition.

Step 3: If the sentence contains more than two atomic propositions, determine which atomic propositions are grouped together and which truth-functional operator connects them.

Step 4: Determine what the main operator of the sentence is (i.e., which truth-functional operator connects the groups of atomic statements together).

Step 5: Once your translation is complete, read it back and see if it accurately captures what the original English sentence conveys. If not,

see if another way of grouping the parts together better captures what the original sentence conveys.

Try using these steps to create your own translations of the sentences in exercise 10 below.

Exercise 10: Translate the following English sentences into our symbolic language using any of the three truth functional operators (i.e., conjunction, negation, and disjunction). Use the constants at the end of each sentence to represent the atomic propositions they are obviously meant for. After you have translated the sentence, identify which truth-functional connective is the main operator of the sentence. (Note: not every sentence requires parentheses; a sentence requires parentheses only if it contains more than two atomic propositions.)

1. Bob does not know how to fly an airplane or pilot a ship, but he does know how to ride a motorcycle. (A, S, M)
2. Tom does not know how to swim or how to ride a horse. (S, H)
3. Theresa writes poems, not novels. (P, N)
4. Bob does not like Sally or Felicia, but he does like Alice. (S, F, A)
5. Cricket is not widely played in the United States, but both football and baseball are. (C, F, B)
6. Tom and Linda are friends, but Tom and Susan aren't—although Linda and Susan are. (T, S, L)
7. Lansing is east of Grand Rapids but west of Detroit. (E, W)
8. Either Tom or Linda brought David home after his surgery; but it wasn't Steve. (T, L, S)
9. Next year, Steve will be living in either Boulder or Flagstaff, but not Phoenix or Denver. (B, F, P, D)
10. Henry VII of England was married to Anne Boleyn and Jane Seymour, but he only executed Anne Boleyn. (A, J, E)
11. Henry VII of England executed either Anne Boleyn and Jane Boleyn or Thomas Cromwell and Thomas More. (A, J, C, M)
12. Children should be seen, but not heard. (S, H)

2.5 "Not both" and "neither nor"

Two common English phrases that can sometimes cause confusion are “not both” and “neither nor.” These two phrases have different meanings and thus are translated with different symbolic logic sentences. Let’s look at an example of each.

Carla will *not* have *both* cake and ice cream.

Carla will have *neither* cake *nor* ice cream.

The first sentence uses the phrase “not both” and the second “neither nor.” One way of figuring out what a sentence means (and thus how to translate it) is by asking the question: What scenarios does this sentence rule out? Let’s apply this to the “not both” statement (which we first saw back in the beginning of section 2.4). There are four possible scenarios, and the statement would be true in every one except the first scenario:

Carla has cake	Carla has ice cream	False
Carla has cake	Carla does not have ice cream	True
Carla does not have cake	Carla has ice cream	True
Carla does not have cake	Carla does not have ice cream	True

To say that Carla will not have both cake and ice cream allows that she can have one or the other (just not both). It also allows that she can have neither (as in the fourth scenario). So the way to think about the “not both” locution is as a negation of a conjunction, since the conjunction is the only scenario that cannot be true if the statement is true. If we use the constant “C” to represent the atomic sentence, “Carla has cake,” and “I” to represent “Carla has ice cream,” then the resulting symbolic translation would be:

$$\sim(C \cdot I)$$

Thus, in general, statements of the form “not both p and q” will be translated as the negation of a conjunction:

$$\sim(p \cdot q)$$

Note that the main operator of the statement is the negation. The negation applies to everything inside the parentheses—i.e., to the conjunction. This is very different from the following sentence (without parentheses):

$$\sim p \cdot q$$

The main operator of this statement is the conjunction and the left conjunct of the conjunction is a negation. In contrast with the “not both” form, this statement asserts that p is not true, while q is true. For example, using our previous example of Carla and the cake, the sentence

$$\sim C \cdot I$$

would assert that Carla will not have cake and will have ice cream. This is a very different statement from $\sim(C \cdot I)$ which, as we have seen, allows the possibility that Carla *will* have cake but *not* ice cream. Thus, again we see the importance of parentheses in our symbolic language.

Earlier (in section 2.3) we made the distinction between what I called an “exclusive or” and an “inclusive or” and I claimed that although we interpret the wedge (\vee) as an inclusive or, we can represent the exclusive or symbolically as well. Since we now know how to translate the “not both,” I can show you how to translate a statement that contains an exclusive or. Recall our example:

Bob placed either first or second in the race.

As we saw, this disjunction contains the two disjuncts, “Bob placed first in the race” (F) and “Bob placed second in the race” (S). Using the wedge, we get:

$$F \vee S$$

However, since the wedge is interpreted as an inclusive or, this statement would allow that Bob got both first and second in the race, which is not possible. So we need to be able to say that although Bob placed either first or second, *he did not place both first and second*. But that is just the “not both” locution. So, to be absolutely clear, we are asserting two things:

Bob placed either first or second.

and

Bob did not place both first and second.

We have already seen that the first sentence is translated: “ $F \vee S$.” The second sentence is simply a “not both F and S ” statement:

C	I	False
C	$\sim I$	False
$\sim C$	I	False

$\sim(F \cdot S)$

Now all we have to do is conjoin the two sentences using the dot:

$$(F \vee S) \cdot \sim(F \cdot S)$$

That is the correct translation of an exclusive or. Notice that when conjoining the “ $F \vee S$ ” to the “ $\sim(F \cdot S)$ ” I needed to put parentheses around the “ $F \vee S$ ” to show that it was grouped together. Thus, it would have been incorrect to write:

$$F \vee S \cdot \sim(F \cdot S)$$

since that is not a well-formed formula. The problem, as before, is that this sentence is ambiguous between two sentences that have different meanings:

$$F \vee (S \cdot \sim(F \cdot S))$$

$$(F \vee S) \cdot \sim(F \cdot S)$$

While both of these sentences are well-formed, only the latter is the correct translation of the exclusive or.

Let’s move on to the English locution “neither...nor” as in:

Carla will eat neither cake nor ice cream.

This statement might be true if, for example, Carla was on a diet (and was sticking to her diet). Using the same method I introduced earlier, we can ask under what conditions the statement would be true or false. As before, there are only four possibilities, which I represent symbolically this time:

$\sim C$	$\sim I$	True
----------	----------	-------------

There is only one circumstance in which this statement is true and that is the one in which it is false that Carla eats cake and false that Carla eats ice cream. That should be obvious from the meaning of the “neither nor” locution. Thus, the correct translation of a “neither nor” statement is as a conjunction of two negations:

$$\sim C \cdot \sim I$$

The main operator of this statement is the dot, which is conjoining the $\sim C$ with the $\sim I$. Thus, the form of any “neither nor” statement can always be translated as a conjunction of two negations:

$$\sim p \cdot \sim q$$

As we will see in a later section (where we will prove it), this statement is also equivalent to a negation of a disjunction:

$$\sim(p \vee q)$$

Thus, the English locution “neither nor” can also be translated using this statement form.

Exercise 11: For each of the following, write out what atomic proposition each constant stands for. Then translate the sentences using the constants you have defined. Finally, after you have translated the sentence, identify which truth-functional connective is the main operator of the sentence.

1. Coral is not both a plant and an animal. (P, A)

2. Although protozoa and chimpanzees are both eukaryotes, they are not both animals. (There are four atomic propositions here; just use A, B, C, and D for each different proposition.)
3. Neither chimpanzees nor protozoa are prokaryotes. (C, P)
4. China has not signed the Kyoto Protocol and neither has the United States. (C, U)
5. Either Chevrolet or McDonald's will support the Olympic team, but they won't both support it. (C, M)
6. Peter Jennings is either a liar or has a really bad memory. (L, M)
7. Peter Jennings is neither a liar nor has a really bad memory. (L, M)
8. Peter Jennings is both a liar and has a really bad memory. (L, M)
9. Peter Jennings is not both a liar and a person with a really bad memory. (L, M)
10. Chevrolet won't support the Olympic team this year, and McDonald's won't either. (C, M)
11. Mother Theresa may be a saint. Even so, she has not been canonized yet by the Catholic Church. (S, C)
12. The best distance runner of the last two decades is either Paul Tergat or Haile Gebrselassie, but it certainly isn't Jim Ryun. (T, G, R)
13. Jim Ryun was the best high school miler of all time, but he ran a slower time than Alan Webb. (R, W)
14. Neither Paul Tergat nor Haile Gebrselassie knows how to play hockey, but they both know how to play soccer. (A, B, C, D)
15. Ethiopians are neither good bobsledders nor tennis players, but they are excellent distance runners. (B, T, D)
16. Before Helen Keller met Annie Sullivan, she could neither speak, read, nor communicate. (S, R, C)
17. Although Helen Keller learned to communicate, she never learned to play soccer or baseball. (C, S, B)
18. Tom is allowed to play football or soccer, but not both. (F, S)
19. Tom will major in either engineering and physics, or business and sociology. (E, P, B, S)
20. Cartman is both xenophobic and racist, but he isn't a murderer or a thief. (X, R, M, T)

2.6 The truth table test of validity

So far, we have learned how to translate certain English sentences into our symbolic language, which consists of a set of constants (i.e., the capital letters that we use to represent different atomic propositions) and the truth-functional connectives. But what is the payoff of doing so? In this section we will learn what the payoff is. In short, the payoff will be that we will have a purely formal method of determining the validity of a certain class of arguments—namely, those arguments whose validity depends on the functioning of the truth-functional connectives. This is what logicians call “propositional logic” or “sentential logic.”

In the first chapter, we learned the informal test of validity, which required us to try to imagine a scenario in which the premises of the argument were true and yet the conclusion false. We saw that if we can imagine such a scenario, then the argument is invalid. On the other hand, if it is not possible to imagine a scenario in which the premises are true and yet the conclusion is false, then the argument is valid. Consider this argument:

1. The convict escaped either by crawling through the sewage pipes or by hiding out in the back of the delivery truck.
2. But the convict did not escape by crawling through the sewage pipes.
3. Therefore, the convict escaped by hiding out in the back of the delivery truck.

Using the informal test of validity, we can see that if we imagine that the first premise and the second premise are true, then the conclusion must follow. However, we can also prove this argument is valid without having to imagine scenarios and ask whether the conclusion would be true in those scenarios. We can do this by a) translating this sentence into our symbolic language and then b) using a truth table to determine whether the argument is valid. Let’s start with the translation. The first premise contains two atomic propositions. Here are the propositions and the constants that I’ll use to stand for them:

S = The convict escaped through the sewage pipes

D = The convict escaped by hiding out in the back of the delivery van

As we can see, the first premise is a disjunction and so, using the constants indicated above, we can translate that first premise as follows:

$$S \vee D$$

The second premise is simply the negation of S:

$$\sim S$$

Finally, the conclusion is simply the atomic sentence, D. Putting this all together in standard form, we have:

1. $S \vee D$
2. $\sim S$
3. $\therefore D$

We will use the symbol “ \therefore ” to denote a conclusion and will read it “therefore.”

The next thing we have to do is to construct a truth table. We have already seen some examples of truth tables when I defined the truth-functional connectives that I have introduced so far (conjunction, disjunction, and negation). A truth table (as we saw in section 2.2) is simply a device we use to represent how the truth value of a complex proposition depends on the truth of the propositions that compose it in every possible scenario. When constructing a truth table, the first thing to ask is how many atomic propositions need to be represented in the truth table. In this case, the answer is “two,” since there are only two atomic propositions contained in this argument (namely, S and D). Given that there are only two atomic propositions, our truth table will contain only four rows—one row for each possible scenario. There will be one row in which both S and D are true, one row in which both S and D are false, one row in which S is true and D is false, and one row in which S is false and D is true.

D	S	$S \vee D$	$\sim S$	D
T	T			
T	F			
F	T			
F	F			

The two furthest left columns are what we call the reference columns of the truth table. **Reference columns** assign every possible arrangement of truth values to the atomic propositions of the argument (in this case, just D and S). The reference columns capture every logically possible scenario. By doing so, we can replace having to use your imagination to imagine different scenarios (as in the informal test of validity) with a mechanical procedure that doesn't require us to imagine or even think very much at all. Thus, you can think of each row of the truth table as specifying one of the possible scenarios. That is, each row is one of the possible assignments of truth values to the atomic propositions. For example, row 1 of the truth table (the first row after the header row) is a scenario in which it is true that the convict escaped by hiding out in the back of the delivery van, and is also true that the convict escaped by crawling through the sewage pipes. In contrast, row 4 is a scenario in which the convict did neither of these things.

The next thing we need to do is figure out what the truth values of the premises and conclusion are for each row of the truth table. We are able to determine what those truth values are because we understand how the truth value of the compound proposition depends on the truth value of the atomic propositions. Given the meanings of the truth functional connectives (discussed in previous sections), we can fill out our truth table like this:

D	S	$S \vee D$	$\sim S$	D
T	T	T	F	T
T	F	T	T	T
F	T	T	F	F
F	F	F	T	F

To determine the truth values for the first premise of the argument (" $S \vee D$ ") we just have to know the truth values of S and D and the meaning of the truth functional connective, the disjunction. The truth table for the disjunction says that a disjunction is true as long as at least one of its disjuncts is true. Thus, every row under the " $S \vee D$ " column should be true, except for the last row since on the last row both D and S are false (whereas in the first three rows at least one or the other is true). The truth values for the second premise ($\sim S$) are easy to determine: we simply look at what we have assigned to "S" in our reference column and then we negate those truth values—the Ts becomes Fs and the Fs becomes Ts. That is just what I've done in the fourth column of the truth table above. Finally, the conclusion in the last column of the truth table will simply

repeat what we have assigned to “D” in our reference column, since the last conclusion simply repeats the atomic proposition “D.”

The above truth table is complete. Now the question is: How do we use this completed truth table to determine whether or not the argument is valid? In order to do so, we must apply what I’ll call the “truth table test of validity.” According to the **truth table test of validity**, an argument is **valid** if and only if for every assignment of truth values to the atomic propositions, if the premises are true then the conclusion is true. An argument is **invalid** if there exists an assignment of truth values to the atomic propositions on which the premises are true and yet the conclusion is false. It is imperative that you understand (and not simply memorize) what these definitions mean. You should see that these definitions of validity and invalidity have a similar structure to the informal definitions of validity and invalidity (discussed in chapter 1). The similarity is that we are looking for the possibility that the premises are true and yet the conclusion is false. If this is possible, then the argument is invalid; if it isn’t possible, then the argument is valid. The difference, as I’ve noted above, is that with the truth table test of validity, we replace having to use your imagination with a mechanical procedure of assigning truth values to atomic propositions and then determining the truth values of the premises and conclusion for each of those assignments.

Applying these definitions to the above truth table, we can see that the argument is valid because there is no assignment of truth values to the atomic propositions (i.e., *no row of our truth table*) on which all the premises are true and yet the conclusion is false. Look at the first row. Is that a row in which all the premises are true and yet the conclusion false? No, it isn’t, because not all the premises are true in that row. In particular, “ $\sim S$ ” is false in that row. Look at the second row. Is that a row in which all the premises are true and yet the conclusion false? No, it isn’t; although both premises are true in that row, the conclusion is also true in that row. Now consider the third row. Is that a row in which all the premises are true and yet the conclusion false? No, because it isn’t a row in which both the premises are true. Finally, consider the last row. Is that a row in which all the premises are true and yet the conclusion false? Again, the answer is “no” because the premises aren’t both true in that row. Thus, we can see that there is no row of the truth table in which the premises are all true and yet the conclusion is false. And that means the argument is valid.

Since the truth table test of validity is a formal method of evaluating an argument's validity, we can determine whether an argument is valid just in virtue of its form, without even knowing what the argument is about! Here is an example:

1. $(A \vee B) \vee C$
2. $\sim A$
3. $\therefore C$

Here is an argument written in our symbolic language. I don't know what A, B, and C mean (i.e., what atomic propositions they stand for), but it doesn't matter because we can determine whether the argument is valid without having to know what A, B, and C mean. A, B, and C could be any atomic propositions whatsoever. If this argument form is invalid then whatever meaning we give to A, B, and C, the argument will always be invalid. On the other hand, if this argument form is valid, then whatever meaning we give to A, B, and C, the argument will always be valid.

The first thing to recognize about this argument is that there are three atomic propositions, A, B, and C. And that means our truth table will have 8 rows instead of only 4 rows like our last truth table. The reason we need 8 rows is that it takes twice as many rows to represent every logically possible scenario when we are working with three different propositions. Here is a simple formula that you can use to determine how many rows your truth table needs:

$$2^n \text{ (where } n \text{ is the number of atomic propositions)}$$

You read this formula "two to the n-th power." So if you have one atomic proposition (as in the truth table for negation), your truth table will have only two rows. If you have two atomic propositions, it will have four rows. If you have three atomic propositions, it will have 8 rows. The number of rows needed grows exponentially as the number of atomic propositions grows linearly. The table below represents the same relationship that the above formula does:

Number of atomic propositions	Number of rows in the truth table
1	2
2	4
3	8
4	16
5	32

So, our truth table for the above argument needs to have 8 rows. Here is how that truth table looks:

A	B	C	$(A \vee B) \vee C$	$\sim A$	C
T	T	T			
T	T	F			
T	F	T			
T	F	F			
F	T	T			
F	T	F			
F	F	T			
F	F	F			

Here is an important point to note about setting up a truth table. You need to make sure that your reference columns capture each distinct possible assignment of truth values. One way to make sure you do this is by following the same pattern each time you construct a truth table. There is no one right way of doing this, but here is how I do it (and recommend that you do it too). Construct the reference columns so that the atomic propositions are arranged alphabetically, from left to right. Then on the right-most reference column (the C column above), alternate true and false each row, all the way to the bottom. On the reference column to the left of that (the B column above), alternate two rows true, two rows false, all the way to the bottom. On the next column to the left (the A column above), alternate 4 true, 4 false, all the way to the bottom.

The next step is to determine the truth values of the premises and conclusion. Note that our first premise is a more complex sentence that consists of two

disjunctions. The main operator is the second disjunction since the two main groupings, denoted by the parentheses, are “ $A \vee B$ ” and “ C ”. Notice, however, that we cannot figure out the truth values of the main operator of the sentence until we figure out the truth values of the left disjunct, “ $A \vee B$.” So that is where we need to start. Thus, in the truth table below, I have filled out the truth values directly underneath the “ $A \vee B$ ” part of the sentence by using the truth values I have assigned to A and B in the reference columns. As you can see in the truth table below, each line is true except for the last two lines, which are false, since a disjunction is only false when both of the disjuncts are false. (If you need to review the truth table for disjunction, please see section 2.3.)

A	B	C	$(A \vee B) \vee C$	$\sim A$	C
T	T	T	T		
T	T	F	T		
T	F	T	T		
T	F	F	T		
F	T	T	T		
F	T	F	T		
F	F	T	F		
F	F	F	F		

Now, since we have figured out the truth values of the left disjunct, we can figure out the truth values under the main operator (which I have emphasized in bold in the truth table below). The two columns you are looking at to determine the truth values of the main operator are the “ $A \vee B$ ” column that we have just figured out above and the “ C ” reference column to the left. It is imperative to understand that the truth values under the “ $A \vee B$ ” are irrelevant once we have figured out the truth values under the main operator of the sentence. That column was only a means to an end (the end of determining the main operator) and so I have grayed those out to emphasize that we are no longer paying any attention to them. (When you are constructing your own truth tables, you may even want to erase these subsidiary columns once you’ve determined the truth values of the main operator of the sentence. Or you may simply want to circle the truth values under the main operator to distinguish them from the rest.)

A	B	C	$(A \vee B) \vee C$	$\sim A$	C
T	T	T	T T		
T	T	F	T T		
T	F	T	T T		
T	F	F	T T		
F	T	T	T T		
F	T	F	T T		
F	F	T	F T		
F	F	F	F F		

Finally, we will fill out the remaining two columns, which is very straightforward. All we have to do for the " $\sim A$ " is negate the truth values that we have assigned to our "A" reference column. And all we have to do for the final column "C" is simply repeat verbatim the truth values that we have assigned to our reference column "C."

A	B	C	$(A \vee B) \vee C$	$\sim A$	C
T	T	T	T T	F	T
T	T	F	T T	F	F
T	F	T	T T	F	T
T	F	F	T T	F	F
F	T	T	T T	T	T
F	T	F	T T	T	F
F	F	T	F T	T	T
F	F	F	F F	T	F

The above truth table is now complete. The next step is to apply the truth table test of validity in order to determine whether the argument is valid or invalid. Remember that what we're looking for is a row in which the premises are true and the conclusion is false. If we find such a row, the argument is invalid. If we do not find such a row, then the argument is valid. Applying this definition to the above truth table, we can see that the argument is invalid because of the 6th row of the table (which I have highlighted). Thus, the explanation of why this argument is invalid is that the sixth row of the table shows a scenario in which the premises are both true and yet the conclusion is false.

Exercise 12: Use the truth table test of validity to determine whether or not the following arguments are valid or invalid.

- | | | |
|---|---|---|
| <p>1.</p> <p>1. $A \vee B$</p> <p>2. B</p> <p>3. $\therefore \sim A$</p> | <p>4.</p> <p>1. $(A \vee B) \cdot (A \vee C)$</p> <p>2. $\sim A$</p> <p>3. $\therefore B \vee C$</p> | <p>7.</p> <p>1. $\sim(A \cdot B)$</p> <p>2. $\therefore \sim A \vee \sim B$</p> |
| <p>2.</p> <p>1. $A \cdot B$</p> <p>2. $\therefore A \vee B$</p> | <p>5.</p> <p>1. $R \cdot (T \vee S)$</p> <p>2. T</p> <p>3. $\therefore \sim S$</p> | <p>8.</p> <p>1. $\sim(A \vee B)$</p> <p>2. $\therefore \sim A \vee \sim B$</p> |
| <p>3.</p> <p>1. $\sim C$</p> <p>2. $\therefore \sim(C \vee A)$</p>
<p>1. $(R \vee S) \cdot \sim D$</p> <p>2. $\sim R$</p> <p>3. $\therefore S \cdot \sim D$</p> | <p>6.</p> <p>1. $A \vee B$</p> <p>2. $\therefore A \cdot B$</p> | <p>9.</p> |

2.7 Conditionals

So far, we have learned how to translate and construct truth tables for three truth functional connectives. However, there is one more truth functional connective that we have not yet learned: the conditional.² The English phrase that is most often used to express conditional statements is “if...then.” For example,

If it is raining then the ground it wet.

Like conjunctions and disjunctions, conditionals connect two atomic propositions. There are two atomic propositions in the above conditional:

It is raining.

² Actually, there is one more truth functional connective that we will not be learning and that is what is called the “biconditional” or “material equivalence.” However, since the biconditional is equivalent to a conjunction of two different conditionals, we don’t actually need it. Although I will discuss material equivalence in section 2.9, we will not be regularly using it.

The ground is wet.

The proposition that follows the “if” is called the **antecedent** of the conditional and the proposition that follows the “then” is called the **consequent** of the conditional. The conditional statement above is not asserting either of these atomic propositions. Rather, it is telling us about the relationship between them. Let’s symbolize “it is raining” as “R” and “the ground is wet” as “G.” Thus, our symbolization of the above conditional would be:

$$R \supset G$$

The “ \supset ” symbol is called the “horseshoe” and it represents what is called the “material conditional.” A material conditional is defined as being true in every case *except* when the antecedent is true and the consequent is false. Below is the truth table for the material conditional. Notice that, as just stated, there is only one scenario in which we count the conditional false: when the antecedent is true and the consequent false.

p	q	$p \supset q$
T	T	T
T	F	F
F	T	T
F	F	T

Let’s see how this applies to the above conditional, “if it is raining, then the ground is wet.” As before, we can think about the meaning of the truth functional connectives by asking whether the sentences containing those connectives would be true or false in the four possible scenarios. The first two are pretty easy. If I assert the above conditional “if it is raining then the ground is wet” when it is both raining and the ground is wet (i.e., the first line of the truth table below), then the conditional statement would be true in that scenario. However, if I assert it and it is raining but the ground isn’t wet (i.e., the second line of the truth table below), then my statement has been shown to be false. Why? Because I’m asserting that any time it is raining, the ground is wet. But if it is raining but the ground isn’t wet, then this scenario is a counterexample to my claim—it shows that my claim is false. Now consider the scenario in which it is not raining but the ground is wet. Would this scenario show that my conditional statement is false? No, it wouldn’t. The reason is that

the conditional statement $R \supset G$ is only asserting something about what is the case when it is raining. So this conditional statement isn't asserting anything about those scenarios in which it *isn't* raining. I'm only saying that *when it is raining*, the ground is wet. But that doesn't mean that the ground couldn't be wet for other reasons (e.g., a sprinkler watering the grass). So the meaning of the material conditional should count a statement true whenever its antecedent is false. Thus, in a scenario in which it is neither raining nor the ground is wet (i.e., the fourth line of the truth table), the conditional statement should still be true. Would the fact of a sunny day and dry ground show that the conditional $R \supset G$ is false? Of course not! Thus, as we've seen, the material conditional is false only when the antecedent is true and the consequent is false.

R	G	$R \supset G$
T	T	T
T	F	F
F	T	T
F	F	T

It is sometimes helpful to think of the material conditional as a rule. For example, suppose that I tell my class:

If you pass all the exams, you will pass the course.

Let's symbolize "you pass all the exams" as "E" and "you pass the course" as "C." We would then symbolize the conditional as:

$E \supset C$

Under what conditions would my statement $E \supset C$ be shown to be false? There are four possible scenarios:

E	C	$E \supset C$
T	T	T
T	F	F
F	T	T
F	F	T

Suppose that you pass all the exams and pass the class (first row). That would confirm my conditional statement $E \supset C$. Suppose, on the other hand, that although you passed all the exams, you did not pass the class (second row). This *would* should my statement is false (and you would have legitimate grounds for complaint!). How about if you don't pass all the exams and yet you do pass the course (third row)? My statement allows this to be true and it is important to see why. When I assert $E \supset C$ I am not asserting anything about the situation in which E is false. I am simply saying that one way of passing the course is by passing all of the exams; but that doesn't mean there aren't other ways of passing the course. Finally, consider the case in which you do not pass all the exams and you also do not pass the course (fourth row). For the same reason, this scenario is compatible with my statement being true. Thus, again, we see that a material conditional is false in only one circumstance: when the antecedent is true and the consequent is false.

There are other English phrases that are commonly used to express conditional statements. Here are some equivalent ways of expressing the conditional, "if it is raining then the ground is wet":

It is raining *only if* the ground is wet

The ground is wet *if* it is raining

Only if the ground is wet is it raining

That it is raining *implies that* the ground is wet

That it is raining *entails that* the ground is wet

As long as it is raining, the ground will be wet

So long as it is raining, the ground will be wet

The ground is wet, *provided that* it is raining

Whenever it is raining, the ground is wet

If it is raining, the ground is wet

All of these conditional statements are symbolized the same way, namely $R \supset G$. The antecedent of a conditional statement always lays down what logicians call a sufficient condition. A **sufficient condition** is a condition that suffices for some other condition to obtain. To say that x is a sufficient condition for y is to say that any time x is present, y will thereby be present. For example, a sufficient condition for dying is being decapitated; a sufficient condition for being a U.S. citizen is being born in the U.S. The consequent of a conditional statement always lays down a necessary condition. A **necessary condition** is a condition that must be in present in order for some other condition to obtain. To say that x is a necessary condition for y is to say that if x were not present, y would not be present either. For example, a necessary condition for being President of the U.S. is being a U.S. citizen; a necessary condition for having a brother is having a sibling. Notice, however, that being a U.S. citizen is not a sufficient condition for being President, and having a sibling is not a sufficient condition for having a brother. Likewise, being born in the U.S. is not a necessary condition for being a U.S. citizen (people can become “naturalized citizens”), and being decapitated is not a necessary condition for dying (one can die without being decapitated).

Exercise 13: Translate the following English sentences into symbolic logic sentences using the constants indicated. Make sure you write out what the atomic propositions are. In some cases this will be straightforward, but not in every case. Remember: atomic propositions never contain any truth functional connectives—and that includes negation! Note: although many of these sentences can be translated using only the horseshoe, others require truth functional connectives other than the horseshoe.

1. The Tigers will win only if the Indians lose their star pitcher. (T, I)
2. Tom will pass the class provided that he does all the homework. (P, H)
3. The car will run only if it has gas. (R, G)
4. The fact that you are asking me about your grade implies that you care about your grade. (A, C)
5. Although Frog will swim without a bathing suit, Toad will swim only if he is wearing a bathing suit. (F, T, B)
6. If Obama isn't a U.S. citizen, then I'm a monkey's uncle. (O, M)
7. If Toad wears his bathing suit, he doesn't want Frog to see him in it. (T, F)
8. If Tom doesn't pass the exam, then he is either stupid or lazy. (P, S, L)

9. Bekele will win the race as long as he stays healthy. (W, H)
10. If Bekele is either sick or injured, he will not win the race. (S, I, W)
11. Bob will become president only if he runs a good campaign and doesn't say anything stupid. (P, C, S)
12. If that plant has three leaves then it is poisonous. (T, P)
13. The fact that the plant is poisonous implies that it has three leaves. (T, P)
14. The plant is poisonous only if it has three leaves. (T, P)
15. The plant has three leaves if it is poisonous. (T, P)
16. Olga will swim in the open water as long as there is a shark net present. (O, N)
17. Olga will swim in the open water only if there is shark net. (O, N)
18. The fact that Olga is swimming implies that she is wearing a bathing suit. (O, B)
19. If Olga is in Nice, she does not wear a bathing suit. (N, B)
20. If Terrence pulls Philip's finger, something bad will happen. (T, B)

2.8 "Unless"

The English term "unless" can be tricky to translate. For example,

The Reds will win unless their starting pitcher is injured.

If we use the constant "R" to stand for the atomic proposition, "the Reds will win" and "S" to stand for the atomic proposition, "the Reds' starting pitcher is injured," how would we translate this sentence using truth functional connectives? Think about what the sentence is saying (think carefully). Is the sentence asserting that the Reds will win? No; it is only saying that

The Reds will win as long as their starting pitcher isn't injured.

"As long as" denotes a conditional statement. In particular, what follows the "as long as" phrase is a sufficient condition, and as we have seen, a sufficient condition is always the antecedent of a conditional. But notice that the sufficient condition also contains a negation. Thus, the correct translation of this sentence is:

$$\sim S \supset R$$

One simple trick you can use to translate sentences which use the term “unless” is just substitute the phrase “if it’s not the case that” for the “unless.” But another trick is just to substitute an “or” for the “unless.” Although it may sound strange in English, a disjunction will always capture the truth functional meaning of “unless.” Thus, we could also correctly translate the sentence like this:

$$S \vee R$$

In the next section we will show how we can prove that these two sentences are equivalent using a truth table.

2.9 Material equivalence

As we saw in the last section, two different symbolic sentences can translate the same English sentence. In the last section I claimed that “ $\sim S \supset R$ ” and “ $S \vee R$ ” are equivalent. More precisely, they are equivalent ways of capturing the truth-functional relationship between propositions. Two propositions are **materially equivalent** if and only if they have the same truth value for every assignment of truth values to the atomic propositions. That is, they have the same truth values on every row of a truth table. The truth table below demonstrates that “ $\sim S \supset R$ ” and “ $S \vee R$ ” are materially equivalent.

R	S	$\sim S \supset R$	$S \vee R$
T	T	F T	T
T	F	T T	T
F	T	F T	T
F	F	T F	F

If you look at the truth values under the main operators of each sentence, you can see that their truth values are identical on every row. That means the two statements are materially equivalent and can be used interchangeably, as far as propositional logic goes.

Let's demonstrate material equivalence with another example. We have seen that we can translate "neither nor" statements as a conjunction of two negations. So, a statement of the form, "neither p nor q" can be translated:

$$\sim p \cdot \sim q$$

But another way of translating statements of this form is as a negation of a disjunction, like this:

$$\sim(p \vee q)$$

We can prove these two statements are materially equivalent with a truth table (below).

p	q	$\sim p \cdot \sim q$	$\sim(p \vee q)$
T	T	F F F	F T
T	F	F F T	F T
F	T	T F F	F T
F	F	T T T	T F

Again, as you can see from the truth table, the truth values under the main operators of each sentence are identical on every row (i.e., for every assignment of truth values to the atomic propositions).

In fact, there is a fifth truth functional connective called "material equivalence" or the "biconditional" that is defined as true when the atomic propositions share the same truth value, and false when the truth values differ. Although we will not be relying on the biconditional, I provide the truth table for it below. The biconditional is represented using the symbol " \equiv " which is called a "tribar."

p	q	$p \equiv q$
T	T	T
T	F	F
F	T	F
F	F	T

Some common ways of expressing the biconditional in English are with the phrases “if and only if” and “just in case.” If you have been paying close attention (or do from now on out) you will see me use the phrase “if and only if” often. It is most commonly used when one is giving a definition, such as the definition of validity and also in defining the “material equivalence” in this very section. It makes sense that the biconditional would be used in this way since when we define something we are laying down an equivalent way of saying it.

Exercise 14: Construct a truth table to determine whether the following pairs of statements are materially equivalent.

1. $A \supset B$ and $\sim A \vee B$
2. $\sim(A \cdot B)$ and $\sim A \vee \sim B$
3. $A \supset B$ and $\sim B \supset \sim A$
4. $A \vee \sim B$ and $B \supset A$
5. $B \supset A$ and $A \supset B$
6. $\sim(A \supset B)$ and $A \cdot \sim B$
7. $A \vee B$ and $\sim A \cdot \sim B$
8. $A \vee (B \cdot C)$ and $(A \vee B) \cdot (A \vee C)$
9. $(A \vee B) \cdot C$ and $A \vee (B \cdot C)$
10. $\sim(A \vee B)$ and $\sim A \vee \sim B$

2.10 Tautologies, contradictions and contingent statements

Can you think of a statement that could never be false? How about a statement that could never be true? It is harder than you think, unless you know how to utilize the truth functional operators to construct a tautology or a contradiction. A **tautology** is a statement that is true in virtue of its form. Thus, we don't even have to know what the statement means to know that it is true. In contrast, a **contradiction** is a statement that is false in virtue of its form. Finally, a contingent statement is a statement whose truth depends on the way the world actually is. Thus, it is a statement that could be either true or false—it just depends on what the facts actually are. In contrast, there is an important sense in which the truth of a tautology or the falsity of a contradiction doesn't depend on how the world is. As philosophers would say, tautologies are true in every

possible world, whereas contradictions are false in every possible world. Consider a statement like:

Matt is either 40 years old or not 40 years old.

That statement is a tautology, and it has a particular form, which can be represented symbolically like this:

$$p \vee \sim p$$

In contrast, consider a statement like:

Matt is both 40 years old and not 40 years old.

That statement is a contradiction, and it has a particular form, which can be represented symbolically like this:

$$p \cdot \sim p$$

Finally, consider a statement like:

Matt is either 39 years old or 40 years old

That statement is a **contingent statement**. It doesn't have to be true (as tautologies do) or false (as contradictions do). Instead, its truth depends on the way the world is. Suppose that Matt is 39 years old. In that case, the statement is true. But suppose he is 37 years old. In that case, the statement is false (since he is neither 39 or 40). We can use truth tables to determine whether a statement is a tautology, contradiction or contingent statement. In a tautology, the truth table will be such that every row of the truth table under the main operator will be true. In a contradiction, the truth table will be such that every row of the truth table under the main operator will be false. And contingent statements will be such that there is mixture of true and false under the main operator of the statement.

The following two truth tables are examples of tautologies and contradictions, respectively.

A	B	$(A \supset B) \vee A$
T	T	T T
T	F	F T
F	T	T T
F	F	T T

A	B	$(A \vee B) \cdot (\sim A \cdot \sim B)$
T	T	T F F F F
T	F	T F F F T
F	T	T F T F F
F	F	F F T F T

Notice that in the second truth table, I had to do quite a lot of work before I could figure out what the truth values of the main operator were. I had to first determine the left conjunct ($A \vee B$) and then the right conjunct ($\sim A \cdot \sim B$), but in order to figure out the truth values of the right conjunct (which is itself a conjunct), I had to determine the negations of A and B. Constructing truth tables can sometimes be a chore, but once you understand what you are doing (and why), it certainly isn't very difficult.

Exercise 15: Construct a truth table to determine whether the following statements are tautologies, contradictions or contingent statements.

1. $A \supset (A \cdot B)$
2. $(A \cdot B) \supset (\sim A \supset \sim B)$
3. $(A \cdot \sim A) \supset B$
4. $(A \supset A) \supset (B \cdot \sim B)$
5. $(A \cdot B) \supset (A \vee B)$
6. $(A \vee B) \supset (A \cdot B)$
7. $(\sim A \supset \sim B) \supset (\sim B \supset \sim A)$
8. $(A \supset B) \supset (\sim B \supset \sim A)$
9. $(B \vee \sim B) \supset A$
10. $(A \vee B) \vee \sim A$

2.11 Proofs and the 8 valid forms of inference

Although truth tables are our only formal method of deciding whether an argument is valid or invalid in propositional logic, there is another formal method of proving that an argument *is* valid: the method of proof. Although you cannot construct a proof to show that an argument is invalid, you can construct proofs to show that an argument is valid. The reason proofs are helpful, is that they allow us to show that certain arguments are valid much more efficiently than do truth tables. For example, consider the following argument:

1. $(R \vee S) \supset (T \supset K)$
2. $\sim K$
3. $R \vee S \quad \therefore \sim T$

(Note: in this section I will be writing the conclusion of the argument to the right of the last premise—in this case premise 3. As before, the conclusion we are trying to derive is denoted by the “therefore” sign, “ \therefore ”.) We could attempt to prove this argument is valid with a truth table, but the truth table would be 16 rows long because there are four different atomic propositions that occur in this argument, R , S , T , and K . If there were 5 or 6 different atomic propositions, the truth table would be 32 or 64 lines long! However, as we will soon see, we could also prove this argument is valid with only two additional lines. That seems a much more efficient way of establishing that this argument is valid. We will do this a little later—after we have introduced the 8 valid forms of inference that you will need in order to do proofs. Each line of the proof will be justified by citing one of these rules, with the last line of the proof being the conclusion that we are trying to ultimately establish. I will introduce the 8 valid forms of inference in groups, starting with the rules that utilize the horseshoe and negation.

The first of the 8 forms of inference is “**modus ponens**” which is Latin for “way that affirms.” Modus ponens has the following form:

1. $p \supset q$
2. p
3. $\therefore q$

What this form says, in words, is that if we have asserted a conditional statement ($p \supset q$) and we have also asserted the antecedent of that conditional statement

(p), then we are entitled to infer the consequent of that conditional statement (q). For example, if I asserted the conditional, “if it is raining, then the ground is wet” and I also asserted “it is raining” (the antecedent of that conditional) then I (or anyone else, for that matter) am entitled to assert the consequent of the conditional, “the ground is wet.”

As with any valid forms of inference in this section, we can prove that modus ponens is valid by constructing a truth table. As you see from the truth table below, this argument form passes the truth table test of validity (since there is no row of the truth table on which the premises are all true and yet the conclusion is false).

p	q	$p \supset q$	p	q
T	T	T	T	T
T	F	F	T	F
F	T	T	F	T
F	F	T	F	F

Thus, any argument that has this same form is valid. For example, the following argument also has this same form (modus ponens):

1. $(A \cdot B) \supset C$
2. $(A \cdot B)$
3. $\therefore C$

In this argument we can assert C according to the rule, modus ponens. This is so even though the antecedent of the conditional is itself complex (i.e., it is a conjunction). That doesn't matter. The first premise is still a conditional statement (since the horseshoe is the main operator) and the second premise is the antecedent of that conditional statement. The rule modus ponens says that if we have that much, we are entitled to infer the consequent of the conditional.

We can actually use modus ponens in the first argument of this section:

1. $(R \vee S) \supset (T \supset K)$
2. $\sim K$
3. $R \vee S \quad \therefore \sim T$

4. $T \supset K$ Modus ponens, lines 1, 3

What I have done here is I have written the valid form of inference (or rule) that justifies the line I am deriving, as well as the lines to which that rule applies, to the right of the new line of the proof that I am deriving. Here I have derived " $T \supset K$ " from lines 1 and 3 of the argument by modus ponens. Notice that line 1 is a conditional statement and line 3 is the antecedent of that conditional statement. This proof isn't finished yet, since we have not yet derived the conclusion we are trying to derive, namely, " $\sim T$." We need a different rule to derive that, which we will introduce next.

The next form of inference is called "**modus tollens**," which is Latin for "the way that denies." Modus tollens has the following form:

1. $p \supset q$
2. $\sim q$
3. $\therefore \sim p$

What this form says, in words, is that if we have asserted a conditional statement ($p \supset q$) and we have also asserted the negated consequent of that conditional ($\sim q$), then we are entitled to infer the negated antecedent of that conditional statement ($\sim p$). For example, if I asserted the conditional, "if it is raining, then the ground is wet" and I also asserted "the ground is not wet" (the negated consequent of that conditional) then I am entitled to assert the negated antecedent of the conditional, "it is not raining." It is important to see that *any* argument that has *this same form* is a valid argument. For example, the following argument is also an argument with this same form:

1. $C \supset (E \vee F)$
2. $\sim(E \vee F)$
3. $\therefore \sim C$

In this argument we can assert $\sim C$ according to the rule, modus tollens. This is so even though the consequent of the conditional is itself complex (i.e., it is a disjunction). That doesn't matter. The first premise is still a conditional statement (since the horseshoe is the main operator) and the second premise is the negated consequent of that conditional statement. The rule modus tollens says that if we have that much, we are entitled to infer the negated antecedent of the conditional.

We can use modus tollens to complete the proof we started above:

1. $(R \vee S) \supset (T \supset K)$
2. $\sim K$
3. $R \vee S \quad \therefore \sim T$
4. $T \supset K \quad \text{Modus ponens, lines 1, 3}$
5. $\sim T \quad \text{Modus tollens, lines 2, 4}$

Notice that the last line of the proof is the conclusion that we are supposed to derive and that each statement that I have derived (i.e., lines 4 and 5) has a rule to the right. That rule cited is the rule that justifies the statement that is being derived and the lines cited are the previous lines of the proof where we can see that the rule applies. This is what is called a proof. A **proof** is a series of statements, starting with the premises and ending with the conclusion, where each additional statement after the premises is derived from some previous line(s) of the proof using one of the valid forms of inference. We will practice this some more in the exercise at the end of this section.

The next form of inference is called "**hypothetical syllogism.**" This is what ancient philosophers called "the chain argument" and it should be obvious why in a moment. Here is the form of the rule:

1. $p \supset q$
2. $q \supset r$
3. $\therefore p \supset r$

As you can see, the conclusion of this argument links p and r together in a conditional statement. We could continue adding conditionals such as " $r \supset s$ " and " $s \supset t$ " and the inferences would be just as valid. And if we lined them all up as I have below, you can see why ancient philosophers referred to this valid argument form as a "chain argument":

$$\begin{array}{l}
 p \supset q \\
 \quad q \supset r \\
 \quad \quad r \supset s \\
 \quad \quad \quad s \supset t \\
 \quad \quad \quad \therefore p \supset t
 \end{array}$$

Notice how the consequent of each preceding conditional statement links up with the antecedent of the next conditional statement in such a way as to create a chain. The chain could be as long as we liked, but the rule that we will cite in our proofs only connects two different conditional statements together. As before, it is important to realize that any argument with this same form is a valid argument. For example,

1. $(A \vee B) \supset \sim D$
2. $\sim D \supset C$
3. $\therefore (A \vee B) \supset C$

Notice that the consequent of the first premise and the antecedent of the second premise are exactly the same term, " $\sim D$ ". That is what allows us to "link" the antecedent of the first premise and the consequent of the second premise together in a "chain" to infer the conclusion. Being able to recognize the forms of these inferences is an important skill that you will have to become proficient at in order to do proofs.

The next four forms of inference we will introduce utilize conjunction, disjunction and negation in different ways. We will start with the rule called "**simplification**," which has the following form:

1. $p \cdot q$
2. $\therefore p$

What this rule says, in words, is that if we have asserted a conjunction then we are entitled to infer either one of the conjuncts. This is the rule that I introduced in the first section of this chapter. It is a pretty "obvious" rule—so obvious, in fact, that we might even wonder why we have to state it. However, every form of inference that we will introduce in this section should be obvious—that is the point of calling them *basic* forms of inference. They are some of the simplest forms of inference, whose validity should be transparently obvious. The idea of a proof is that although the inference being made in the argument is *not* obvious, we can break that inference down in steps, each of which *is* obvious. Thus, the obvious inferences ultimately justify the non-obvious inference being made in the argument. Those obvious inferences thus function as *rules* that we use to justify each step of the proof. Simplification is a prime example of one of the more obvious rules.

As before, it is important to realize that any inference that has the same form as simplification is a valid inference. For example,

1. $(A \vee B) \cdot \sim(C \cdot D)$
2. $\therefore (A \vee B)$

is a valid inference because it has the same form as simplification. That is, line 1 is a conjunction (since the dot is the main operator of the sentence) and line 2 is inferring one of the conjuncts of that conjunction in line 1. (Just think of the “ $A \vee B$ ” as the “ p ” and the “ $\sim(C \cdot D)$ ” as the “ q ”.)

The next rule we will introduce is called “**conjunction**” and is like the reverse of simplification. (Don’t confuse the *rule* called conjunction with the *type of complex proposition* called a conjunction.) Conjunction has the following form:

1. p
2. q
3. $\therefore p \cdot q$

What this rule says, in words, is that if you have asserted two different propositions, then you are entitled to assert the conjunction of those two propositions. As before, it is important to realize that any inference that has the same form as conjunction is a valid inference. For example,

1. $A \supset B$
2. $C \vee D$
3. $\therefore (A \supset B) \cdot (C \vee D)$

is a valid inference because it has the same form as conjunction. We are simply conjoining two propositions together; it doesn’t matter whether those propositions are atomic or complex. In this case, of course, the propositions we are conjoining together are complex, but as long as those propositions have already been asserted as premises in the argument (or derived by some other valid form of inference), we can conjoin them together into a conjunction.

The next form of inference we will introduce is called “**disjunctive syllogism**” and it has the following form:

1. $p \vee q$

2. $\sim p$
3. $\therefore q$

In words, this rule states that if we have asserted a disjunction and we have asserted the negation of one of the disjuncts, then we are entitled to assert the other disjunct. Once you think about it, this inference should be pretty obvious. If we are taking for granted the truth of the premises—that either p or q is true; and that p is not true—then it *has to follow* that q is true in order for the original disjunction to be true. (Remember that we *must* assume the premises are true when evaluating whether an argument is valid.) If I assert that it is true that either Bob or Linda stole the diamond, and I assert that Bob did not steal the diamond, then it has to follow that Linda did. That is a disjunctive syllogism. As before, any argument that has this same form is a valid argument. For example,

1. $\sim A \vee (B \cdot C)$
2. $\sim \sim A$
3. $\therefore B \cdot C$

is a valid inference because it has the same form as disjunctive syllogism. The first premise is a disjunction (since the wedge is the main operator), the second premise is simply the negation of the left disjunct, " $\sim A$ ", and the conclusion is the right disjunct of the original disjunction. It may help you to see the form of the argument if you treat " $\sim A$ " as the p and " $B \cdot C$ " as the q . Also notice that the second premise contains a double negation. Your English teacher may tell you never to use double negatives, but as far as logic is concerned, there is absolutely nothing wrong with a double negation. In this case, our left disjunct in premise 1 is itself a negation, while premise 2 is simply a negation of that left disjunct.

The next rule we'll introduce is called "**addition**." It is not quite as "obvious" a rule as the ones we've introduced above. However, once you understand the conditions under which a disjunction is true, then it should be obvious why this form of inference is valid. Addition has the following form:

1. p
2. $\therefore p \vee q$

What this rule says, in words, is that if we have asserted some proposition, p , then we are entitled to assert the disjunction of that proposition p and *any*

other proposition q we wish. Here's the simple justification of the rule. If we know that p is true, and a disjunction is true if at least one of the disjuncts is true, then we know that the disjunction $p \vee q$ is true even if we don't know whether q is true or false. Why? Because *it doesn't matter whether q is true or false, since we already know that p is true.* The hardest thing to understand about this rule is why we would ever want to use it. The best answer I can give you for that right now is that it can help us out when doing proofs.³

As before, is it important to realize that any argument that has this same form, is a valid argument. For example,

1. $A \vee B$
2. $\therefore (A \vee B) \vee (\sim C \vee D)$

is a valid inference because it has the same form as addition. The first premise asserts a statement (which in this case is complex—a disjunction) and the conclusion is a disjunction of that statement and *some other statement*. In this case, that other statement is itself complex (a disjunction). But an argument or inference can have the same form, regardless of whether the components of those sentences are atomic or complex. That is the important lesson that I have been trying to drill in in this section.

The final of our 8 valid forms of inference is called "**constructive dilemma**" and is the most complicated of them all. It may be most helpful to introduce it using an example. Suppose I reasoned thus:

The killer is either in the attic or the basement. If the killer is in the attic then he is above me. If the killer is in the basement then his is below me. Therefore, the killer is either _____ or _____.

Can you fill in the blanks with the phrases that would make this argument valid? I'm guessing that you can. It should be pretty obvious. The conclusion of the argument is the following:

³ A better answer is that we need this rule in order to make this set of rules that I am presenting a sound a complete set of rules. That is, without it there would be arguments that are valid but that we aren't able to show are valid using this set of rules. In more advanced areas of logic, such as metalogic, logicians attempt to prove things about a particular system of logic, such as proving that the system is sound and complete.

The killer is either above me or below me.

That this argument is valid should be obvious (can you imagine a scenario where all the premises are true and yet the conclusion is false?). What might not be as obvious is the form that this argument has. However, you should be able to identify that form if you utilize the tools that you have learned so far. The first premise is a disjunction. The second premise is a conditional statement whose antecedent is the left disjunct of the disjunction in the first premise. And the third premise is a conditional statement whose antecedent is the right disjunct of the disjunction in the first premise. The conclusion is the disjunction of the consequents of the conditionals in premises 2 and 3. Here is this form of inference using symbols:

1. $p \vee q$
2. $p \supset r$
3. $q \supset s$
4. $\therefore r \vee s$

We have now introduced each of the 8 forms of inference. In the next section I will walk you through some basic proofs that utilize these 8 rules.

Exercise 16: Fill in the blanks with the valid form of inference that is being used and the lines the inference follows from. Note: the conclusion is written to the right of the last premise, following the “/∴” symbols.

Example 1:

1. $M \supset \sim N$
2. M
3. $H \supset N \quad \quad \quad \therefore \sim H$
4. $\sim N \quad \underline{\text{Modus ponens, 1, 2}}$
5. $\sim H \quad \underline{\text{Modus tollens, 3, 4}}$

Example 2:

1. $A \vee B$
2. $C \supset D$
3. $A \supset C$
4. $\sim D \quad \quad \quad \therefore B$
5. $A \supset D \quad \underline{\text{Hypothetical syllogism, 3, 2}}$

Chapter 2: Formal methods of evaluating arguments

6. $\sim A$ Modus tollens, 5, 4
 7. B Disjunctive syllogism, 1, 6

1

1. $A \cdot C \quad \therefore (A \vee E) \cdot (C \vee D)$
2. A _____
3. C _____
4. $A \vee E$ _____
5. $C \vee D$ _____
6. $(A \vee E) \cdot (C \vee D)$ _____

4. $D \supset E \quad \therefore E \vee B$

5. $C \supset E$ _____
6. $C \vee A$ _____
7. $E \vee B$ _____

2

1. $A \supset (B \supset D)$
2. $\sim D$
3. $D \vee A \quad \therefore \sim B$
4. A _____
5. $B \supset D$ _____
6. $\sim B$ _____

#6

1. $(A \vee M) \supset R$
2. $(L \supset R) \cdot \sim R$
3. $\sim(C \cdot D) \vee (A \vee M) \quad \therefore \sim(C \cdot D)$
4. $\sim R$ _____
5. $\sim(A \vee M)$ _____
6. $\sim(C \cdot D)$ _____

3

1. $A \supset \sim B$
2. $A \vee C$
3. $\sim\sim B \cdot D \quad \therefore C$
4. $\sim\sim B$ _____
5. $\sim A$ _____
6. C _____

#7

1. $(H \cdot K) \supset L$
2. $\sim R \cdot K$
3. $K \supset (H \vee R) \quad \therefore L$
4. K _____
5. $H \vee R$ _____
6. $\sim R$ _____
7. H _____
8. $H \cdot K$ _____
9. L _____

#4

1. $A \supset B$
2. $A \cdot \sim D$
3. $B \supset C \quad \therefore C \cdot \sim D$
4. A _____
5. $A \supset C$ _____
6. C _____
7. $\sim D$ _____
8. $C \cdot \sim D$ _____

#8

1. $C \supset B$
2. $\sim D \cdot \sim B$
3. $(A \supset (B \supset C)) \vee D$
4. $A \vee C \quad \therefore B \supset C$
5. $\sim D$ _____
6. $A \supset (B \supset C)$ _____
7. $\sim B$ _____
8. $\sim C$ _____
9. A _____
10. $B \supset C$ _____
11. $(B \supset C) \vee B$ _____

#5

1. C
2. $A \supset B$
3. $C \supset D$

2.12 How to construct proofs

You can think of constructing proofs as a game. The goal of the game is to derive the conclusion from the given premises using only the 8 valid rules of inference that we have introduced. Not every proof requires you to use every rule, of course. But you may use any of the rules—as long as your use of the rule is correct. Like most games, people can be better or worse at the “game” of constructing proofs. Better players will be able to a) make fewer mistakes, b) construct the proofs more quickly, and c) construct the proofs more efficiently. In order to construct proofs, it is imperative that you internalize the 8 valid forms of inference introduced in the previous section. You will be citing these forms of inference as rules that will justify each new line of your proof that you add. By “internalize” I mean that you have memorized them so well that you can see those forms manifest in various sentences almost without even thinking about it. If you internalize the rules in this way, constructing proofs will be a pleasant diversion, rather than a frustrating activity. In addition to internalizing the 8 valid forms of inference, there are a couple of different strategies that can help when you’re stuck and can’t figure out what to do next. The first is the **strategy of working backwards**. When we work backwards in a proof, we ask ourselves what rule we can use to derive the sentence(s) we need to derive. Here is an example:

1. $R \cdot S$
2. $T \quad \therefore (T \vee L) \cdot (R \cdot S)$

The conclusion, which is to the right of the second premise and follows the “ \therefore ” symbol, is a conjunction (since the dot is the main operator). If we are trying to “work backwards,” the relevant question to ask is: What rule can we use to derive a conjunction? If you know the rules, you should know the answer to that question. There is only one rule that allows us to derive (infer) a sentence that is a conjunction. That rule is called “conjunction.” The form of the rule conjunction says that in order to derive a conjunction, we need to have each conjunct on a separate line. So, what are the two conjuncts that we would need in order to derive the conjunction that is the conclusion (i.e., “ $(T \vee L) \cdot (R \cdot S)$ ”). We would need both “ $T \vee L$ ” on a line and “ $R \cdot S$ ” on a separate line. But look at premise 1—we already have “ $R \cdot S$ ” on its own line! So the only other thing we need to derive is the sentence “ $T \vee L$ ”. Once we have that on a separate

line, then we can use the rule conjunction to conjoin those two sentences to get the conclusion! So the next question we have to ask is: How can I derive the sentence " $T \vee L$ "? Again, if we are working backwards, the relevant question to ask here is: What rule allows me to derive a disjunction? There are only two: constructive dilemma and addition. However, we know that we won't be using constructive dilemma since none of the premises are conditional statements, and constructive dilemma requires conditional statements as premises. That leaves addition. Addition allows us to disjoin any statement we like to an existing statement. Since we have " T " as the second premise, the rule addition allows us to disjoin " L " to that statement. The first new line of the proof should thus look like this:

3. $T \vee L$ Addition 2

What I have done is number a new line of the proof (continuing the numbering from the premises) and then have written the rule that justifies that new line as well as the line(s) from which that line was derived via that rule. In this case, since addition is a rule that allows you to derive a sentence directly from just one line, I have cited only one line. The next step of the proof should be clear since we have already talked through it above. All we have to do now is go directly to the conclusion, since the conclusion is a conjunction and we now have (on separate lines of the proof) each conjunct. Thus, the final line of this (quite simple) proof should look like this:

4. $(T \vee L) \cdot (R \cdot S)$ Conjunction 1, 3

Again, all I've done is the write the new line of the proof (continuing the numbering from the previous line) and then have written the rule that justifies that new line as well as the line(s) from which that line was derived via that rule. In this case, the rule conjunction requires that we cite two lines (i.e., each conjunct that we are conjoining). So, I have to find the lines that contained " $T \vee L$ " and " $R \cdot S$ " and cite those lines. It does not matter the order in which you cite the lines as long as you have cited the correct lines (e.g., I could have equally well have written, "Conjunction 3, 1" as the justification). Thus the complete proof should look like this:

1. $R \cdot S$
 2. T $\therefore (T \vee L) \cdot (R \cdot S)$
 3. $T \vee L$ Addition 2

4. $(T \vee L) \cdot (R \cdot S)$ Conjunction 1, 3

That's it. That is all there is to constructing a proof. The last line of the proof is the conclusion to be derived: check. Each line of the proof follows by the rule and the line(s) cited: check. Since both of those requirements check out, our proof is complete and correct.

I have just walked you through a simple proof using the strategy of working backwards. This strategy works well as long as the conclusion we are trying to derive is complex—that is, if it contains truth functional connectives. However, sometimes our conclusion will simply be an atomic statement. In that case, we will not as easily be able to utilize the strategy of working backwards. But there is another strategy that we can utilize: the **strategy of working forward**. To utilize the strategy of working forward, we simply ask ourselves what rules we can apply to the existing premises to derive *something*, even if it isn't the conclusion we are ultimately trying to derive. As a part of this strategy, we should typically break apart a conjunction whenever we have one as a premise of our argument. Doing this can help to see where to go next. (If you've ever played Scrabble, then you can think of this as rearranging your Scrabble tiles in order to see what words you can build.) Here is an example of a proof where we should utilize the strategy of working forward:

1. $A \cdot B$
2. $B \supset C \quad \therefore C$

Notice that since the conclusion is atomic, we cannot utilize the strategy of working backwards. Instead, we should try working forward. As part of this strategy, we should break apart conjunctions by using the rule "simplification." That will be the first step of our proof:

1. $A \cdot B$
2. $B \supset C \quad \therefore C$
3. $A \quad$ Simplification 1
4. $B \quad$ Simplification 1

The first two lines of the proof is just breaking down the conjunction in line 1, where line 3 is just the left conjunct and line 4 is just the right conjunct. Both lines 3 and 4 follow by the same rule and the same line, in this case. The next question we ask when utilizing the strategy of working forward is: what lines of

the proof we can apply some rule to in order to derive *something or other*? Look at the conditional on line 2. We haven't used that yet. So what rule can we apply to that line? You should be thinking of the rules that utilize conditional statements (modus ponens, modus tollens, and hypothetical syllogism). We can rule out hypothetical syllogism since here we have only *one* conditional and the rule hypothetical syllogism requires that we have two. If you look at line 4 (that we have just derived) you should see that it is the antecedent of the conditional statement on line 2. And you should know that that means we can apply the rule, modus ponens. So our next step is to do that:

1. $A \cdot B$
2. $B \supset C \quad \therefore C$
3. $A \quad$ Simplification 1
4. $B \quad$ Simplification 1
5. $C \quad$ Modus ponens 2, 4

But now also notice that the line that we have just derived is in fact the conclusion of the argument. So our proof is finished.

Before the close of this section, let's work through a bit longer proof. Remember: any proof, long or short, is the same process and utilizes the same strategy. It is just a matter of keeping track of where you are in the proof and what you're ultimately trying to derive. So here is a bit more complex proof:

1. $(\sim A \vee B) \supset L$
2. $\sim B$
3. $A \supset B$
4. $L \supset (\sim R \vee D)$
5. $\sim D \cdot (R \vee F) \quad \therefore (L \vee G) \cdot \sim R$

The conclusion is a conjunction of " $L \vee G$ " and " $\sim R$ " so we know that if we can get each of those sentences on a separate line, then we can use the rule conjunction to derive the conclusion. That will be our long range goal here (and this is utilizing the strategy of working backwards). However, we cannot see how to directly get there from here at this point, so we will begin utilizing the strategy of working forward. The first thing we'll do is simplify the conjunction on line 5:

6. $\sim D \quad$ Simplification 5

7. $R \vee F$ Simplification 5

Look at lines 2 and 6: they are both negated atomic propositions. Another part of the strategy of working forward is to utilize either atomic or negated atomic sentences. We should look for how we can utilize modus tollens or disjunctive syllogism by plugging these negated atomic sentences into other lines of the proof. Look at lines 2 and 3. You should see a modus tollens there. That will be our next step:

8. $\sim A$ Modus tollens 2, 3

The next step of this proof can be a bit tricky. There are a couple different ways we could go. One would be to utilize the rule "addition." Can you see how we might helpfully utilize this rule using either line 6 or 8? If not, I'll give you a hint: what if we were to use addition on line 8 to derive " $\sim A \vee B$ "? Can you see how we could then plug that into line 1? In fact, " $\sim A \vee B$ " is the antecedent of the conditional in line 1, so we could then use modus ponens to derive the consequent. Thus, let's try starting with the addition on line 8:

9. $\sim A \vee B$ Addition 8

Next, we'll utilize line 9 and line 1 with modus ponens to derive the next line:

10. L Modus ponens 1, 9

Notice at this point that what we have derived on line 10 is " L " and what we earlier said we needed as one of the conjuncts was " $L \vee G$ ". You should recognize that we have a rule that will allow us to infer directly from " L " to " $L \vee G$ ". That rule is addition (again). That will be the next line of the proof:

11. $L \vee G$ Addition 10

At this point, our strategy should be to try to derive the other conjunct, " $\sim R$ ". Notice that " $\sim R$ " is contained within the sentence on line 4, but it is embedded. How can we "get it free"? Start by noticing that the $\sim R$ is a part of a disjunction, which is itself a consequent of a conditional statement. Also notice that we have already derived the antecedent of that conditional statement, which means that we can use modus ponens to derive the consequent:

12. $\sim R \vee D$ Modus ponens 4, 10

The penultimate step is to use a disjunctive syllogism to derive " $\sim R$ ".

13. $\sim R$ Disjunctive syllogism 6, 12

The final step is simply to conjoin lines 11 and 13 to get the conclusion:

14. $(L \vee G) \cdot \sim R$ Conjunction 11, 13

Thus, here is the completed proof:

1. $(\sim A \vee B) \supset L$
2. $\sim B$
3. $A \supset B$
4. $L \supset (\sim R \vee D)$
5. $\sim D \cdot (R \vee F) \quad \therefore (L \vee G) \cdot \sim R$
6. $\sim D$ Simplification 5
7. $R \vee F$ Simplification 5
8. $\sim A$ Modus tollens 2, 3
9. $\sim A \vee B$ Addition 8
10. L Modus ponens 1, 9
11. $L \vee G$ Addition 10
12. $\sim R \vee D$ Modus ponens 4, 10
13. $\sim R$ Disjunctive syllogism
14. $(L \vee G) \cdot \sim R$ Conjunction 11, 13

Constructing proofs is a skill that takes practice. The following exercises will give you some practice with constructing proofs.

Exercise 17: Construct proofs for the following valid arguments. The first fifteen proofs *can* be complete in three or less additional lines. The next five proofs will be a bit longer. It is important to note that there is always more than one way to construct a proof. If your proof differs from the answer key, that doesn't mean it is wrong.

#1

1. $A \cdot B$

2. $(A \vee C) \supset D \quad \therefore A \cdot D$

Chapter 2: Formal methods of evaluating arguments

- #2
 1. A
 2. B $\therefore (A \vee C) \cdot B$
- #3
 1. $D \supset E$
 2. $D \cdot F \therefore E$
- #4
 1. $J \supset K$
 2. J $\therefore K \vee L$
- #5
 1. $A \vee B$
 2. $\sim A \cdot \sim C \therefore B$
- #6
 1. $A \supset B$
 2. $\sim B \cdot \sim C \therefore \sim A$
- #7
 1. $D \supset E$
 2. $(E \supset F) \cdot (F \supset D) \therefore D \supset F$
- #8
 1. $(T \supset U) \cdot (T \supset V)$
 2. T $\therefore U \vee V$
- #9
 1. $(E \cdot F) \vee (G \supset H)$
 2. $I \supset G$
 3. $\sim(E \cdot F) \therefore I \supset H$
- #10
 1. $M \supset N$
 2. $O \supset P$
 3. $N \supset P$
 4. $(N \supset P) \supset (M \vee O) \therefore N \vee P$
- #11
 1. $A \vee (B \supset A)$
2. $\sim A \cdot C \therefore \sim B$
- #12
 1. $(D \vee E) \supset (F \cdot G)$
 2. D $\therefore F$
- #13
 1. $T \supset U$
 2. $\forall \vee \sim U$
 3. $\sim V \cdot \sim W \therefore \sim T$
- #14
 1. $(A \vee B) \supset \sim C$
 2. $C \vee D$
 3. A $\therefore D$
- #15
 1. $L \vee (M \supset N)$
 2. $\sim L \supset (N \supset O)$
 3. $\sim L \therefore M \supset O$
- #16
 1. $A \supset B$
 2. $A \vee (C \cdot D)$
 3. $\sim B \cdot \sim E \therefore C$
- #17
 1. $(F \supset G) \cdot (H \supset I)$
 2. $J \supset K$
 3. $(F \vee J) \cdot (H \vee L) \therefore G \vee K$
- #18
 1. $(E \vee F) \supset (G \cdot H)$
 2. $(G \vee H) \supset I$
 3. E $\therefore I$
- #19
 1. $(N \vee O) \supset P$
 2. $(P \vee Q) \supset R$
 3. $Q \vee N$
 4. $\sim Q \therefore R$

#20

1. $J \supset K$
2. $K \vee L$
3. $(L \cdot \sim J) \supset (M \cdot \sim J)$

4. $\sim K$ $\therefore M$

2.13 Short review of propositional logic

So far in this chapter we have learned a formal method for determining whether a certain class of arguments (i.e., those that utilize only truth functional operators) are valid or invalid. That method is the truth table test of validity. We have also learned a formal method for proving arguments are valid or invalid (the method of proof). The other important skill we have learned in this chapter so far is translating sentences into propositional logic. Thus, there are three different skills that you should know how to do:

1. Translate sentences from English into propositional logic
2. Construct truth tables in order to determine whether an argument is valid or invalid
3. Construct proofs to prove an argument is valid

It is important to reiterate that truth tables are the only formal method that allow us to determine *whether* an argument is valid or invalid; proofs can only show that an argument is valid, but not that it is invalid. You might think that you can use proofs to show that an argument is invalid—for example, if you are unable to construct a proof for an argument, that means that the argument is invalid. However, this doesn't follow. There could be many reasons why you are unable to construct a proof, including that you just aren't skilled enough to construct proofs. But the fact that you aren't skilled enough to find a proof for an argument wouldn't mean that the argument is invalid, it would just mean that you weren't skilled enough to show that it is valid! So we cannot use one's inability to construct a proof for an argument to establish that the argument is invalid. Again, only the truth table test of validity can establish that an argument is invalid.

The study of propositional logic has given us a way of understanding what "formal" means in the phrase, "formal logic." We can see this clearly with the truth table test of validity. After we translate an argument into propositional

logic using constants and the truth functional connectives, we don't need to know what the constants mean in order to know whether the argument is valid or invalid. We simply have to fill out the truth table in the mechanical way we have learned and then apply the truth table test of validity (which is also a mechanical procedure). Thus, once an argument has been translated into propositional logic, determining whether an argument passes the truth table test of validity is something a computer could easily do. The translation from English to symbolic format is not as easy for a computer to do because successfully doing so depends on *understanding* the nuances of English. Although today there are computer programs that are pretty good at doing this, it has taken many years to get there. In contrast, any simple computer program from half a century ago could easily construct and evaluate a truth table using the truth table test of validity because this doesn't take any *understanding*—it is simply a mechanical procedure. There are many different programs, many of which are readily available on the web, that allow you to construct and evaluate truth tables.

In contrast, the informal test of validity (from chapter 1) requires that we understand the meaning of the statements involved in the argument in order for us to be try to imagine the premises as true and the conclusion as false. Since this test requires the use of our imagination, it clearly also requires that we understand the meanings of the statements in the argument. The truth table test of validity does not require any of this. Since the truth table method does not require understanding of the meaning of the statements involved in the argument, but only an awareness of their logical form, we refer to it as a formal logic. Formal logic is a kind of logic that looks only at the form, rather than the content (meaning) of the statements. We can easily see this by constructing an argument where the atomic propositions use silly, made-up words, such as those from Lewis Carroll's "Jabberwocky":

1. If toves are slithy, then the borogoves are mimsy
2. Borogoves are not mimsy
3. Therefore, toves are not slithy

If we translate "toves are slithy" at "T" and "borogoves are mimsy" as "B" then the form of this argument is clearly modus tollens, which is one of the 8 valid forms of inference:

1. $T \supset B$

2. $\sim B$
3. $\therefore \sim T$

We can thus see that this argument is valid even though we have no idea what “toves” or “borogoves” are or what “slithy” and “mimsy” mean. Thus, propositional logic, which includes the truth table test of validity, is a kind of formal logic, whereas the informal test of validity is not. There are other kinds of formal logic besides propositional logic. In the next section I will introduce another kind of formal logic: categorical logic.

2.14 Categorical logic

Consider the following argument:

1. All humans are mortal
2. All mortal things die
3. Therefore, all humans die

If we were to apply the informal test of validity (from chapter 1) to this argument, we would see that the argument is valid because it is not possible to imagine a scenario in which the premises are true and yet the conclusion is false. However, look at what happens if we try to translate it using propositional logic. Since “all humans are mortal” is atomic, (i.e., it does not contain any truth functional operators) we can translate it using the constant “H.” The second premise, “all mortal things die,” is also atomic, so we can translate it using the constant, “M.” Finally, the conclusion, is yet another atomic statement, “All humans die,” which we can translate “D.” But then the form of our argument is just this:

1. H
2. M
3. $\therefore D$

The problem is that this argument is not valid, which we can clearly see by constructing a truth table. Since there are three different atomic components, our truth table will be 8 rows. (In the following truth table, since the reference columns would just be identical to the premise and conclusion columns, I just collapsed the two in order to make the truth table less redundant.)

H	M	D
T	T	T
T	T	F
T	F	T
T	F	F
F	T	T
F	T	F
F	F	T
F	F	F

Notice the second row of the truth table (which I have bolded). The premises are both true on that row and yet the conclusion is false. That means that this argument does not pass the truth table test of validity and so is invalid. But clearly this argument *is* valid. If it is true that all humans are mortal and that all mortal things die, then it must be true that all humans die. What this argument reveals is one of the limitations of propositional logic. There are some arguments that are intuitively valid (such as this one) but that cannot be shown to be valid using the methods of propositional logic. This shows that we need other kinds of formal logic to be able to capture a wider range of logically valid inferences. Categorical logic allows us to supplement propositional logic with a formal method that will handle arguments like this that propositional logic is unable to handle.

Categorical logic is the logic that deals with the logical relationship between categorical statements. A **categorical statement** is simply a statement about a category or type of thing. For example, the first premise of the above argument is a statement about the categories of humans and things that are mortal. The second premise is a statement about the categories of things that are mortal and things that die. Finally, the conclusion is a statement about humans and things that die. Although you may think that this argument has a similar form as a hypothetical syllogism, it is distinct from a hypothetical syllogism because the premises are not composed of two different atomic propositions. Rather, each premise contains only one atomic proposition.

In categorical logic, the logical terms (analogous to the truth functional operators of propositional logic) are the terms “all” and “some.” In contrast with propositional logic, in categorical logic we will use capital letters to stand

for categories of things in the world, rather than for atomic propositions. Thus, we can represent the statement:

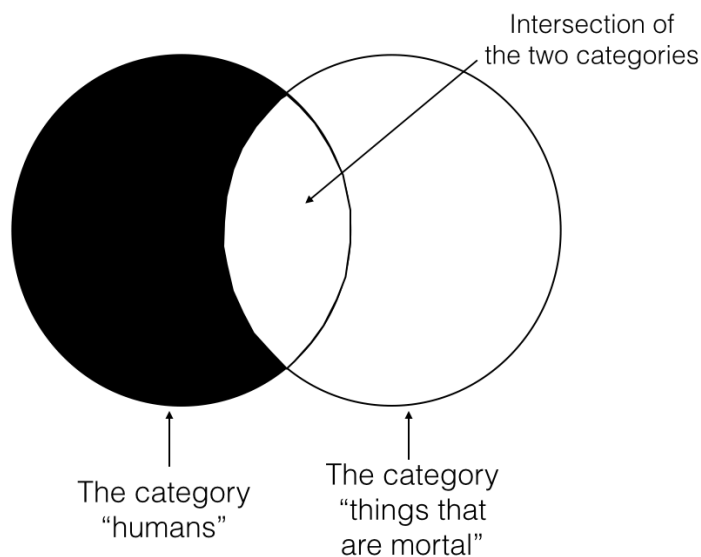
All humans are mortal

as

All H are M

where “H” stands for the category of “humans” and “M” stands for the category, “things that are mortal.” Notice that the categories are nouns or noun phrases. Thus, instead of saying that the category is “mortal” I said the category is “things that are mortal.” It is important to recognize the difference between how the capital letters are being used in categorical logic and how they were used in propositional logic. In categorical logic, the capital letters stand for noun phrases that denote categories of things in the world—for example, “cars” or “things that are man-made” or “mammals” or “things that are red.”

In categorical logic, we will use what are called **Venn diagrams** to represent the logical relationships between the different kinds of categorical statements. A Venn diagram is simply a way of graphically representing the logical relationship between two different categorical statements. Below is a Venn diagram that represents the statement, “all humans are mortal.”



Here is how to understand this Venn. There are two circles that represent the two categories, "humans" and "things that are mortal." These two categories are overlapping so that the intersection of those two categories (i.e. the place where the two circles overlap) represents things that are both human and mortal. Any shaded portions of the Venn diagram (by "shaded" I will mean "blacked out") represent that there is nothing in that area of the category. So the above Venn says that there is nothing in the category "humans" that is not also in the category "things that are mortal." The above Venn also allows that there are things that are in the category "things that are mortal" but that aren't in the category "humans" (which is as it should be since, of course, dogs are mortal and yet not human). So the reason the category "things that are mortal" is left unshaded is that in saying "all humans are mortal" I leave open the possibility that there are things that are not human and yet mortal.

As noted above, the statement, "all humans are mortal," has a particular form:

All H are M.

This is one of the **four categorical forms**. The way we will represent these categorical forms generally are with an "S" (which stands for "subject term") and a "P" (which stands for "predicate term"). Thus, the categorical statement, "all humans are mortal," has the following categorical form:

All S are P

The way we interpret statements of this form are as follows: everything in the category S is also in the category P. This statement form is what we call a "**universal affirmative**," since it is a universal statement that does not contain a negation. There are three other categorical statement forms that you will have to become familiar with in order to do categorical logic. Here they are (with the name of the type of statement in parentheses to the right:

No S are P (**universal negative**)

Some S are P (**particular affirmative**)

Some S are not P (**particular negative**)

Here are three examples of statements that have these three forms (respectively):

No reptiles give live birth

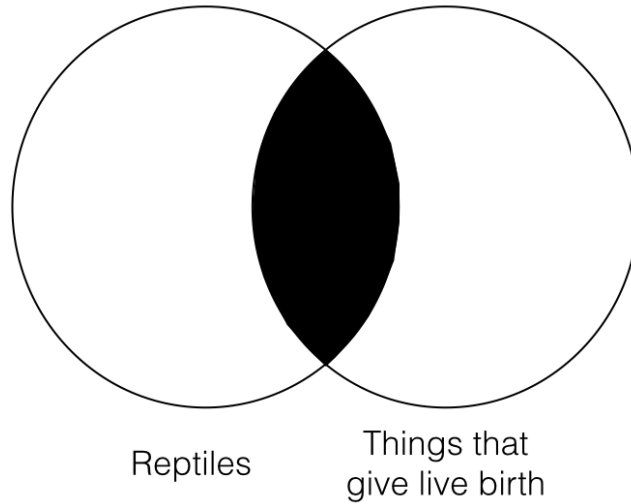
Some birds are taller than President Obama

Some birds don't fly

Notice that although these three statements don't have exactly the same form as the statement forms above, they can be translated into those same forms. All we have to do is figure out the noun phrase that describes each category that the statement is referring to. Let's start with "no reptiles give live birth." This categorical statement refers to two different categories: the category of "reptiles" and the category of "things that give live birth." Notice, again, that I added "things that..." to the predicate of the sentence ("give live birth") because "give live birth" is not a description of a category. Rather, the way of describing the category is with the noun phrase, "things that give live birth." Using these two category descriptions, we can translate this sentence to have the same form as its categorical form. All we have to do is substitute in the name of the subject category (i.e., the "S" term) and the description of the predicate category (i.e., the "P" term). Doing that will yield the following sentence:

No reptiles are things that give live birth

Although this sentence sounds strange in English, it has the same form as the categorical form, no S are P, and this translation allows us to clearly see that it does and thus to see what the two categories are. Here is what the Venn diagram for this statement looks like:

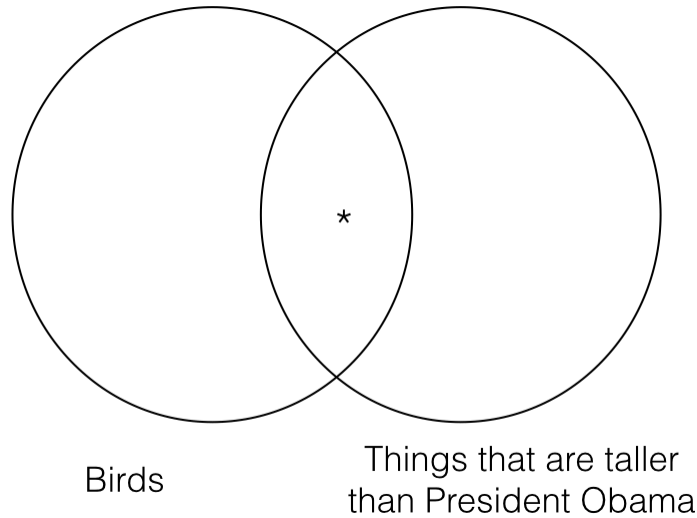


This Venn diagram represents that there is nothing in the intersection of the two categories, “reptiles” and “things that give live birth.” If you think about it, this is exactly what our original statement was saying: there isn’t anything that is both a reptile and gives live birth.

Let’s look at the next statement, “some birds are taller than President Obama.” This is a statement not about all birds, but about some birds. What are the two categories? One category is clearly “birds.” The other category is “things that are taller than President Obama.” That may sound like a strange category, but it is perfectly legitimate category. It includes things like adult ostriches, large grizzly bears standing on their hind legs, giraffes, the Flatiron Building, a school bus, etc. Here is how we’d translate this sentence using our two categories:

Some birds are things that are taller than President Obama.

Again, although this sentence sounds strange in English, it has the same form as the categorical form, some S are P, and it allows us to clearly see what the two categories are. Below is the Venn diagram for this statement:

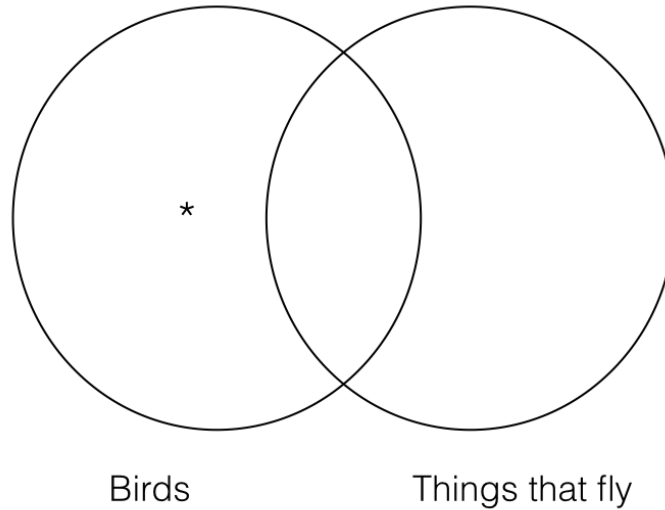


By convention, an asterisk on the Venn diagram means that there is at least one thing in that category. By putting the asterisk in the intersection of the two categories, we are saying that there is at least one thing that is a bird and is taller than President Obama, which is exactly what our original sentence was saying.

Finally, let's consider the statement, "some birds don't fly." How would we translate this sentence to have the "some S are not P" form? The first step is to get the descriptions of the two categories using either nouns or noun phrases. The "S" term is easy; it is just "birds" again. But we have to be a bit more careful with the "P" term, since its predicate contains a negation. We do not want any of our categories to contain a negation. Rather, the negation is contained in the form (i.e., the "not"). The category cannot be simply "fly" or even "flies" since neither of these are a category of *thing*. We have to use our trick of turning the predicate into a noun phrase, i.e., "things that fly." Given these two category descriptions, we can then translate the sentence to have the categorical form, some S are not P:

Some *birds* are not *things that fly*

Again, although the English sounds clunky here, it has the same form as the categorical form, some S are not P, and it allows us to clearly see what the two categories are. Below is the Venn diagram for this statement:



By convention, an asterisk on the Venn diagram means that there is at least one thing in that category. By putting the asterisk inside the “birds” category, but outside the “things that fly” category, we are representing that at least one thing that is a bird isn’t a thing that flies. This is exactly what our original sentence was saying.

Translating categorical statements into their categorical form can be tricky. In fact, it is probably one of the trickier things you’ll do in formal logic. There is no simple way of doing it other than asking yourself whether your translation accurately captures the meaning of the original English sentence. Here is an example of a tricky categorical statement:

Nobody loves me but my mother.

This is a categorical statement, but which of the four categorical forms does it have? The first step is to ask what two categories are being referred to in this sentence. Here are the two categories: “things that love me” and “things that are my mother.” Notice that the category couldn’t just be “my mother” since that isn’t a category; it’s a particular thing. Again, this sounds strange, but it is important to remember that we are describing *categories of things*. The next question is: what is this sentence saying is the relationship between these two categories? Hint: it has to be one of the four categorical forms (since any categorical statement can be translated into one of these four forms). The sentence is saying that the only things that love me are things that are my

mother. The categorical form of the statement is the “all S are P” form. Thus, the sentence, translated into the correct categorical form would be:

All things that love me are things that are my mother.

We will end this section with one last example. Consider the following categorical statement:

The baboon is a fearsome beast.

Which of the four categorical forms does this statement have? Although the article “the,” which often denotes particulars, may lead one to think that this is a particular affirmative form (some S are P), it is actually a universal affirmative form (all S are P). This English sentence has the sense of “baboons are fearsome beasts” rather than of “that (particular) baboon is a fearsome beast.” English is strange, which is what makes translation one of the trickiest parts of logic. So, the two categories are: “baboons” and “fearsome beasts.” Notice that since “fearsome beasts” is already a noun phrase, we don’t have to add “things that are...” to it. Using the two category descriptions, the translation into the “all S are P” categorical form is thus:

All baboons are fearsome beasts.

In this section we have learned what categorical statements are, how to translate categorical statements into one of the four categorical forms, and how to construct Venn diagrams for each of the four categorical forms. The following exercises will give you some practice with the translation part; in subsequent sections we will learn how to use Venn diagrams as a formal method of evaluating a certain class of arguments.

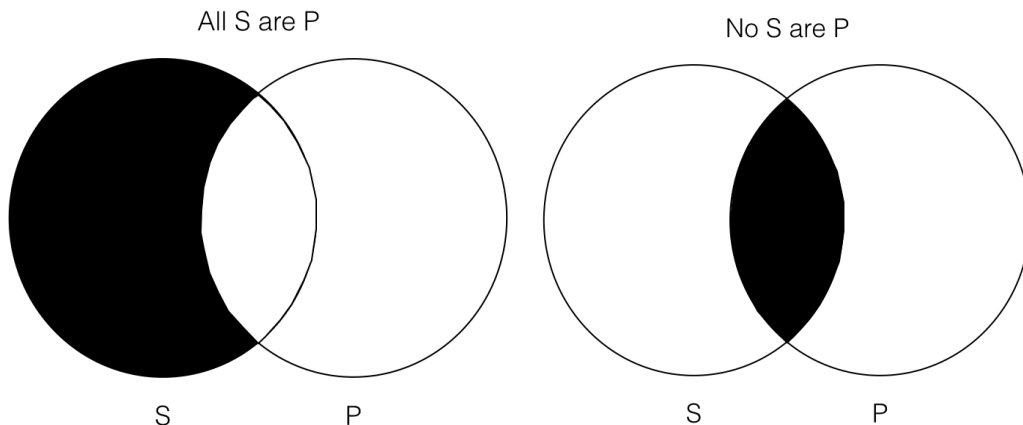
Exercise 18: Translate each of the following sentences into one of the four categorical forms (universal affirmative, universal negative, particular affirmative, particular negative). Make sure that the descriptions of the two categories are nouns or noun phrases (rather than adjectives or verbs).

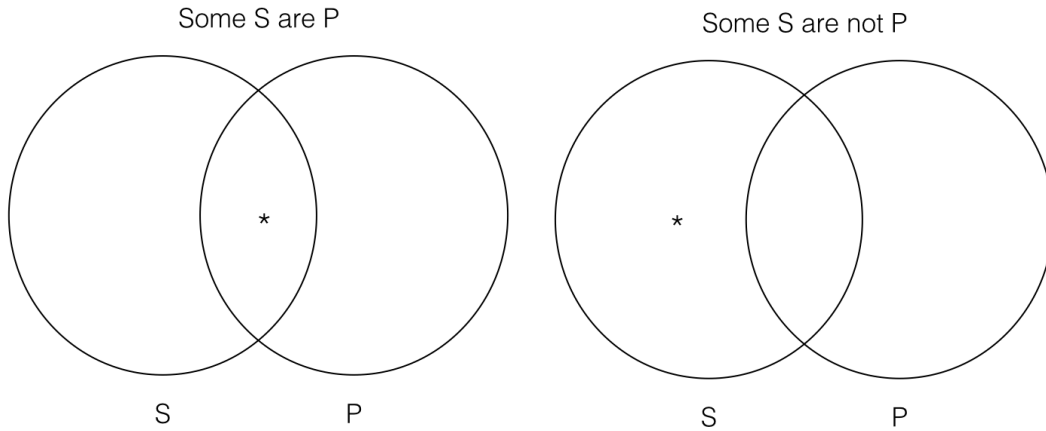
1. Real men wear pink.
2. Dinosaurs are not birds.
3. Birds evolved from dinosaurs.

4. Some mammals are not predators.
5. Some predators are not mammals.
6. Not all who wander are lost.
7. All presidents are not women.
8. Boxers aren't rich.
9. If someone is sleeping then they aren't conscious.
10. If someone is conscious then they aren't sleeping.
11. All's well that ends well.
12. My friends are the only ones that care.
13. Someone loves you.
14. Jesus loves everyone.
15. Jesus loves the little children.
16. Some people don't love Jesus.
17. Only pedestrians may use the Appalachian Trail.
18. Only citizens can be president.
19. Anyone who is a Hindu believes in God.
20. Anything that is cheap is no good.
21. Some expensive things are no good.
22. Not all mammals have legs.
23. There are couples without children.
24. There are no people who hate chocolate.
25. There are people who hate cats.
26. Nothing that is sharp is safe.
27. No poodle could run faster than a cheetah.
28. No professional runner is slow.
29. Baboons aren't friendly.
30. Pigs will eat anything.

2.15 The Venn test of validity for immediate categorical inferences

In the last section, we introduced the four categorical forms. Those forms are below.

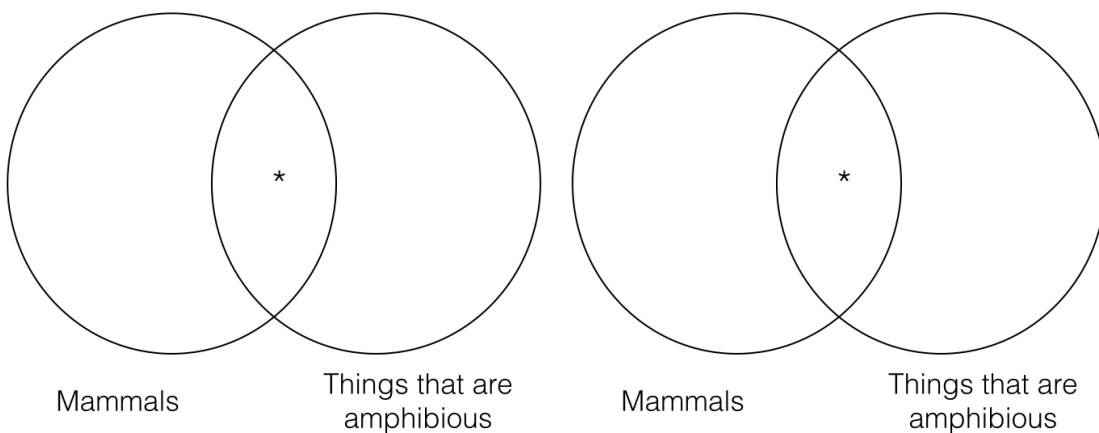




We can use Venn diagrams in order to determine whether certain kinds of arguments are valid or invalid. One such type of argument is what we will call “immediate categorical inferences.” An immediate categorical inference is simply an argument with one premise and one conclusion. For example:

1. Some mammals are amphibious.
2. Therefore, some amphibious things are mammals.

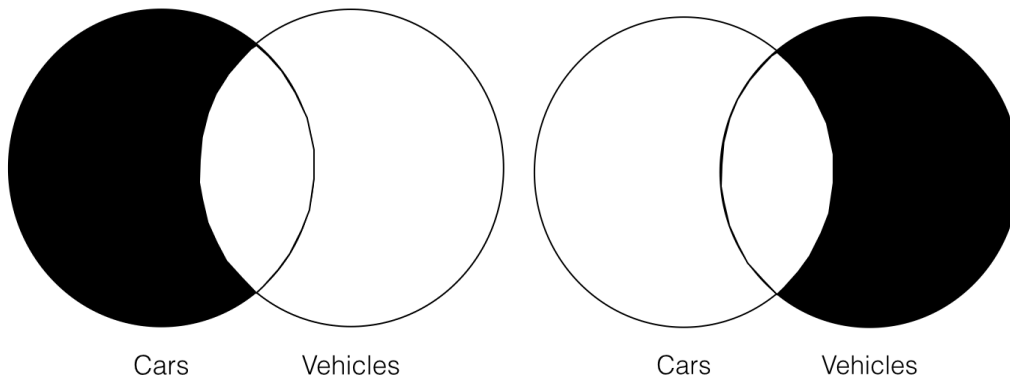
If we construct a Venn diagram for the premise and another Venn diagram for the conclusion, we will see that the Venn diagrams are identical to each other.



That is, the information that is represented in the Venn for the premise, is exactly the same information represented in the Venn for the conclusion. This argument passes the **Venn test of validity** because the conclusion Venn contains no additional information that is not already contained in the premise Venn. Thus, this argument is valid. Let’s now turn to an example of an invalid argument.

1. All cars are vehicles.
2. Therefore, all vehicles are cars.

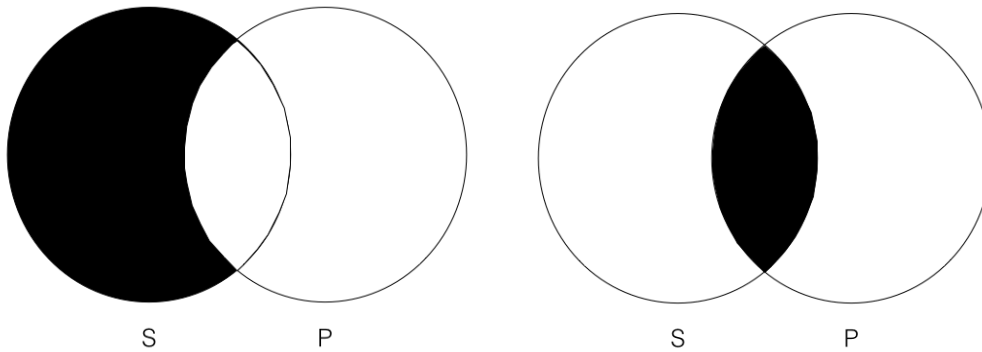
Here are the Venns for the premise and the conclusion, respectively:



In this case, the Venns are clearly not the same. More importantly, we can see that the conclusion Venn (on the right) contains additional information that is not already contained in the premise Venn. In particular, the conclusion Venn allows that a) there could be things in the “car” category that aren’t in the “vehicle” category and b) that there cannot be anything in the “vehicle” category that isn’t also in the “car” category. That is not information that is contained in the premise Venn, which says that a) there isn’t anything in the category “car” that isn’t also in the category “vehicle” and b) that there could be things in the category “vehicle” that aren’t in the category “car.” Thus, this argument does not pass the Venn test of validity since there is information contained in the conclusion Venn that is not already contained in the premise Venn. Thus, this argument is invalid.

The Venn test of validity is a formal method, because we can apply it even if we only know the form of the categorical statements, but don’t know what the categories referred to in the statements represent. For example, we can simply use “S” and “P” for the categories—and we clearly don’t know what these represent. For example:

1. All S are P
2. No P are S



The conclusion (on the right) contains information that is not contained in the premise (on the left). In particular, the conclusion Venn explicitly rules out that there is anything that is both in the category "S" and in the category "P" while the premise Venn allows that this is the case (but does not require it). Thus, we can say that this argument fails the Venn test of validity and thus is invalid. We know this even though we have no idea what the categories "S" and "P" are. This is the mark of a formal method of evaluation.

Exercise 19: Apply the Venn test of validity in order to determine whether the following categorical inferences are valid or invalid.

1. All S are P; therefore, all P are S
2. Some S are P; therefore, some P are S
3. Some S are P; therefore, some P are not S
4. Some S are P; therefore, all P are S
5. No S are P; therefore, no P are S
6. No P are S; therefore, some S are P
7. Some S are not P; therefore, some P are not S
8. All S are P; therefore some P are not S

2.16 Universal statements and existential commitment

Consider the following inference:

1. All S are P
2. Therefore, some S are P

Is this inference valid or invalid? As it turns out, this is an issue on which there has been much philosophical debate. On the one hand, it seems that many

times when we make a universal statement, such as “all dogs are mammals,” we imply that there *are* dogs—i.e., that dogs exist. Thus, if we assert that *all* dogs are mammals, that implies that *some* dogs are mammals (just as if I say that *everyone* at the party was drunk, this implies that at least *someone* at the party was drunk). In general, it may seem that “all” implies “some” (since some is encompassed by all). This reasoning would support the idea that the above inference is valid: universal statements imply certain particular statements. Thus, statements of the form “all S are P” would imply that statements of the form “some S are P.” This is what is called “**existential commitment.**”

In contrast to the reasoning just laid out, modern logicians *reject* existential commitment; they do *not* take statements of the form “all S are P” to imply that there exists anything in the “S” category. Why would they think this? One way of understanding why universal statements are interpreted in this way in modern logic is by considering laws such as the following:

All trespassers will be fined.

All bodies that are not acted on by any force are at rest.

All passenger cars that can travel 770 mph are supersonic.

The “S” terms in the above categorical statements are “trespassers,” “bodies that are not acted on by any force,” and “passenger cars that can travel 770 mph.” Now ask yourself: do these statements commit us to the existence of either trespassers or bodies not acted on by any force? No, they don’t. Just because we assert the rule that all trespassers will be fined, we do not necessarily commit ourselves to the claim that there are trespassers. Rather, what we are saying is *anything that is a trespasser* will be fined. But this can be true, even if there are no trespassers! Likewise, when Isaac Newton asserted that all bodies that are not acted on by any force remain at rest, he was not committing himself to the existence of “bodies not acted on by any force.” Rather, he was saying that *anything that is a body not acted on by any force* will remain in motion. But this can be true, even if there are no bodies not acted on by any force! (And there aren’t any such bodies, since even things that are stationary like your house or your car parked in the driveway are still acted on by forces such as gravity and friction.) Finally, in asserting that all passenger cars that can travel 770 mph are supersonic, we are not committing ourselves to the

existence of any such car. Rather, we are only saying that *were there any such car*, it would be supersonic (i.e., it would travel faster than the speed of sound).

For various reasons (that we will not discuss here), modern logic treats a universal categorical statement as a kind of conditional statement. Thus, a statement like,

All passenger cars that can travel 770 mph are supersonic

is interpreted as follows:

For any x , *if x is a passenger car that can travel 770 mph then x is supersonic.*

But since conditional statements do not assert either the antecedent or the consequent, the universal statement is not asserting the existence of passenger cars that can travel 770 mph. Rather, it is just saying that *if* there were passenger cars that could travel that fast, then those things would be supersonic.

We will follow modern logic in denying existential commitment. That is, we will not interpret universal affirmative statements of the form "All S are P" as implying particular affirmative statements of the form "some S are P." Likewise, we will not interpret universal negative statements of the form "no S are P" as implying particular negative statements of the form "some S are not P." Thus, when constructing Venn diagrams, you can always rely on the fact that if there is no particular represented in the premise Venn (i.e., there is no asterisk), then if the conclusion Venn represents a particular (i.e., there is an asterisk), the argument will be invalid. This is so since no universal statement logically implies the existence of any particular. Conversely, if the premise Venn does represent a particular statement (i.e., it contains an asterisk), then if the conclusion doesn't contain particular statement (i.e., doesn't contain an asterisk), the argument will be invalid.

Exercise 20: Construct Venn diagrams to determine which of the following immediate categorical inferences are valid and which are invalid. Make sure you remember that we are not interpreting universal statements to imply existential commitment.

1. All S are P; therefore, some S are P
2. No S are P; therefore, some S are not P
3. All S are P; therefore, some P are S
4. No S are P; therefore, some P are not S

2.17 Venn validity for categorical syllogisms

A **categorical syllogism** is just an argument with two premises and a conclusion, where every statement of the argument is a categorical statement. As we have seen, there are four different types (forms) of categorical statement:

All S are P	(universal affirmative)
No S are P	(universal negative)
Some S are P	(particular affirmative)
Some S are not P	(particular negative)

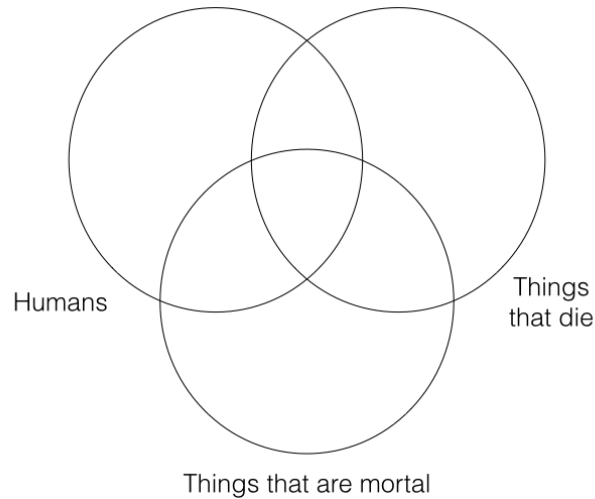
Thus, any categorical syllogism's premises and conclusion will be some mixture of these different types of statement. The argument I gave at the beginning of section 2.13 was a categorical syllogism. Here, again, is that argument:

1. All humans are mortal
2. All mortal things die
3. Therefore, all humans die

As we can see now that we have learned the four categorical forms, each one of the statements in this syllogism is a "universal affirmative" statement of the form, "all S are P." Let's first translate each statement of this argument to have the "all S are P" form:

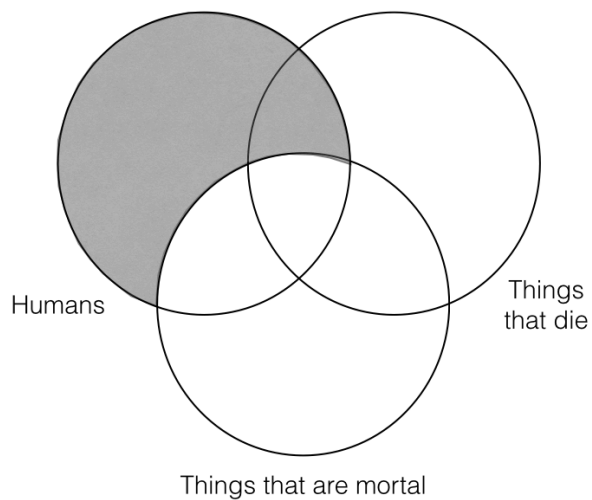
1. All humans are things that are mortal.
2. All things that are mortal are things that die.
3. All humans are things that die.

In determining the validity of categorical syllogisms, we must construct a three category Venn diagram for the premises and a two category Venn diagram for the conclusion. Here is what the three category Venn looks like for the premises:

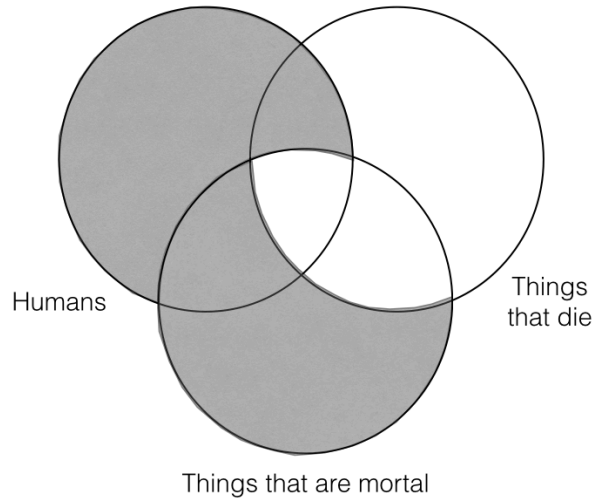


We need a three category Venn for the premises since the two premises refer to three different categories. The way you should construct the Venn is with the circle that represents the "S" category of the conclusion (i.e., the category "humans") on left, the circle that represents the "P" category of the conclusion (i.e., the category "things that die") on the right, and the remaining category ("things that are mortal") in the middle, as I have done above. Constructing your three category Venn in this way will allow you to easily determine whether the argument is valid.

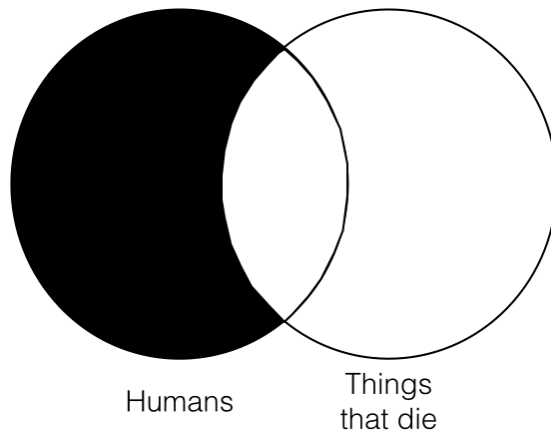
The next thing we must do is represent the information from the first two premises in our three category Venn. We'll start with the first premise, which says "all humans are things that are mortal." That means that we must shade out anything that is in the "human" category, but that isn't in the "things that are mortal" category, like this:



The next thing we have to do is fill in the information for the second premise, all things that are mortal are things that die. That means that there isn't anything that is in the category "things that are mortal" but that isn't in the "things that die" category. So we must shade out all of the parts of the "things that are mortal" category that lie outside the "things that die" category, like this:



The next thing we have to do is construct a two category Venn for the conclusion and then compare the information represented by the three category Venn for the premises to the two category Venn for the conclusion.



The conclusion represents the information that there is nothing in the "humans" category that isn't also in the "things that die" category. It also allows that there are things that die, but that aren't humans. The premise Venn also includes this same information, since every part of the "humans" category that is outside the

“things that die” category is shaded out. Thus, this argument passes the Venn test of validity and is thus valid since there is no information represented in the conclusion Venn that is not also represented in the premise Venn. Notice that it doesn’t matter that the premise Venn contains more information than the conclusion Venn. That is to be expected, since the premise Venn is representing a whole other category that the conclusion Venn isn’t. This is perfectly allowable. What isn’t allowable (and thus would make an argument fail the Venn test of validity) is if the conclusion Venn contained information that wasn’t already contained in the premise Venn. However, since this argument does not do that, it is valid.

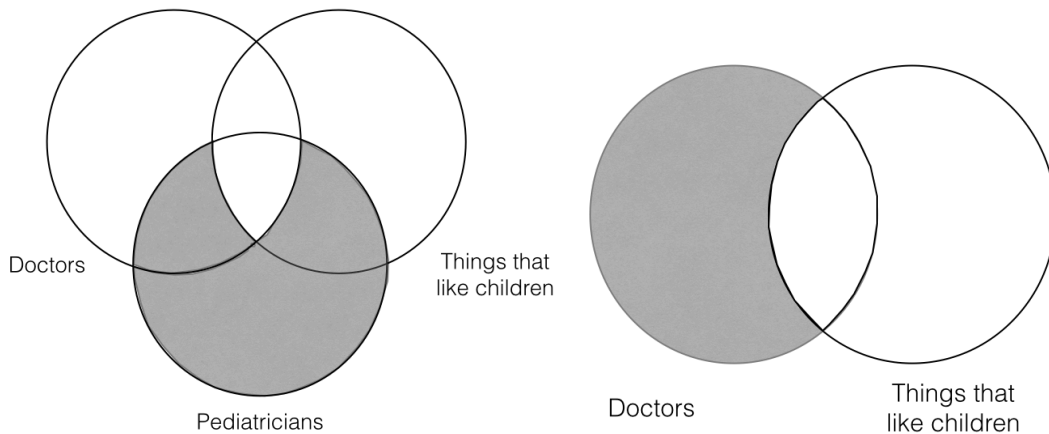
Let’s try another one.

1. All pediatricians are doctors
2. All pediatricians like children
3. Therefore, all doctors like children.

The first step is to identify the three categories referred to in this categorical syllogism. They are:

- Pediatricians
- Doctors
- Things that like children

The next step is to fill out the three category Venn for the premises and the two category Venn for the conclusion.



This argument does not pass the Venn test of validity because there is information contained in the conclusion Venn that is not contained in the premise Venn. In particular, the conclusion says that there is nothing in the “doctors” category that is outside the “things that like children category.” However, the premises do not represent that information, since the section of the category “doctors” that lies outside of the intersection of the category “things that like children” is unshaded, thus representing that there can be things there.

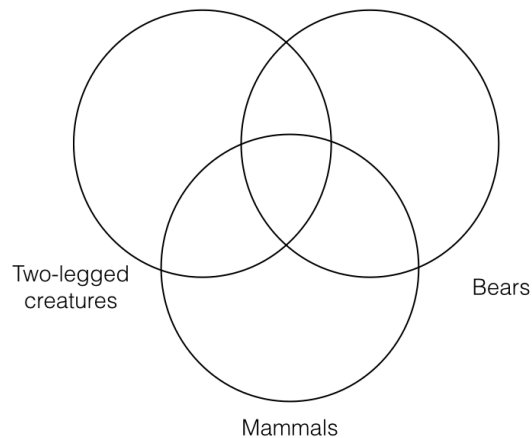
Sometimes when filling in particular statements on a three category for the premises, you will encounter a problem that requires another convention in order to accurately represent the information in the Venn. Here is an example where this happens:

1. Some mammals are bears
2. Some two-legged creatures are mammals
3. Therefore, some two-legged creatures are bears

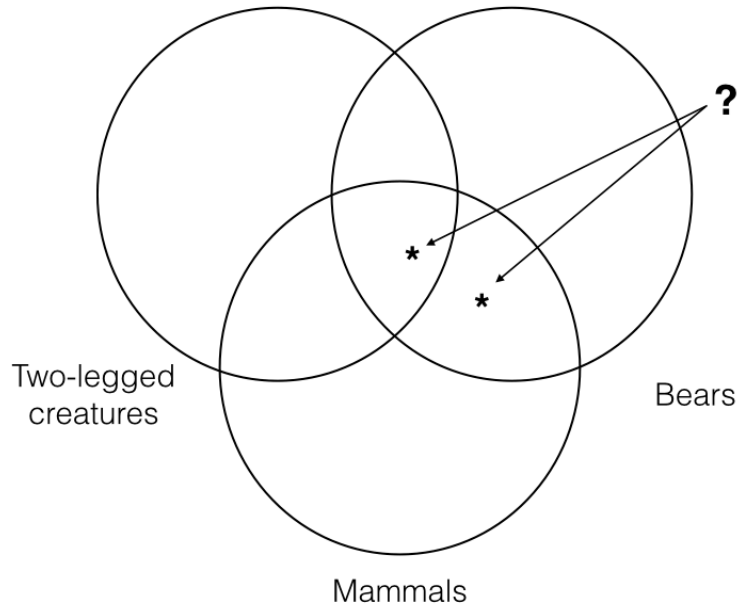
There are three categories referred to in this categorical syllogism:

- Mammals
- Bears
- Two-legged creatures

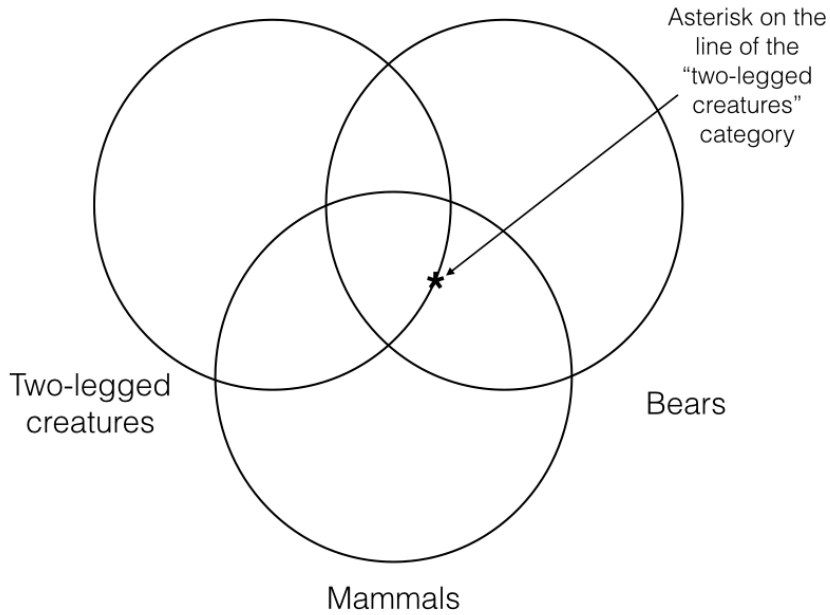
As always, we will put the “S” term of the conclusion on the left of our three category Venn, the “P” term on the right, and the remaining term in the middle, as follows:



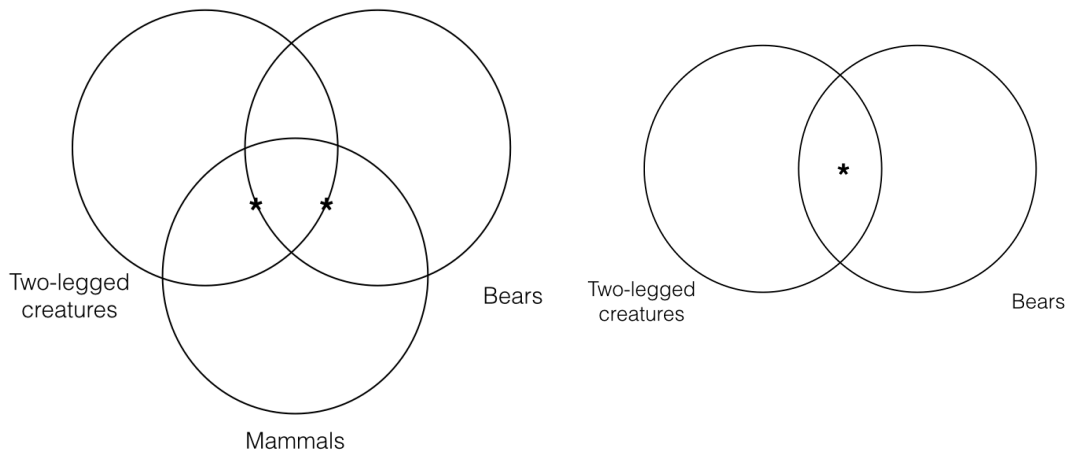
Now we need to represent the first premise, which means we need to put an asterisk in the intersection of the “mammals” and “bears” categories. However, here we have a choice to make. Since the intersection of the “bears” and “mammals” categories contains a section that is outside the “two-legged creatures” category and a section that is inside the “two-legged creatures” category, we must choose between representing the particular as part of the “two-legged creatures” category or not.



But *neither* of these can be right, since the first premise says *nothing at all* about whether the thing that is both a bear and a mammal is two-legged! Thus, in order to accurately represent the information contained in this premise, we must adopt a new convention. That convention says that when we encounter a situation where we must represent a particular on our three category Venn, but the premise says nothing about a particular category, then we must put the asterisk *on the line of that category* as I have done below. When we do this, it will represent that the particular is neither inside the category or outside the category.



We must do this same thing for the second premise, since we encounter the same problem there. Thus, when putting the asterisk in the intersection of the "two-legged creatures" and "mammals" categories, we cannot put the asterisk either inside or outside the "bears" category. Instead, we must put the asterisk on the line of the "bears" category. Thus, using this convention, we can represent the premise Venn and conclusion Venn as follows:



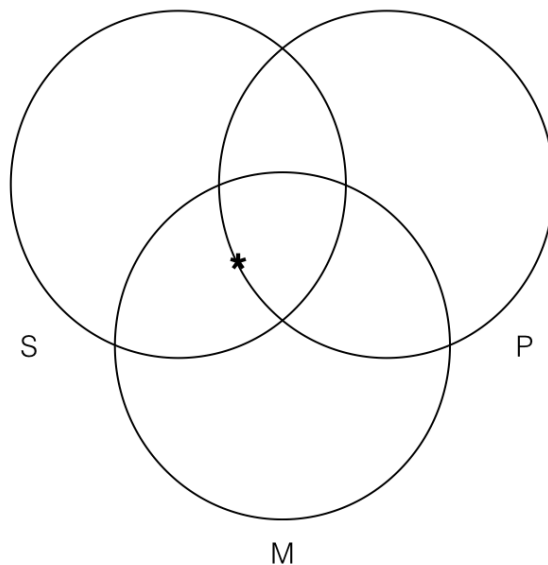
Keeping in mind the convention we have just introduced, we can see that this argument fails the Venn test of validity and is thus invalid. The reason is that the conclusion Venn clearly represents an individual in the intersection of the "two-legged creatures" and "bears" categories, whereas the premise Venn contains

no such information. Thus, the conclusion Venn contains information that is not contained in the premise Venn, which means the argument is invalid.

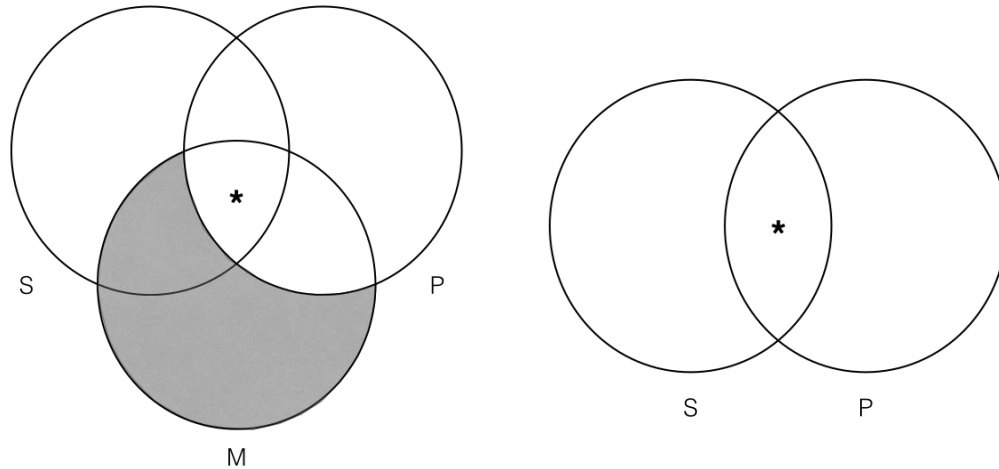
We will close this section with one last example that will illustrate an important strategy. The strategy is that we should always map universal statements before mapping particular statements. Here is a categorical syllogism that illustrates this point. This time I am going to switch to just using the capital letters S, P, and M to represent the categories. Recall that we can do this because the Venn test of validity is a formal evaluation method where we don't have to actually understand what the categories represent in the world in order to determine whether the argument is valid.

1. Some S are M
2. All M are P
3. \therefore Some S are P

If we think about mapping the first premise on our three category Venn, it seems that we will have to utilize the convention we just introduced, since the first premise is a particular categorical statement that mentions only the categories S and M and nothing about the category P:



However, as it turns out, we don't have to use this convention because when we map premise 2, which is a universal statement, this clears up where the asterisk has to go:



We can see that once we've mapped the universal statement onto the premise Venn (on the left), there is only one section where the asterisk can go that is in the intersection of S and M. The reason is that once we have mapped the "all M are P" premise, and have thus shaded out any portion of the M category that is outside the P category, we know that that asterisk cannot belong inside the M category, given that it has to be inside the P category. When we apply the Venn test of validity to the above argument, we can see that it is valid since the conclusion Venn does not contain any information that isn't already contained in the premise Venn. The conclusion simply says that there is some thing that is both S and P, and that information is already represented in our premise Venn. Thus, the argument is valid. The point of strategy here is that we should always map our universal statements onto our three category Venns before mapping our particular statements. The reason is that the universal can determine how we map our particular statements (but not vice versa).

Exercise 21: Use the Venn test of validity to determine whether the following syllogisms are valid or invalid.

1. All M is P
All M is S
 \therefore All S is P

2. All P is M
All M is S
 \therefore All S is P

3. All M is P
Some M is S
 \therefore Some S is P

4. All P is M
Some M is S
 \therefore Some S is P

5. All P is M
Some S is M
∴ Some S is P
6. All P is M
Some S is not M
∴ Some S is not P
7. All M is P
Some S is not M
∴ Some S is not P
8. All M is P
Some M is not S
∴ Some S is not P
9. No M is P
Some S is M
∴ Some S is not P
10. No P is M
Some S is M
∴ Some S is not P
11. No P is M
Some S is not M
∴ Some S is not P
12. No M is P
Some S is not M
∴ Some S is not P
13. No P is M
Some M is not S
∴ Some S is not P
14. No P is M
No M is S
∴ No S is P
15. No P is M
All M is S
∴ No S is P
16. No P is M
All S is M
∴ No S is P
17. All P is M
No S is M
∴ No S is P
18. All M is P
No S is M
∴ No S is P
19. Some M is P
Some M is not S
∴ Some S is not P
20. Some P is M
Some S is not M
∴ Some S is P

3.1 Inductive arguments and statistical generalizations

As we saw in chapter 1 (section 1.8), an inductive argument is an argument whose conclusion is supposed to follow from its premises with a high level of probability, rather than with certainty. This means that although it is possible that the conclusion doesn't follow from its premises, it is unlikely that this is the case. We said that inductive arguments are "defeasible," meaning that we could turn a strong inductive argument into a weak inductive argument simply by adding further premises to the argument. In contrast, deductive arguments that are valid can never be made invalid by adding further premises. Recall our "Tweets" argument:

1. Tweets is a healthy, normally functioning bird
2. Most healthy, normally functioning birds fly
3. Therefore, Tweets probably flies

Without knowing anything else about Tweets, it is a good bet that Tweets flies. However, if we were to add that Tweets is 6 ft. tall and can run 30 mph, then it is no longer a good bet that Tweets can fly (since in this case Tweets is likely an ostrich and therefore can't fly). The second premise, "most healthy, normally functioning birds fly," is a statistical generalization. **Statistical generalizations** are generalizations arrived at by empirical observations of certain regularities. Statistical generalizations can be either universal or partial. Universal generalizations assert that *all* members (i.e., 100%) of a certain class have a certain feature, whereas partial generalizations assert that *most* or *some percentage* of members of a class have a certain feature. For example, the claim that "67.5% of all prisoners released from prison are rearrested within three years" is a partial generalization that is much more precise than simply saying that "most prisoners released from prison are rearrested within three years." In contrast, the claim that "all prisoners released from prison are rearrested within three years" is a universal generalization. As we can see from these examples, deductive arguments typically use universal statistical generalizations whereas inductive arguments typically use partial statistical generalizations. Since statistical generalizations are often crucial premises in both deductive and inductive arguments, being able to evaluate when a statistical generalization is good or bad is crucial for being able to evaluate arguments. What we are doing in evaluating statistical generalizations is determining whether the premise in our argument is true (or at least well-

supported by the evidence). For example, consider the following inductive argument, whose premise is a (partial) statistical generalization:

1. 70% of voters say they will vote for candidate X
2. Therefore, candidate X will probably win the election

This is an inductive argument because even if the premise is true, the conclusion could still be false (for example, an opponent of candidate X could systematically kill or intimidate those voters who intend to vote for candidate X so that very few of them will actually vote). Furthermore, it is clear that the argument is intended to be inductive because the conclusion contains the word “probably,” which clearly indicates that an inductive, rather than deductive, inference is intended. Remember that in evaluating arguments we want to know about the strength of the inference from the premises to the conclusion, *but we also want to know whether the premise is true!* We can assess whether or not a statistical generalization is true by considering whether the statistical generalization meets certain conditions. There are two conditions that any statistical generalization must meet in order for the generalization to be deemed “good.”

1. **Adequate sample size:** the sample size must be large enough to support the generalization.
2. **Non-biased sample:** the sample must not be biased.

A **sample** is simply a portion of a population. A **population** is the totality of members of some specified set of objects or events. For example, if I were determining the relative proportion of cars to trucks that drive down my street on a given day, the population would be the total number of cars and trucks that drive down my street on a given day. If I were to sit on my front porch from 12-2 pm and count all the cars and trucks that drove down my street, that would be a sample. A good statistical generalization is one in which the sample is **representative** of the population. When a sample is representative, the characteristics of the sample match the characteristics of the population at large. For example, my method of sampling cars and trucks that drive down my street would be a good method as long as the proportion of trucks to cars that drove down my street between 12-2 pm matched the proportion of trucks to cars that drove down my street during the whole day. If for some reason the number of trucks that drove down my street from 12-2 pm was much higher than the average for the whole day, my sample would not be representative of the

population I was trying to generalize about (i.e., the total number of cars and trucks that drove down my street in a day). The “adequate sample size” condition and the “non-biased sample” condition are ways of making sure that a sample is representative. In the rest of this section, we will explain each of these conditions in turn.

It is perhaps easiest to illustrate these two conditions by considering what is wrong with statistical generalizations that fail to meet one or more of these conditions. First, consider a case in which the sample size is too small (and thus the adequate sample size condition is not met). If I were to sit in front of my house for only fifteen minutes from 12:00-12:15 and saw only one car, then my sample would consist of only 1 automobile, which happened to be a car. If I were to try to generalize from that sample, then I would have to say that only cars (and no trucks) drive down my street. But the evidence for this universal statistical generalization (i.e., “every automobile that drives down my street is a car”) is extremely poor since I have sampled only a very small portion of the total population (i.e., the total number of automobiles that drive down my street). Taking this sample to be representative would be like going to Flagstaff, AZ for one day and saying that since it rained there on that day, it must rain every day in Flagstaff. Inferring to such a generalization is an informal fallacy called “hasty generalization.” One commits the fallacy of **hasty generalization** when one infers a statistical generalization (either universal or partial) about a population from too few instances of that population. Hasty generalization fallacies are very common in everyday discourse, as when a person gives just one example of a phenomenon occurring and implicitly treats that one case as sufficient evidence for a generalization. This works especially well when fear or practical interests are involved. For example, Jones and Smith are talking about the relative quality of Fords versus Chevys and Jones tells Smith about his uncle’s Ford, which broke down numerous times within the first year of owning it. Jones then says that Fords are just unreliable and that that is why he would never buy one. The generalization, which is here ambiguous between a universal generalization (i.e., all Fords are unreliable) and a partial generalization (i.e., most/many Fords are unreliable), is not supported by just one case, however convinced Smith might be after hearing the anecdote about Jones’s uncle’s Ford.

The non-biased sample condition may not be met even when the adequate sample size condition is met. For example, suppose that I count all the cars on my street for a three hour period from 11-2 pm during a weekday. Let’s assume

that counting for three hours straight give us an adequate sample size. However, suppose that during those hours (lunch hours) there is a much higher proportion of trucks to cars, since (let's suppose) many work trucks are coming to and from worksites during those lunch hours. If that were the case, then my sample, although large enough, would not be representative because it would be biased. In particular, the number of trucks to cars in the sample would be higher than in the overall population, which would make the sample unrepresentative of the population (and hence biased).

Another good way of illustrating sampling bias is by considering polls. So consider candidate X who is running for elected office and who strongly supports gun rights and is the candidate of choice of the NRA. Suppose an organization runs a poll to determine how candidate X is faring against candidate Y, who is actively anti gun rights. But suppose that the way the organization administers the poll is by polling subscribers to the magazine, *Field and Stream*. Suppose the poll returned over 5000 responses, which, let's suppose, is an adequate sample size and out of those responses, 89% favored candidate X. If the organization were to take that sample to support the statistical generalization that "most voters are in favor of candidate X" then they would have made a mistake. If you know anything about the magazine *Field and Stream*, it should be obvious why. *Field and Stream* is a magazine whose subscribers would tend to own guns and support gun rights. Thus we would expect that subscribers to that magazine would have a much higher percentage of gun rights activists than would the general population, to which the poll is attempting to generalize. But in this case, the sample would be unrepresentative and biased and thus the poll would be useless. Although the sample would allow us to generalize to the population, "*Field and Stream* subscribers," it would not allow us to generalize to the population at large.

Let's consider one more example of a sampling bias. Suppose candidate X were running in a district in which there was a high proportion of elderly voters. Suppose that candidate X favored policies that elderly voters were against. For example, suppose candidate X favors slashing Medicare funding to reduce the budget deficit, whereas candidate Y favored maintaining or increasing support to Medicare. Along comes an organization who is interested in polling voters to determine which candidate is favored in the district. Suppose that the organization chooses to administer the poll via text message and that the results of the poll show that 75% of the voters favor candidate X. Can you see what's wrong with the poll—why it is biased? You probably recognize that this polling

method will not produce a representative sample because elderly voters are much less likely to use cell phones and text messaging and so the poll will leave out the responses of these elderly voters (who, we've assumed make up a large segment of the population). Thus, the sample will be biased and unrepresentative of the target population. As a result, any attempt to generalize to the general population would be extremely ill-advised.

Before ending this section, we should consider one other source of bias, which is a bias in the polling questionnaire itself (what statisticians call the "instrument"). Suppose that a poll is trying to determine how much a population favors organic food products. We can imagine the questionnaire containing a choice like the following:

Which do you prefer?

- a. products that are expensive and have no FDA proven advantage over the less expensive products
- b. products that are inexpensive and have no FDA proven disadvantage over more expensive products

Because of the phrasing of the options, it seems clear that many people will choose option "b." Although the two options do accurately describe the difference between organic and non-organic products, option "b" sounds much more desirable than option "a." The phrasing of the options is biased insofar as "a" is a stand-in for "organic" and "b" is stand-in for "non-organic." Even people who favor organic products may be more inclined to choose option "b" here. Thus, the poll would not be representative because the responses would be skewed by the biased phrasing of the options. Here is another example with the same point:

Which do you favor?

- a. Preserving a citizen's constitutional right to bear arms
- b. Leaving honest citizens defenseless against armed criminals

Again, because option "b" sounds so bad and "a" sounds more attractive, those responding to a poll with this question might be inclined to choose "a" even if they don't really support gun rights. This is another example of how bias can creep into a statistical generalization through a biased way of asking a question.

Random sampling is a common sampling method that attempts to avoid any kinds of sampling bias by making selection of individuals for the sample a matter of random chance (i.e., anyone in the population is as likely as anyone else to be chosen for the sample). The basic justification behind the method of random sampling is that if the sample is truly random (i.e., anyone in the population is as likely as anyone else to be chosen for the sample), then the sample will be representative. The trick for any random sampling technique is to find a way of selecting individuals for the sample that doesn't create any kind of bias. A common method used to select individuals for a random sample (for example, by Gallup polls) is to call people on either their landline or cell phones. Since most voting Americans have either a landline or a cell phone, this is a good way of ensuring that every American has an equal chance of being included in the sample. Next, a random number generating computer program selects numbers to dial. In this way, organizations like Gallup are able to get something close to a random sample and are able to represent the whole U.S. population with a sample size as small as 1000 (with a margin of error of +/- 4). As technology and social factors change, random sampling techniques have to be updated. For example, although Gallup used to call only landlines, eventually this method became biased because many people no longer owned landlines, but only cell phones. If some new kind of technology replaces cell phones and landlines, then Gallup will have to adjust the way it obtains a sample in order to reflect the changing social reality.

Exercise 22: What kinds of problems, if any, do the following statistical generalizations have? If there is a problem with the generalization, specify which of the two conditions (adequate sample size, non-biased sample) are not met. Some generalizations may have multiple problems. If so, specify all of the problems you see with the generalization.

1. Bob, from Silverton, CO drives a 4x4 pickup truck, so most people from Silverton, CO drive 4x4 pickup trucks.
2. Tom counts and categorizes birds that land in the tree in his backyard every morning from 5:00-5:20 am. He counts mostly Mourning Doves and generalizes, "most birds that land in my tree in the morning are Mourning Doves."
3. Tom counts and categorizes birds that land in the tree in his backyard every morning from 5:00-6:00 am. He counts mostly Mourning Doves and generalizes, "most birds that land in my tree during the 24-hour day are Mourning Doves."

4. Tom counts and categorizes birds that land in the tree in his backyard every day from 5:00-6:00 am, from 11:00-12:00 pm, and from 5:00-6:00 pm. He counts mostly Mourning Doves and generalizes, "most birds that land in my tree during the 24-hour day are Mourning Doves."
5. Tom counts and categorizes birds that land in the tree in his backyard every evening from 10:00-11:00 pm. He counts mostly owls and generalizes, "most birds that land in my tree throughout the 24-hour day are owls."
6. Tom counts and categorizes birds that land in the tree in his backyard every evening from 10:00-11:00 pm and from 2:00-3:00 am. He counts mostly owls and generalizes, "most birds that land in my tree throughout the night are owls."
7. A poll administered to 10,000 registered voters who were homeowners showed that 90% supported a policy to slash Medicaid funding and decrease property taxes. Therefore, 90% of voters support a policy to slash Medicaid funding.
8. A telephone poll administered by a computer randomly generating numbers to call, found that 68% of Americans in the sample of 2000 were in favor of legalizing recreational marijuana use. Thus, almost 70% of Americans favor legalizing recreation marijuana use.
9. A randomized telephone poll in the United States asked respondents whether they supported a) a policy that allows killing innocent children in the womb or b) a policy that saves the lives of innocent children in the womb. The results showed that 69% of respondents choose option "b" over option "a." The generalization was made that "most Americans favor a policy that disallows abortion."
10. Steve's first rock and roll concert was an Ani DiFranco concert, in which most of the concert-goers were women with feminist political slogans written on their t-shirts. Steve makes the generalization that "most rock and roll concert-goers are women who are feminists." He then applies this generalization to the next concert he attends (Tom Petty) and is greatly surprised by what he finds.
11. A high school principal conducts a survey of how satisfied students are with his high school by asking students in detention to fill out a satisfaction survey. Generalizing from that sample, he infers that 79% of students are dissatisfied with their high school experience. He is surprised and saddened by the result.

12. After having attended numerous Pistons home games over 20 years, Alice cannot remember a time when she didn't see ticket scalpers selling tickets outside the stadium. She generalizes that there are always scalpers at every Pistons home game.
13. After having attended numerous Pistons home games over 20 years, Alice cannot remember a time when she didn't see ticket scalpers selling tickets outside the stadium. She generalizes that there are ticket scalpers at every NBA game.
14. After having attended numerous Pistons home games over 20 years, Alice cannot remember a time when she didn't see ticket scalpers selling tickets outside the stadium. She generalizes that there are ticket scalpers at every sporting event.
15. Bob once ordered a hamburger from Burger King and got violently ill shortly after he ate it. From now on, he never eats at Burger King because he fears he will get food poisoning.

3.2 Inference to the best explanation and the seven explanatory virtues

Explanations help us to understand *why* something happened, not simply convince us *that* something happened (see chapter 1, section 1.3). However, there is a common kind of inductive argument that takes the best explanation of *why* x occurred as an argument for the claim *that* x occurred. For example, suppose that your car window is broken and your iPod (which you left visible in the front seat) is missing. The immediate inference you would probably make is that someone broke the window of your car and stole your iPod. What makes this a reasonable inference? What makes it a reasonable inference is that this explanation explains all the relevant facts (broken window, missing iPod) and *does so better than any other competing explanation*. In this case, it is perhaps possible that a stray baseball broke your window, but since (let us suppose) there is no baseball diamond close by, and people don't play catch in the parking garage you are parked in, this seems unlikely. Moreover, the baseball scenario doesn't explain why the iPod is gone. Of course, it could be that some inanimate object broke your window and then someone saw the iPod and took it. Or perhaps a dog jumped into the window that was broken by a stray baseball and ate your iPod. These are all possibilities, but they are remote and thus much less likely explanations of the facts at hand. The much better explanation is that a thief both broke the window and took the iPod. This explanation explains all the relevant facts in a simple way (i.e., it was the thief

responsible for both things) and this kind of thing is (unfortunately) not uncommon—it happens to other people at other times and places. The baseball-dog scenario is not as plausible because it doesn't happen in contexts like this one (i.e., in a parking garage) nearly as often and it is not as simple (i.e., we need to posit two different events that are unconnected to each other—stray baseball, stray dog—rather than just one—the thief). **Inference to the best explanation** is a form of inductive argument whose premises are a set of observed facts, a hypothesis that explains those observed facts, and a comparison of competing explanations, and whose conclusion is that the hypothesis is true. The example we've just been discussing is an inference to the best explanation. Here is its form:

1. Observed facts: Your car window is broken and your iPod is gone.
2. Explanation: The hypothesis that a thief broke the window and stole your iPod provides a reasonable explanation of the observed facts.
3. Comparison: No other hypothesis provides as reasonable an explanation.
4. Conclusion: Therefore, a thief broke your car window and stole your iPod.

Notice that this is an inductive argument because the premises could all be true and yet the conclusion false. Just because something is reasonable, doesn't mean it is true. After all, sometimes things happen in the world that defy our reason. So perhaps the baseball-dog hypothesis was actually true. In that case, the premises of the argument would still be true (after all, the thief hypothesis is still *more reasonable* than the baseball-dog hypothesis) and yet the conclusion would be false. But the fact that the argument is not a deductive argument isn't a defect of the argument, because inference to the best explanation arguments are not intended to be deductive arguments, but inductive arguments. As we saw in chapter 1, inductive arguments can be strong even if the premises don't entail the conclusion. That isn't a defect of an inductive argument, it is simply a definition of what an inductive argument is!

As we've seen, in order to make a strong inference to the best explanation, the favored explanation must be the best (or the most reasonable). But what makes an explanation reasonable? There are certain conditions that any good explanation must meet. The more of these conditions are met, the better the explanation. The first, and perhaps most obvious condition, is that the hypothesis proposed must actually explain all the observed facts. For example,

if, in order to explain the facts that your car window was broken and your iPod was missing, someone were to say offer the hypothesis that a rock thrown up from a lawnmower broke the window of your car, then this hypothesis wouldn't account for all the facts because it wouldn't explain the disappearance of your iPod. It would lack the explanatory virtue of explaining all the observed facts. The baseball-dog hypothesis would explain all the observed facts, but it would lack certain other explanatory virtues, such as "power" and "simplicity." In the remainder of this section, I will list the seven explanatory virtues and then I will discuss each one in turn. The seven explanatory virtues are:

1. **Explanatoriness:** Explanations must explain all the observed facts.
2. **Depth:** Explanations should not raise more questions than they answer.
3. **Power:** Explanations should apply in a range of similar contexts, not just the current situation in which the explanation is being offered.
4. **Falsifiability:** Explanations should be falsifiable—it must be possible for there to be evidence that would show that the explanation is incorrect.
5. **Modesty:** Explanations should not claim any more than is needed to explain the observed facts. Any details in the explanation must relate to explaining one of the observed facts.
6. **Simplicity:** Explanations that posit fewer entities or processes are preferable to explanations that posit more entities or processes. All other things being equal, the simplest explanation is the best. This is sometimes referred to as "**Ockham's razor**" after William of Ockham (1287-1347), the medieval philosopher and logician.
7. **Conservativeness:** Explanations that force us to give up fewer well-established beliefs are better than explanations that force us to give up more well-established beliefs.

Suppose that when confronted with the observed facts of my car window being broken and my iPod missing, my colleague Jeff hypothesizes that my colleague, Paul Jurczak did it. However, given that I am friends with Paul, that Paul could easily buy an iPod if he wanted one, and that I know Paul to be the kind of person who has probably never stolen anything in his life (much less broken a car window), this explanation would raise many more questions than it answers. Why would Paul want to steal my iPod? Why would he break my car window to do so? Etc. This explanation raises as many questions as it answers and thus it lacks the explanatory virtue of "depth."

Consider now an explanation that lacks the explanatory virtue of “power.” A good example would be the stray baseball scenario which is supposed to explain, specifically, the breaking of the car window. Although it is possible that a stray baseball broke my car window, that explanation would not apply in a range of similar contexts since people don’t play baseball in or around parking garages. So not many windows broken in parking garages can be explained by stray baseballs. In contrast, many windows broken in parking garages can be explained by thieves. Thus, the thief explanation would be a more powerful explanation, whereas the stray baseball explanation would lack the explanatory virtue of power.

Falsifiability can be a confusing concept to grasp. How can anything having to do with being *false* be a *virtue* of an explanation? An example will illustrate why the *possibility* of being false is actually a necessary condition for any good empirical explanation. Consider the following explanation. My socks regularly disappear and then sometime reappear in various places in the house. Suppose I were to explain this fact as follows. There is an invisible sock gnome that lives in our house. He steals my socks and sometimes he brings them back and sometimes he doesn’t. This explanation sounds silly and absurd, but how would you show that it is false? It seems that the hypothesis of the sock gnome is designed such that it cannot be shown to be false—it cannot be falsified. The gnome is invisible, so you can never see it do its thing. Since there is no way to observe it, it seems you can never prove nor disprove the existence of the sock gnome. Thus, you can neither confirm nor disconfirm the hypothesis. But such a hypothesis is a defective hypothesis. Any empirical hypothesis (i.e., a hypothesis that is supposed to explain a set of observed facts) must at least *be able to be shown false*. The sock gnome hypothesis lacks this virtue—that is, it lacks the explanatory virtue of being falsifiable. In contrast, if I were to hypothesize that our dog, Violet, ate the sock, then this hypothesis is falsifiable. For example, I could perform surgery on Violet and see if I found remnants of a sock. If I didn’t, then I would have shown that the hypothesis is false. If I did, then I would thereby have confirmed the hypothesis. So the “dog ate the sock” hypothesis is falsifiable, and this is a *good* thing. The difference between a true hypothesis and a false one is simply that the true hypothesis has not yet been shown to be false, whereas the false one has. Falsifiability requires only that it be *possible* to show that the hypothesis is false. If we look for evidence that would show that the hypothesis is false, but we won’t find that evidence, then we have confirmed that hypothesis. In contrast, an unfalsifiable hypothesis

cannot be confirmed because we cannot specify any evidence that would show it was false, so we can't try to look for such evidence (which is what a rigorous scientific methodology requires).

Suppose, to return to my broken window/missing iPod scenario, that my friend Chris hypothesized that a 24 year old Chinese man with a Tweety Bird tattoo on his left shoulder broke the window of my car and stole the iPod. This explanation would lack the explanatory virtue of "modesty." The problem is that the hypothesis is far more specific than it needs to be in order to explain the relevant observed facts. The details in any explanation should be relevant to explaining the observed facts. However, there is no reason to include the details that the thief was 24 years old, Chinese, and had a Tweety Bird tattoo on his left shoulder. How do those details help us to understand why the observed facts occurred? They don't. It would be just as explanatory to say, simply, that it was a thief rather than to include all those details about the thief, which don't help us to understand or explain any of the observed facts.

The explanatory virtue of "simplicity" tells us that all other things being equal, the simplest explanation is the better explanation. More precisely, an explanation that posits fewer entities or processes in order to explain the observed facts is better than an explanation that posits more entities and processes to explain that same set of observed facts. Here is an example of an explanation that would lack the virtue of simplicity. Suppose that all three of our cars in our driveway were broken into one night and that the next morning the passenger's side rear windows of each car were broken out. If I were to hypothesize that three separate, unrelated thieves at three different times of the night broke into each of the cars, then this would be an explanation that lacks the virtue of simplicity. The far simpler explanation is that it was one thief (or one related group of thieves) that broke into the three cars at roughly the same time. In the domain of science, upholding simplicity is often a matter of not positing new entities or laws when we can explain the observed facts in terms of existing entities and laws. My earlier example of the sock gnome stealing the socks vs. our dog Violet taking the socks is a good example to illustrate this. Sock gnomes would be a new kind of entity that we don't have any independent reason to think exists, but our dog Violet clearly already exists and since the observed facts can be explained by Violet's actions rather than that of a sock gnome, the Violet explanation possesses the explanatory virtue of simplicity, whereas the sock gnome explanation lacks the explanatory virtue of simplicity. However, sometimes science requires that we posit new kinds of entities or

processes, as when Copernicus and Galileo suggested that the sun, rather than the earth, was at the center of the “solar system” in order to explain certain astronomical observations. In physics new entities are often posited in order to explain the observations that physicists make. For example, the elementary particle dubbed “the Higgs boson” was hypothesized by Peter Higgs (and others) in 1964 and was confirmed in 2012. Much earlier, in 1897, J.J. Thompson and his collaborators, drawing on the work of earlier German physicists, discovered the electron—one of the first elementary particles to be discovered. So there is nothing wrong with positing new laws or entities—that is how science progresses. Simplicity doesn’t say that one should never posit new entities; that would be absurd. Rather, it tells us that *if* the observed facts can be explained without having to posit new entities, then that explanation is preferable to an explanation that does posit new entities (all other things being equal). Of course, sometimes the observations cannot be explained without having to change the way we understand that world. This is when it is legitimate to posit new entities or scientific laws.

The last explanatory virtue—conservativeness—tells us that better explanations are ones that force us to give up fewer well-established beliefs. Like simplicity, conservativeness is an explanatory virtue only when we are considering two explanations that each explain all the observed facts, but where one conflicts with well-established beliefs and the other doesn’t. In such a case, the former explanation would lack the explanatory virtue of conservativeness, whereas the latter explanation would possess the virtue of conservativeness. Here is an example to illustrate the virtue of conservativeness. Suppose that there are some photographs that vaguely seem to indicate a furry, bipedal humanoid creature that does not look human. My friend Chris offers the following explanation: the creature in those photos is Bigfoot, or Sasquatch. In contrast, I maintain that the creature in the photos is a person in a Bigfoot suit. Given just this evidence (the blurry photos), Chris’s explanation lacks the virtue of conservativeness since his explanation requires the existence of Bigfoot, which is contrary to well-established beliefs that Bigfoot is merely folklore, not a real creature. In contrast, my explanation possesses the virtue of conservativeness since there is nothing about someone dressing up in a costume and being caught on camera (or even someone doing so to play a practical joke or to perpetuate a false belief in a certain population) that conflicts with well-established beliefs. My explanation doesn’t require the existence of Bigfoot, but just the existence of human beings dressed up to look like Bigfoot.

It should be stated that some of the examples I have given could illustrate more than one explanatory virtue. For example, the example of the invisible sock gnome hypothesis could illustrate either lack of falsifiability or lack of simplicity. In identifying which explanatory virtues a particular explanation may lack, what is important is that you give the correct reasoning for why the explanation lacks that particular virtue. For example, if you say that the explanation isn't falsifiable, then you need to make sure you give the right explanation of why it isn't falsifiable (i.e., that there is no evidence that could ever show that the hypothesis is false). In contrast, if the explanation lacks simplicity, you'd have to say that there is another explanation that can equally explain all the observed facts but that posits fewer entities or processes.

Exercise 23: Identify which explanatory virtues, if any, the following explanations lack and explain why it lacks that particular virtue. If there is a better explanation, suggest what it might be.

1. Bob explains the fact that he can't remember what happened yesterday by saying that he must have been kidnapped by aliens, who performed surgery on him and then erased his memory of everything that happened the day before returning him to his house.
2. Mrs. Jones hears strange noises at night such as the creaking of the floor downstairs and rattling of windows. She explains these phenomena by hypothesizing that there is a 37-pound badger that inhabits the house and that emerges at night in search of Wheat Thins and Oreos.
3. Edward saw his friend Tom at the store in their hometown of Lincoln, Nebraska just an hour ago. Then, while watching the World Cup on television, he saw someone that looked just like Tom in the crowd at the game in Brazil. He hypothesizes that his friend Tom must have an identical twin that Tom has never told him about.
4. Edward's friend Tom died two years ago. But just yesterday Tom saw someone who looked and spoke exactly like Tom. Edward hypothesizes that Tom must have come back to life.
5. Edward's friend Tom died twenty years ago when Tom was just 18. But just yesterday Edward saw someone who looked and spoke exactly like Tom. Edward hypothesizes that Tom must have had a son that he did not know about and that this person must have been Tom's son.

6. Elise has the uncanny feeling that although her family members look exactly the same, something just isn't right about them. She hypothesizes that her family members have been replaced with imposters who look and act exactly like her real family members and that no one can prove that this didn't happen.
7. John thinks that since something cannot come from nothing and since we know there was a Big Bang, an all-powerful but invisible and undetectable being must have been the cause of the Big Bang.
8. Erin feels that she is being followed. Every time she looks over her shoulder, she sees someone duck behind an object to avoid being seen. She hypothesizes that it must be her 5th grade teacher, Mr. Sanchez.
9. While walking through the forest at night, Claudia hears some rustling in the bushes. It is clear to her that it isn't just the wind, because she can hear sticks cracking on the ground. She hypothesizes that it must be an escaped zoo animal.
10. While driving on the freeway, Bill sees the flashing lights of a cop car in his rear view mirror. He hypothesizes that the cops must have finally found out about the library book that he never returned when he was in fifth grade and are coming to get him.
11. While driving on the freeway, Bill sees the flashing lights of a cop car in his rear view mirror. He hypothesizes that the cops are going to pull someone over for speeding.
12. While driving on the freeway, Bill sees the flashing lights of a cop car in his rear view mirror. He hypothesizes that the cops are going to pull someone over for going 13.74 mph over the speed limit.
13. Stacy cannot figure out why the rat poison she is using is not killing the rats in her apartment. She hypothesizes that the rats must be a new breed of rats that are resistant to any kind of poison and that evolved in the environment of her apartment.
14. Stacy cannot figure out why the rat poison she is using is not killing the rats in her apartment. She hypothesizes that the rats must be a new breed of rats that are immortal and that evolved in the environment of her apartment.
15. Bob is fed up with his life. He intends to kill himself so he gets his gun, puts bullets into it and pull the trigger. Miraculously, he is not killed. Bob hypothesizes that he must be immortal.

3.3 Analogical arguments

Another kind of common inductive argument is an argument from analogy. In an **argument from analogy**, we note that since some thing x shares similar properties to some thing y , then since y has characteristic A , x probably has characteristic A as well. For example, suppose that I have always owned Subaru cars in the past and that they have always been reliable and I argue that the new car I've just purchased will also be reliable because it is a Subaru. The two things in the analogy are 1) the Subarus I have owned in the past and 2) the current Subaru I have just purchased. The similarity between these two things is just that they are both Subarus. Finally, the conclusion of the argument is that *this* Subaru will share the characteristic of *being reliable* with the past Subarus I have owned. Is this argument a strong or weak inductive argument? Partly it depends on how many Subarus I've owned in the past. If I've only owned one, then the inference seems fairly weak (perhaps I was just lucky in that one Subaru I've owned). If I've owned ten Subarus then the inference seems much stronger. Thus, the reference class that I'm drawing on (in this case, the number of Subarus I've previously owned) must be large enough to generalize from (otherwise we would be committing the fallacy of "hasty generalization"). However, even if our reference class was large enough, what would make the inference even stronger is knowing not simply that the new car is a Subaru, but also specific things about its origin. For example, if I know that this particular model has the same engine and same transmission as the previous model I owned and that nothing significant has changed in how Subarus are made in the intervening time, then my argument is strengthened. In contrast, if this new Subaru was made after Subaru was bought by some other car company, and if the engine and transmission were actually made by this new car company, then my argument is weakened. It should be obvious why: the fact that the car is still called "Subaru" is not relevant establishing that it will have the same characteristics as the other cars that I've owned that were called "Subarus." Clearly, what the car is called has no inherent relevance to whether the car is reliable. Rather, what is relevant to whether the car is reliable is the quality of the parts and assembly of the car. Since it is possible that car companies can retain their name and yet drastically alter the quality of the parts and assembly of the car, it is clear that the name of the car isn't *itself* what establishes the quality of the car. Thus, the original argument, which invoked merely that the new car was a Subaru is not as strong as the argument that the car was

constructed with the same quality parts and quality assembly as the other cars I'd owned (and that had been reliable for me). What this illustrates is that better arguments from analogy will invoke more *relevant similarities* between the things being compared in the analogy. This is a key condition for any good argument from analogy: the similar characteristics between the two things cited in the premises must be *relevant* to the characteristic cited in the conclusion.

Here is an ethical argument that is an argument from analogy.¹ Suppose that Bob uses his life savings to buy an expensive sports car. One day Bob parks his car and takes a walk along a set of train tracks. As he walks, he sees in the distance a small child whose leg has become caught in the train tracks. Much to his alarm, he sees a train coming towards the child. Unfortunately, the train will reach the child before he can (since it is moving very fast) and he knows it will be unable to stop in time and will kill the child. At just that moment, he sees a switch near him that he can throw to change the direction of the tracks and divert the train onto another set of tracks so that it won't hit the child. Unfortunately, Bob sees that he has unwittingly parked his car on that other set of tracks and that if he throws the switch, his expensive car will be destroyed. Realizing this, Bob decides not to throw the switch and the train strikes and kills the child, leaving his car unharmed. What should we say of Bob? Clearly, that was a horrible thing for Bob to do and we would rightly judge him harshly for doing it. In fact, given the situation described, Bob would likely be criminally liable. Now consider the following situation in which you, my reader, likely find yourself (whether you know it or not—well, now you *do* know it). Each week you spend money on things that you do not need. For example, I sometimes buy \$5 espressos from Biggby's or Starbuck's. I do not need to have them and I could get a much cheaper caffeine fix, if I chose to (for example, I could make a strong cup of coffee at my office and put sweetened hazelnut creamer in it). In any case, I really don't need the caffeine at all! And yet I regularly purchase these \$5 drinks. (If \$5 drinks aren't the thing you spend money on, but in no way need, then fill in the example with whatever it is that fits your own life.) With the money that you could save from forgoing these luxuries, you could, quite literally, save a child's life. Suppose (to use myself as an example) I were to buy two \$5 coffees a week (a conservative estimate). That is \$10 a week, roughly \$43 a month and \$520 a year. Were I to donate that amount (just \$40/month) to an organization such as the Against Malaria Foundation, I could save a child's

¹ This argument comes (with interpretive liberties on my part) from Peter Singer's, "The Singer Solution to World Poverty" published in the NY Times Magazine, September 5, 1999.

life in just six years.² Given these facts, and comparing these two scenarios (Bob's and your own), the argument from analogy proceeds like this:

1. Bob chose to have a luxury item for himself rather than to save the life of a child.
2. "We" regularly choose having luxury items rather than saving the life of a child.
3. What Bob did was morally wrong.
4. Therefore, what we are doing is morally wrong as well.

The two things being compared here are Bob's situation and our own. The argument then proceeds by claiming that since we judge what Bob did to be morally wrong, and since our situation is analogous to Bob's in relevant respects (i.e., choosing to have luxury items for ourselves rather than saving the lives of dying children), then our actions of purchasing luxury items for ourselves must be morally wrong for the same reason.

One way of arguing against the conclusion of this argument is by trying to argue that there are relevant disanalogies between Bob's situation and our own. For example, one might claim that in Bob's situation, there was something much more immediate he could do to save the child's life right then and there. In contrast, our own situation is not one in which a child that is physically proximate to us is in imminent danger of death, where there is something we can immediately do about it. One might argue that this disanalogy is enough to show that the two situations are not analogous and that, therefore, the conclusion does not follow. Whether or not this response to the argument is adequate, we can see that the way of objecting to an argument from analogy is by trying to show that there are relevant differences between the two things being compared in the analogy. For example, to return to my car example, even if the new car was a Subaru and was made under the same conditions as all of my other Subarus, if I purchased the current Subaru used, whereas all the other Subarus had been purchased new, then that could be a relevant difference that would weaken the conclusion that this Subaru will be reliable.

So we've seen that an argument from analogy is strong only if the following two conditions are met:

² <http://www.givewell.org/giving101/Your-dollar-goes-further-overseas>

1. The characteristics of the two things being compared must be *similar in relevant respects* to the characteristic cited in the conclusion.
2. There must not be any *relevant disanalogies* between the two things being compared.

Arguments from analogy that meet these two conditions will tend to be stronger inductive arguments.

Exercise 24: Evaluate the following arguments from analogy as either strong or weak. If the argument is weak, cite what you think would be a relevant disanalogy.

1. Every painting by Rembrandt contains dark colors and illuminated faces, therefore the original painting that hangs in my high school is probably by Rembrandt, since it contains dark colors and illuminated faces.
2. I was once bitten by a poodle. Therefore, this poodle will probably bite me too.
3. Every poodle I've ever met has bitten me (and I've met over 300 poodles). Therefore, this poodle will probably bite me too.
4. My friend took Dr. Van Cleave's logic class last semester and got an A. Since Dr. Van Cleave's class is essentially the same this semester and since my friend is no better a student than I am, I will probably get an A as well.
5. Bill Cosby used his power and position to seduce and rape women. Therefore, Bill Cosby probably also used his power to rob banks.
6. Every car I've ever owned had seats, wheels and brakes and was also safe to drive. This used car that I am contemplating buying has seats, wheels and brakes. Therefore, this used car is probably safe to drive.
7. Every Volvo I've ever owned was a safe car to drive. My new car is a Volvo. Therefore, my new car is probably safe to drive.
8. Dr. Van Cleave did not give Jones an excused absence when Jones missed class for his grandmother's funeral. Mary will have to miss class to attend her aunt's funeral. Therefore, Dr. Van Cleave should not give Mary an excused absence either.
9. Dr. Van Cleave did not give Jones an excused absence when Jones missed class for his brother's birthday party. Mary will have to miss class to attend her aunt's funeral. Therefore, Dr. Van Cleave should not give Mary an excused absence either.

10. If health insurance companies pay for heart surgery and brain surgery, which can both increase an individual's happiness, then they should also pay for cosmetic surgery, which can also increase an individual's happiness.
11. A knife is an eating utensil that can cut things. A spoon is also an eating utensil. So a spoon can probably cut things as well.
12. Any artificial, complex object like a watch or a telescope has been designed by some intelligent human designer. But naturally occurring objects like eyes and brains are also very complex objects. Therefore, complex naturally occurring objects must have been designed by some intelligent non-human designer.
13. The world record holding runner, Kenenisa Bekele ran 100 miles per week and twice a week did workouts comprised of ten mile repeats on the track in the weeks leading up to his 10,000 meter world record. I have run 100 miles per week and have been doing ten mile repeats twice a week. Therefore, the next race I will run will probably be a world record.
14. I feel pain when someone hits me in the face with a hockey puck. We are both human beings, so you also probably feel pain when you are hit in the face with a hockey puck.
15. The color I experience when I see something as "green" has a particular quality (that is difficult to describe). You and I are both human beings, so the color you experience when you see something green probably has the exact same quality. (That is, what you and I experience when we see something green is the exact same experiential color.)

3.4 Causal reasoning

When I strike a match, it will produce a flame. It is natural to take the striking of the match as the cause that produces the effect of a flame. But what if the matchbook is wet? Or what if I happen to be in a vacuum in which there is no oxygen (such as in outer space)? If either of those things is the case, then the striking of the match will not produce a flame. So it isn't simply the striking of the match that produces the flame, but a combination of the striking of the match together with a number of other conditions that must be in place in order for the striking of the match to create a flame. Which of those conditions we call the "cause" depends in part on the context. Suppose that I'm in outer space

striking a match (suppose I'm wearing a space suit that supplies me with oxygen but that I'm striking the match in space, where there is no oxygen). I continuously strike it but no flame appears (of course). But then someone (also in a space suit) brings out a can of compressed oxygen that they spray on the match while I strike it. All of a sudden a flame is produced. In this context, it looks like it is the spraying of oxygen that causes flame, not the striking of the match. Just as in the case of the striking of the match, any cause is more complex than just a simple event that produces some other event. Rather, there are always multiple conditions that must be in place for any cause to occur. These conditions are called **background conditions**. That said, we often take for granted the background conditions in normal contexts and just refer to one particular event as the cause. Thus, we call the striking of the match the cause of the flame. We don't go on to specify all the other conditions that conspired to create the flame (such as the presence of oxygen and the absence of water). But this is more for convenience than correctness. For just about any cause, there are a number of conditions that must be in place in order for the effect to occur. These are called necessary conditions (recall the discussion of necessary and sufficient conditions from chapter 2, section 2.7). For example, a necessary condition of the match lighting is that there is oxygen present. A necessary condition of a car running is that there is gas in the tank. We can use necessary conditions to diagnose what has gone wrong in cases of malfunction. That is, we can consider each condition in turn in order to determine what caused the malfunction. For example, if the match doesn't light, we can check to see whether the matches are wet. If we find that the matches are wet then we can explain the lack of the flame by saying something like, "dropping the matches in the water caused the matches not to light." In contrast, a sufficient condition is one which, if present, will always bring about the effect. For example, a person being fed through an operating wood chipper is sufficient for causing that person's death (as was the fate of Steve Buscemi's character in the movie *Fargo*).

Because the natural world functions in accordance with natural laws (such as the laws of physics), causes can be generalized. For example, any object near the surface of the earth will fall towards the earth at 9.8 m/s^2 unless impeded by some contrary force (such as the propulsion of a rocket). This generalization applies to apples, rocks, people, wood chippers and every other object. Such causal generalizations are often parts of explanations. For example, we can explain why the airplane crashed to the ground by citing the causal generalization that *all* unsupported objects fall to the ground and by noting that the airplane had lost any method of propelling itself because the engines had

died. So we invoke the causal generalization in explaining why the airplane crashed. Causal generalizations have a particular form:

For any x , if x has the feature(s) F , then x has the feature G

For example:

For any *human*, if *that human* has been fed through an operating wood chipper, then *that human* is dead.

For any *engine*, if *that engine* has no fuel, then *that engine* will not operate.

For any *object near the surface of the earth*, if *that object* is unsupported and not impeded by some contrary force, then *that object* will fall towards the earth at 9.8 m/s^2 .

Being able to determine when causal generalizations are true is an important part of becoming a critical thinker. Since in both scientific and every day contexts we rely on causal generalizations in explaining and understanding our world, the ability to assess when a causal generalization is true is an important skill. For example, suppose that we are trying to figure out what causes our dog, Charlie, to have seizures. To simplify, let's suppose that we have a set of potential candidates for what causes his seizures. It could be either:

- A) eating human food,
- B) the shampoo we use to wash him,
- C) his flea treatment,
- D) not eating at regular intervals,

or some combination of these things. Suppose we keep a log of when these things occur each day and when his seizures (S) occur. In the table below, I will represent the absence of the feature by a negation. So in the table below, " $\sim A$ " represents that Charlie did not eat human food on that day; " $\sim B$ " represents that he did not get a bath and shampoo that day; " $\sim S$ " represents that he did not have a seizure that day. In contrast, " B " represents that he did have a bath and shampoo, whereas " C " represents that he was given a flea treatment that day. Here is how the log looks:

Day 1	~A	B	C	D	S
Day 2	A	~B	C	D	~S
Day 3	A	B	~C	D	~S
Day 4	A	B	C	~D	S
Day 5	A	B	~C	D	~S
Day 6	A	~B	C	D	~S

How can we use this information to determine what might be causing Charlie to have seizures? The first thing we'd want to know is what feature is present every time he has a seizure. This would be a necessary (but not sufficient) condition. And that can tell us something important about the cause. The **necessary condition test** says that any candidate feature (here A, B, C, or D) that is absent when the target feature (S) is present is eliminated as a possible necessary condition of S.³ In the table above, A is absent when S is present, so A can't be a necessary condition (i.e., day 1). D is also absent when S is present (day 4) so D can't be a necessary condition either. In contrast, B is never absent when S is present—that is every time S is present, B is also present. That means B is a necessary condition, based on the data that we have gathered so far. The same applies to C since it is never absent when S is present. Notice that there are times when both B and C are absent, but on those days the target feature (S) is absent as well, so it doesn't matter.

The next thing we'd want to know is which feature is such that every time it is present, Charlie has a seizure. The test that is relevant to determining this is called the sufficient condition test. The **sufficient condition test** says that any candidate that is present when the target feature (S) is absent is eliminated as a possible sufficient condition of S. In the table above, we can see that no one candidate feature is a sufficient condition for causing the seizures since for each candidate (A, B, C, D) there is a case (i.e. day) where it is present but that no seizure occurred. Although no one feature is sufficient for causing the seizures (according to the data we have gathered so far), it is still possible that certain features are **jointly sufficient**. Two candidate features are jointly sufficient for a target feature if and only if there is no case in which both candidates are present

³ This discussion draws heavily on chapter 10, pp. 220-224 of Sinnott-Armstrong and Fogelin's *Understanding Arguments*, 9th edition (Cengage Learning).

and yet the target is absent. Applying this test, we can see that B and C are jointly sufficient for the target feature since any time both are present, the target feature is always present. Thus, from the data we have gathered so far, we can say that the likely cause of Charlie’s seizures are when we both give him a bath and then follow that bath up with a flea treatment. Every time those two things occur, he has a seizure (sufficient condition); and every time he has a seizure, those two things occur (necessary condition). Thus, the data gathered so far supports the following causal conditional:

Any time Charlie is given a shampoo bath and a flea treatment, he has a seizure.

Although in the above case, the necessary and sufficient conditions were the same, this needn’t always be the case. Sometimes sufficient conditions are not necessary conditions. For example, being fed through a wood chipper is a sufficient condition for death, but it certainly isn’t necessary! (Lot’s of people die without being fed through a wood chipper, so it can’t be a necessary condition of dying.) In any case, determining necessary and sufficient conditions is a key part of determining a cause.

When analyzing data to find a cause it is important that we rigorously test each candidate. Here is an example to illustrate rigorous testing. Suppose that on every day we collected data about Charlie he ate human food but that on none of the days was he given a bath and shampoo, as the table below indicates.

Day 1	A	~B	C	D	~S
Day 2	A	~B	C	D	~S
Day 3	A	~B	~C	D	~S
Day 4	A	~B	C	~D	S
Day 5	A	~B	~C	D	~S
Day 6	A	~B	C	D	S

Given this data, A trivially passes the necessary condition test since it is always present (thus, there can never be a case where A is absent when S is present). However, in order to rigorously test A as a necessary condition, we have to look for cases in which A is not present and then see if our target condition S is present. We have rigorously tested A as a necessary condition only if we have collected data in which A was not present. Otherwise, we don’t really know whether A is a necessary condition. Similarly, B trivially passes the sufficient

condition test since it is never present (thus, there can never be a case where B is present but S is absent). However, in order to rigorously test B as a sufficient condition, we have to look for cases in which B is present and then see if our target condition S is absent. We have rigorously tested B as a sufficient condition only if we have collected data in which B is present. Otherwise, we don't really know whether B is a sufficient condition or not.

In **rigorous testing**, we are actively looking for (or trying to create) situations in which a candidate feature *fails* one of the tests. That is why when rigorously testing a candidate for the necessary condition test, we must seek out cases in which the candidate is not present, whereas when rigorously testing a candidate for the sufficient condition test, we must seek out cases in which the candidate is present. In the example above, A is not rigorously tested as a necessary condition and B is not rigorously tested as a sufficient condition. If we are interested in finding a cause, we should always rigorously test each candidate. This means that we should always have a mix of different situations where the candidates and targets are sometimes present and sometimes absent.

The necessary and sufficient conditions tests can be applied when features of the environment are wholly present or wholly absent. However, in situations where features of the environment are always present in some degree, these tests will not work (since there will never be cases where the features are absent and so rigorous testing cannot be applied). For example, suppose we are trying to figure out whether CO₂ is a contributing cause to higher global temperatures. In this case, we can't very well look for cases in which CO₂ is present but high global temperatures aren't (sufficient condition test), since CO₂ and high temperatures are always present to some degree. Nor can we look for cases in which CO₂ is absent when high global temperatures are present (necessary condition test), since, again, CO₂ and high global temperatures are always present to some degree. Rather, we must use a different method, the method that J.S. Mill called **the method of concomitant variation**. In concomitant variation we look for how things vary vis-à-vis each other. For example, if we see that as CO₂ levels rise, global temperatures also rise, then this is evidence that CO₂ and higher temperatures are positively correlated. When two things are **positively correlated**, as one increases, the other also increases at a similar rate (or as one decreases, the other decreases at a similar rate). In contrast, when two things are **negatively correlated**, as one increases, the other decreases at similar rate (or vice versa). For example, if as a police department increased the number of police officers on the street, the number of crimes reported

decreases, then number of police on the street and number of crimes reported would be negatively correlated. In each of these examples, we may think we can directly infer the cause from the correlation—the rising CO₂ levels are causing the rising global temperatures and the increasing number of police on the street is causing the crime rate to drop. However, we cannot directly infer causation from correlation. Correlation is not causation. If A and B are positively correlated, then there are four distinct possibilities regarding what the cause is:

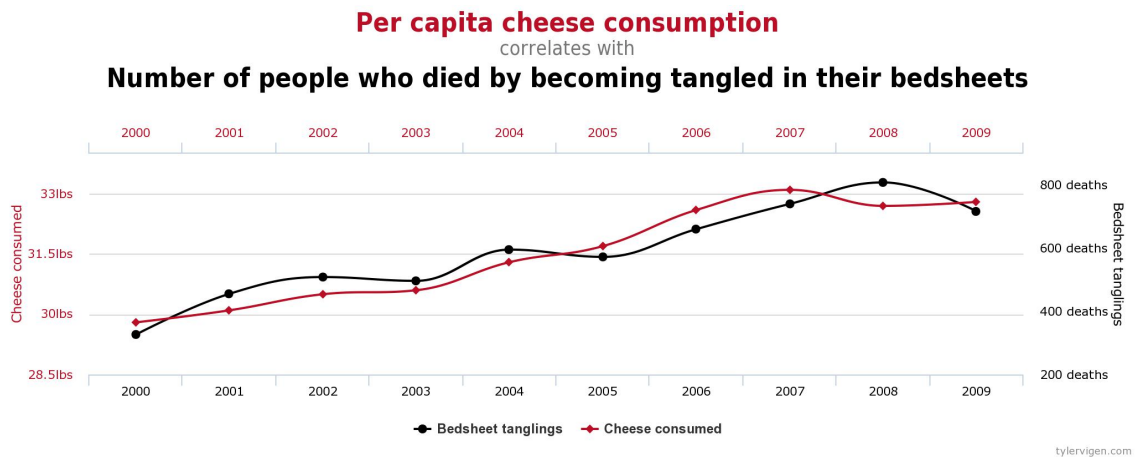
1. A is the cause of B
2. B is the cause of A
3. Some third thing, C, is the cause of both A and B increasing
4. The correlation is accidental

In order to infer what causes what in a correlation, we must rely on our general **background knowledge** (i.e., things we know to be true about the world), our scientific knowledge, and possibly further scientific testing. For example, in the global warming case, there is no scientific theory that explains how rising global temperatures could cause rising levels of CO₂ but there is a scientific theory that enables us to understand how rising levels of CO₂ could increase average global temperatures. This knowledge makes it plausible to infer that the rising CO₂ levels are causing the rising average global temperatures. In the police/crime case, drawing on our background knowledge we can easily come up with an inference to the best explanation argument for why increased police presence on the streets would lower the crime rate—the more police on the street, the harder it is for criminals to get away with crimes because there are fewer places where those crimes could take place without the criminal being caught. Since criminals don't want to risk getting caught when they commit a crime, seeing more police around will make them less likely to commit a crime. In contrast, there is no good explanation for why decreased crime would cause there to be more police on the street. In fact, it would seem to be just the opposite: if the crime rate is low, the city should cut back, or at least remain stable, on the number of police officers and put those resources somewhere else. This makes it plausible to infer that it is the increased police officers on the street that is causing the decrease in crime.

Sometimes two things can be correlated without either one causing the other. Rather, some third thing is causing them both. For example, suppose that Bob discovers a correlation between waking up with all his clothes on and waking up

with a headache. Bob might try to infer that sleeping with all his clothes on causes headaches, but there is probably a better explanation than that. It is more likely that Bob's drinking too much the night before caused him to pass out in his bed with all his clothes on, as well as his headache. In this scenario, Bob's inebriation is the **common cause** of both his headache and his clothes being on in bed.

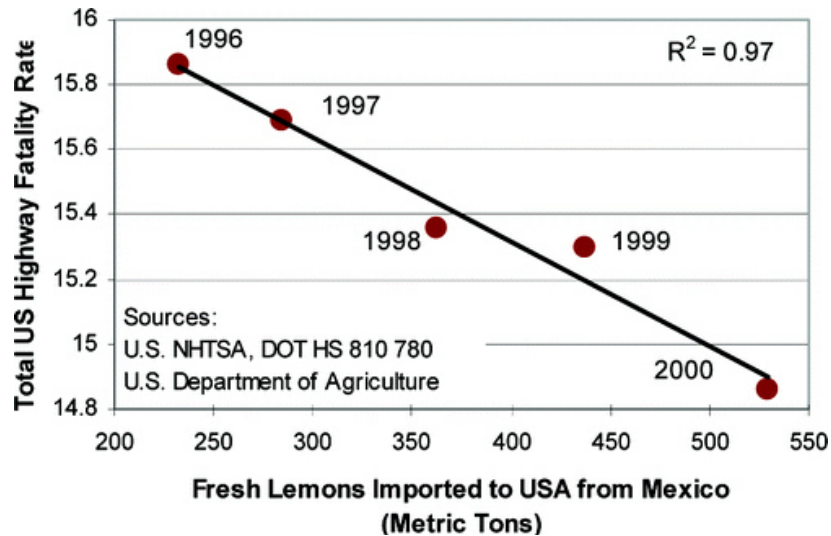
Sometimes correlations are merely accidental, meaning that there is no causal relationship between them at all. For example, Tyler Vigen⁴ reports that the per capita consumption of cheese in the U.S. correlates with the number of people who die by becoming entangled in their bedsheets:



And the number of Mexican lemons imported to the U.S. correlates with the number of traffic fatalities⁵:

⁴ <http://tylervigen.com/spurious-correlations>

⁵ Stephen R. Johnson, The Trouble with QSAR (or How I Learned To Stop Worrying and Embrace Fallacy). *J. Chem. Inf. Model.*, 2008, 48 (1), pp. 25–26.



Clearly neither of these correlations are causally related at all—they are **accidental correlations**. What makes them accidental is that we have no theory that would make sense of how they could be causally related. This just goes to show that it isn't simply the correlation that allows us to infer a cause, but, rather, some additional background theory, scientific theory, or other evidence that establishes one thing as causing another. We can explain the relationship between correlation and causation using the concepts of necessary and sufficient conditions (first introduced in chapter 2): correlation is a necessary condition for causation, but it is not a sufficient condition for causation.

Our discussion of causes has shown that we cannot say that just because A precedes B or is correlated with B, that A caused B. To claim that since A precedes or correlates with B, A must therefore be the cause of B is to commit what is called the **false cause fallacy**. The false cause fallacy is sometimes called the "post hoc" fallacy. "Post hoc" is short for the Latin phrase, "post hoc ergo propter hoc," which means "after this therefore because of this." As we've seen, false cause fallacies occur any time someone assumes that two events that are correlated must be in a causal relationship, or that since one event precedes another, it must cause the other. To avoid the false cause fallacy, one must look more carefully into the relationship between A and B to determine whether there is a true cause or just a common cause or accidental correlation. Common causes and accidental correlations are more common than one might think.

Exercise 25: Determine which of the candidates (A, B, C, D) in the following examples pass the necessary condition test or the sufficient

condition test relative to the target (G). In addition, note whether there are any candidates that aren't rigorously tested as either necessary or sufficient conditions.

1.

Case 1	A	B	~C	D	~G
Case 2	~A	B	C	D	G
Case 3	A	~B	C	D	G

2.

Case 1	A	B	C	D	G
Case 2	~A	B	~C	D	~G
Case 3	A	~B	C	~D	G

3.

Case 1	A	B	C	D	G
Case 2	~A	B	C	D	G
Case 3	A	~B	C	D	G

4.

Case 1	A	B	C	D	~G
Case 2	~A	B	C	D	G
Case 3	A	~B	C	~D	G

5.

Case 1	A	B	~C	D	~G
Case 2	~A	B	C	D	G
Case 3	A	~B	~C	~D	~G

6.

Case 1	A	B	C	D	~G
Case 2	~A	B	C	~D	~G
Case 3	A	~B	~C	D	G

7.

Case 1	A	B	~C	D	G
--------	---	---	----	---	---

Case 2	~A	~B	C	D	~G
Case 3	A	~B	C	~D	~G

8.

Case 1	A	B	~C	D	~G
Case 2	~A	~B	~C	D	~G
Case 3	A	~B	~C	~D	~G

9.

Case 1	A	B	C	D	G
Case 2	~A	~B	C	D	G
Case 3	A	~B	~C	D	~G

10.

Case 1	~A	B	~C	D	~G
Case 2	~A	B	C	D	G
Case 3	~A	~B	~C	D	G

Exercise 26: For each of the following correlations, use your background knowledge to determine whether A causes B, B causes A, a common cause C is the cause of both A and B, or the correlations is accidental.

1. There is a positive correlation between U.S. spending on science, space, and technology (A) and suicides by hanging, strangulation, and suffocation (B).
2. There is a positive correlation between our dog Charlie's weight (A) and the amount of time we spend away from home (B). That is, the more time we spend away from home, the heavier Charlie gets (and the more we are at home, the lighter Charlie is).
3. The height of the tree in our front yard (A) positively correlates with the height of the shrub in our backyard (B).
4. There is a negative correlation between the number of suicide bombings in the U.S. (A) and the number of hairs on a particular U.S. President's head (B).
5. There is a high positive correlation between the number of fire engines in a particular borough of New York City (A) and the number of fires that occur there (B).

6. At one point in history, there was a negative correlation between the number of mules in the state (A) and the salaries paid to professors at the state university (B). That is, the more mules, the lower the professors' salaries.
7. There is a strong positive correlation between the number of traffic accidents on a particular highway (A) and the number of billboards featuring scantily-clad models (B).
8. The girth of an adult's waist (A) is negatively correlated with the height of their vertical leap (B).
9. Olympic marathon times (A) are positively correlated with the temperature during the marathon (B). That is, the more time it takes an Olympic marathoner to complete the race, the higher the temperature.
10. The number gray hairs on an individual's head (A) is positively correlated with the number of children or grandchildren they have (B).

3.5 Probability

As we have seen, a strong inductive argument is one in which the truth of the premises makes the conclusion highly *probable*. The distinction between strong inductive arguments and valid (deductive) arguments is that whereas the premises of strong inductive arguments make their conclusions *highly probable*, the premises of valid arguments make their conclusions *certain*. We can think of probability as how likely it is that something is (or will be) true, given a particular body of evidence. Using numbers between 0 and 1, we can express probabilities numerically. For example, if I have a full deck of cards and pick one at random, what is the probability that the card I pick is a queen? Since there are 52 cards in the deck, and only four of them are queens, the probability of picking a queen is $4/52$, or .077. That is, I have about a 7.7% chance of picking a queen at random. In comparison, my chances of picking any "face" card would be much higher. There are three face cards in each suit and four different suits, which means there are 12 face cards total. So, $12/52 = .23$ or 23%. In any case, the important thing here is that probabilities can be expressed numerically. In using a numerical scheme to represent probabilities, we take 0 to represent an impossible event (such as a contradiction) and 1 to represent an event that is certain (such as a tautology).

Probability is important to understand because it provides the basis for formal methods of evaluating inductive arguments. While there is no universally agreed upon method of evaluating inductive arguments in the way there is with deductive arguments, there are some basic laws of probability that it is important to keep in mind. As we will see in the next few sections, although these laws of probability are seemingly simple, we misapply them all the time.

We can think of the rules of probability in terms of some of the truth functional operators, introduced in chapter 2: the probability of conjunctions, the probability of negations, and the probability of disjunctions. The probability of conjunctions is the probability that two, independent events will both occur. For example, what is the probability that you randomly draw a queen and then (after returning it to the pile and reshuffling the deck) you draw another queen? Since we are asking what is the probability that these two events both occur, this is a matter of calculating the probability of a joint occurrence. In the following, "a" and "b" will refer to independent events, and the locution "P(a)" stands for "the probability of a." Here is how we calculate the **probability of conjunctions**:

$$P(a \text{ and } b) = P(a) \times P(b)$$

So, to apply this to my example of drawing two queens, we have to multiply the probability of drawing one queen, "P(a)" by the probability of drawing yet another queen, "P(b)." Since we have already calculated the probability of drawing a queen at .077, the math is quite simple:

$$.077 \times .077 = .0059$$

That is, there a less than 1% chance (.59% to be precise) of drawing two queens in this scenario. So, obviously, you'd not be wise to place a bet on that happening! Let's try another example where we have to calculate the probability of a conjunction. Suppose I want to know what the probability that both my father and mother will die of brain cancer. (Macabre, I know.) I'd have to know the probability of dying of brain cancer, which is about 5/100,000. That is, 5 out of every 100,000 people die of brain cancer. That is a very small number: .00005. But the chance of both of them dying of brain cancer is going to be an even smaller number:

$$.00005 \times .00005 = .0000000025$$

That is almost 1 in a *billion* chance. So not very likely. Let's consider a final example with more manageable numbers. Suppose I wanted to know the probability of rolling a 12 when rolling two, six-sided dice. Since the only way to roll a 12 is when I roll a 6 on each die, I can compute the probability of rolling a 6 and then the independent probability of rolling another 6 on the other die. The probability of rolling a six on 1 die is just $1/6 = .166$. Thus,

$$.166 \times .166 = .028$$

Thus, you have a 2.8% chance of rolling a 12. We could have also calculated this using fractions instead of decimals:

$$1/6 \times 1/6 = 1/36$$

Calculating the **probability of negations** is simply a matter of subtracting the probability that some event, say event a, will occur from 1. The result is the probability that event a will not occur:

$$P(\text{not-a}) = 1 - P(a)$$

For example, suppose I am playing monopoly and I wanted to determine the probability that I do *not* roll a 12 (since if I roll a 12 I will land on Boardwalk, which my opponent owns with hotels). Since we have already determined that the probability of rolling a 12 is .028, we can calculate the probability of not rolling a 12 thus:

$$1 - .028 = .972$$

Thus, I have 97.2% chance of not rolling a 12. So it is highly likely that I won't (thank goodness).

Here's another example. What are the chances that my daughter doesn't get into Harvard? Since the acceptance rate at Harvard is about 6% (or .06), I simply subtract that from 1, which yields .94, or 94%. So my daughter has a 94% chance of not getting into Harvard.

We should pause here to make some comments about probability. The probability of an event occurring is relative to some reference class. So, for example, the probability of getting osteoporosis is much higher if you are a

woman over 50 (16%) than if you are a man over 50 (4%). So if you want accurate data concerning probability, you have to take into account all the relevant factors. In the case of osteoporosis, that means knowing whether you are a woman or a man and are over or under 50. The same kind of point applies to my example of getting into Harvard. Here's an anecdote that will illustrate the point. Some years ago, I agreed to be a part of an interviewing process for candidates for the "presidential scholarship" at the college at which I was teaching at the time. The interviewees were high school students and we could have calculated the probability that any one of them would win the scholarship simply by noting the number of scholarships available and the number of applicants for them. But after having interviewed the candidates I was given to interview, it was very clear to me that one of them easily outshined all the rest. Thus, given the new information I had, it would have been silly for me to assign the same, generic probability to this student winning the award. This student was extremely well-spoken, well-put-together, and answered even my hardest questions (with which other candidates struggled) with an ease and confidence that stunned me. On top of all of that, she was a Hispanic woman, which I knew would only help her in the process (since colleges value diversity in their student population). I recommended her highly for the scholarship, but I also knew that she would end up at a much better institution (and probably with one of their most competitive scholarships). Some time later, I was wondering where she did end up going to college, so I did a quick search on her name and, sure enough, she was a freshman at Harvard. No surprise to me. The point of the story is that although we could have said that this woman's chances of not getting into Harvard are about 94%, this would neglect all the other things about her which in fact drastically increase her chances of getting into Harvard (and thus drastically decrease her chances of *not* getting in). So our assessments of probability are only as good as the information we use to assess them. If we were omniscient (i.e., all-knowing), then arguably we could know every detail and would be able to predict with 100% accuracy any event. Since we aren't, we have to rely on the best information we do have and use that information to determine the *chances* that an event will occur.

Calculating the probability of disjunctions is simply a matter of figuring out the probability that either one event or another will occur. To calculate the **probability of a disjunction** we simply add the probability of the two events together:

$$P(a \text{ or } b) = P(a) + P(b)$$

For example, suppose I wanted to calculate the probability of drawing randomly from a shuffled deck either a spade or a club. Since there are four suits (spades, clubs, diamonds, hearts) each with an equal number of cards, the probability of drawing a spade is $\frac{1}{4}$ or .25. Likewise the probability of drawing a club is .25. Thus, the probability of drawing either a spade or club is:

$$.25 + .25 = .50$$

So you have a 50% chance of drawing either a spade or a club.

Sometimes events are not independent. For example, suppose you wanted to know the probability of drawing 5 clubs from the deck (which in poker is called a “flush”). This time you are holding on to the cards after you draw them rather than replacing them back into the deck. The probability of drawing the first club is simply $\frac{13}{52}$ (or $\frac{1}{4}$). However, each of the remaining four draws will be affected by the previous draws. If one were to successfully draw all clubs then after the first draw, there would be only 51 cards left, 12 of which were clubs; after the second draw, there would be only 50 cards left, 11 of which were clubs, and so on, like this:

$$\frac{13}{52} \times \frac{12}{51} \times \frac{11}{50} \times \frac{10}{49} \times \frac{9}{48} = \frac{33}{66,640}$$

As you can see, we’ve had to determine the probability of a conjunction, since we want card 1 *and* card 2 *and* card 3 etc. to all be clubs. That is a conjunction of different events. As you can also see, the probability of drawing such a hand is extremely low—about .0005 or .05%. A flush is indeed a rare hand.

But suppose we wanted to know, not the chances of drawing a flush in a specific suit, but just the chances of drawing a flush in *any* suit. In that case, we’d have to calculate the probability of a disjunction of drawing *either* a flush in clubs *or* a flush in spades *or* a flush in diamonds *or* a flush in hearts. Recall that in order to calculate a disjunction we must add together the probabilities:

$$.0005 + .0005 + .0005 + .0005 = .002$$

So the probability of drawing a flush in any suit is still only about .2% or one fifth of one percent—i.e., very low.

Let's examine another example before closing this section on probability. Suppose we want to know the chances of flipping at least 1 head in 6 flips of a fair coin. You might reason as follows: There is a 50% chance I flip heads on the first flip, a 50% chance on the second, etc. Since I want to know the chance of flipping at least one head, then perhaps I should simply calculate the probability of the disjunction like this:

$$.5 + .5 + .5 + .5 + .5 + .5 = 3 \text{ (or 300\%)}$$

However, this cannot be right, because the probability of any event is between 1 and 0 (including 0 and 1 for events that are impossible and absolutely certain). However, this way of calculating the probability leaves us with an event that is three times more than certain. And nothing is more than 100% certain—100% certainty is the limit. So something is wrong with the calculation. Another way of seeing that something must be wrong with the calculation is that it isn't impossible that I flip 6 tails in a row (and thus no heads). Since that is a real possibility (however improbable), it cannot be 100% certain that I flip at least one head. Here is the way to think about this problem. What is the probability that I flip *all tails*? That is simply the probability of the conjunction of 6 events, each of which has the probability of .5 (or 50%):

$$.5 \times .5 \times .5 \times .5 \times .5 \times .5 = .015 \text{ (or 1.5\%)}$$

Then we simply use the rule for calculating the probability of a negation, since we want to know the chances that we *don't* flip 6 tails in a row (i.e., we flip at least one head):

$$1 - .015 = .985$$

So the probability of flipping at least one head in 6 flips of the coin is 98.5%. (It would be exactly the same probability of flipping at least one tails in 6 flips.)

Exercise 27: Use the three different rules of calculating probabilities (conjunctions, negations, disjunctions) to calculate the following probabilities, which all related to fair, six-sided dice.

1. What is the probability of rolling a five on one throw one die?
2. What is the probability of not rolling a five on one throw of one die?

3. What is the probability of rolling a five on your first throw and another five on the second throw of that die?
4. If you roll two dice at one time, what are the chances of rolling two twos?
5. If you roll two dice at one time, what are the chances of rolling at least one two?
6. If you roll two dice at once, what are the chances that at most one of the dice will come up a two?
7. If you roll two dice at once, what are the chances that the total of the dice adds up to 8?
8. If you roll two dice at once, what are the chances that there will be no fours?
9. If you roll two dice at once, what are the chances of rolling double fives?
10. If you roll two dice at once, what are the chances of rolling doubles (of any number)?

3.6 The conjunction fallacy

In this and the remaining sections of this chapter, we will consider some formal fallacies of probability. These fallacies are easy to spot once you see them, but they can be difficult to detect because of the way our minds mislead us— analogous to the way our minds can be misled when watching a magic trick. In addition to introducing the fallacies, I will suggest some psychological explanations for why these fallacies are so common, despite how easy they are to see once we've spotted them.

The conjunction fallacy is best introduced with an example.⁶

Linda is 31 years old, single, outspoken, and very bright. She majored in philosophy. As a student, she was deeply concerned with issues of discrimination and social justice, and also participated in anti-nuclear demonstrations.

⁶ The following famous example comes from Tversky, A. and Kahneman, D. (1983). Extension versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90(4), 293–315.

Given this information about Linda, which of the following is more probable?

- a. Linda is a bank teller.
- b. Linda is a bank teller and is active in the feminist movement.

If you are like most people who answer this question, you will answer “b.” But that cannot be correct because it violates the basic rules of probability. In particular, notice that option b contains option a (i.e., Linda is a bank teller). But option b also contains more information—that Linda is also active in the feminist movement. The problem is that *a conjunction can never be more probable than either one of its conjuncts*. Suppose we say it is very probable that Linda a bank teller (how boring, given the description of Linda which makes her sound interesting!). Let’s set the probability low, say .4. Then what is the probability of her being active in the feminist movement? Let’s set that high, say .9. However, the probability that she is both a bank teller and active in the feminist movement must be computed as the probability of a conjunction, like this:

$$.4 \times .9 = .36$$

So given these probability assignments (which I’ve just made up but seem fairly plausible), the probability of Linda being both a bank teller and active in the feminist movement is .36. But .36 is a lower probability than .4, which was the probability that she is bank teller. So option b cannot be more probable than option a. Notice that even if we say it is absolutely certain that Linda is active in the feminist movement (i.e., we set the probability of her being active in the feminist movement at 1), option b is still only equal to the probability of option a, since $(.4)(1) = .4$.

Sometimes it is easy to avoid committing conjunction fallacies. Here is an example that illustrates that we can in fact easily see that a conjunction is not more probable than either of its conjuncts.

Which is more probable?

- a. Mark has hair
- b. Mark has blonde hair

In this case, it is clear which of the options is more probable. Clearly a is more probable since it requires less to be true. Option a would be true even if option

b is true. But option a could also be true even if option b were false (i.e., Mark could have brown hair or red hair, etc.).

Thus there are cases where we can easily avoid committing the conjunction fallacy. So what is the difference between this case and the Linda case? The Nobel Prize-winning psychologist, Daniel Kahneman (and his long-time collaborator, Amos Tversky), has for many years suggested a psychological explanation for this difference. The explanation is complex, but I can give you the gist of it quite simply. Kahneman suggests that our minds are wired to find patterns and many of these patterns we find are based on what he calls “representativeness.” In the Linda case, the idea of Linda being active in the feminist movement fits better with the description of Linda as a philosophy major, as being active in social justice movements, and, perhaps, as being single. We build up a picture of Linda and then we try to match the descriptions to her. “Bank teller” doesn’t really match anything in the description of Linda. That is, the description of Linda is not representative of a bank teller. However, for many people, it is representative of a feminist. Thus, our minds more or less automatically see the match between representativeness of the description of Linda and option b, which mentions she is a feminist. Kahneman thinks that in cases like these, our minds substitute a question of representativeness for the question of probability, thus answering the probability question incorrectly.⁷ We are distracted from the probability question by seeking representativeness, which our minds more automatically look for and think about than probability. For Kahneman, the psychological explanation is needed to explain why even trained mathematicians and those who deal regularly with probability still commit the conjunction fallacy. The psychological explanation that our brains are wired to look for representativeness, and that we unwittingly substitute the question of representativeness for the question of probability, explains why even experts make these kinds of mistakes.

3.7 The base rate fallacy

Consider the following scenario. You go in for some testing for some health problems you’ve been having and after a number of tests, you test positive for colon cancer. What are the chances that you really do have colon cancer? Let’s

⁷ Kahneman gives this explanation numerous places, including, most exhaustively (and for a general audience) in his 2011 book, *Thinking Fast and Slow*. New York, NY: Farrar, Straus and Giroux.

suppose that the test is not perfect, but it is 95% accurate. That is, in the case of those who really do have colon cancer, the test will detect the cancer 95% of the time (and thus miss it 5% of the time). (The test will also misdiagnose those who don't actually have colon cancer 5% of the time.) Many people would be inclined to say that, given the test and its accuracy, there is a 95% chance that you have colon cancer. However, if you are like most people and are inclined to answer this way, you are wrong. In fact, you have committed the fallacy of ignoring the base rate (i.e., the base rate fallacy).

The **base rate** in this example is the rate of those who have colon cancer in a population. There is very small percentage of the population that actually has colon cancer (let's suppose it is .005 or .5%), so the probability that you have it must take into account the very low probability that you are one of the few that have it. That is, prior to the test (and not taking into account any other details about you), there was a very low probability that you have it—that is, a half of one percent chance (.5%). Yes, the test is 95% accurate, but *given the very low prior probability that you have colon cancer*, we cannot simply now say that there is a 95% chance that you have it. Rather, we must temper that figure with the very low base rate. The general point is this: when a condition (x) is very rare, then even if a highly accurate test identifies condition (x) as being present, we should still suspect that condition (x) is not present. In the above scenario, the prior probability (i.e., before the test) that you have colon cancer is really, really low. And that means that *even after the test* we should suspect that the probability is still fairly low. That's the logic of the matter and we can understand that conceptually without actually even getting into the math.

But since we are given numbers to work with, we can actually use math to figure the actual probability that you have colon cancer. Here is how we do it. Let's suppose that our population is 100,000 people. The base rate tells us that .5% of the population has colon cancer. That means that of the 100,000 people only 500 of them have colon cancer. If we were to apply the 95% accurate test to 500 people, the test would correctly diagnose 475 of them. That is, the test would deliver 475 correct identifications. These are called **true positives**. However, the test will also mistakenly tell us that some of the people who don't have colon cancer actually do have it. When this happens, it is called a **false positive**. Our base rate tells us that most of our population (99,500) do not have colon cancer and the 95% accurate test will misdiagnose 5% of those has having colon cancer. This comes out to 4975 false positives!

So what are the chances that you are true positive rather than a false positive? It is simply the number of true positives (475) divided by the total number of positive identifications that the test would make. This latter number includes those the test would misidentify (4975) as well as the number it would accurately identify (475)—thus the total number the test would identify as having colon cancer would be 5450. So the probability that you have it, given the positive test = $475/5450 = .087$ or 8.7%. So the probability that you have cancer, given the evidence of the positive test is 8.7%. Thus, contrary to our initial reasoning that there was a 95% chance that you have colon cancer, the chance is only a tenth of that—it is less than 10%! In thinking that the probability that you have cancer is closer to 95% you would be ignoring the extremely low probability of having the disease in the first place (i.e., the low base rate). Neglecting to account for low base rates in determining the probability of some event is the signature of any base rate fallacy.

The general lesson here is that the number of false positives will be quite high (even when the identification method is fairly accurate) as long as the base rate of the phenomenon we're looking for is very low. And if the number of false positives is high, then this will significantly lower the probability that the identification method has correctly identified the phenomenon in question.

From the above example we can see that the general method for determining probabilities when base rates are in play is the following.

1. Determine the number of false positives (i.e., the number of instances that are *incorrectly* identified by the method)
2. Determine the number of true positives (i.e., the number of instances that are *correctly* determined by the method)
3. Use the following equation to figure probability:

$$\frac{\textit{true positives}}{\textit{true positives} + \textit{false positives}}$$

Before closing this section, let's look at a couple more examples of a base rate fallacy. Suppose that the government has developed a machine that is able to detect terrorist intent with an accuracy of 90%. During a joint meeting of congress, a highly trustworthy source says that there is a terrorist in the building. (Let's suppose, for the sake of simplifying this example, that there is in fact a terrorist in the building.) In order to determine who the terrorist is, the building

security seals all the exits, rounds up all 3000 people in the building and uses the machine to test each person. The first 30 people pass without triggering a positive identification from the machine, but on the very next person, the machine triggers a positive identification of terrorist intent. The question is: what are the chances that the person who set off the machine really is a terrorist?⁸ Consider the following three possibilities: a) 90%, b) 10%, or c) .3%. If you answered 90%, then you committed the base rate fallacy again. The actually answer is "c"—less than 1%! Here's why. The base rate is the likelihood that any given individual is a terrorist and this is exceedingly low since there is only one terrorist in the building and there are 3000 people in the building. That means the probability of any one person being a terrorist, before any results of the test, is exceedingly low: $1/3000$. Since the test is 90% accurate, that means that out of the 2999 non-terrorists, it will misidentify 10% of them as terrorists = ~ 300 false positives. Assuming the machine doesn't misidentify the one actual terrorist, the machine will identify a total of 301 individuals as those "possessing terrorist intent." The probability that any one of them actually possesses terrorist intent is $1/301 = .3\%$. So the probability is drastically lower than 90%. It's not even close. This is another good illustration of how far off probabilities can be when the base rate is ignored.

Last one. Suppose that Bob is a super-eyewitness. When he observes a crime he is 99% accurate in identifying the suspect. Suppose that Bob identifies Nancy as the person who robbed the American Apparel store. In a population where .5% (half of one percent) of the population are robbers, what is the probability that Nancy really is a robber, given that Bob's eyewitness testimony identified her as a robber?

At this point, having been sensitized to the base rate fallacy, you should suspect that the probability is nowhere near as high as the accuracy of Bob's eyewitness skills. Here's the math. Suppose our population is 1000 people: 995 non-robbers and 5 robbers (based on the above base rate of robbers in the population). Of 995 non-robbers, Bob will misidentify 9.95 as robbers (false positives) and accurately identify 4.95 as robbers for a totally of 14.9 robber-identifications. So the chances that Nancy really is a robber, given Bob's eyewitness evidence is:

⁸ This example is taken (with certain alterations) from:
http://news.bbc.co.uk/2/hi/uk_news/magazine/8153539.stm

$$\frac{\text{\# of robbers Bob correctly identifies}}{\text{total \# of Bob's robber identifications}}$$

= 4.95/14.9 = 33%. Thus, it is more likely that Nancy is *not* the robber.

3.8 The small numbers fallacy

Suppose a study showed that of the 3,141 counties of the United States, the incidence of kidney cancer was lowest in those counties which are mostly rural, sparsely populated, and located in traditionally Republican states. In fact, this is true.⁹ What accounts for this interesting finding? Most people would be tempted to look for a causal explanation—to look for features of the rural environment that account for the lower incidence of cancer. However, they would be wrong (in this case) to do so. It is easy to see why once we consider the counties that have the highest incidence of kidney cancer: they are counties that are mostly rural, sparsely populated, and located in traditionally Republican states! So whatever it was you thought might account for the lower cancer rates in rural counties can't be the right explanation, since these counties also have the highest rates of cancer. It is important to understand that it isn't the same counties that have the highest and lowest rates—for example, county X doesn't have both a high and a low cancer rate (relative to other U.S. counties). That would be a contradiction (and so can't possibly be true). Rather, what is the case is that counties that have the highest kidney cancer rates are "mostly rural, sparsely populated, and located in traditionally Republican states" but also counties that have the lowest kidney cancer rates are "mostly rural, sparsely populated, and located in traditionally Republican states." How could this be? Before giving you the explanation, I'll give you a simpler example and see if you can figure it out from that example.

Suppose that a jar contains equal amounts of red and white marbles. Jack and Jill are taking turns drawing marbles from the jar. However, they draw marbles at different rates. Jill draws 10 marbles at a time while Jack draws 3 marbles at a time. Who is more likely to draw either all red or all white marbles more often: Jack or Jill?¹⁰

⁹ This example taken from Kahneman (2011), op. cit., p. 109.

¹⁰ This example is also taken (with minor modifications) from Kahneman (2011), p. 110.

The answer here should be obvious: Jack is more likely to draw marbles of all the same color more often, since Jack is only drawing 3 marbles at a time. Since Jill is drawing 10 marbles at a time, it will be less likely that her draws will yield marbles of all the same color. This is simply a fact of sampling and is related to the sampling errors discussed in section 3.1. A sample that is too small will tend not to be representative of the population. In the marbles case, if we view Jack's draws as samples, then his samples will be far less representative of the ratio of marbles in the jar. Jill, on the other hand, will tend to get far more representative samples. Since Jill is drawing a larger number of marbles, it is less likely that her samples would be drastically off in the way Jack's would be. The general point to be taken from this example is that smaller samples tend to the extremes—both in terms of overrepresenting some feature (e.g., red marbles) and in underrepresenting that same feature.

Can you see how this might apply to the case of kidney cancer rates in rural, sparsely populated counties? There is a national kidney cancer rate which is an average of all the kidney cancer rates of the 3,141 counties in the U.S. Imagine ranking each county in terms of the cancer rates from highest to lowest. The finding is that there is a relatively larger proportion of the sparsely populated counties at the top of this list, but also a relatively larger proportion of the sparsely populated counties at the bottom of the list. But why would it be that the more sparsely populated counties would be overrepresented at both ends of the list? The reason is that these counties have smaller populations, so they will tend to have more extreme results (of either the higher or lower rates). Just as Jack is more likely to get either all white marbles or all red marbles (an extreme result), the less populated counties will tend to have cancer rates that are at the extreme, relative to the national average. And this is a purely statistical fact; it has nothing to do with features of those environments causing the cancer rate to be higher or lower. Just as Jack's extreme draws have nothing to do with the way he is drawing (but are simply the result of statistical, mathematical facts), the extremes of the smaller counties have nothing to do with features of those counties, but only with the fact that they are smaller and so will tend to have more extreme results (i.e., cancer rates that are either higher or lower than the national average).

The first take home lesson here is that smaller groups will tend to be less representative than larger groups (thus, smaller groups will tend to inhabit the extremes of a spectrum). We can call this the **law of small numbers**. The second take home message is that our brains are wired to look for causal

explanations rather than mathematical explanations, and because of this we are prone to ignore the law of small numbers and look for a causal explanation of phenomena instead. The **small numbers fallacy** is our tendency to seek a causal explanation for some phenomenon when only the law of small numbers is needed to explain that phenomenon.

We will end this section with a somewhat humorous and incredible example of a small numbers bias that, presumably, wasted billions of dollars. This example, too, comes from Kahneman, who in turn heard the anecdote from some of his colleagues who are statisticians.¹¹ Some time ago, the Gates foundation (which is the charitable foundation of Microsoft founder, Bill Gates) donated 1.7 billion to research a curious finding: smaller schools tend to be more successful than larger schools. That is, if you consider a rank ordering of the most successful schools, the smaller schools will tend to be overrepresented near the top (i.e., there is a higher proportion of them near the top of the list compared to the proportion of larger schools at the top of the list). This is the finding that the Gates Foundation invested 1.7 billion dollars to help understand. In order to do so, they created smaller schools, sometimes splitting larger schools in half. However, none of this was necessary. Had the Gates Foundation (or those advising them) looked that the characteristics of the worst schools, they would have found that those schools also tended to be smaller! The “finding” is merely a result of the law of small numbers: smaller groups tend towards the extremes (on both ends of a spectrum) more so than larger groups. In this case, the fact that smaller schools tend to be both more successful and less successful is explained in the same way as we explain why Jack tends to get either all red or all white marbles more often than Jill.

3.9 Regression to the mean fallacy

Humans are prone to see causes even when no such cause is present. For example, if I have just committed some wrong and then immediately after the thunder cracks, I may think that my wrong action caused the lightning (e.g., because the gods were angry with me). The term “snake oil” refers to a product that promises certain (e.g., health) benefits but is actually fraudulent and has no benefits whatsoever. For example, consider a product that is supposed to help you recover from a common cold. You take the medicine and then within a few

¹¹ Kahneman (2011), pp. 117-118.

days, you are all better! No cold! It must have been the medicine. Or maybe you just regressed to the mean. **Regression to the mean** describes the tendency of things to go back to normal or to return to something close to the relevant statistical average. In the case of a cold, when you have a cold, you are outside of the average in terms of health. But you will naturally return to the state of health, with or without the “medicine.” If anyone were to try to convince you to buy such a medicine, you shouldn’t. Because the fact that you got better from your cold more likely has to do with the fact that you will naturally regress to the mean (return to normal) than it has to do with the special medicine.

Another example. Suppose you live in Lansing and it has been over 100 degrees for two weeks straight. Someone says that if you pay tribute and do a special dance to Baal, the temperature will drop. Suppose you do this and the temperature does drop. Was it Baal or just regression to the mean? Probably regression to the mean, unless we have some special reason for thinking it is Baal. The point is, extreme situations tend to regress towards less extreme, more average situations. Since it is very rare for it to ever be over 100 degrees in Lansing, the fact that the temperature drops is to be expected, regardless of one’s prayers to Baal.

Suppose that a professional golfer has been on a hot streak. She has been winning every tournament she enters by ten strokes—she’s beating the competition like they were middle school golfers. She is just playing so much better than them. Then something happens. The golfer all of a sudden starts playing like average. What explains her fall from greatness? The sports commentators speculate: could it be that she switched her caddy, or that it is warmer now than it was when she was on her streak, or perhaps it was fame that went to her head once she had started winning all those tournaments? Chances are, none of these are the right explanation because no such explanation is needed. Most likely she just regressed to the mean and is now playing like everyone else—still like a pro, just not like a golfer who is out of this world good. Even those who are skilled can get lucky (or unlucky) and when they do, we should expect that eventually that luck will end and they will regress to the mean.

As these examples illustrate, one commits the **regression to the mean fallacy** when one tries to give a causal explanation of a phenomenon that is merely statistical or probabilistic in nature. The best way to rule out that something is

not to be explained as regression to the mean is by doing a study where one compares two groups. For example, suppose we could get our snake oil salesman to agree to a study in which a group of people who had colds took the medicine (experimental group) and another group of people didn't take the medicine or took a placebo (control group). In this situation, if we found that the experimental group got better and the control group didn't, or if the experimental group got better more quickly than the control group, then perhaps we'd have to say that maybe there is something to this snake oil medicine. But without the evidence of a control for comparison, even if lots of people took the snake oil medicine and got better from their colds, it wouldn't prove anything about the efficacy of the medicine.

3.10 Gambler's fallacy

The gambler's fallacy occurs when one thinks that independent, random events can be influenced by each other. For example, suppose I have a fair coin and I have just flipped 4 heads in a row. Erik, on the other hand, has a fair coin that he has flipped 4 times and gotten tails. We are each taking bets that the next coin flipped is heads. Who should you bet flips the head? If you are inclined to say that you should place the bet with Erik since he has been flipping all tails and since the coin is fair, the flips must even out soon, then you have committed the gambler's fallacy. The fact is, each flip is independent of the next, so the fact that I have just flipped 4 heads in a row does not increase or decrease my chances of flipping a head. Likewise for Erik. It is true that as long as the coin is fair, then over a large number of flips we should expect that the proportion of heads to tails will be about 50/50. But there is no reason to expect that a particular flip will be more likely to be one or the other. Since the coin is fair, each flip has the same probability of being heads and the same probability of being tails—50%.

4.1 Formal vs. informal fallacies

A **fallacy** is simply a mistake in reasoning. Some fallacies are formal and some are informal. In chapter 2, we saw that we could define validity formally and thus could determine whether an argument was valid or invalid without even having to know or understand what the argument was about. We saw that we could define certain valid rules of inference, such as modus ponens and modus tollens. These inference patterns are valid in virtue of their form, not their content. That is, any argument that has the same form as modus ponens or modus tollens will automatically be valid. A formal fallacy is simply an argument whose form is invalid. Thus, any argument that has that form will automatically be invalid, regardless of the meaning of the sentences. Two formal fallacies that are similar to, but should never be confused with, modus ponens and modus tollens are **denying the antecedent** and **affirming the consequent**. Here are the forms of those invalid inferences:

Denying the antecedent

$$\begin{array}{l} p \supset q \\ \sim p \\ \therefore \sim q \end{array}$$

Affirming the consequent

$$\begin{array}{l} p \supset q \\ q \\ \therefore p \end{array}$$

Any argument that has either of these forms is an invalid argument. For example:

1. If Kant was a deontologist, then he was a non-consequentialist.
2. Kant was not a deontologist.
3. Therefore, Kant was a not a non-consequentialist.

The form of this argument is:

1. $D \supset C$
2. $\sim D$
3. $\therefore \sim C$

As you can see, this argument has the form of the fallacy, denying the antecedent. Thus, we know that this argument is invalid even if we don't know what "Kant" or "deontologist" or "non-consequentialist" means. ("Kant" was a famous German philosopher from the early 1800s, whereas "deontology" and "non-consequentialist" are terms that come from ethical theory.) It is mark of a formal fallacy that we can identify it even if we don't really understand the meanings of the sentences in the argument. Recall our Jabberwocky argument from chapter 2. Here's an argument which uses silly, made-up words from Lewis Carroll's "Jabberwocky." See if you can determine whether the argument's form is valid or invalid:

1. If toves are brillig then toves are slithy.
2. Toves are slithy
3. Therefore, toves are brillig.

You should be able to see that this argument has the form of affirming the consequent:

1. $B \supset S$
2. S
3. $\therefore B$

As such, we know that the argument is invalid, even though we haven't got a clue what "toves" are or what "slithy" or "brillig" means. The point is that we can identify formal fallacies without having to know what they mean.

In contrast, **informal fallacies** are those which cannot be identified without understanding the concepts involved in the argument. A paradigm example of an informal fallacy is the fallacy of composition. We will consider this fallacy in the next sub-section. In the remaining subsections, we will consider a number of other informal logical fallacies.

4.1.1 Composition fallacy

Consider the following argument:

- Each member on the gymnastics team weighs less than 110 lbs.
Therefore, the whole gymnastics team weighs less than 110 lbs.

This argument commits the composition fallacy. In the composition fallacy one argues that since each part of the whole has a certain feature, it follows that the whole has that same feature. However, you cannot generally identify any argument that moves from statements about parts to statements about wholes as committing the composition fallacy because whether or not there is a fallacy depends on what feature we are attributing to the parts and wholes. Here is an example of an argument that moves from claims about the parts possessing a feature to a claim about the whole possessing that same feature, but doesn't commit the composition fallacy:

Every part of the car is made of plastic. Therefore, the whole car is made of plastic.

This conclusion *does* follow from the premises; there is no fallacy here. The difference between this argument and the preceding argument (about the gymnastics team) isn't their form. In fact both arguments have the same form:

Every part of X has the feature f. Therefore, the whole X has the feature f.

And yet one of the arguments is clearly fallacious, while the other isn't. The difference between the two arguments is not their form, but their content. That is, the difference is what feature is being attributed to the parts and wholes. Some features (like *weighing a certain amount*) are such that if they belong to each part, then it does not follow that they belong to the whole. Other features (such as *being made of plastic*) are such that if they belong to each part, it follows that they belong to the whole.

Here is another example:

Every member of the team has been to Paris. Therefore the team has been to Paris.

The conclusion of this argument does not follow. Just because each member of the team has been to Paris, it doesn't follow that the whole team has been to Paris, since it may not have been the case that each individual was there at the same time and was there in their capacity as a member of the team. Thus, even though it is plausible to say that the team is composed of every member of the team, it doesn't follow that since every member of the team has been to Paris, the whole team has been to Paris. Contrast that example with this one:

Every member of the team was on the plane. Therefore, the whole team was on the plane.

This argument, in contrast to the last one, contains no fallacy. It is true that if every member is on the plane then the whole team is on the plane. And yet these two arguments have almost exactly the same form. The only difference is that the first argument is talking about the property, *having been to Paris*, whereas the second argument is talking about the property, *being on the plane*. The only reason we are able to identify the first argument as committing the composition fallacy and the second argument as not committing a fallacy is that we understand the relationship between the concepts involved. In the first case, we understand that it is possible that every member could have been to Paris without the team ever having been; in the second case we understand that as long as every member of the team is on the plane, it has to be true that the whole team is on the plane. The take home point here is that in order to identify whether an argument has committed the composition fallacy, one must understand the concepts involved in the argument. This is the mark of an informal fallacy: we have to rely on our understanding of the meanings of the words or concepts involved, rather than simply being able to identify the fallacy from its form.

4.1.2 Division fallacy

The division fallacy is like the composition fallacy and they are easy to confuse. The difference is that the division fallacy argues that since the whole has some feature, each part must also have that feature. The composition fallacy, as we have just seen, goes in the opposite direction: since each part has some feature, the whole must have that same feature. Here is an example of a division fallacy:

The house costs 1 million dollars. Therefore, each part of the house costs 1 million dollars.

This is clearly a fallacy. Just because the whole house costs 1 million dollars, it doesn't follow that each part of the house costs 1 million dollars. However, here is an argument that has the same form, but that doesn't commit the division fallacy:

The whole team died in the plane crash. Therefore, each individual on the team died in the plane crash.

In this example, since we seem to be referring to one plane crash in which all the members of the team died ("the" plane crash), it follows that if the whole team died in the crash, then every individual on the team died in the crash. So this argument does not commit the division fallacy. In contrast, the following argument has exactly the same form, but does commit the division fallacy:

The team played its worst game ever tonight. Therefore, each individual on the team played their worst game ever tonight.

It can be true that the whole team played its worst game ever even if it is true that no individual on the team played their worst game ever. Thus, this argument does commit the fallacy of division even though it has the same form as the previous argument, which doesn't commit the fallacy of division. This shows (again) that in order to identify informal fallacies (like composition and division), we must rely on our understanding of the concepts involved in the argument. Some concepts (like "team" and "dying in a plane crash") are such that if they apply to the whole, they also apply to all the parts. Other concepts (like "team" and "worst game played") are such that they can apply to the whole even if they do not apply to all the parts.

4.1.3 Begging the question

Consider the following argument:

Capital punishment is justified for crimes such as rape and murder because it is quite legitimate and appropriate for the state to put to death someone who has committed such heinous and inhuman acts.

The premise indicator, "because" denotes the premise and (derivatively) the conclusion of this argument. In standard form, the argument is this:

1. It is legitimate and appropriate for the state to put to death someone who commits rape or murder.
2. Therefore, capital punishment is justified for crimes such as rape and murder.

You should notice something peculiar about this argument: the premise is essentially the same claim as the conclusion. The only difference is that the premise spells out what capital punishment means (the state putting criminals to death) whereas the conclusion just refers to capital punishment by name, and the premise uses terms like “legitimate” and “appropriate” whereas the conclusion uses the related term, “justified.” But these differences don’t add up to any real differences in meaning. Thus, the premise is essentially saying the same thing as the conclusion. This is a problem: we want our premise to provide a reason for accepting the conclusion. But if the premise is the same claim as the conclusion, then it can’t possibly provide a reason for accepting the conclusion! **Begging the question** occurs when one (either explicitly or implicitly) assumes the truth of the conclusion in one or more of the premises. Begging the question is thus a kind of circular reasoning.

One interesting feature of this fallacy is that *formally* there is nothing wrong with arguments of this form. Here is what I mean. Consider an argument that explicitly commits the fallacy of begging the question. For example,

1. Capital punishment is morally permissible
2. Therefore, capital punishment is morally permissible

Now, apply any method of assessing validity to this argument and you will see that it is valid by any method. If we use the informal test (by trying to imagine that the premises are true while the conclusion is false), then the argument passes the test, since any time the premise is true, the conclusion will have to be true as well (since it is the exact same statement). Likewise, the argument is valid by our formal test of validity, truth tables. But while this argument is technically valid, it is still a really bad argument. Why? Because the point of giving an argument in the first place is to provide some reason for thinking the conclusion is true *for those who don’t already accept the conclusion*. But if one doesn’t already accept the conclusion, then simply restating the conclusion in a different way isn’t going to convince them. Rather, a good argument will provide some reason for accepting the conclusion that is sufficiently independent of that conclusion itself. Begging the question utterly fails to do this and this is why it counts as an informal fallacy. What is interesting about begging the question is that there is absolutely nothing wrong with the argument formally.

Whether or not an argument begs the question is not always an easy matter to sort out. As with all informal fallacies, detecting it requires a careful understanding of the meaning of the statements involved in the argument. Here is an example of an argument where it is not as clear whether there is a fallacy of begging the question:

Christian belief is warranted because according to Christianity there exists a being called “the Holy Spirit” which reliably guides Christians towards the truth regarding the central claims of Christianity.¹

One might think that there is a kind of circularity (or begging the question) involved in this argument since the argument appears to assume the truth of Christianity in justifying the claim that Christianity is true. But whether or not this argument really does beg the question is something on which there is much debate within the sub-field of philosophy called epistemology (“study of knowledge”). The philosopher Alvin Plantinga argues persuasively that the argument does not beg the question, but being able to assess that argument takes patient years of study in the field of epistemology (not to mention a careful engagement with Plantinga’s work). As this example illustrates, the issue of whether an argument begs the question requires us to draw on our general knowledge of the world. This is the mark of an informal, rather than formal, fallacy.

4.1.4 False dichotomy

Suppose I were to argue as follows:

Raising taxes on the wealthy will either hurt the economy or it will help it.
But it won’t help the economy. Therefore, it will hurt the economy.

The standard form of this argument is:

1. Either raising taxes on the wealthy will hurt the economy or it will help it.
2. Raising taxes on the wealthy won’t help the economy.
3. Therefore, raising taxes on the wealthy will hurt the economy.

¹ This is a much simplified version of the view defended by Christian philosophers such as Alvin Plantinga. Plantinga defends (something like) this claim in: Plantinga, A. 2000. *Warranted Christian Belief*. Oxford, UK: Oxford University Press.

This argument contains a fallacy called a “false dichotomy.” A **false dichotomy** is simply a disjunction that does not exhaust all of the possible options. In this case, the problematic disjunction is the first premise: either raising the taxes on the wealthy will hurt the economy or it will help it. But these aren’t the only options. Another option is that raising taxes on the wealthy will have no effect on the economy. Notice that the argument above has the form of a disjunctive syllogism:

$$\begin{array}{l} A \vee B \\ \sim A \\ \therefore B \end{array}$$

However, since the first premise presents two options as if they were the only two options, when in fact they aren’t, the first premise is false and the argument fails. Notice that the form of the argument is perfectly good—the argument is valid. The problem is that this argument isn’t sound because the first premise of the argument commits the false dichotomy fallacy. False dichotomies are commonly encountered in the context of a disjunctive syllogism or constructive dilemma (see chapter 2).

In a speech made on April 5, 2004, President Bush made the following remarks about the causes of the Iraq war:

Saddam Hussein once again defied the demands of the world. And so I had a choice: Do I take the word of a madman, do I trust a person who had used weapons of mass destruction on his own people, plus people in the neighborhood, or do I take the steps necessary to defend the country? Given that choice, I will defend America every time.

The false dichotomy here is the claim that:

Either I trust the word of a madman or I defend America (by going to war against Saddam Hussein’s regime).

The problem is that these aren’t the only options. Other options include ongoing diplomacy and economic sanctions. Thus, even if it true that Bush shouldn’t have trusted the word of Hussein, it doesn’t follow that the only other

option is going to war against Hussein's regime. (Furthermore, it isn't clear in what sense this was needed to defend *America*.) That is a false dichotomy.

As with all the previous informal fallacies we've considered, the false dichotomy fallacy requires an understanding of the concepts involved. Thus, we have to use our understanding of world in order to assess whether a false dichotomy fallacy is being committed or not.

4.1.5 Equivocation

Consider the following argument:

Children are a headache. Aspirin will make headaches go away.
Therefore, aspirin will make children go away.

This is a silly argument, but it illustrates the fallacy of equivocation. The problem is that the word "headache" is used equivocally—that is, in two different senses. In the first premise, "headache" is used figuratively, whereas in the second premise "headache" is used literally. The argument is only successful if the meaning of "headache" is the same in both premises. But it isn't and this is what makes this argument an instance of the fallacy of equivocation.

Here's another example:

Taking a logic class helps you learn how to argue. But there is already too much hostility in the world today, and the fewer arguments the better. Therefore, you shouldn't take a logic class.

In this example, the word "argue" and "argument" are used equivocally. Hopefully, at this point in the text, you recognize the difference. (If not, go back and reread section 1.1.)

The fallacy of equivocation is not always so easy to spot. Here is a trickier example:

The existence of laws depends on the existence of intelligent beings like humans who create the laws. However, some laws existed before there were any humans (e.g., laws of physics). Therefore, there must be some non-human, intelligent being that created these laws of nature.

The term “law” is used equivocally here. In the first premise it is used to refer to societal laws, such as criminal law; in the second premise it is used to refer to laws of nature. Although we use the term “law” to apply to both cases, they are importantly different. Societal laws, such as the criminal law of a society, are enforced by people and there are punishments for breaking the laws. Natural laws, such as laws of physics, cannot be broken and thus there are no punishments for breaking them. (Does it make sense to scold the electron for not doing what the law says it will do?)

As with every informal fallacy we have examined in this section, equivocation can only be identified by understanding the meanings of the words involved. In fact, the definition of the fallacy of equivocation refers to this very fact: the same word is being used in two different senses (i.e., with two different meanings). So, unlike formal fallacies, identifying the fallacy of equivocation requires that we draw on our understanding of the meaning of words and of our understanding of the world, generally.

4.2 Slippery slope fallacies

Slippery slope fallacies depend on the concept of **vagueness**. When a concept or claim is vague, it means that we don’t know precisely what claim is being made, or what the boundaries of the concept are. The classic example used to illustrate vagueness is the “**sorites paradox**.” The term “sorites” is the Greek term for “heap” and the paradox comes from ancient Greek philosophy. Here is the paradox. I will give you two claims that each sound very plausible, but in fact lead to a paradox. Here are the two claims:

1. One grain of sand is not a heap of sand.
2. If I start with something that is not a heap of sand, then adding one grain of sand to that will not create a heap of sand.

For example, two grains of sand is not a heap, thus (by the second claim) neither is three grains of sand. But since three grains of sand is not a heap then (by the second claim again) neither is four grains of sand. You can probably see where this is going. By continuing to add one grain of sand over and over, I will eventually end up with something that is clearly a heap of sand, but that won’t be counted as a heap of sand if we accept both claims 1 and 2 above.

Philosophers continue to argue and debate about how to resolve the sorites paradox, but the point for us is just to illustrate the concept of vagueness. The concept “heap” is a vague concept in this example. But so are so many other concepts, such as color concepts (red, yellow, green, etc.), moral concepts (right, wrong, good, bad), and just about any other concept you can think of. The one domain that seems to be unaffected by vagueness is mathematical and logical concepts. There are two fallacies related to vagueness: the causal slippery slope and the conceptual slippery slope. We’ll cover the conceptual slippery slope first since it relates most closely to the concept of vagueness I’ve explained above.

4.2.1 Conceptual slippery slope

It may be true that there is no essential difference between 499 grains of sand and 500 grains of sand. But even if that is so, it doesn’t follow that there is no difference between 1 grain of sand and 5 billion grains of sand. In general, just because we cannot draw a distinction between A and B, and we cannot draw a distinction between B and C, it doesn’t mean we cannot draw a distinction between A and C. Here is an example of a conceptual slippery slope fallacy.

It is illegal for anyone under 21 to drink alcohol. But there is no difference between someone who is 21 and someone who is 20 years 11 months old. So there is nothing wrong with someone who is 20 years and 11 months old drinking. But since there is no real distinction between being one month older and one month younger, there shouldn’t be anything wrong with drinking at any age. Therefore, there is nothing wrong with allowing a 10 year old to drink alcohol.

Imagine the life of an individual in stages of 1 month intervals. Even if it is true that there is no distinction in kind between any one of those stages, it doesn’t follow that there isn’t a distinction to be drawn at the extremes of either end. Clearly there is a difference between a 5 year old and a 25 year old—a distinction in kind that is relevant to whether they should be allowed to drink alcohol. The conceptual slippery slope fallacy assumes that because we cannot draw a distinction between adjacent stages, we cannot draw a distinction at all between any stages. One clear way of illustrating this is with color. Think of a color spectrum from purple to red to orange to yellow to green to blue. Each color grades into the next without there being any distinguishable boundaries

between the colors—a continuous spectrum. Even if it is true that for any two adjacent hues on the color wheel, we cannot distinguish between the two, it doesn't follow from this that there is no distinction to be drawn between any two portions of the color wheel, because then we'd be committed to saying that there is no distinguishable difference between purple and yellow! The example of the color spectrum illustrates the general point that just because the boundaries between very similar things on a spectrum are vague, it doesn't follow that there are no differences between any two things on that spectrum.

Whether or not one will identify an argument as committing a conceptual slippery slope fallacy, depends on the other things one believes about the world. Thus, whether or not a conceptual slippery slope fallacy has been committed will often be a matter of some debate. It will itself be vague. Here is a good example that illustrates this point.

People are found not guilty by reason of insanity when they cannot avoid breaking the law. But people who are brought up in certain deprived social circumstances are not much more able than the legally insane to avoid breaking the law. So we should not find such individuals guilty any more than those who are legally insane.

Whether there is conceptual slippery slope fallacy here depends on what you think about a host of other things, including individual responsibility, free will, the psychological and social effects of deprived social circumstances such as poverty, lack of opportunity, abuse, etc. Some people may think that there are big differences between those who are legally insane and those who grow up in deprived social circumstances. Others may not think the differences are so great. The issues here are subtle, sensitive, and complex, which is why it is difficult to determine whether there is any fallacy here or not. If the differences between those who are insane and those who are the product of deprived social circumstances turn out to be like the differences between one shade of yellow and an adjacent shade of yellow, then there is no fallacy here. But if the differences turn out to be analogous to those between yellow and green (i.e., with many distinguishable stages of difference between) then there would indeed be a conceptual slippery slope fallacy here. The difficulty of distinguishing instances of the conceptual slippery slope fallacy, and the fact that distinguishing it requires us to draw on our knowledge about the world, shows that the conceptual slippery slope fallacy is an informal fallacy.

4.2.2 Causal slippery slope fallacy

The causal slippery slope fallacy is committed when one event is said to lead to some other (usually disastrous) event via a chain of intermediary events. If you have ever seen Direct TV's "get rid of cable" commercials, you will know exactly what I'm talking about. (If you don't know what I'm talking about you should Google it right now and find out. They're quite funny.) Here is an example of a causal slippery slope fallacy (it is adapted from one of the Direct TV commercials):

If you use cable, your cable will probably go on the fritz. If your cable is on the fritz, you will probably get frustrated. When you get frustrated you will probably hit the table. When you hit the table, your young daughter will probably imitate you. When your daughter imitates you, she will probably get thrown out of school. When she gets thrown out of school, she will probably meet undesirables. When she meets undesirables, she will probably marry undesirables. When she marries undesirables, you will probably have a grandson with a dog collar. Therefore, if you use cable, you will probably have a grandson with a dog collar.

This example is silly and absurd, yes. But it illustrates the causal slippery slope fallacy. Slippery slope fallacies are always made up of a series of conjunctions of probabilistic conditional statements that link the first event to the last event. A causal slippery slope fallacy is committed when one assumes that just because each individual conditional statement is probable, the conditional that links the first event to the last event is also probable. Even if we grant that each "link" in the chain is individually probable, it doesn't follow that the whole chain (or the conditional that links the first event to the last event) is probable. Suppose, for the sake of the argument, we assign probabilities to each "link" or conditional statement, like this. (I have italicized the consequents of the conditionals and assigned high conditional probabilities to them. The high probability is for the sake of the argument; I don't actually think these things are as probable as I've assumed here.)

If you use cable, then *your cable will probably go on the fritz* (.9)

If your cable is on the fritz, then *you will probably get angry* (.9)

If you get angry, then *you will probably hit the table* (.9)

If you hit the table, *your daughter will probably imitate you* (.8)

If your daughter imitates you, *she will probably be kicked out of school* (.8)

If she is kicked out of school, *she will probably meet undesirables* (.9)

If she meets undesirables, *she will probably marry undesirables* (.8)

If she marries undesirables, *you will probably have a grandson with a dog collar* (.8)

However, even if we grant the probabilities of each link in the chain is high (80-90% probable), the conclusion doesn't even reach a probability higher than chance. Recall that in order to figure the probability of a conjunction, we must multiply the probability of each conjunct:

$$(.9) \times (.9) \times (.9) \times (.8) \times (.8) \times (.9) \times (.8) \times (.8) = .27$$

That means the probability of the conclusion (i.e., that if you use cable, you will have a grandson with a dog collar) is only 27%, despite the fact that each conditional has a relatively high probability! The causal slippery slope fallacy is actually a formal probabilistic fallacy and so could have been discussed in chapter 3 with the other formal probabilistic fallacies. What makes it a formal rather than informal fallacy is that we can identify it without even having to know what the sentences of the argument mean. I could just have easily written out a nonsense argument comprised of series of probabilistic conditional statements. But I would still have been able to identify the causal slippery slope fallacy because I would have seen that there was a series of probabilistic conditional statements leading to a claim that the conclusion of the series was also probable. That is enough to tell me that there is a causal slippery slope fallacy, even if I don't really understand the meanings of the conditional statements.

It is helpful to contrast the causal slippery slope fallacy with the valid form of inference, hypothetical syllogism. Recall that a hypothetical syllogism has the following kind of form:

A \supset B

B \supset C

C \supset D

D \supset E

\therefore A \supset E

The only difference between this and the causal slippery slope fallacy is that whereas in the hypothetical syllogism, the link between each component is *certain*, in a causal slippery slope fallacy, the link between each event is *probabilistic*. It is the fact that each link is probabilistic that accounts for the fallacy. One way of putting this is point is that *probability is not transitive*. Just because A makes B probable and B makes C probable and C makes X probable, it doesn't follow that A makes X probable. In contrast, when the links are certain rather than probable, then if A always leads to B and B always leads to C and C always leads to X, then it has to be the case that A always leads to X.

4.3 Fallacies of relevance

What all fallacies of relevance have in common is that they make an argument or response to an argument that is irrelevant to that argument. Fallacies of relevance can be *psychologically* compelling, but it is important to distinguish between rhetorical techniques that are psychologically compelling, on the one hand, and *rationally* compelling arguments, on the other. What makes something a fallacy is that it fails to be rationally compelling, once we have carefully considered it. That said, arguments that fail to be rationally compelling may still be psychologically or emotionally compelling. The first fallacy of relevance that we will consider, the ad hominem fallacy, is an excellent example of a fallacy that can be psychologically compelling.

4.3.1 Ad hominem

"Ad hominem" is a Latin phrase that can be translated into English as the phrase, "against the man." In an ad hominem fallacy, instead of responding to (or attacking) the *argument* a person has made, one attacks the *person* him or herself. In short, one attacks the person making the argument rather than the argument itself. Here is an anecdote that reveals an ad hominem fallacy (and that has actually occurred in my ethics class before).

A philosopher named Peter Singer had made an argument that it is morally wrong to spend money on luxuries for oneself rather than give all of your money that you don't strictly need away to charity. The argument is actually an argument from analogy (whose details I discussed in section 3.3), but the essence of the argument is that there are every day in this world children who die from preventable deaths, and there are charities

who could save the lives of these children if they are funded by individuals from wealthy countries like our own. Since there are things that we all regularly buy that we don't need (e.g., Starbucks lattes, beer, movie tickets, or extra clothes or shoes we don't really need), if we continue to purchase those things rather than using that money to save the lives of children, then we are essentially contributing to the deaths of those children if we choose to continue to live our lifestyle of buying things we don't need, rather than donating the money to a charity that will save lives of children in need. In response to Singer's argument, one student in the class asked: "Does Peter Singer give his money to charity? Does he do what he says we are all morally required to do?"

The implication of this student's question (which I confirmed by following up with her) was that if Peter Singer himself doesn't donate all his extra money to charities, then his argument isn't any good and can be dismissed. But that would be to commit an *ad hominem* fallacy. Instead of responding to the *argument* that Singer had made, this student attacked *Singer himself*. That is, they wanted to know how Singer lived and whether he was a hypocrite or not. Was he the kind of person who would tell us all that we had to live a certain way but fail to live that way himself? But all of this is irrelevant to assessing Singer's *argument*. Suppose that Singer didn't donate his excess money to charity and instead spent it on luxurious things for himself. Still, the argument that Singer has given can be assessed on its own merits. Even if it were true that Peter Singer was a total hypocrite, his argument may nevertheless be rationally compelling. And it is the quality of the argument that we are interested in, not Peter Singer's personal life and whether or not he is hypocritical. Whether Singer is or isn't a hypocrite, is irrelevant to whether the argument he has put forward is strong or weak, valid or invalid. The argument stands on its own and it is that argument rather than Peter Singer himself that we need to assess.

Nonetheless, there is something psychologically compelling about the question: Does Peter Singer practice what he preaches? I think what makes this question seem compelling is that humans are very interested in finding "cheaters" or hypocrites—those who say one thing and then do another. Evolutionarily, our concern with cheaters makes sense because cheaters can't be trusted and it is essential for us (as a group) to be able to pick out those who can't be trusted. That said, whether or not a person giving an argument is a hypocrite is irrelevant to whether that person's argument is good or bad. So there may be psychological reasons why humans are prone to find certain kinds of ad

hominem fallacies *psychologically compelling*, even though ad hominem fallacies are not *rationally compelling*.

Not every instance in which someone attacks a person's character is an ad hominem fallacy. Suppose a witness is on the stand testifying against a defendant in a court of law. When the witness is cross examined by the defense lawyer, the defense lawyer tries to go for the witness's credibility, perhaps by digging up things about the witness's past. For example, the defense lawyer may find out that the witness cheated on her taxes five years ago or that the witness failed to pay her parking tickets. The reason this isn't an ad hominem fallacy is that in this case the lawyer is trying to establish whether what the witness is saying is true or false and in order to determine that we have to know whether the witness is trustworthy. These facts about the witness's past may be relevant to determining whether we can trust the witness's word. In this case, the witness is making claims that are either true or false rather than giving an argument. In contrast, when we are assessing someone's argument, the argument stands on its own in a way the witness's testimony doesn't. In assessing an argument, we want to know whether the argument is strong or weak and we can evaluate the argument using the logical techniques surveyed in this text. In contrast, when a witness is giving testimony, they aren't trying to argue anything. Rather, they are simply making a claim about what did or didn't happen. So although it may seem that a lawyer is committing an ad hominem fallacy in bringing up things about the witness's past, these things are actually relevant to establishing the witness's credibility. In contrast, when considering an argument that has been given, we don't have to establish the arguer's credibility because we can assess the argument they have given on its own merits. The arguer's personal life is irrelevant.

4.3.2 Straw man

Suppose that my opponent has argued for a position, call it position A, and in response to his argument, I give a rationally compelling argument against position B, which is related to position A, but is much less plausible (and thus much easier to refute). What I have just done is attacked a straw man—a position that “looks like” the target position, but is actually not that position. When one attacks a straw man, one commits the straw man fallacy. The straw man fallacy misrepresents one's opponent's argument and is thus a kind of irrelevance. Here is an example.

Two candidates for political office in Colorado, Tom and Fred, are having an exchange in a debate in which Tom has laid out his plan for putting more money into health care and education and Fred has laid out his plan which includes earmarking more state money for building more prisons which will create more jobs and, thus, strengthen Colorado's economy. Fred responds to Tom's argument that we need to increase funding to health care and education as follows: "I am surprised, Tom, that you are willing to put our state's economic future at risk by sinking money into these programs that do not help to create jobs. You see, folks, Tom's plan will risk sending our economy into a tailspin, risking harm to thousands of Coloradans. On the other hand, my plan supports a healthy and strong Colorado and would never bet our state's economic security on idealistic notions that simply don't work when the rubber meets the road."

Fred has committed the straw man fallacy. Just because Tom wants to increase funding to health care and education does not mean he does not want to help the economy. Furthermore, increasing funding to health care and education does not entail that fewer jobs will be created. Fred has attacked a position that is not the position that Tom holds, but is in fact a much less plausible, easier to refute position. However, it would be silly for any political candidate to run on a platform that included "harming the economy." Presumably no political candidate would run on such a platform. Nonetheless, this exact kind of straw man is ubiquitous in political discourse in our country.

Here is another example.

Nancy has just argued that we should provide middle schoolers with sex education classes, including how to use contraceptives so that they can practice safe sex should they end up in the situation where they are having sex. Fran responds: "proponents of sex education try to encourage our children to a sex-with-no-strings-attached mentality, which is harmful to our children and to our society."

Fran has committed the straw man (or straw woman) fallacy by misrepresenting Nancy's position. Nancy's position is not that we should encourage children to have sex, but that we should make sure that they are fully informed about sex so that if they do have sex, they go into it at least a little less blindly and are able to make better decisions regarding sex.

As with other fallacies of relevance, straw man fallacies can be compelling on some level, even though they are irrelevant. It may be that part of the reason we are taken in by straw man fallacies is that humans are prone to “demonize” the “other”—including those who hold a moral or political position different from our own. It is easy to think bad things about those with whom we do not regularly interact. And it is easy to forget that people who are different than us are still people just like us in all the important respects. Many years ago, atheists were commonly thought of as highly immoral people and stories about the horrible things that atheists did in secret circulated widely. People believed that these strange “others” were capable of the most horrible savagery. After all, they may have reasoned, if you don’t believe there is a God holding us accountable, why be moral? The Jewish philosopher, Baruch Spinoza, was an atheist who lived in the Netherlands in the 17th century. He was accused of all sorts of things that were commonly believed about atheists. But he was in fact as upstanding and moral as any person you could imagine. The people who knew Spinoza knew better, but how could so many people be so wrong about Spinoza? I suspect that part of the reason is that since at that time there were very few atheists (or at least very few people actually admitted to it), very few people ever knowingly encountered an atheist. Because of this, the stories about atheists could proliferate without being put in check by the facts. I suspect the same kind of phenomenon explains why certain kinds of straw man fallacies proliferate. If you are a conservative and mostly only interact with other conservatives, you might be prone to holding lots of false beliefs about liberals. And so maybe you are less prone to notice straw man fallacies targeted at liberals because the false beliefs you hold about them incline you to see the straw man fallacies as true.

4.3.3 Tu quoque

“Tu quoque” is a Latin phrase that can be translated into English as “you too” or “you, also.” The tu quoque fallacy is a way of avoiding answering a criticism by bringing up a criticism of your opponent rather than answer the criticism. For example, suppose that two political candidates, A and B, are discussing their policies and A brings up a criticism of B’s policy. In response, B brings up her own criticism of A’s policy rather than respond to A’s criticism of her policy. B has here committed the tu quoque fallacy. The fallacy is best understood as a way of avoiding having to answer a tough criticism that one may not have a good answer to. This kind of thing happens all the time in political discourse.

Tu quoque, as I have presented it, is fallacious when the criticism one raises is simply in order to avoid having to answer a difficult objection to one's argument or view. However, there are circumstances in which a tu quoque kind of response is not fallacious. If the criticism that A brings toward B is a criticism that equally applies not only to A's position but to any position, then B is right to point this fact out. For example, suppose that A criticizes B for taking money from special interest groups. In this case, B would be totally right (and there would be no tu quoque fallacy committed) to respond that not only does A take money from special interest groups, but every political candidate running for office does. That is just a fact of life in American politics today. So A really has no criticism at all to B since everyone does what B is doing and it is in many ways unavoidable. Thus, B could (and should) respond with a "you too" rebuttal and in this case that rebuttal is not a tu quoque fallacy.

4.3.4 Genetic fallacy

The genetic fallacy occurs when one argues (or, more commonly, implies) that the origin of something (e.g., a theory, idea, policy, etc.) is a reason for rejecting (or accepting) it. For example, suppose that Jack is arguing that we should allow physician assisted suicide and Jill responds that that idea first was used in Nazi Germany. Jill has just committed a genetic fallacy because she is implying that because the idea is associated with Nazi Germany, there must be something wrong with the idea itself. What she should have done instead is explain what, exactly, is wrong with the idea rather than simply assuming that there must be something wrong with it since it has a negative origin. The origin of an idea has nothing inherently to do with its truth or plausibility. Suppose that Hitler constructed a mathematical proof in his early adulthood (he didn't, but just suppose). The validity of that mathematical proof stands on its own; the fact that Hitler was a horrible person has nothing to do with whether the proof is good. Likewise with any other idea: ideas must be assessed on their own merits and the origin of an idea is neither a merit nor demerit of the idea.

Although genetic fallacies are most often committed when one associates an idea with a negative origin, it can also go the other way: one can imply that because the idea has a positive origin, the idea must be true or more plausible. For example, suppose that Jill argues that the Golden Rule is a good way to live one's life because the Golden Rule originated with Jesus in the Sermon on the Mount (it didn't, actually, even though Jesus does state a version of the Golden Rule). Jill has committed the genetic fallacy in assuming that the (presumed)

fact that Jesus is the origin of the Golden Rule has anything to do with whether the Golden Rule is a good idea.

I'll end with an example from William James's seminal work, *The Varieties of Religious Experience*. In that book (originally a set of lectures), James considers the idea that if religious experiences could be explained in terms of neurological causes, then the legitimacy of the religious experience is undermined. James, being a materialist who thinks that all mental states are physical states—ultimately a matter of complex brain chemistry, says that the fact that any religious experience has a physical cause does not undermine that veracity of that experience. Although he doesn't use the term explicitly, James claims that the claim that the physical origin of some experience undermines the veracity of that experience is a genetic fallacy. Origin is irrelevant for assessing the veracity of an experience, James thinks. In fact, he thinks that religious dogmatists who take the origin of the Bible to be the word of God are making exactly the same mistake as those who think that a physical explanation of a religious experience would undermine its veracity. We must assess ideas for their merits, James thinks, not their origins.

4.3.5 Appeal to consequences

The appeal to consequences fallacy is like the reverse of the genetic fallacy: whereas the genetic fallacy consists in the mistake of trying to assess the truth or reasonableness of an idea based on the *origin* of the idea, the appeal to consequences fallacy consists in the mistake of trying to assess the truth or reasonableness of an idea based on the (typically negative) *consequences* of accepting that idea. For example, suppose that the results of a study revealed that there are IQ differences between different races (this is a fictitious example, there is no such study that I know of). In debating the results of this study, one researcher claims that if we were to accept these results, it would lead to increased racism in our society, which is not tolerable. Therefore, these results must not be right since if they were accepted, it would lead to increased racism. The researcher who responded in this way has committed the appeal to consequences fallacy. Again, we must assess the study on its own merits. If there is something wrong with the study, some flaw in its design, for example, then that would be a relevant criticism of the study. However, the fact that the results of the study, if widely circulated, would have a negative effect on society is not a reason for rejecting these results as false. The consequences of some idea (good or bad) are irrelevant to the truth or reasonableness of that idea.

Notice that the researchers, being convinced of the negative consequences of the study on society, might rationally choose *not to publish* the study (for fear of the negative consequences). This is totally fine and is not a fallacy. The fallacy consists not in choosing not to publish something that could have adverse consequences, but in claiming that the *results themselves* are undermined by the negative consequences they could have. The fact is, sometimes truth can have negative consequences and falsehoods can have positive consequences. This just goes to show that the consequences of an idea are irrelevant to the truth or reasonableness of an idea.

4.3.6 Appeal to authority

In a society like ours, we have to rely on authorities to get on in life. For example, the things I believe about electrons are not things that I have ever verified for myself. Rather, I have to rely on the testimony and authority of physicists to tell me what electrons are like. Likewise, when there is something wrong with my car, I have to rely on a mechanic (since I lack that expertise) to tell me what is wrong with it. Such is modern life. So there is nothing wrong with needing to rely on authority figures in certain fields (people with the relevant expertise in that field)—it is inescapable. The problem comes when we invoke someone whose expertise is not relevant to the issue for which we are invoking it. For example, suppose that a group of doctors sign a petition to prohibit abortions, claiming that abortions are morally wrong. If Bob cites that fact that these doctors are against abortion, therefore abortion must be morally wrong, then Bob has committed the appeal to authority fallacy. The problem is that doctors are not authorities on what is morally right or wrong. Even if they are authorities on how the body works and how to perform certain procedures (such as abortion), it doesn't follow that they are authorities on *whether or not these procedures should be performed*—the *ethical* status of these procedures. It would be just as much an appeal to consequences fallacy if Melissa were to argue that since some other group of doctors supported abortion, that shows that it must be morally acceptable. In either case, since doctors are not authorities on moral issues, their opinions on a moral issue like abortion is irrelevant. In general, an appeal to authority fallacy occurs when someone takes what an individual says as evidence for some claim, when that individual has no particular expertise in the relevant domain (even if they do have expertise in some other, unrelated, domain).

Exercise 1

- | | |
|---|-----------------------------------|
| 1. Statement | 8. Statement |
| 2. Statement | 9. Statement |
| 3. Not a statement (question) | 10. Statement |
| 4. Statement | 11. Not a statement (question) |
| 5. Not a statement (command) | 12. Not a statement (exclamation) |
| 6. Not a statement
(command/request) | 13. Not a statement (command) |
| 7. Statement | 14. Statement |
| | 15. Statement |

Exercise 2

1. Argument. Conclusion: The woman in the hat is not a witch.
2. Not an argument
3. Argument. Conclusion: Albert won't be willing to help me wash the dishes.
4. Not an argument
5. Not an argument
6. Not an argument
7. Not an argument
8. Argument. Conclusion: Obesity has become a problem in the U.S.
9. Not an argument
10. Argument. Conclusion: Albert isn't a fireman.
11. Argument. Conclusion: Charlie and Violet don't sweat.
12. Argument (explanation). Conclusion: I forgot to lock the door.
13. Not an argument
14. Argument (explanation). Conclusion: Apple stole some of Samsung's patents.
15. Argument. Conclusion: No one who gets frostbite while on K2 will ever survive.

Exercise 3

1. Explanation. Conclusion: Wanda rode the bus today.
2. Explanation. Conclusion: Wanda has not picked up her car from the shop.
3. Argument. Conclusion: Bob rode the bus to work today.
4. Argument. Conclusion: It can't be snowing right now.
5. Explanation. Conclusion: Some people with schizophrenia hear voices in their head.
6. Argument. Conclusion: Fracking should be allowed.
7. Argument. Conclusion: Wanda did not ride the bus today.

8. Argument. Conclusion: The Tigers will not win their game against the Pirates.
9. Argument. Conclusion: No one living in Pompeii could have escaped before the lava from Mt. Vesuvius hit.
10. Explanation: When a person moves to Cincinnati, their allergies worsen.

Exercise 4

1.
 1. There is nothing wrong with consensual sexual and economic relations between adults.
 2. There is no difference between a man paying directly for sex and a man taking a woman on a blind date, paying for it, and then having sex with her.
 3. Therefore, there is nothing wrong with prostitution. (from 1, 2 independently)
2.
 1. Multiple surveys done with prostitutes show that a high percentage of them report having been sexually abused as children.
 2. Therefore, prostitution involves women who were typically abused as children. (from 1)
 3. Therefore, prostitution is wrong. (from 2)
3.
 1. There was warm water in the cabin's tea kettle.
 2. There was wood still smoldering in the fireplace.
 3. Therefore, someone was in this cabin recently. (from 1-2)
 4. Tim has been with me the whole time.
 5. Therefore, the person in the cabin couldn't have been Tim. (from 3-4)
 6. Therefore, there is someone else in these woods. (from 6)
4.
 1. Marla Runyan ran in the Sydney 2000 Olympics when she was blind.
 2. Therefore, it is possible to be blind and yet run in the Olympics. (from 1)
5.
 1. The bridge was out.
 2. Therefore, the train had to take a longer, alternate route. (from 1)
 3. Therefore, the train was late. (from 2)
6.
 1. Iran has threatened to destroy Israel multiple times.
 2. If Iran had a nuclear missile, it could destroy Israel.
 3. Therefore, if Iran had a nuclear missile, Israel is not safe. (from 1-2)

4. Iran has been developing enriched uranium.
 5. Enriched uranium is the most difficult part of the process of building a nuclear weapon—every other part of the process is easy compared to that.
 6. Therefore, Israel is not safe. (from 3-5)
- 7.
1. All professional hockey players are missing front teeth.
 2. Martin is a professional hockey player.
 3. Therefore, Martin is missing his front teeth. (from 1-2)
 4. Almost all professional athletes who are missing their front teeth have false teeth.
 5. Therefore, Martin probably has false teeth. (from 3-4)
- 8.
1. Every time I have eaten crab Rangoon at China Food restaurant, I have had stomach troubles afterward.
 2. Therefore, probably anyone who eats crab Rangoon at China Food restaurant will have stomach troubles afterward. (from 1)
 3. Bob ate the crab Rangoon at China Food restaurant.
 4. Therefore, Bob will probably have stomach troubles afterward. (from 2-3)
- 9.
1. Albert and Caroline like to run together.
 2. Albert never runs without Caroline.
 3. Therefore, any time Albert is running, so is Caroline. (from 1-2)
 4. Albert is panting hard.
 5. Therefore, Albert looks like he has just run (from 4)
 6. Therefore, Caroline has probably just ran as well. (from 3, 5)
- 10.
1. Jeremy's prints would be expected to be on his own gun.
 2. Someone could have stolen Jeremy's gun and used it to kill Tim.
 3. Therefore, just because Jeremy's prints were on the gun that killed Tim and the gun was registered to Tim, it doesn't follow that Jeremy killed Tim. (from 1-2)

Exercise 5 (Note: there are many possible counterexamples to the arguments that are invalid; you don't have to have the counterexamples I provide to be correct.)

1. Invalid. Counterexample: Katie is severely mentally handicapped and so is not smarter than a chimpanzee.

2. Invalid. Counterexample: suppose Bob just became a fireman, so he has never actually put out any fires.
3. Invalid. Counterexample: Although Gerald probably knows how to teach mathematics, suppose he has just had a traumatic brain injury and no longer knows how to teach mathematics. And suppose the injury is recent enough for him not to have lost his job as a mathematics professor yet.
4. Invalid. Counterexample: A similar counterexample as #3 would work equally as well here.
5. Valid.
6. Invalid. Counterexample: Perhaps Craig and Linda are married, but Monique is Linda's secret lover and Craig just finds about it and is angry and hates Monique (but still loves Linda).
7. Invalid. Counterexample: suppose that although Orel Hershizer believes that God exists, in fact God doesn't exist. In that case, Orel can't communicate with God, since you can't communicate with something that doesn't exist (i.e., communication is a two-way interaction).
8. Valid.
9. Valid.
10. Valid.

Exercise 6 (Note: there is more than one possible correct answer to some of these.)

1. Anyone who rides horses is a cowboy.
2. Driving over the speed limit is wrong.
3. It is raining.
4. Olaf is an elf.
5. Any time a person has a choice of who to take to homecoming, they will take the person they like the most.
6. I have looked at the watch in frequent intervals—much more often than every 12 hours.
7. Only those who have drunk too much fall out of apartment windows.
8. Mark is on Earth and is unassisted by any devices that help him overcome the Earth's gravity.
9. Any nation in which there is a large discrepancy between net worths of different races is a racist nation.
10. The water is at sea level.
11. First missing premise: We should not allow policies that have the potential of taking innocent lives, unless there is a very good reason to do

- so. Second missing premise: there is no very good reason to allow capital punishment.
12. First missing premise: We should not allow policies that take working class jobs away from working class folks, unless there is some very good reason to do so. Second missing premise: there is no very good reason to allow immigration that would offset the harm done to working class folks.
 13. Any fair economic exchange between consenting adults should be allowed.
 14. Anything that privileges using a student-athlete to make money for the college over that student-athlete's education should be banned.
 15. Any student who receives an F in a course should have studied more for that course.

Exercise 7

- | | | |
|----------------|----------------|----------------|
| 1. Discounting | 5. Discounting | 9. Discounting |
| 2. Assuring | 6. Guarding | 10. Assuring |
| 3. Discounting | 7. Assuring | |
| 4. Assuring | 8. Discounting | |

Exercise 8

1. Not truth functional.
2. Truth functional: Tom is a fireman. Tom is a father.
3. Not truth functional.
4. Not truth functional.
5. Truth functional: Cameron Dias has had several relationships. Cameron Dias has never married.
6. Not truth functional.
7. Not truth functional.
8. Not truth functional.
9. Truth functional: Jack is a cowboy. Jill is a cowboy.
10. Truth functional: Josiah is Amish. Josiah is a drug dealer.
11. Truth functional: The Tigers are the best baseball team in the state. The Tigers are not as good as the Yankees.
12. Truth functional: Bob went to the beach to enjoy some rest. Bob went to the beach to enjoy some relaxation.
13. Truth functional: Lauren isn't the fastest runner on the team. Lauren is fast enough to have made it to the national championship.
14. Truth functional: The ring is beautiful. The ring is expensive.

15. Truth functional: It is sad that many Americans do not know where their next meal will come from. It is true that many Americans do not know where their next meal will come from.

Exercise 9

- | | | |
|---------------|-------------------------|--------------|
| 1. $B \vee T$ | 5. $\sim G$ | 9. $\sim B$ |
| 2. $\sim S$ | 6. $\sim M \vee \sim T$ | 10. $\sim A$ |
| 3. $\sim B$ | 7. $\sim T$ | |
| 4. $A \vee B$ | 8. $C \cdot D$ | |

Exercise 10

1. $(\sim A \cdot \sim S) \cdot M$ (The main operator is the second dot—in this case it doesn't actually matter which dot since the sentence has the same meaning whichever of the conjuncts you treat as the main operator.)
2. $\sim S \cdot \sim H$ (The main operator is the conjunction.)
3. $P \cdot \sim N$ (The main operator is the conjunction.)
4. $(\sim S \cdot \sim F) \cdot A$ (The main operator is the second dot—in this case it doesn't actually matter which dot since the sentence has the same meaning whichever of the conjuncts you treat as the main operator.)
5. $\sim C \cdot (F \cdot B)$ (The main operator is the first dot—in this case it doesn't actually matter which dot since the sentence has the same meaning whichever of the conjuncts you treat as the main operator.)
6. $(T \cdot \sim S) \cdot L$ (The main operator is the second dot—in this case it doesn't actually matter which dot since the sentence has the same meaning whichever of the conjuncts you treat as the main operator.)
7. $E \cdot W$ (There is only one truth functional operator, the conjunction. So that is by default the main operator!)
8. $(T \vee L) \cdot \sim S$ (The main operator is the conjunction.)
9. $(B \vee F) \cdot \sim(P \vee D)$ (The main operator is the conjunction.)
10. $(A \cdot J) \cdot E$ (The main operator is the second dot—in this case it doesn't actually matter which dot since the sentence has the same meaning whichever of the conjuncts you treat as the main operator.)
11. $(A \cdot J) \vee (C \cdot M)$ (The wedge is the main operator.)
12. $S \cdot \sim H$ (The main operator is the conjunction.)

Exercise 11

1.
 - a. $P = \text{Coral is a plant}; A = \text{Coral is an animal}$
 - b. $\sim(P \cdot A)$

- c. Main operator is the negation
- 2.
- a. A = Protozoa are eukaryotes; B = Chimpanzees are eukaryotes; C = Protozoa are animals; D = Chimpanzees are animals
- b. $(A \cdot B) \cdot \sim(C \cdot D)$
- c. Main operator is the second dot
- 3.
- a. C = Chimpanzees are prokaryotes; P = Protozoa are prokaryotes
- b. $\sim(C \vee P)$
- c. Main operator is the negation
- 4.
- a. C = China has not signed the Kyoto Protocol; U = The United States has not signed the Kyoto Protocol
- b. $\sim(C \vee U)$
- c. Main operator is the negation
- 5.
- a. C = Chevrolet will support the Olympic team; M = McDonald's will support the Olympic team
- b. $(C \vee M) \cdot \sim(C \cdot M)$
- c. Main operator is the first dot
- 6.
- a. L = Peter Jennings is a liar; M = Peter Jennings has a really bad memory
- b. $L \vee M$
- c. Main (and only) operator is the wedge
- 7.
- a. L = Peter Jennings is a liar; M = Peter Jennings has a really bad memory
- b. $\sim(L \vee M)$
- c. Main operator is the negation
- 8.
- a. L = Peter Jennings is a liar; M = Peter Jennings has a really bad memory
- b. $L \cdot M$
- c. Main (and only) operator is the dot
- 9.
- a. L = Peter Jennings is a liar; M = Peter Jennings has a really bad memory
- b. $\sim(L \cdot M)$

10. c. Main operator is the negation
10. a. C = Chevrolet will support the Olympic team this year; M = McDonald's will support the Olympic team this year
10. b. $\sim(C \vee M)$
10. c. Main operator is the negation
11. a. S = Mother Theresa is a saint; C = Mother Theresa has been canonized by the Catholic Church
11. b. $S \cdot \sim C$
11. c. Main operator is the dot
12. a. T = Paul Tergat was the best distance runner of the last two decades; G = Haile Gebrselassie was the best distance runner of the last two decades; R = Jim Ryun was the best distance runner of the last two decades
12. b. $(T \vee G) \cdot \sim R$
12. c. Main operator is the dot
13. a. R = Jim Ryun was the best high school miler of all time; W = Jim Ryun ran a slower time than Alan Webb
13. b. $R \cdot W$
13. c. Main (and only) operator is the dot
14. a. A = Paul Tergat knows how to play hockey; B = Haile Gebrselassie knows how to play hockey; C = Paul Tergat knows how to play soccer; D = Haile Gebrselassie knows how to play soccer
14. b. $\sim(A \vee B) \cdot (C \cdot D)$
14. c. Main operator is the first dot
15. a. B = Ethiopians are good bobsledders; T = Ethiopians are good tennis players; D = Ethiopians are good distance runners
15. b. $\sim(B \vee T) \cdot D$
15. c. Main operator is the dot
16. a. S = Before Helen Keller met Annie Sullivan she could speak; R = Before Helen Keller met Annie Sullivan she could read; C = Before Helen Keller met Annie Sullivan she could communicate
16. b. $(\sim S \cdot \sim R) \cdot \sim C$

- c. Main operator is the second dot
- 17.
- a. C = Helen Keller learned to communicate; S = Helen Keller learned to play soccer; B = Helen Keller learned to play baseball
- b. $C \cdot (\sim S \cdot \sim B)$
- c. Main operator is the first dot
- 18.
- a. F = Tom is allowed to play football; S = Tom is allowed to play soccer
- b. $(F \vee S) \cdot \sim(F \cdot S)$
- c. Main operator is the first dot
- 19.
- a. E = Tom will major in engineering; P = Tom will major in physics; B = Tom will major in business; S = Tom will major in sociology
- b. $(E \cdot P) \vee (B \cdot S)$
- c. Main operator is the wedge
- 20.
- a. X = Cartman is xenophobic; R = Cartman is a racist; M = Cartman is a murderer; T = Cartman is the thief
- b. $(X \cdot R) \cdot (\sim M \cdot \sim T)$
- c. Main operator is the second dot

Exercise 12

1. Invalid
2. Valid
3. Invalid
4. Valid

A	B	C	$(A \vee B) \cdot (A \vee C)$	$\sim A$	$B \vee C$
T	T	T	T T T	F	T
T	T	F	T T T	F	T
T	F	T	T T T	F	T
T	F	F	T T T	F	F
F	T	T	T T T	T	T
F	T	F	T F F	T	T
F	F	T	F F T	T	T
F	F	F	F F F	T	F

5. Invalid

R	T	S	$R \cdot (T \vee S)$	T	$\sim S$
T	T	T	T T	T	F
T	T	F	T T	T	T
T	F	T	T T	F	F
T	F	F	F F	F	T
F	T	T	F T	T	F
F	T	F	F T	T	T
F	F	T	F T	F	F
F	F	F	F F	F	T

6. Invalid

7. Valid

8. Valid

A	B	$\sim(A \vee B)$	$\sim A \vee \sim B$
T	T	F T	F F F
T	F	F T	F T T
F	T	F T	T T F
F	F	T F	T T T

9. Valid

D	R	S	$(R \vee S) \cdot \sim D$	$\sim R$	$S \cdot \sim D$
T	T	T	T F F	F	F
T	T	F	T F F	F	F
T	F	T	T F F	T	F
T	F	F	F F F	T	F
F	T	T	T T T	F	T
F	T	F	T T T	F	F
F	F	T	T T T	T	T
F	F	F	F F T	F	F

Exercise 13

1. $T \supset I$ (T = The Tigers will win; I = The Indians will lose their star pitcher)
2. $H \supset P$ (P = Tom will pass the class; H = Tom does all of his homework)
3. $R \supset G$ (R = The car will run; G = The car has gas)

4. $A \supset C$ (A = You are asking me about your grade; C = You care about your grade)
5. $F \cdot (T \supset B)$ (F = Frog will swim without his bathing suit; T = Toad will swim; B = Toad is wearing a bathing suit)
6. $\sim O \supset M$ (O = Obama is not a U.S. citizen; M = I am a monkey's uncle)
7. $T \supset \sim F$ (T = Toad wears his bathing suit; F = Toad wants Frog to see him in his bathing suit)
8. $\sim P \supset (S \vee L)$ (P = Tom passes his exam; S = Tom is stupid; L = Tom is lazy)
9. $H \supset W$ (W = Bekele will win the race; H = Bekele stays healthy)
10. $(S \vee I) \supset \sim W$ (S = Bekele is sick; I = Bekele is injured; W = Bekele will win the race)
11. $P \supset (C \cdot \sim S)$ (P = Bob will become president; C = Bob runs a good campaign; S = Bob says something stupid)
12. $T \supset P$ (T = That plant has three leaves; P = That plant is poisonous)
13. $P \supset T$ (T = That plant has three leaves; P = That plant is poisonous)
14. $P \supset T$ (T = That plant has three leaves; P = That plant is poisonous)
15. $P \supset T$ (T = That plant has three leaves; P = That plant is poisonous)
16. $N \supset O$ (O = Olga will swim in the open water; N = There is a shark net present)
17. $O \supset N$ (O = Olga will swim in the open water; N = There is a shark net present)
18. $O \supset B$ (O = Olga is swimming; B = Olga is wearing a bathing suit)
19. $N \supset \sim B$ (N = Olga is in Nice; B = Olga wears a bathing suit)
20. $T \supset B$ (T = Terrence pulls Philip's finger; B = Something bad will happen)

Exercise 14

- | | |
|-------------------|--------------------|
| 1. Equivalent | 6. Equivalent |
| 2. Equivalent | 7. Not equivalent |
| 3. Equivalent | 8. Equivalent |
| 4. Equivalent | 9. Not equivalent |
| 5. Not equivalent | 10. Not equivalent |

Exercise 15

1. Contingent

A	B	$A \supset (A \cdot B)$
T	T	T T
T	F	F F
F	T	T F
F	F	T F

2. Tautology

A	B	$(A \cdot B) \supset (\sim A \supset \sim B)$
T	T	T T T
T	F	F T T
F	T	F T F
F	F	F T T

3. Tautology

A	B	$(A \cdot \sim A) \supset B$
T	T	F T T
T	F	F T F
F	T	F T T
F	F	F T F

4. Contradiction

A	B	$(A \supset A) \supset (B \cdot \sim B)$
T	T	T F F
T	F	T F F
F	T	T F F
F	F	T F F

5. Tautology

A	B	$(A \cdot B) \supset (A \vee B)$
T	T	T T T
T	F	F T T
F	T	F T T
F	F	F T F

6. Contingent

A	B	$(A \vee B) \supset (A \cdot B)$
T	T	T T T
T	F	T F F
F	T	T F F
F	F	F T F

7. Contingent

A	B	$(\sim A \supset \sim B) \supset (\sim B \supset \sim A)$
T	T	T T T
T	F	T F F
F	T	F T T
F	F	T T T

8. Tautology

A	B	$(A \supset B) \supset (\sim B \supset \sim A)$
T	T	T T T
T	F	F T F
F	T	T T T
F	F	T T T

9. Contingent

A	B	$(B \vee \sim B) \supset A$
T	T	T T
T	F	T T
F	T	T F
F	F	T F

10. Tautology

A	B	$(A \vee B) \vee \sim A$
T	T	T T F
T	F	T T F
F	T	T T T
F	F	F T T

Exercise 16

- | | |
|---|---|
| <p>1.</p> <p style="padding-left: 20px;">Simplification 1
Simplification 1
Addition 2
Addition 3
Conjunction 4, 5</p> <p>2.</p> <p style="padding-left: 20px;">Disjunctive syllogism 2, 3
Modus ponens 1, 4
Modus tollens 2, 5</p> <p>3.</p> <p style="padding-left: 20px;">Simplification 3
Modus tollens 1, 4
Disjunctive syllogism 2, 5</p> <p>4.</p> <p style="padding-left: 20px;">Simplification 2
Hypothetical syllogism 1, 3
Modus ponens 4, 5
Simplification 2
Conjunction 6, 7</p> <p>5.</p> <p style="padding-left: 20px;">Hypothetical syllogism 3, 4</p> | <p>Addition 1
Constructive dilemma 2, 5, 6</p> <p>6.</p> <p style="padding-left: 20px;">Simplification 2
Modus tollens 1, 4
Disjunctive syllogism 3, 5</p> <p>7.</p> <p style="padding-left: 20px;">Simplification 2
Modus ponens 3, 4
Simplification 2
Disjunctive syllogism 5, 6
Conjunction 4, 7
Modus ponens 1, 8</p> <p>8.</p> <p style="padding-left: 20px;">Simplification 2
Disjunctive syllogism 3, 5
Simplification 2
Modus tollens 1, 7
Disjunctive syllogism 4, 8
Modus ponens 6, 9
Addition 10</p> |
|---|---|

Exercise 17

- 1.
1. $A \cdot B$
 2. $(A \vee C) \supset D \quad \therefore A \cdot D$
 3. A Simplification 1
 4. $A \vee C$ Addition 3
 5. D Modus ponens 2, 4
 6. $A \cdot D$ Conjunction 3, 5
- 2.
1. A
 2. $B \quad \therefore (A \vee C) \cdot B$
 3. $A \vee C$ Addition 1
 4. $(A \vee C) \cdot B$ Conjunction 2, 3

3.

1. $D \supset E$
2. $D \cdot F \quad \therefore E$
3. $D \quad$ Simplification 2
4. $E \quad$ Modus ponens 1, 3

4.

1. $J \supset K$
2. $J \quad \therefore K \vee L$
3. $K \quad$ Modus ponens 1, 2
4. $K \vee L \quad$ Addition 3

5.

1. $A \vee B$
2. $\sim A \cdot \sim C \quad \therefore B$
3. $\sim A \quad$ Simplification 2
4. $B \quad$ Disjunctive syllogism 1, 3

6.

1. $A \supset B$
2. $\sim B \cdot \sim C \quad \therefore \sim A$
3. $\sim B \quad$ Simplification 2
4. $\sim A \quad$ Modus tollens, 1, 3

7.

1. $D \supset E$
2. $(E \supset F) \cdot (F \supset D) \quad \therefore D \supset F$
3. $E \supset F \quad$ Simplification 2
4. $D \supset F \quad$ Hypothetical syllogism 1, 3

8.

1. $(T \supset U) \cdot (T \supset V)$
2. $T \quad \therefore U \vee V$
3. $T \vee U \quad$ Addition 2
4. $U \vee V \quad$ Constructive dilemma 1, 2, 3

9.

1. $(E \cdot F) \vee (G \supset H)$
2. $I \supset G$

3. $\sim(E \cdot F) \therefore I \supset H$
4. $G \supset H$ Disjunctive syllogism 1, 3
5. $I \supset H$ Hypothetical syllogism 2, 4

10.

1. $M \supset N$
2. $O \supset P$
3. $N \supset P$
4. $(N \supset P) \supset (M \vee O) \therefore N \vee P$
5. $M \vee O$ Modus ponens 3, 4
6. $N \vee P$ Constructive dilemma 1, 2, 5

11.

1. $A \vee (B \supset A)$
2. $\sim A \cdot C \therefore \sim B$
3. $\sim A$ Simplification 2
4. $B \supset A$ Disjunctive syllogism 1, 3
5. $\sim B$ Modus tollens 3, 4

12.

1. $(D \vee E) \supset (F \cdot G)$
2. $D \therefore F$
3. $D \vee E$ Addition 2
4. $F \cdot G$ Modus ponens 1, 3
5. F Simplification 4

13.

1. $T \supset U$
2. $V \vee \sim U$
3. $\sim V \cdot \sim W \therefore \sim T$
4. $\sim V$ Simplification 3
5. $\sim U$ Disjunctive syllogism 2, 4
6. $\sim T$ Modus tollens 1, 5

14.

1. $(A \vee B) \supset \sim C$
2. $C \vee D$
3. $A \therefore D$
4. $A \vee B$ Addition 3

- 5. $\sim C$ Modus ponens 1, 4
- 6. D Disjunctive syllogism 2, 5

15.

- 1. $L \vee (M \supset N)$
- 2. $\sim L \supset (N \supset O)$
- 3. $\sim L$ $\therefore M \supset O$
- 4. $N \supset O$ Modus ponens 2, 3
- 5. $M \supset N$ Disjunctive syllogism 1, 3
- 6. $M \supset O$ Hypothetical syllogism 4, 5

16.

- 1. $A \supset B$
- 2. $A \vee (C \cdot D)$
- 3. $\sim B \cdot \sim E$ $\therefore C$
- 4. $\sim B$ Simplification 3
- 5. $\sim A$ Modus tollens 1, 4
- 6. $C \cdot D$ Disjunctive syllogism 2, 5
- 7. C Simplification 6

17.

- 1. $(F \supset G) \cdot (H \supset I)$
- 2. $J \supset K$
- 3. $(F \vee J) \cdot (H \vee L)$ $\therefore G \vee K$
- 4. $F \supset G$ Simplification 1
- 5. $F \vee J$ Simplification 3
- 6. $G \vee K$ Constructive dilemma 2, 4, 5

18.

- 1. $(E \vee F) \supset (G \cdot H)$
- 2. $(G \vee H) \supset I$
- 3. E $\therefore I$
- 4. $E \vee F$ Addition 3
- 5. $G \cdot H$ Modus ponens 1, 4
- 6. G Simplification 5
- 7. $G \vee H$ Addition 6
- 8. I Modus ponens 2, 7

19.

1. $(N \vee O) \supset P$
2. $(P \vee Q) \supset R$
3. $Q \vee N$
4. $\sim Q \quad \therefore R$
5. N Disjunctive syllogism 3, 4
6. $N \vee O$ Addition 5
7. P Modus ponens 1, 6
8. $P \vee Q$ Addition 7
9. R Modus ponens 2, 8

20.

1. $J \supset K$
2. $K \vee L$
3. $(L \cdot \sim J) \supset (M \cdot \sim J)$
4. $\sim K \quad \therefore M$
5. L Disjunctive syllogism 2, 4
6. $\sim J$ Modus tollens 1, 4
7. $L \cdot \sim J$ Conjunction 5, 6
8. $M \cdot \sim J$ Modus ponens 3, 7
9. M Simplification 8

Exercise 18

1. All real men are things that wear pink.
2. No dinosaurs are birds.
3. All birds are things that evolved from dinosaurs.
4. Some mammals are not predators. [Already in "Some S are not P" categorical form.]
5. Some predators are not mammals. [Already in "Some S are not P" categorical form.]
6. Some things that wander are not things that are lost.
7. No presidents are women.
8. No boxers are rich people.
9. No things that are sleeping are things that are conscious.
10. No things that are conscious are things that are sleeping.
11. All things that end well are things that are well. [Note that this is in a different order than the sentence states. However, if you think about the meaning of the cliché, it should be clear that this is the correct order.]

- What is being said is that as long as something ends well, it is well. It is not saying that the *only* things that can be well are things that end well.]
12. All things that care are things that are my friends.
 13. Some person is a person who loves you.
 14. All people are people who are loved by Jesus.
 15. All little children are people who are loved by Jesus.
 16. Some people are people who don't love Jesus.
 17. All things that use the Appalachian Trail are pedestrians.
 18. All presidents are citizens.
 19. All Hindus are people who believe in God.
 20. No cheap things are good things.
 21. Some expensive things are not good things.
 22. Some mammals are not things that have legs.
 23. Some couples are couples without children.
 24. No people are chocolate-haters.
 25. Some people are cat-haters.
 26. No sharp things are safe things.
 27. No poodles are things that can run faster than a cheetah.
 28. No professional runner is a person that runs slowly.
 29. No baboons are friendly creatures.
 30. All things are things that would be eaten by a pig.

Exercise 19

- | | |
|------------|------------|
| 1. Invalid | 5. Valid |
| 2. Valid | 6. Invalid |
| 3. Invalid | 7. Invalid |
| 4. Invalid | 8. Invalid |

Exercise 20

1. Invalid
2. Invalid
3. Invalid
4. Invalid

Exercise 21

- | | | |
|------------|------------|------------|
| 1. Invalid | 4. Invalid | 7. Invalid |
| 2. Invalid | 5. Invalid | 8. Invalid |
| 3. Valid | 6. Valid | 9. Valid |

- | | | |
|-------------|-------------|-------------|
| 10. Valid | 14. Invalid | 18. Invalid |
| 11. Invalid | 15. Invalid | 19. Invalid |
| 12. Invalid | 16. Valid | 20. Invalid |
| 13. Invalid | 17. Valid | |

Exercise 22

1. Hasty generalization (you can't infer something general from just one case here—the sample size is way too small). There is also a sampling bias present: even if many others people from Silverton, CO drove pickups, it doesn't follow that people generally do. There is a high percentage of trucks in Silverton because the rough roads there almost require trucks.
2. Biased sample: even if he has an adequate sample size, Tom needs to sample from different times during the morning to be sure that he has a representative sample. If morning doves are disproportionately represented during the early morning hours, then his sample will be biased.
3. Even more clearly than the previous example, this one is a biased sample: even if he has an adequate sample size, Tom needs to sample from different times of the day. It is likely that morning doves will be disproportionately represented in the morning, since they are more likely to be out in the morning than other kinds of birds.
4. This example corrects the problems of the previous two: Tom has sampled from different times during the day. As long as he has taken these samples on multiple different days (preferably in different seasons too), then his sample is representative and his generalization is good.
5. Biased sample. Same problem, mutatis mutandis, as #3.
6. This seems to be a good generalization, assuming that he keeps up this regimen on multiple days. The difference, of course, is that instead of making his generalization cover the whole day, his generalization is only about the birds that land in his tree during the night.
7. Biased sample. Of course the home owners will be likely to support a policy that slashes property taxes. Most likely, those on Medicaid (governmental health care support for the elderly) will not be homeowners but will be in nursing home facilities. If the poll had been administered to Medicaid recipients (who are less likely to own homes), the results would likely have been different.
8. This seems a good generalization. Telephone polls are a good way of getting a random sample, and the sample size is large enough if a good random sampling technique is used.

9. Sampling bias because of the biased way the question is asked: “killing innocent children” uses strong, evaluative language and may influence how people answer, making them more likely to choose option b over option a (who wants to say they support “killing innocent children?”).
10. Steve’s problem is that he has gotten a biased sample. Ani Difrancos concert-goers are not representative of concert-goers tout court. Since Ani Difrancos is very political (and from a feminist perspective), we should expect to see a much higher proportion of such speech at an Ani Difrancos concert. In contrast, Tom Petty is about as apolitical as any musician.
11. Biased sample. We should expect students in detention to be less satisfied, on average, than students generally. Thus, since the principal’s survey was only administered to students in detention, the rate of dissatisfaction will be much higher, which will make the sample unrepresentative and the generalization bad.
12. This seems to be a good generalization. Her generalization only covers “all Pistons games” (rather than all NBA games or all professional sports games, more generally) and she has attended many games over many years. Thus the sample seems to be both representative (i.e., non-biased) and large enough.
13. Unlike the last example, Alice’s generalization now applies to all NBA games, but still uses only her experience at Pistons games. But unless we are given some reason for thinking that Pistons games are representative of all NBA games, we should not assume that Pistons games are representative of all NBA games. Thus, the sample is probably biased (although we do not know for sure that it is, we cannot assume it isn’t without further investigation).
14. Even more than the last example, this one is a biased sample. Unless we have a good reason for thinking that Pistons games are representative of all professional sporting events, we cannot assume that they are.
15. Although we can understand Bob’s fear, this is clearly a hasty generalization since he is generalizing from only one case at one Burger King to all Burger Kings, all the time.

Exercise 23 (Note: for many of these, there is more than one correct answer. The important thing to do is to give the correct explanation for why the explanation lacks the virtue you have chosen.)

1. This could be any number of them, including: depth (why would the aliens have kidnapped him and then returned him to his home?), power (this explanation cannot be used in a range of different circumstances—a

- better explanation is simply that he has some kind of amnesia), or simplicity (if we don't have any other reason to admit there are aliens, then we should simply chalk it up to some kind of amnesia).
2. Modesty. There is no reason she should posit all of those specific details about the badger, even if it was a badger. However, even just saying it's a badger or a large rodent is an explanation that seems to lack simplicity. If houses naturally creak and windows rattle from the wind, then positing a large rodent seems unnecessary. A better explanation would simply be that the house creaks naturally as it slightly shifts and the wind is rattling the windows.
 3. Simplicity and modesty. It is simpler to simply assume that there is someone who looks like Bob, whether or not he is Bob's identical twin. It is also more modest since positing someone who looks like Bob could include someone that is Bob's identical twin, but also leaves open the possibility that it's just an unrelated person who happens to look like Bob. The explanation might also lack power insofar as it raises more questions than it answers. For example, why did Bob never tell you about his identical twin?
 4. Conservativeness: people don't die and come back to life, as far as we know. Thus, we could also say it lacks power since this kind of explanation doesn't apply in any other cases we know. A better explanation is that there is someone who looks just like Tom.
 5. Modesty. Like #3, a more modest explanation is that this is someone who looks like Tom, whether or not it is Tom's son. The explanation might also lack depth since we would want to know why you had never seen or heard of Tom's son for 20 years.
 6. The last line is the giveaway: this explanation lacks falsifiability. The reason is that Elise says that there is no way to prove that this happened (she just knows it). The explanation also lacks depth since we would want to know why and how this replacement was done!
 7. If this explanation lacks an explanatory virtue, it is probably falsifiability: there is no way (within current science) to show that there wasn't such a being. Furthermore, it might also lack depth since it raises the question: where did this all-powerful being come from?
 8. Modesty. Why think that it is her 5th grade teacher rather than just some person following her? The explanation is far more specific than it needs to be in order to explain the observations she has made. Thus, it lacks modesty.

9. Again, this explanation lacks modesty. Why not just say that it is “an animal” rather than “an escaped zoo animal.” Unless she has some evidence relevant to the escaped zoo animal hypothesis, she should just leave it at the more general “animal” hypothesis. Furthermore, the explanation may be said to lack power, as well. Since most such noises are made by creatures in the wild, not escaped zoo animals, the “creatures in the wild” explanation is more powerful, since it is used to explain a much wider range of similar observations (i.e., hearing rustling in the bushes and sticks cracking on the ground while in the woods).
10. Simplicity. The simpler explanation is that Bill was speeding, not that they had tracked his overdue library book. It also lacks power since most of the time when people are pulled over on the highway it is for speeding, not unreturned library books.
11. This is a good explanation and seems to lack no explanatory virtue.
12. This explanation clearly lacks modesty. Why say that someone was going precisely 13.74 mph over the speed limit rather than saying that they were going over the speed limit (without specifying how far)? That specificity is not justified by the observed facts.
13. Conservativeness. We have no good reason for positing some whole new breed of rats—especially if the claim is that they evolved in her apartment only. This would violate what we know about how evolution works (i.e., we probably need a much larger population for this to happen than the population of rats that are contained in only her apartment). Furthermore, the explanation lacks power since a better explanation that applied to a wider range of circumstances is simply that the rats were not taking the bait.
14. Even more clearly than #13, this one lacks conservativeness. There are no known cases of anything being immortal and this idea violates our understanding of the basic laws of nature. Nothing is immortal.
15. Again, this explanation lacks conservativeness (i.e., it violates our understanding of nature which says that nothing is immortal). A better explanation is that the bullets Bob put in his gun were blanks (cf. the movie, *Crash*).

1. Weak: if the painting is hanging in your high school, it probably isn't a Rembrandt. That is the disanalogy: even if the colors are very similar, almost all Rembrandts hang in galleries, not in high schools.
2. Weak. Although the similarity is that they are both poodles, there is probably some other characteristic that accounted for me being bitten. That is, it probably wasn't the fact that the dog that bit me was a poodle, but more likely that I was invading its space or it felt threatened, etc. It could have likely been some other breed in the same circumstances. So it isn't "poodleness" that accounts for the biting. That said, if we had evidence that poodles are much more likely to bite than other breeds then this argument would be stronger.
3. Strong. Unlike, the last one, this argument delivers a much stronger analogy between past events (poodle-encounters and poodle-bitings) and the current event (poodle-encounter).
4. Strong. The relevant similarities are: 1) Van Cleave's class doesn't change much from semester to semester, 2) the person has the same abilities as their friend who got the A.
5. Weak. Although both are crimes, there are many relevant differences between committing rape and robbing a bank.
6. Weak. There is no particular relationship between having seats, wheels, and brakes, on the one hand, and being safe to drive, on the other. So having seats, wheels and brakes is not a relevant similarity between the two cars, if what we are interested in is how safe they are.
7. Strong. The car company (Volvo) is a relevant similarity between the old cars and the new car. We can expect similar quality between cars from the same company. In contrast, knowing that a car has wheels, brakes and seats tells us essentially nothing about its quality, including its safety.
8. Strong.
9. Weak. A birthday party and a funeral are not relevantly similar in this case. A funeral is a much more important family event than a birthday party (typically). So we should not expect similarity with respect to a professor's absence policy when comparing birthday parties to funerals.
10. Weak. Although both may influence happiness, the relevant difference is that whereas heart and brain surgery are typically a matter of life and death (and hence much more likely to be paid for by insurance), cosmetic surgery is not a matter of life and death.
11. Weak. Although a knife and spoon share the property of being eating utensils, that is not a relevant similarity on which we can expect that they will share functional properties like cutting.

12. Whether this famous argument for the existence of God is strong or weak is a matter of some debate. One reason for saying it is a weak argument is that there is a disanalogy between artificial objects and natural objects, since complex natural objects may evolve without being designed by an intelligent designer, whereas no artificial objects (yet) can evolve on their own.
13. Weak. Running the same number of miles as an elite runner is not a relevant similarity for determining how fast one will run a race. The relevant dissimilarity here is that although Bekele runs his mile repeats at close to 4:00 flat, I can only run mine at 5:30. So it is the pace at which one runs, rather than the number of miles one runs, that is the better predictor of how fast one can run a race.
14. Strong. The fact that we are both humans is relevant to determining whether someone will feel pain. Humans all have similar physiology, which is why we should expect that if x causes one person physical pain, then x will also cause anyone else a similar pain. (However, this argument also raises a famous problem in philosophy of mind called “the problem of other minds.” The issue is whether or not we can ever know that people have mental states, such as pain, like my own. Even if you exhibit pain behavior in similar instances in which I experience pain, how do I know that you are actually feeling what I am feeling—that you are having the experience of pain, rather than simply exhibiting pain behavior without have the mental experience of it? Many philosophers have argued that we cannot overcome this problem and must admit that we cannot know whether people other than ourselves actually have mental states like ours.)
15. Again, the common sense answer would be that this is a strong argument based on a strong analogy. Since you and I are both human and share similar perceptual systems, we should expect that we will perceive the world very similarly (even if not exactly the same). (However, we can raise the same “problem of other minds” problem here as I did in #14 above. Suppose we both point at the grass and say that it is green. However, how do I know that your experience of green is like my experience of green? Maybe your experience of green is more like my experience of red and vice versa.)

Exercise 25

1. C is sufficient since any time it is present, the target G is present. Both C and D are necessary, since any time the target G is present, they are present.
2. A and C are sufficient; A and C are also necessary.
3. C and D are necessary. All candidates are sufficient, since there is never a case where a candidate is present but the target is absent. However, the test for sufficient conditions is not rigorous, since there is no case in which the target is absent (hence, it is trivially true that there is no case in which a candidate is present when the target is absent).
4. No one candidate is sufficient (case 1 rules them all out). Only C is necessary.
5. C is sufficient; B, C, and D are necessary.
6. None are sufficient (case 1 rules them all out); A and D are necessary.
7. B is sufficient; A, B and D are necessary.
8. C is sufficient; A, B, C, and D are all necessary. However, the test for necessary conditions is not rigorous, since there is no case in which the target is present (hence, it is trivially true that there is no case in which a candidate is absent when the target is present).
9. B and C are sufficient; C and D are necessary.
10. A and C are sufficient; D is necessary.

Exercise 26

1. Accidental
2. B causes A. Perhaps when we are away from home, Charlie gets less exercise, hence putting on more weight. In that case B is (indirectly) causing A, since our being away from home results in Charlie getting less exercise, which results in him gaining more weight.
3. Common cause. The common cause is simply the factors that make plants grow, such a sunlight, water, and good soil. In this case, A and B are not causing each other, but there is something else (i.e., growth factors of plants) that is causing each one, independently to grow.
4. A causes B. The more bombing, the more stress for the president. And the more stress for the president, the more of his hairs fall out.
5. B causes A. Arguably, the average number of fires per year will influence the number of fire engines needed. Fewer fires would require fewer engines (in which case they'd likely retire some); more fires would require more engines (in which case they'd likely acquire some).
6. Common cause: agrarian societies will have more mules and will also probably pay professors less, since higher education is not as important in

- an agrarian society. So the common cause is being a (more or less) agrarian society.
7. B causes A. Drivers are distracted by the scantily-clad models on the billboards and more distracted drivers causes more accidents.
 8. A causes B. The wider one's waist, the higher the weight due to fat. The higher the weight due to fat, the lower one's vertical leap.
 9. B causes A. The heat causes slower marathon times.
 10. Common cause: Aging. The older one is, the more likely one will have gray hair and the more likely one will have more children or grandchildren. However, neither of these factors are causing the other. Rather, both are being caused independently by a common factor: age.

Exercise 27

- | | |
|-----------------------------------|-----------------------------------|
| 1. $1/6$ | 6. $35/36$ |
| 2. $5/6$ | 7. $1 - (5/6 \times 5/6) = 11/36$ |
| 3. $1/36$ | 8. $1 - 11/36 = 25/36$ |
| 4. $1/36$ | 9. $1/36$ |
| 5. $1 - (5/6 \times 5/6) = 11/36$ | 10. $6/36$ |

Glossary/Index

- "A but B" locution, 33
- Accidental correlation, 163
- Ad hominem fallacy, 196
- Addition, 98
- Adequate sample size condition, 137
- Affirming the consequent, 182
- Antecedent, 82
- Appeal to authority fallacy, 203
- Appeal to consequences fallacy, 202
- Argument, 3
- Argument from analogy, 151
- Assuring, 30
- Atomic proposition, 50
- Background conditions, 156
- Background knowledge, 161
- Base rate fallacy, 175
- Begging the question fallacy, 187
- Categorical logic, 112
- Categorical statement, 112
- Categorical syllogism, 126
- Causal slippery slope fallacy, 194
- Common cause, 162
- Complex proposition, 50
- Composition fallacy, 183
- Conceptual slippery slope fallacy, 192
- Conclusion, 1
- Conclusion indicator, 4, 5
- Conditional statement, 17, 81
- Conjunct, 50
- Conjunction (sentence type), 50
- Conjunction (inference rule), 97
- Conjunction fallacy, 172
- Consequent, 82
- Conservativeness, 145
- Constant, 57
- Constructive dilemma, 99
- Contingent statement, 90
- Contradiction, 89
- Counterexample, 19
- Deductive argument, 23
- Defeasible argument, 24
- Denying the antecedent, 182
- Depth, 145
- Descriptive statement, 27
- Descriptive language, 34
- Discounting, 30
- Disjunct, 59
- Disjunction, 59
- Disjunctive syllogism, 97
- Division fallacy, 185
- Empirical science, 37
- Equivocation, 190
- Evaluative language, 34
- Exclusive or, 60
- Existential commitment, 124
- Explanation, 8
- Explanatoriness, 145
- Fallacies of relevance, 196
- Fallacy, 182
- False cause fallacy, 163
- False dichotomy fallacy, 189
- False positive, 175
- Falsifiability, 145
- Four categorical forms, 114
- Gambler's fallacy, 181
- Genetic fallacy, 201
- Guarding, 30
- Hasty generalization fallacy, 138
- Hypothetical syllogism, 95
- Inclusive or, 60
- Inductive argument, 23
- Inference to the best explanation, 144
- Informal fallacy, 183
- Informal test of validity, 19
- Intermediate conclusion, 11
- Invalid argument, 19
- Is-ought gap, 28
- Joint sufficiency, 158
- Law of small numbers, 178
- Main argument, 12
- Main conclusion, 10

Glossary/Index

- Main operator, 64
- Material equivalence, 87
- Method of concomitant variation, 160
- Missing premise, 27
- Modesty, 145
- Modus ponens, 92
- Modus tollens, 94
- Necessary condition, 85
- Necessary condition test, 158
- Negation, 55
- Negative correlation, 160
- Non-biased sample condition, 137
- Normative statement, 27
- Normative concept, 37
- Normative science, 37
- Ockham's razor, 145
- Paraphrase, 15
- Particular affirmative, 114
- Particular negative, 114
- Population, 137
- Positive correlation, 160
- Power, 145
- Premise, 1
- Premise indicator, 4, 5
- Principle of charity, 27
- Probability of conjunctions, 167
- Probability of disjunctions, 170
- Probability of negations, 168
- Proof, 95.
- Proposition, 49
- Propositional logic, 49
- Purely descriptive term, 34
- Purely evaluative term, 34
- Random sampling, 141
- Reference column, 75
- Regression to the mean, 180
- Regression to the mean fallacy, 181
- Representativeness, 137
- Rigorous testing, 160
- Sample, 137
- Simplicity, 145
- Simplification, 96
- Small numbers fallacy, 179
- Sorites paradox, 191
- Sound argument, 22
- Standard argument form, 3
- Statement, 1
- Statistical generalization, 136
- Strategy of working backwards, 102
- Strategy of working forward, 104
- Straw man fallacy, 198
- Strong claim, 32
- Strong inductive argument, 24
- Subargument, 12
- Substitution test, 6
- Sufficient condition, 85
- Sufficient condition test, 158
- Tautology, 89
- Truth table, 51
- Truth table test of validity, 76
- Truth value, 50
- Truth-functional connective, 50
- Tu quoque fallacy, 200
- Universal affirmative, 114
- Universal generalization, 25
- Universal negative, 114
- Vagueness, 191
- Valid argument, 18
- Venn diagram, 113
- Venn test of validity, 121
- Weak claim, 32
- Weak inductive argument, 24
- Well-formed formula, 64