

Arguments in Context

Arguments in Context

AN INTRODUCTION TO CRITICAL THINKING

THADDEUS ROBINSON

Muhlenberg College
Allentown, PA



Arguments in Context by Thaddeus Robinson is licensed under a [Creative Commons Attribution-NonCommercial 4.0 International License](https://creativecommons.org/licenses/by-nc/4.0/), except where otherwise noted.

Contents

Acknowledgements	vii
Introduction	ix
UNIT I. <u>AN INTRODUCTION TO REASONING</u>	
1. Reasoning and Argument	3
2. Reasoning as a Social Process	15
Unit #1 Summary	25
UNIT II. <u>ARGUMENT ANALYSIS</u>	
3. Analysis, Standardization, and Diagramming	29
4. Complications to Argument Analysis	39
5. Missing Premises	49
6. Argument Analysis--Some Final Notes	61
Unit #2 Summary	75
UNIT III. <u>AN INTRODUCTION TO EVALUATION</u>	
7. Evaluation and Logical Strength	79
8. Factual Correctness	89
9. Bias and Motivated Reasoning	99
Unit #3 Summary	111
UNIT IV. <u>AN INTRODUCTION TO DEDUCTIVE ARGUMENTS</u>	
10. Deduction and Argument Form	115
11. The Language of Deduction	123
Unit #4 Summary	133

UNIT V. COMMON INDUCTIVE ARGUMENTS

12. Arguments from Analogy	137
13. Inference to the Best Explanation	147
14. Inductive Generalizations	155
15. Inductive Applications	167
Unit #5 Summary	175

UNIT VI. SOCIAL ARGUMENTS

16. Testimony and Trust	179
17. Trust and the Ad Hominem	191
18. Agreement, Disagreement, and Popularity	199
Unit #6 Summary	211

UNIT VII. SCIENTIFIC REASONING

19. Dealing with Statistics	215
20. Correlation and Cause	227
21. Significant Correlations and Controlled Studies	237
Unit #7 Summary	245
Concluding Thoughts	247

Acknowledgements

In my view, courses in “Critical Thinking” are difficult to teach well, and this text is my effort to make this task a little easier. In developing this text I’ve tried to build on what I take to be the best books out there, and I want to start by acknowledging three works in particular: Trudy Govier’s *A Practical Study of Argument*, Richard Feldman’s *Reason and Argument*, and Merrilee Salmon’s *Introduction to Logic and Critical Thinking*. More importantly, many individuals have helped me along the way. In particular, I would like to extend my thanks to several of my colleagues at Muhlenberg: to Dan Doviak for many informal conversations about pedagogy in critical thinking, to Steve Coutinho for helping me recognize the subtle importance of names for inclusive pedagogy, and to Cathy Ouelette for reading drafts of Chapter 2. In addition, I’d like to thank Kevin Sharpe (St. Cloud State) in particular who used an early version of this text in his course. The book has changed significantly since then, due in part, to Kevin’s helpful feedback. I also owe a special thanks to Andrew Leahy (Muhlenberg Class of ’21) who worked with me on a full overhaul of the text during the summer of 2019, and to Hallie Hoffman (Muhlenberg Class of ’22) who worked with me to develop the discussion in Chapter 16 of testimony, the web, and social media. My decision to create an OER using Pressbooks was guided and informed by a series of seminars developed by Trexler Library and Muhlenberg’s Digital Learning team, and I would like to thank Tim Clarke in particular for his help and encouragement putting together this project. Many thanks are owed to my wife, Katie Robinson, for her support, and for being a patient sounding board during the development of this text. Finally, I am grateful to the many students whose questions and comments helped me to improve and refine the text along the way.

Thaddeus “Tad” Robinson

Associate Professor of Philosophy, Muhlenberg College

tadrobinson@muhlenberg.edu

Introduction

This is what is commonly called a “Critical Thinking” textbook, and is intended to be used as a resource in a semester-long course. The goal of this text is to help students reflectively enhance their general reasoning skills. Before unpacking this goal, it is important to say something about how a course in Critical Thinking differs from the many courses students have already taken on the subject of reasoning. Our society takes the reasoning skills of its citizens very seriously, and we expend vast resources teaching our children the quantitative reasoning skills that we generically refer to as ‘math’. In kindergarten (if not before) we start with basic arithmetic, and over at least the next decade of their lives we move on to teach children fractions, ratios, exponents, algebra, geometry, and so forth. What unifies these subjects is that they are largely focused on methods for determining precise numerical values (albeit in different contexts and/or situations). Society’s commitment to learning these skills makes perfect sense: knowing how to manipulate and determine numerical values is crucial in the modern world, and most people wouldn’t learn these skills without instruction.

However, not all of our reasoning is quantitative. Many of the conclusions we come to, and decisions we make, are not about numerical values. Moreover, not all the information that we use to draw these conclusions and make these decisions is quantitative. For example, consider a jury deliberation. Each member of the jury is supposed to consider the evidence presented by the prosecution and defense to determine the guilt or innocence of the defendant. Of course, quantitative information might be presented to the jury, but a lot of the information presented to the jury will not be quantitative. For example, witness testimony requires the jurors to determine whether the witness is trustworthy, and to decide to what extent their testimony is relevant to the case at hand. Moreover, the jury will need to weigh the competing explanations of the crime presented by the opposing attorneys, and think about the plausibility of alternative explanations. Finally, of course, the jury will have to determine the defendant’s guilt or innocence using the non-quantitative standard of ‘beyond a reasonable doubt’. More broadly, as we navigate the world we are constantly engaged in reasoning that is not primarily quantitative, and this is the kind of reasoning that we will focus on in this text.

Given this, let us return to the text’s goal of helping readers reflectively enhance their general reasoning skills. What does this mean? Let us start with the term *reflectively enhance*. Quantitative reasoning does not come naturally to most people, and we need to be taught these skills through repeated and prolonged instruction. In contrast, we seem to have natural aptitude for other kinds of everyday reasoning, and this makes it easy to take these skills for granted. However, the fact that a skill comes naturally to us, doesn’t mean we cannot improve on it, and this is precisely the aim of this text. Just as systematically reflecting on the mechanics of running can help a person improve their natural ability to run, so too reflecting on reasoning can enhance our existing skills, and that’s what we’ll do here.

This brings us to the second part of this text's goal: to reflectively enhance our *general reasoning skills*. As we've seen, we will be focusing on non-quantitative reasoning skills. In addition, we will be focused on skills that are general insofar as they are applicable in a wide array of different circumstances and situations. Overall, I take our general reasoning skills to involve identifying, analyzing, creating, questioning, evaluating, and refining arguments. Despite this generality, sharpening these skills can have dramatic effects on our thinking and decision-making by helping to limit the options we take seriously, identify irrelevant issues or poor sources, know what questions to ask, and to reject inadequate arguments offered for or against some course of action (to name just a few).

To these ends, the text features (i) an extensive unit that teaches students to spot and visually represent arguments in multiple ways, (ii) a careful explanation of the standards for argument evaluation, (iii) a detailed consideration of common argument types, along with common mistakes and the questions we need to ask to avoid making them. In addition, a distinctive element of the text is its emphasis on the contexts in which we reason. Specifically, this book emphasizes the social and psychological contexts for reasoning, and it is important to say something about each of these contexts up front.

Let us start with the social context. Although reasoning occurs within the private confines of individual minds, it is almost always a consequence of our interactions with other people. This shouldn't be surprising, after all, we are social animals. We get a great deal of our information from other people, we regularly turn to other people to check our thinking, and we use what other people are saying or doing as a guide for our own lives. Given this, social factors play a significant role in what we know and what we do. Social factors can be a good guide to the truth, but they can lead us astray in some cases, and knowing how social factors influence us can put us in a position to avoid or mitigate the influence of misleading factors. The text emphasizes interpersonal influences on our thinking in a number of ways. First, it places upfront a discussion of the importance of 'cooperative dialogue' and 'cooperative disagreement'. Second, the book includes a unit on 'social arguments' focusing in particular on the role of trust in thinking. Third, the text emphasizes a variety of influences from the social psychology literature, from false consensus to fundamental attribution error and beyond.

This text also focuses on psychological context for reasoning. Psychologists have been studying reasoning for decades, and they have discovered that we are subject to a variety of biases and cognitive illusions. Just as we have to learn that the pencil in the glass of water is not really broken despite the way it looks, so too we need to learn that in some contexts arguments that look good—aren't. This is complicated by the fact that there seem to be different types of reasoning processes going on within us—from slow and consciously controlled thinking to semi-autonomous and completely autonomous processes. A special focus on these factors comes in a discussion of dual-process theories in Unit #1, and bias and motivated reasoning in Unit #3. The text is also focused on fostering the transfer of these skills from the classroom into students' own thinking. Thus, the text draws upon examples and exercises from a wide range of everyday situations in an effort to show the wide relevance and applicability of these ideas. Finally, the text illustrates terms and concepts through conversation between named individuals to both encourage transfer, as well as highlight the social context for much of our thinking.

To put it briefly, then, this book aims to improve its readers' reasoning by teaching the basics of argument identification and evaluation within the social and psychological contexts in which it normally occurs, and ultimately, to give readers' the confidence to think through the challenging problems and decisions we all face.

UNIT I

AN INTRODUCTION TO REASONING

Reasoning and Argument

SECTION 1: INTRODUCTION

We spend a lot of time trying to figure things out, and we do so primarily by means of reason. For example, we try to predict what will happen, explain what has already occurred, generalize from our experiences, and extrapolate from what we know. While we are sometimes interested in abstract questions, most of the time we put our reasoning skills to work on entirely practical matters. We have goals, plans, and interests, and being able to accurately predict, explain, generalize on, and extrapolate from, our experiences are crucial skills for achieving our goals. Indeed, we all know how to reason. We do it all the time, and it is something we are relatively good at doing. Nevertheless, reasoning is also something that we can improve on, and sharpening our skills can have dramatic effects on our beliefs and decisions. In this chapter we will start by defining reasoning and explaining how it relates to arguments. We will then briefly introduce two important skills: argument analysis and argument evaluation. Along the way, we will begin building a vocabulary for thinking about and developing these skills.

SECTION 2: WHAT IS REASONING?

We reason all the time, but what are we doing exactly? This answer might not be immediately obvious, so let's begin with some straightforward cases to see what they have in common.

- Solving a math problem
- Figuring out why your phone won't work correctly
- Deciding who to vote for
- Working out why your friend is angry with you
- Determining whether you can afford to buy a new car

What do these have in common? We could pick out a number of features, but we will focus on two in particular. First, in each case our thinking is driving toward a specific outcome or conclusion (e.g. "my phone won't work because...", or "the answer is...", etc.). Second, in each one of these cases this conclusion will be based on reasons. That is, we will arrive at a specific conclusion because we think we have good reasons for doing so. Let us take a closer look at each of these features.

Reasoning is a mental process that ends with a conclusion. Sometimes this conclusion is a newly formed belief. You might, for example, be asked to find the average of 88, 69, 94, and 77 and arrive at the new belief

that the average of these numbers is 82. Alternatively, you might troubleshoot your phone and arrive at the new belief that the operating system wasn't properly installed. While reasoning always leads to a conclusion, that conclusion need not be a new belief. In some cases, these processes lead us to be more (or less) confident in beliefs we already hold. We can see both kinds of conclusion at work in the following example.

Ex. 1:

Talia wakes up one morning to discover that her car is missing. As she thinks about it she quickly concludes that her brother has probably borrowed it. Her reasons for drawing this conclusion are that he i) knows where her spare set of keys are, ii) has borrowed it without asking in the past, and iii) is supposed to pick up a cake at a bakery across town today. Just after she has come to this conclusion the phone rings. It is one of Talia's friends who mentions that she saw Talia's car parked at the bakery across town.

In the first part of Ex. 1 Talia is reasoning to a new belief, namely that her brother has borrowed her car. When her friend calls, she gets a new piece of information. However, this information does not lead her to any new belief; after all, this information is an additional reason for thinking that her brother has borrowed the car, and she already believes that. Instead, as a result of this news, she is even more confident that her brother has borrowed the car. In light of this distinction, we will say that reasoning is a process that leads to a change in a person's system of belief, and we will understand a person's **system of belief** to include not only their beliefs, but also the relationships between those beliefs, and the confidence with which they are held.

The second defining feature of reasoning is that it is a process whereby we change our system of beliefs because we have *reasons* for doing so. In general, to have reasons for drawing some conclusion is to have some group of existing beliefs that indicate in one way or another that the conclusion in question is true. In the example above, Talia has a variety of existing beliefs about her brother and his circumstances, and she takes this information to point toward the fact that her brother has borrowed the car.

Importantly, our reasons can indicate the truth of the conclusion to different degrees. When we take ourselves to have solved a math problem correctly, for example, we take ourselves to have *shown* that the conclusion is correct. This is a bit different from Talia's reasoning—she probably wouldn't say that her reasons *show* or *prove* that her brother borrowed the car. Nonetheless, she thinks that her reasons are good enough to draw the conclusion. In general, reasons can support conclusions with different degrees of strength, and it should be no surprise that we have many different ways of talking about this support. We can say, for example, that when we reason we take our existing beliefs to *indicate*, *give good reason for*, *offer evidence on behalf of*, *establish*, *warrant*, or *demonstrate* a change to our system of beliefs. These differences will be important later, but for now we will simply say that when we reason, we take our existing beliefs to **justify** a particular change in our system of belief.

Now that we have taken a brief look at the reasoning process, we can return to the question we began with and define reasoning as follows:

Reasoning is the process whereby a person changes their system of belief on the basis of reasons which they take to justify this change.

SECTION 3: TYPES OF REASONING

We have just arrived at a definition of reasoning, but it is important to note that this definition captures only one kind of reasoning. Over the last 20-30 years, psychologists and cognitive scientists have come to the

conclusion that we use a variety of different methods and mechanisms to update our system of beliefs. Some of these methods involve conscious directed attention, but many do not. Here is the basic idea: many of the things that humans do are automatic and do not involve conscious control. We do not, for example, need to tell ourselves to breathe or blink (although we can). Moreover, it is not just physical activity that can be automatic; mental activities can occur automatically as well. Take recognition for example. You do not decide to recognize people—it is something you automatically do. The same goes for laughing. You do not need to think about laughing—it is often an automatic reaction (though sometimes events can become funny as you think about them). Cognitive scientists and psychologists now think that, in addition, there are a variety of automatic and semi-automatic *reasoning* processes. These processes update our systems of belief in ways that are often outside of any conscious effort or awareness.

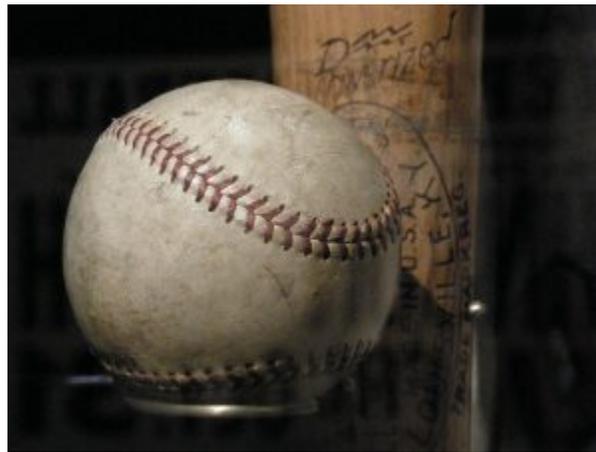
In order to get a sense for different kinds of reasoning processes, consider the following example.¹ In answering this question make sure to note the first answer that comes to mind. Then stop and think about it a little more.

Ex. 2:

A bat and ball cost \$1.10 in total. The bat costs \$1.00 more than the ball. How much does the ball cost?

First Answer that Occurs to You: _____.

Considered Answer: _____.



"National Baseball Hall of Fame: P. Waner's 3000th Career Hit Ball and Bat" by IslesPunkFan CC BY-NC 2.0

Most people read this question and a particular answer just pops into their heads. This answer appears immediately to them, without any particular effort. However, if you take a closer look at this automatic answer, you will see that it can't be right, and figuring out the actual answer to this question takes some work (follow this note if you need some help).² The way the question in Ex. 2 is structured triggers in many people an automatic reasoning mechanism which gives an answer that our conscious reasoning system subsequently recognizes as wrong.

In this case, a largely automatic process leads us to make a mistake. But this is not always the case; in fact, we rely on reasoning processes like this all of the time. As we navigate through the world we automatically and reliably respond to our environment, and we will see that these kinds of reasoning processes naturally generate intuitive reactions to, and impressions of, people, places, circumstances, and ideas (among other

things). These intuitions and impressions contribute to our system of belief, and can thereby inform successive conscious forms of reasoning.

In this book, we will primarily address consciously directed reasoning, since this is the kind of reasoning over which we have the most control. Nonetheless, we cannot ignore other types of reasoning processes. These processes are pervasive features of our thinking and inform our conscious judgments and decisions both positively and negatively. Given this, it is important to know when to be skeptical of the impressions and intuitions that these processes give us. As such, if we want to think more clearly and make better choices, we will have to take these kinds of processes into account. Since we will be primarily discussing directed conscious reasoning, let us simply refer to this as “reasoning” in line with the definition given at the end of section 2. When we need to talk about automatic reasoning processes we will explicitly identify them.

SECTION 4: FROM REASONING TO ARGUMENT

In thinking about reasoning more generally, we will focus our attention primarily on the investigation of arguments. In order to illustrate the difference between reasoning and argument, we will need to start by talking in more detail about beliefs. Normally, we would agree that two people can share a belief. Consider, for example, the fact that both Maria and Jackson are enrolled in Economics 101, and so share the belief that the course is taught on Tuesdays and Thursdays from 2-3:15. In what sense do they share this belief? As we’ve seen, a belief is part of an individual’s system of belief, and so is a psychological state of that individual. Maria’s beliefs are hers, and Jackson’s are his, and this means that in an important sense, Maria and Jackson do not have the same belief. Rather, they share the belief insofar as their distinct beliefs are about the same thing, or have the same content. The sharable content of a belief is called a **proposition**, and we will say that Maria and Jackson believe the same proposition, namely that *Economics 101 is on Tuesdays and Thursdays from 2-3:15*. Making the distinction between a belief and its content is useful, at least in part, because it allows us to talk about the truth of the proposition without talking about any individual person.

This is relevant because there is a similar distinction between reasoning and argument. Suppose that Maria and Jackson are sitting in their Economics class; at 2:05 they both look out the window and see their friend Logan, who is also enrolled in the course, speeding away from the building on his bike. They both conclude that Logan is not coming to class. Let’s compare their thinking. On the one hand, Maria and Jackson have engaged in distinct reasoning processes. After all, Maria’s reasoning had led her to update her system of belief, and Jackson’s reasoning has led him to update his. Nevertheless, there is something common to their distinct reasoning processes. They have both updated their systems of belief to include the same proposition, namely ‘Logan is not coming to class,’ on the basis of the same reasons, namely their respective beliefs in the propositions that ‘It is 2:05’ and ‘Logan is speeding away from the building’. These distinct reasoning processes share the same content, and we will call the content of a process of reasoning, an argument. That is, we will say that **an argument** is a collection of propositions in which one is purportedly justified by the others. When we reason, we take our existing belief that one or more propositions are true (our reasons), to justify our belief that some other proposition is also true (our conclusion). That is, we reason by means of arguments. Like the distinction between beliefs and propositions, drawing the distinction between reasoning and argument is useful because it allows us to evaluate a reasoning processes independently of who is engaging in it.

SECTION 5: ARGUMENTS AND THEIR PARTS

When it comes to arguments, it is important to make two terminological distinctions. First, the term ‘argument’ as it is defined above differs from another common sense of the term. We often use the term ‘an argument’ to refer to a disagreement or a dispute. This is *not* how we will be using the term. Here is an example of a dispute that is not an argument as this book uses the term.

Ex. 3:

Maria: Eating meat is irresponsible and unnecessary.

Jackson: Are you crazy? No it is not.

Why isn’t this disagreement an example of an argument? The answer, in short, is because neither person has tried to justify what they are saying. Presumably each person has reasons for thinking they are right, but as conversation stands all that has been publically expressed is a disagreement. Compare Ex. 3 to the following.

Ex. 4:

I bet the Phillies will win their game tonight since they are on a hot streak.

Ex. 5:

I needed to get at least 90% of the points in this class to earn an A-. Because I got 84%, I didn’t earn an A-.

Ex. 6:

The choice for dinner is either lasagna or pizza. The pizza is too gross to even consider eating. I guess it’s lasagna for me!

Examples 4-6 are all arguments since, in each case, a reason is offered on behalf of a conclusion. As these examples of argumentation show, arguments are common in everyday thinking and need not concern abstract or theoretical topics (although they certainly may).

A second terminological note is that all arguments have two parts—the premises and the conclusion. The **premises** of an argument give reasons or evidence on behalf of the conclusion; put otherwise, premises are the pieces of information that back-up or justify the conclusion. The **conclusion**, on the other hand, is the proposition for which reasons or evidence are given, it is that proposition which is backed-up or justified. We can label the parts of the arguments above accordingly:

Ex. 4:

Premise—The Phillies are on a hot streak

Conclusion—The Phillies will win their game tonight.

Ex. 5:

Premise—I needed to have gotten at least 90% of the points in this class to get an A-

Premise—I got 84%

Conclusion—I won’t get an A-

Ex. 6:

Premise—My choices are either lasagna or pizza.

Premise—There is no way I can eat the pizza—it is always gross.

Conclusion—I will be having lasagna.

With these terminological issues out of the way we can focus on arguments themselves. First, although we reason by means of arguments, it is important to recognize that we often think and reason *about* arguments as well. Here is an example:

Ex. 7:

Maria says “I don’t buy the argument that since voting is a restricted activity, we should require IDs at the polls.”

Maria is considering the following argument:

Premise: Voting is a restricted activity.

Conclusion: We should require IDs at the polls.

In this instance Maria is talking about an argument, but not making one herself. Maria is claiming that one reason people sometimes give for thinking that we should require IDs at the polls is not, to her mind, a good one.



“Vote for Nacho” by rikkis_refuge CC BY 2.0

This raises a second point. Arguments can be good or bad. When we come to believe a conclusion on the basis of the premises we do so because we have judged that the premises justify or establish the conclusion. But we can be wrong about this. Sometimes arguments that we take to be good, are not. In general terms, good arguments are arguments in which the premises establish their conclusion, whereas bad arguments are those in which the premises do not. Correspondingly, we reason well when our beliefs are based on good arguments and we reason poorly when they are not. Thus, in order to improve our reasoning, we will have to learn how to properly evaluate arguments. Doing so is a two-step process. Most obviously, we will need to learn how to distinguish between good arguments and bad ones. This information is useless, however, if we cannot accurately identify and analyze arguments. Put otherwise, you cannot accurately assess whether an argument is good or bad, if you don’t know what the premises are, and how they are related to the conclusion. Let us take an introductory look at these two steps.

SECTION 6: AN INTRODUCTION TO SPOTTING ARGUMENTS

There are a number of words that authors and speakers use to indicate that they are making an argument. We will call these words *indicator words*, since they typically indicate the presence of an argument. Some words and phrases, like 'since', 'because', 'for', 'on account of', and 'given that...', specifically indicate the presence of a premise.

Ex. 8:

There is no way the Spartans will make the playoffs this year, *since* they are 6 games back with less than two weeks to go.

Ex. 9:

Given the suspect's blood/alcohol level at the time of the accident, it is clear that she was driving over the legal limit.

Other words and phrases specifically indicate the presence of an argument's conclusion: 'thus', 'therefore', 'hence', 'so', 'consequently'.

Ex. 10:

Malik doesn't have any brothers or sisters; *hence* he is an only child.

Ex. 11:

The number 8 is even; *consequently*, it is not a prime number.

We need to keep two qualifications in mind. First, these brief lists include the most common indicator words, but there are many ways that authors and speakers can indicate the presence of an argument without using these terms. Thus, we cannot merely memorize the terms above and be done with it. Second, unfortunately indicator words do not *always* indicate the presence of an argument. Consider the following:

Ex. 12:

The marching band hasn't gotten anything less than a #1 at contest *since* 2015.

In this case, the term 'since' is being used to refer to time, not to a premise. Given that indicator words are not 100% reliable as indications of argumentation, we cannot infer that we have an argument merely because one of our indicator words shows up. As we will see, we need to pay attention to the contexts in which these words are used.

There are a couple of common obstacles to spotting arguments we should note at the outset. The first has to do with conditional claims. As noted, an argument amounts to a set of propositions in which one proposition is supported by the others. Propositions report a fact about the world and are either true or false, e.g. 'Dogs are mammals', 'Yosemite is in California', 'Jordan's mom is a lawyer'. Presumably you get the point, but there is one kind of proposition that people often find confusing: conditionals. What is a conditional? To start off, let's look at some examples:

Ex. 13:

You are legally eligible to purchase alcoholic beverages in this state only if you are 21 years of age or older.

Ex. 14:

If the candle is lit, then oxygen is present.

Ex. 15:

You are not permitted to play poker at the Platinum table if you are not willing to bet at least \$20 per hand.

What do all of these propositions have in common? While we might identify a number of features, the most important for our purposes is that each expresses a relationship between two things. More specifically, each tells us that one thing is dependent *or conditional* in some way on another thing: a burning candle depends on oxygen, playing at the Platinum table is conditional on betting \$20 per hand, and so forth. In very general terms, this is what a conditional is—a proposition that says that one thing is dependent on another (more on this later though). Part of what makes thinking about conditionals difficult is that we are used to thinking about objects and their characteristics, but less so about the *relations* between them. Nevertheless, relations are just as much a part of the world as anything else.

There is a lot to say about conditionals, but for the time being we will focus on two noteworthy features of conditionals in particular. First, as you can see there are many ways that conditionals can be expressed, though it is very common to use the word ‘if’. Second, because conditional claims express a relation between two or more things, they commonly appear in arguments about the things they relate. You might conclude, for example, that you are not permitted to play poker at the Platinum table, since you aren’t willing to bet \$20 per hand, and this is a requirement for playing at that table.

A second obstacle to spotting arguments has to do with opinions. Suppose somebody says “teens really shouldn’t be watching R-rated movies, since for the most part they are not mature enough to handle the psychological and emotional effects of mature content.” Is this an argument? Many people are tempted to say ‘no’—this is just an opinion. When we use ‘opinion’ in this way, we are identifying an idea or claim as particularly controversial, uncertain, or debatable. Understood in this sense the claim that “teens really shouldn’t be watching R-rated movies” is an opinion, whereas something like “The bookshelf weighs 80 lbs.” is not. After all, the weight of the bookshelf should not be controversial or debatable—we can use objective and commonly agreed upon methods to determine its weight.



“Opinionated” by iceman9294 CC BY-NC 2.0

We need to be aware, however, that whether a statement is controversial or debatable is not relevant to whether there is an argument present. Recall that anytime a speaker or author gives a reason to believe a conclusion, they have given an argument—regardless whether anything the author has said or written is controversial, uncertain, or debatable. In the example above, the speaker uses the word ‘since’ to indicate

the presence of a premise, and consequently, the presence of an argument. So: whether somebody's claim is an opinion won't tell us anything about whether they have offered an argument. In fact, often controversial, uncertain, and debatable claims are precisely the sort of thing that people offer arguments for!

SECTION 7: EVALUATING ARGUMENTS

Last, let us turn to the most important topic this book will take up: argument evaluation. To evaluate an argument is to decide whether it is good or bad. We have an intuitive ability to evaluate arguments—we can usually distinguish good arguments from bad ones just by looking at them. This native ability is not, however, infallible; in fact, there are certain contexts and kinds of cases where we tend to make mistakes. Thus it is important to ask: what is the difference between a good argument and a bad one? Put otherwise, we need to know what makes a good argument, good, and a bad argument, bad. Let us start with an example.

Ex. 16:

Premise: The largest city in the U.S. is located in Nebraska.

Premise: New York City is the largest city in the U.S.

Conclusion: So, New York City is located in Nebraska.

Clearly this is a bad argument. The problem is that one of its premises is false—we know that the largest city in the U.S. is not in the state of Nebraska. This example shows one way in which an argument can be bad: when it has false premises. Let us say the following:

An argument is **factually correct** when (and only when) all of its premises are true. It is **factually incorrect** otherwise.

Thus, the argument in Ex. 16 is factually incorrect because not all of its premises are true. In addition, whether an argument is factually correct or not is solely a matter of whether the premises are true—an argument with all true premises but a false conclusion is still factually correct. Factual correctness is not, however, the only feature of an argument relevant to its evaluation. Consider the following case.

Ex. 17:

Premise: Selena passed her driver's license exam.

Conclusion: So, Selena will pass her calculus exam.

Let us say that it is true that Selena passed her driver's license exam. Even so, clearly this is a bad argument. This is a poor argument because the premise does not support the truth of the conclusion. That is, the premise, though true, does not give us good or sufficient reason to believe the conclusion is true. We will refer to this feature of arguments as logical strength and say:

An argument is **logically strong** when (and only when) the premises—if true—provide strong support for the truth of the conclusion. An argument is **logically weak** otherwise.

It is crucial to see that these features of arguments are independent of one another. An argument may be factually incorrect, but logically strong (see, e.g. Ex. 16), factually correct, but logically weak (Ex. 17), both correct and strong, or both incorrect and weak. In light of these distinctions we can distinguish good arguments from bad ones in the following way. Let us say that:

An argument is good, henceforth **sound**, when (and only when) it is both factually correct and logically strong. An argument is **unsound** otherwise.

Ideally, all of our reasoning would proceed by means of sound arguments. However, we are always working with limited information, and this means that we all sometimes endorse unsound arguments. Nevertheless, as we will see in this text, we can take steps to limit these kinds of cases.

EXERCISES

Exercise Set 1A:

Directions: For each of the following passages, determine whether either speaker is giving an argument or whether the speakers are merely having a dispute.

#1:

A: We should go to Drew's place tonight.

B: Um, no. We should avoid Drew's place like the plague.

#2:

A: Obama is our worst president!

B: What are you talking about? Our "worst president"? Do you even know what that means?

#3:

A: Many peer-reviewed studies have shown that asbestos is a carcinogen.

B: These studies can't be right, since my uncle worked in an asbestos mine for 40 years and didn't get cancer.

#4:

A: Minneapolis has about the same population as Indianapolis.

B: No it doesn't, my dad told me that Indianapolis is much bigger.

Exercise Set 1B:

Directions: For each of the following passages, determine whether there is an indicator word(s) present and if so, identify it.

#1:

People who can't speak Russian are excluded from this opportunity, so you can't come.

#2:

The first inhabitants of the island were Dutch settlers. The next wave of settlers were mainly from Italy and Greece.

#3:

There are two main reasons for rejecting this option—first, it is immoral; second, it doesn't achieve what we are actually trying to do!

#4:

Seeing as it has air conditioning, I think you should count yourself lucky to live in that dorm.

Exercise Set 1C:

Directions: Determine whether each of the following arguments is sound or not. If unsound, then say whether it is factually incorrect, logically weak, or both.

#1:

Premise: There are exactly 52 states that make up the United States.

Conclusion: So, if I've only been to 50 states, then I haven't been to them all.

#2:

Premise: All squares are rectangles.

Premise: No rectangles are circles.

Conclusion: So some squares are circles.

#3:

Premise: The word 'since' always indicates the presence of an argument.

Conclusion: So, the claim "*the marching band hasn't gotten anything less than a #1 at contest since 2015*" is an argument.

#4:

Premise: If an argument is sound, then it is logically strong.

Premise: If an argument is logically strong, then the premises—if true—provide strong support for the truth of the conclusion.

Conclusion: So, if an argument is sound, then the premises—if true—provide strong support for the truth of the conclusion.

Exercise Set 1D:

#1:

Our systems of belief change all the time. Give a brief example of a time your system of belief changed.

#2:

What is the difference between reasoning and argument as we are using the terms in this book.

#3:

List 1 true conditional claim. List 1 false conditional claim.

#4:

Describe the experience of playing scrabble, or bananagrams, or some other similar word game. How might automatic reasoning processes be at work in this kind of experience?

Notes

1. Frederick, Shane. (2005) "Cognitive Reflection and Decision Making." *Journal of Economic Perspectives* 19 (4), 25-42.
2. On first glance most people think the answer is that the ball costs \$.10. However, that can't be right. The bat costs \$1.00 more than the ball, so if the ball costs \$.10, the bat costs \$1.10, and together they cost \$1.20.

CHAPTER 2

Reasoning as a Social Process

SECTION 1: INTRODUCTION

As we saw in the last chapter, reasoning is a process we use to update our system of beliefs. This process is often, in addition, a social one. In fact, other people are probably the most important influence on our thinking. For example, we learn a lot from what we are told, or read in the news, or see online. Moreover, we seek the advice and guidance of others, we bounce our ideas off of other people to see if they make sense, and we work together to solve problems or answer questions. We are constantly looking to, and engaging with, other people in order to achieve our goals and follow our interests. In light of this, we will focus on a particularly important way that we engage with one another that we will call *cooperative dialogue*. Cooperative dialogue is a fundamental form of communication that is especially conducive to learning from others. It turns out, however, that in contexts of disagreement or controversy we are liable to unnecessarily abandon cooperative dialogue. In this chapter, we take a look at this form of communication, explain why we tend to abandon it, and offer some guidelines for sustaining cooperation through difficult social contexts.

SECTION 2: COOPERATIVE DIALOGUE

Cooperative dialogue, as we will understand it, is a form of communication between two or more people that has at its heart a shared effort to learn, understand, figure out, or decide something. In order to illustrate the idea, let's start with a simple case of communication.

Ex. 1: You are in a coffee shop, and are waiting for an important email. You'd like to access the shop's wifi, and ask for the password. The employee gives it to you.

This everyday interaction is a cooperative dialogue—though a brief one. What makes this a cooperative dialogue is that you made a request for information and the employee cooperated with that request. Had the employee ignored you, responded sarcastically, or lied to you this would not have been a case of cooperative dialogue.

The inquiry in Ex. 1 has a clear and straightforward answer which employee knows and can easily share. In many instances of cooperative dialogue, however, this is not the case.

Ex. 2: You are a college student trying to decide whether you should major in business and minor in art or to double major. You can see costs and benefits of both options, and go to talk to your academic advisor about it. You sit down with your advisor and she raises a number of relevant issues and shares her experience with you.

This is a complicated decision, and what you are looking for in this case is a deeper understanding of the options, not for your advisor to tell you what to do. The cooperative dialogue in this case is a sustained conversation with the shared goal of an informed decision.

In the examples above, one person is looking directly to another for answers or guidance. However, we often cooperatively engage with other people who do not necessarily know or understand what we are inquiring about, but are either themselves interested or are willing to talk with us about the issue. That is, in many cases of cooperative dialogue, nobody involved can authoritatively answer the question at hand. Consider the following example:

Ex. 3:

Sofia: Did you hear that Ian and Isabel broke up?

Alex: No. What happened?

Sofia: I don't know, but I'd bet Isabel broke it off; I mean she never seemed all that into Ian.

Alex: Maybe you are right, I saw her ignore his phone calls a couple of times last week. But on the other hand, she was just telling me the other day about her plans to throw him a surprise party for his birthday next week.

Sofia: Hmm...maybe something happened between them recently.

In this case, Alex and Sofia are engaging in a shared inquiry, and each person is bringing what they know to the conversation in an effort to answer the question. Again, this is an everyday example, but the same will go for more significant or controversial kinds of inquiry. Imagine Alex and Sofia raising questions about whether Grandpa needs to move into assisted-living, what to think about a local ballot measure, or whether to take the job-offer. In each case, Alex and Sofia will cooperatively engage with one another to the extent that they bring their ideas to together in order to make progress with the inquiry.

The examples of cooperative dialogue above are all in-person conversations, but we can cooperatively engage with people by text, email, or on social media as well. In addition, it is important to consider one other kind of cooperative dialogue. Consider a typical reading assignment for a college course. In the kinds of texts normally assigned, the author is trying to share what they know, and thereby help the reader understand something in particular. A cooperative dialogue with the author can occur in this context if, in reading the assignment, *we share in the effort to understand*. Indeed, this is exactly what your teacher is hoping for in giving you the assignment in the first place! That is, your teacher is hoping that you'll use the text as a source of information for deeper understanding, or treat the text as an interlocutor in a shared inquiry.

Admittedly, a cooperative dialogue with a text is very different from an in-person conversation; after all, the author can't answer your questions or objections. Nevertheless, a reader can engage with a text by seeking to understand, evaluate, and question the author's claims just as we could in conversation. As long as the reader's effort is to better understand or decide something, the conditions for cooperative dialogue have been met.

SECTION 3: COOPERATIVE DISAGREEMENT

Let us look at another kind of cooperative dialogue: cooperative disagreement. This might strike you as an oxymoron: how can two people cooperate *and* disagree at the same time? In order to clarify, consider the fact that disagreements raise all kinds of questions. For example:

“Why don’t we agree? I believe this is true; how could somebody deny this? What are my reasons for thinking this in the first place? Are those reasons really as good as I think they are? Do we disagree because I know something he doesn’t? Wait...what if he knows something I don’t!?”

These are valuable questions because answering them can (i) reveal that we’ve missed something or otherwise made a mistake, and (ii) thereby put us in a position to correct ourselves. That is, we can learn a lot from thinking about disagreement, and often the best way to do this is to engage with those who disagree (as opposed to, say, speculating about their reasons). Thus, in a **cooperative disagreement** disputants work together for the sake of understanding the disagreement, and ultimately for the sake of deeper understanding of the disputed issue overall. This means, among other things, that the disputants in a cooperative disagreement make the effort to articulate why they disagree, and then discuss those reasons. Not all cases of disagreement are instances of cooperative dialogue; in fact, many are not. As we will see, cooperative disagreements can be difficult to sustain. This is too bad, because, again, disagreement is an especially good opportunity for increased understanding and better decision making. Let’s take a look at an example.

Ex. 4:

Liam: The new policy of keeping the doors to the dorm locked at all times is totally ridiculous. It is a real pain... Yesterday, I forgot my keys and had to call a friend to let me in, and I ended up being late for class. I mean, it is not like there have been many problems with non-students getting into the dorm anyway.

Minh: I think this policy is great, and I will tell you why. According to my RA there have been at least 5 incidents already this semester. Keeping the doors locked makes me feel safer. Also, the dorm is not a public building and it shouldn’t be open to the public.

In this case, Minh not only disagrees with Liam, but is willing to seek a cooperative dialogue with him. She explains why she disagrees, and it turns out that Minh knows something that Liam does not. Does the fact that there have been more incidents than Liam knew about change his view? It might, but it surely wouldn’t have if Minh hadn’t been willing to engage with Liam in this way.

SECTION 4: COOPERATIVE PERSUASION

We often engage with each other in order to persuade. Our aim in persuasive contexts is to get someone else to believe or act in a particular way. So, what is the relationship between cooperative dialogue and persuasion? Suppose that a friend is trying to persuade you to vote for a particular candidate. Can they do so cooperatively? Sure. Again, the heart of cooperative dialogue has to do with a shared effort to figure out or decide something, and doing so is consistent with also trying to persuade another person. Suppose that your friend thinks that Sandy Berners is the best candidate in an upcoming election, and has a clear set of reasons for this belief. If, in sharing those reasons with you she is honestly trying to give you what she regards as good reasons for supporting this candidate, then she really is trying to help you understand and decide. If, however, her goal is to get you to support Sandy Berners regardless of whether you do so for what she would regard as good reasons, then she isn’t cooperating with you, she is trying to manipulate you.

Here is another case: the honest realtor. Imagine that you are shopping for a new house and have talked with a realtor about your budget and what you are looking for. Taking your interests, budget, and available inventory into account, he identifies several houses you might like. If, however, the realtor’s primary goal is to sell a particular house and tries to talk you into it knowing full well that it isn’t what you are looking for,

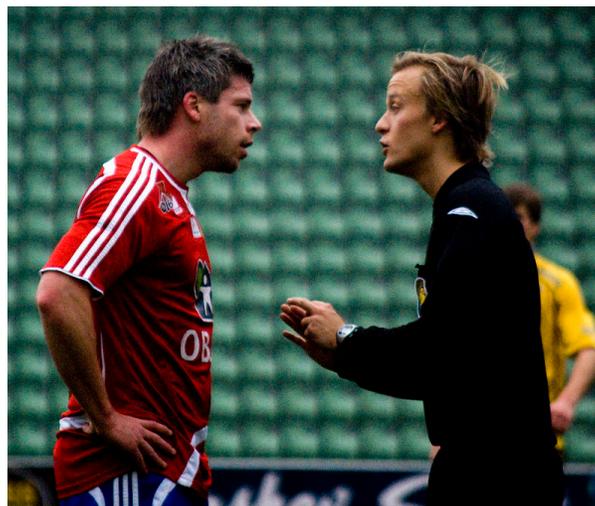
then this dialogue is not a cooperative one. After all, in this case, he is not trying to help you make a good decision, but is trying to help himself by getting you to make a decision that is in his interest.

These examples show that there are cases of **cooperative persuasion**. Again, not all cases of persuasion are cooperative dialogues. The difference has to do with whether their goal is to persuade you by means of helping you to understand something in order to make a good decision, or merely to get you to believe or act in accordance with their wishes.

SECTION 5: OBSTACLES TO COOPERATIVE DIALOGUE

Cooperative dialogue is an easy and natural form of communication. When other people ask us questions, we typically do our best to help them out if we can, and are usually happy to weigh-in with our view when an interesting question or issue is raised by friends, relatives, or colleagues. In addition to being easy and natural, cooperative dialogue is also extremely useful. After all, through cooperative dialogue we are able to bring the knowledge and experience of multiple people to bear on understanding, figuring out, or deciding something.

However, we do not always seek cooperative dialogue—even when we could and when it would be in our interest to do so. Why not? The short answer is that we have competing interests and goals besides improved understanding or decision-making. Think for example, of the honest realtor who really needs a sale. As his need for a sale grows, it can become more and more difficult for him to engage with homebuyers in a cooperative way. It is not difficult to think of other cases like this, but it is important to note that not all competing interests are financial. Another significant obstacle to cooperative dialogue, and the one we will focus on here, has to do with the social value of ideas. We care what other people think about us, and this can affect how we decide to engage with them. We want other people to think highly of us, or at least not think poorly of us, and so we sometimes avoid disagreeing with others or even raising topics that might be controversial. The worry is that if we disagree or express an alternative viewpoint others will think less of us. As a result, we often “go along to get along” instead of contributing to the conversation. This worry is enshrined in the following piece of advice: in polite conversation its best to avoid talking about money, politics, or religion.



“JoA in an argument” by Anders V CC BY-NC-SA 2.0

Let's dig a little deeper into this, and start by focusing on the claim that getting along with others means avoiding conversations about money, religion, and politics. Lying behind this piece of advice is the fact that people often have conflicting and emotionally charged views about these topics. Our beliefs about these subjects can be especially important to us because they tend to be tied to how we see ourselves, and how we see our place in the broader community. Let us call these kinds of beliefs, **identity-beliefs**. Identity-beliefs need not be tied exclusively to our views about money, or politics, or religion. They might be tied to facts about where we live (e.g. "I'm a New Yorker through and through), what we do for a job (e.g. physician, farmer, police officer), the kinds of food we eat (e.g. vegetarian), or even to the kinds of products we purchase (e.g. "I don't buy PCs; I am an Apple-person"), among many others. Because they are tied to our self-image, a criticism, objection, or challenge to one of our identity-beliefs can feel like a criticism, objection, or challenge, to who we are, and a disagreement about one of these beliefs can become a defense of one's character or identity.

There is another side to this as well. We have a natural tendency to tie a person's ideas or beliefs to who they are as a person, and a consequence of this tendency is a predisposition to vilify those who object to, or challenge, our identity-beliefs. We vilify another person when we take the fact that they disagree with us as evidence that there is something wrong with *them*—normally in terms of some defect in their intelligence or character.¹ Thus, we might think that people who disagree with us are ignorant, naïve, or malicious, in which case we'll think there is nothing to learn from them and no reason to talk with them in the first place. We are left with the ironic fact that we tend to do to others precisely what we fear they will do to us!

In light of all this, it isn't surprising that we often avoid disagreement and controversial topics. After all, nobody likes to be criticized or be called on to defend who they are, especially if they think there is no point to engaging anyway (since they've judged that the only reason that they disagree is because they are ignorant or mean). Indeed, in many contexts it makes perfect sense to "go along to get along." Not always, however. Money, politics, and religion, while divisive, are also important elements of our world—elements that we need to deal with as we navigate our daily lives. We *should* be thinking about these issues, and considering other's opinions can help us see things we've missed and refine our own views. But how do we do that? That is, how can we cooperatively dialogue with others when controversial topics are on the table or when people have strong disagreements?

SECTION 6: GUIDELINES FOR ENGAGING IN AND SUSTAINING COOPERATIVE DIALOGUE

Cooperative dialogue is useful to us, so let us briefly outline some guidelines for engaging in and sustaining cooperative dialogue. As noted above, cooperative dialogue is particularly difficult in cases of disagreement or possible disagreement (for example, conversations about controversial issues). So, we will limit our discussion of guidelines that are especially relevant in these contexts.

Guideline #1: Respect your interlocutor

Cooperation is difficult, if not impossible, without some degree of mutual respect, and this is probably the single most important guideline. To respect your interlocutor is to treat them as you would like to be treated in the same circumstance. Put in different terms, Guideline #1 urges us to follow what we might call *The Golden Rule of Cooperative Dialogue*. In a cooperative dialogue, and especially in a disagreement, we want our interlocutor to listen intently, to make an effort to understand what we are saying, and to interpret us in a fair and plausible way (among other things). Consequently, we should make the effort to engage with our interlocutors similarly, to do our best to charitably interpret their viewpoint, and to seek clarification when

what is said is ambiguous or sounds implausible to you. It is very important to note that to respect somebody in this way does not mean that you have to like them or agree with them. In this way, cooperative dialogue is like a lot of other cooperative activities. A person does not have to be friends with or share the political or religious beliefs with their teammates in an athletic competition, for example. Successful teams can put aside differences that are not relevant to achieving their goals, and the same goes for sustaining cooperative dialogue through disagreement.

Guideline #2: Remember that we are all fallible

We all make mistakes. Our experience is limited, we are almost always working with incomplete information, and we can be inattentive to or misremember relevant facts. Moreover, it can be easy to misunderstand or misinterpret others, we are prone to bias, and sometimes we simply reason poorly. Further, our ability to discuss ideas with one another is limited: some subjects or topics are difficult to talk about, and sometimes it is difficult to isolate why we believe what we believe. The fact that we are fallible in these ways doesn't mean that some people aren't more or less knowledgeable than others, and it doesn't mean that we can't have strongly held beliefs. It does mean, however, that even decent, well-meaning, and intelligent people can end up being wrong, or making poor decisions, or having a hard time discussing certain ideas. It also means that people with conflicting beliefs can both be reasonable in their beliefs. Consider the following case.

Ex. 5:

Carlos: You want to get some lunch? The cafeteria is still open.

Nick: I would, but, uhh, it is closed. It is 1:45 and it closes at 1:30.

Carlos: Right, but my roommate works over there, and told me that since there is a special event this afternoon, the cafeteria would be open an extra hour.

Nick thinks that Carlos is mistaken in thinking that they can still get some lunch, and Nick has good reason—namely that it is almost always closed at this time of day. Nevertheless, in this case it is Nick who is mistaken because he is missing relevant information, information that Carlos has. Both Carlos and Nick make reasonable inferences, but since they are reasoning from different information, they end up with conflicting beliefs. This is a simple case, and it important to emphasize that people can have widely divergent histories, experiences, and background knowledge, and consequently, widely divergent beliefs and ways of understanding the world. That is, controversy and disagreement are a natural consequence of the inherent limitations of our knowledge and experience. In fact, it is precisely because we are fallible in these ways that cooperative dialogue is so important! By drawing on other people's thinking and experience and comparing it with our own we can fill in some of the gaps in our knowledge and experience, check our reasoning, and more broadly correct for our intellectual liabilities.

Guideline #3: Separate people from their ideas, and focus the discussion on the reasons behind disputed ideas

The chief obstacle to cooperative dialogue noted above lies in the fact that we have a natural tendency to connect a person's beliefs or ideas to their character. This is what leads us to vilify those who disagree with us and, at the same time, to fear others will vilify us if we disagree. According to Guideline #3, we should focus our attention on the ideas and arguments in question, and not a person's motives, background, or character. Here is an example:

Ex. 6:

Jacob: I saw that the school decided to fund a new study-abroad scholarship instead of refurbishing the stands for the soccer fields.

Anna: What a waste. The school shouldn't have to sponsor study-abroad. Anybody can study abroad if they really want to. I mean, if you really want it, then you'll find a way.

Jacob: Are you serious? You only think that because you are rich!

Jacob clearly disagrees with Anna, but instead of objecting to her reason for thinking that funding the study-abroad scholarship is wasteful, he makes the objection about her. This puts Anna in a position to be defensive, and it wouldn't be surprising if she decided to disengage from this conversation. It is important to see, however, that this would be a missed opportunity. After all, from Jacob's perspective Anna is reasoning on the basis of a mistake, and pointing this out is an opportunity to improve Anna's understanding and perhaps change her mind.

It is important to be clear that who we are, the experiences we've had, and where we've grown up are factors that shape the ways that we think about the world. Maybe Anna really does believe this because she comes from a wealthy family, and so doesn't understand the struggles of those less well-off. Nevertheless, if we want to cooperatively engage with other people we should, in most cases, set aside these kinds of considerations for the sake of sustaining the dialogue. The point is not to set them aside because they are irrelevant, but because drawing on them easily transforms a conversation about ideas into a challenge to who a person is. In this case, Jacob misses out on the opportunity to share his reasons for disagreeing with Anna—reasons that might persuade Anna to change her mind. This is not to say that in cooperative dialogue we can never talk about a person's motives or experience or character. We can; but it can be difficult to do so without creating an atmosphere of defensiveness that undermines the dialogue.

This guideline applies to criticisms or objections leveled at our beliefs as well. If you are trying to engage cooperatively, then you need to be open to hearing criticism without taking it personally. When we find ourselves taking something personally, and getting irritated, insulted, or offended, we need to give ourselves a moment to think about whether the criticism really is personal or not. We miss out if we break off an otherwise promising conversation because we've misinterpreted our interlocutor's criticism. Of course, sometimes it really is personal, and we may want to break off the conversation. It is worth, adding, however, that even if the criticism is personal, it might be worth trying to continue to cooperatively engage. After all, some conversations and disagreements are really important to get out into the open, and by refusing to take criticism personally, we can show our interlocutor that we are serious about having a cooperative dialogue.

Guideline #4: Choose your words carefully

Separating people from their ideas means more than just setting aside our interlocutor's motives, background, or character, it also means being careful about how we express our disagreement. Because people have a natural tendency to hear criticism as a personal attack, we should try to communicate and reinforce that our disagreement is only at the level of ideas in the way that we speak. There are all kinds of ways that you might do this. Interpersonal dynamics are complicated, however, and there is no rulebook or simple formula here. To give just one example of the complications involved, think about the way you can express a disagreement with your best friend as compared to the way you can express it with a parent or teacher or boss. That said, here are a few concrete suggestions to illustrate some of the techniques you might use.

First, and perhaps most obviously, in expressing your disagreement you can explicitly direct it toward an idea. There can be a big difference in the mind of your interlocutor between hearing “*I disagree with you*”, “*I disagree with your claim that...*”, and “*I disagree with the claim that...*”. The first example suggests the disagreement is about the person. The second is less suggestive, but still identifies the claim as ‘yours’ thereby leaving open a connection between the disagreement and the person. The third claim, in contrast, connects the disagreement to an idea—an idea that could have been expressed by anybody.

Second, try to frame your disagreement within a broader context of agreement. By contextualizing your disagreement in this way you not only clarify the locus of disagreement, but show that you and your interlocutor have common beliefs or values. Thus, you might add: “*I think you are right that X is a serious problem, but I disagree with the claim that it is the most pressing problem we have right now.*” Third, if you sincerely think your interlocutor has done something well, say so. By identifying things that your interlocutor has done well, you show that you’ve been listening, and are being fair-minded in your engagement with them. To continue with the example, you might say: “*You’ve made a good case, and I think you are right that X is a serious problem, but...*” Admittedly, these are small changes, and they might not be appropriate for every situation, but they are changes that can make a difference.

Guideline #5: Don’t expect closure

It is important to keep your expectations for the conversation in line with its cooperative character. A good cooperative disagreement can leave the dispute unsettled. Indeed, what makes cooperative disagreement successful is not that it ends with agreement, but that the parties involved have a better sense for the roots of the disagreement, and have new perspectives and arguments to consider. This can be difficult to keep in mind since, as we have seen, we have values and goals that can compete with our interest in increased understanding, and it can be easy to shift from one goal to another without noticing it. As soon as our conversational goal becomes looking smart, or saving face, or getting the other person to agree or capitulate for example, we have slipped from cooperation to some other kind of dialogue.



“The Purpose of Argument” by *ImNotQuiteJack* CC BY-SA 2.0

SECTION 7: FINAL THOUGHTS

Unfortunately, substantive cooperative dialogue and disagreement can be hard to find. After all, it is not solely up to us—it takes two to cooperate, and sometimes we are surrounded by people who, for whatever reason, are not open to cooperating with us. What we can do, however, is work to foster a social environ-

ment that allows for cooperative dialogue and disagreement. We can do so by practicing *The Golden Rule of Cooperative Dialogue* in our interactions with others. By treating others in the way we'd like to be treated, we implicitly invite them to engage with us cooperatively.

It is important to take up one final question, namely: when should we seek cooperative dialogue? There is no one-size-fits-all answer to this question. It will depend on the issue at hand, the situation, and the people involved. There are certainly cases in which an individual is obligated to seek cooperative dialogue. Think of a general creating a battle plan. Lives are on the line, and the general should draw on all the resources he or she can to come up with the best plan. Of course, few cases are as serious as this one. Moreover, there are many cases in which it doesn't make sense to seek cooperative dialogue at all: trying to cooperatively engage with a bunch of internet trolls is a fool's errand, for example. More broadly speaking, if you've got good reasons for thinking that a person is not, or will not engage with you sincerely, then you've got good reason to break off or refrain from cooperative efforts. Many situations fall between these two extremes, however, and thereby provide opportunities for cooperative dialogue.

Where does this leave us? Cooperative dialogue can be a particularly effective way of learning about our world, and making better decisions. We do not always pursue cooperative dialogue, however, and it can be especially difficult to maintain in cases of disagreement. This is due, in large part, to the fact that we care what other people think about us, and fear that they will think less of us as people *because* we disagree. As we've seen there are a number of techniques we can use to combat our natural inclination to connect a person's ideas to their character, and to thereby sustain cooperation through disagreement and foster an environment that allows for greater cooperative engagement.

EXERCISES

#1: Turn on CNN, MSNBC or Fox News and watch a few conversations. In a paragraph or two describe at least one of the conversations you watched and say whether it was a cooperative dialogue or not. Make sure to explain your thinking.

#2: Is cooperative dialogue or disagreement different on social media? If so, how? More generally, is cooperative dialogue or disagreement different or more complicated when the dialogue is public and other people can see it? If so, how?

#3: Explain, in your own words, why cooperative disagreement is valuable.

#4: Think back to some of the difficult conversations you've had (or tried to have). In general terms, what allowed the conversation to proceed? If it didn't work, why not?

#5: In Ex. 6 Jacob shuts down the conversation by making his objection personal. What could Jacob have said in reply to Anna instead?

#6: According to Guideline #4, we need to choose our words carefully if we want to sustain a cooperative dialogue in a context of disagreement or controversy. In addition to the practical suggestions listed above, what are some other specific phrases or techniques you can use to clearly separate people and ideas? Give at least three examples.

Notes

1. See Ichheiser, G. (1949). "Misunderstandings in human relations: A study in false social perception." *American Journal of Sociology* 55 (Suppl.), 39. More recent studies trace this phenomenon to our natural tendency to think that we see, understand, and interpret the world as it really is. Given this, we reason that other intelligent people should see the world as we do, and when they don't, it is their fault. For a summary see Ross, L., Lepper, M. and Ward, A. (2010). "History of Social Psychology: Insights, Challenges, and Contributions to Theory and Application." *Handbook of Social Psychology*. Hoboken, NJ: Wiley.

Unit #1 Summary

We reason all the time, and are pretty good at it. That said, we can improve our skills, and in chapter 1 we began to develop a vocabulary for talking about reasoning and argument more clearly. In addition, we introduced two key skills: argument identification and evaluation. If we are going to talk about reasoning and argument, then we will need to know how to recognize cases of reasoning in our own thinking, speaking, and writing, and in the speaking and writing of others. Moreover, we want to reason well, and we closed the chapter by introducing basic standards of argument evaluation. Reasoning is often a social process, and other people are one of the most important influences on our thinking. In chapter 2 we took a look at the communicative context that is, in general, most conducive to learning from other people—cooperative dialogue. We looked at a number of obstacles to cooperative dialogue, and ended the chapter with some guidelines for sustaining cooperative dialogue through difficult social contexts.

KEY TERMS

- Reasoning
- System of Belief
- Justify
- Automatic Reasoning
- Argument
- Proposition
- Premise
- Conclusion
- Indicator Words
- Opinion
- Logical Strength
- Factual Correctness
- Soundness
- Cooperative Dialogue
- Cooperative Persuasion
- Cooperative Disagreement
- Identity-beliefs
- 5 Guidelines

FOR FURTHER READING

For an accessible introduction to automatic reasoning processes see Malcom Gladwell's *Blink*. For a readable but more sophisticated account see Daniel Kahneman's *Thinking Fast and Slow*. For a critical, persuasive, and well-written approach to the psychology of reasoning see *The Enigma of Reason* by Hugo Mercier and Dan Sperber. For additional practical advice about sustaining conversation see *Crucial Conversations* by Patterson, Grenny, McMillan, and Switzler.

UNIT II

ARGUMENT ANALYSIS

Analysis, Standardization, and Diagramming

SECTION 1: INTRODUCTION

In this unit, we will focus on argument identification and analysis. Argument identification involves spotting arguments and distinguishing them from their surroundings. We briefly discussed argument identification in Chapter 1, and saw that we need to be on the lookout for particular words (e.g. 'since', 'so', 'because', and 'hence') that frequently indicate the presence of an argument. The other skill we will focus on in this unit is argument analysis. In general, to **analyze** an argument is to determine its intended structure.

As we've seen, the crucial elements or "bones" of any argument are its premises and its conclusion, and we determine an argument's structure by identifying these elements and figuring out how they fit together. We are beginning with the skills of identification and analysis in particular because recognizing and accurately determining argumentative structure are prerequisites for accurate evaluation. After all, we can't assess whether an argument is good or bad until we know what the argument is supposed to be! More precisely, we can't say whether an argument is sound or not, unless we know what its conclusion is, and what premises have been offered on its behalf.

Before we get started, it is important to make two notes. First, you are already skilled at argument identification and analysis. Again, these are prerequisites for evaluation, and you have been evaluating arguments almost your whole life. However, these skills are largely intuitive. The goal here, then, is not to develop new skills, but instead to sharpen or enhance existing ones. To this end, we will slow down and think through the steps involved in identification and analysis. By sensitizing ourselves to argumentative language in this way, we will make what is mostly implicit in everyday practice, explicit. More specifically, this will involve giving names to a variety of argumentative moves and situations, identifying a method for seeing behind an author's words, and learning some techniques for representing arguments. Proceeding in this way will allow us to see in detail how arguments are expressed and structured, and will ultimately put us in a better position to evaluate and engage with arguments.

The second note has to do with reading. We will be focusing on arguments in text, and one of the things that we will find is that in all but the simplest cases, argument analysis requires active engagement with the text. We actively engage with a text when we ask questions of it: *what is the author trying to say? What reasons does the author give for the conclusion? Why do they believe the premises are true?* We do not always read actively. When we read a magazine or a novel, we normally sit back and let the author provide us with information or tell us a story. However, when it comes to argument analysis passive reading will not do. In most cases, argument analysis will involve actively looking to the text for clues, and using these clues to

piece together the author's intended argument. Active engagement with a text can be demanding—especially at first, but it is an important skill, and something you will likely find becomes easier with practice.

SECTION 2: ANALYZING SIMPLE ARGUMENTS

Let's start with a simple argument.

Ex. 1:

Since *Banana Republic's* sale items will go fast, I think we should go there first.

The first thing we should notice in this case is the word 'since'. This word typically indicates the presence of an argument, and a look back at the context confirms this. Now that we have an argument, how do we analyze it? We need to break the argument into its components, i.e. its premises and conclusion. Because the word 'since' is not only an argument indicator, but a premise indicator, we know that what follows it—'*Banana Republic's* sale items will go fast'—is being used as a premise. This premise supports the proposition 'we should go there first' which is the conclusion. We have captured all there is to say about this argument's structure, and so our analysis is complete.

Importantly, what makes this argument simple, as we will use the term, is not that it is short or easy to understand. Instead, this argument is **simple** because it draws only one conclusion. As we will see later in the chapter, simple arguments can be chained together to form complex arguments. For now, however, let us take a look at another simple argument.

Ex. 2:

The coach will likely be fired, given the allegations of unethical recruiting practices and the team's sub-par performance the last three years.

The term 'given' and a quick look at the context tells us that we have an argument on our hands. 'Given' is a premise indicator, and in what immediately follows the author makes two claims. The author is claiming that there are allegations of unethical recruiting practices *and* that the team's performance over the last three years has been sub-par. The author is offering these two premises as evidence for the proposition that the coach will likely be fired, and so this latter proposition is the conclusion.



"Coaches Shout Instructions" by Chris Hunkeler CC BY-SA 2.0

This second example is similar to the first; however, we should note two important differences. First, in this example there are two premises, whereas in Ex. 1 there was only one. This does not change the fact that

like Ex. 1, Ex. 2 is a simple argument. Again, what makes an argument simple is that it draws only one conclusion, and this means that simple arguments can have many premises. The second point has to do with order. In Ex. 1 the premise came first and the conclusion last, whereas in Ex. 2 the order is reversed. This teaches us an important lesson: *premises are not always presented first, nor are conclusions always presented last*. As it turns out, in everyday speech and writing, arguments are presented in many different ways. Sometimes, in fact, the conclusion is placed between premises (as we will see).

SECTION 3: REPRESENTING ARGUMENTATIVE STRUCTURE

The point of analyzing an argument is to uncover its structure, and it will be useful to have a uniform or standard way of representing the bones of an argument. For simple arguments we will use the following process of representation. For each distinct part of an argument (each of its premises and its conclusion), we will assign a unique number, assigning the highest number to the conclusion. We will then stack the propositions in numerical order, and add a conclusion indicator to the conclusion for clarity. We will call this way of representing an argument, a **standardization**.

In Ex.1, the argument consists in two propositions: the premise and conclusion, and so in our standardization we will stack the numbers 1) and 2). We will assign 2) to the conclusion, and then add the conclusion indicator 'so' to the argument. This gives us the following standardization of Ex. 1:

1. *Banana Republic's* sale items will go fast.
2. **So**, we should go there first.

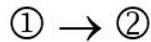
The argument in Ex. 2, consists of three propositions: two premises and the conclusion. As a consequence, our standardization will stack the numbers 1, 2, and 3, we will make 3 the conclusion, and add the indicator word 'so' to it. This gives us the following standardization for Ex. 2:

1. There have been allegations of unethical recruiting practices.
2. The team's performance over the last three years has been sub-par.
3. **So**, the coach will likely be fired.

Standardizing arguments is a useful way of representing an argument's structure, but it is not the only way. A different way of representing the bones of an argument is called **diagramming**. Diagramming is a more abstract way of representing arguments that allows us to see an argument's structure independently of its subject matter.

Just as we would in a standardization, we start a diagram by assigning each part of an argument a number. We start numbering at one and reserve the highest number for the conclusion. This is where the similarity ends, however. First, when diagramming, we let the assigned number stand in for the whole proposition; we do not rewrite the premises and conclusions as we do in standardizations. Second, when we diagram we will use numbers with circles around them to indicate propositions. Third, we will not stack the numbers, but work horizontally to capture the relation between the premise and conclusion. Doing so, requires the use of two additional symbols: the arrow ('→') represents the relationship between an argument's premises and its conclusion. The arrow points away from the premises, and toward the conclusion the premises purport to establish. Further we will use the plus ('+'). The + represents the idea that an author intends two or more distinct propositions to be taken together as evidence.

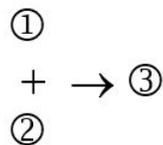
The diagram of Ex. 1 looks like this:



Thus, 1 represents the proposition '*Banana Republic's* sale items will go fast', while 2 represents the conclusion that 'we should go there first'. Moreover, the arrow points from 1 to 2 because 1 is the evidence that purports to establish 2.

In the case of Ex. 2, the author offers two pieces of evidence on behalf of his conclusion, and this is reflected here by the conjunction of 1 and 2 (which represent the propositions that there are allegations of unethical recruiting practices *and* that the team's performance over the last three years has been sub-par).

Diagram of Ex. 2:



It is important to note that standardization and diagramming are distinct ways of representing argumentative structure, and we can use one without the other. That is, we might standardize an argument without diagramming it, and vice versa (though if you choose to diagram an argument without also standardizing it, you'll have to find a way connect your numbers to the propositions they represent). In this unit, however, we will always standardize and diagram examples, and we will use the numbers assigned in the standardization of an argument for the diagram. Now that we have discussed simple arguments, let's see how analysis works in more complex cases.

SECTION 4: ANALYZING COMPLEX ARGUMENTS

A simple argument draws only one conclusion. However, simple arguments can be put together to create complex arguments. As we will see, **complex arguments** draw one or more sub-conclusions on the way to the main conclusion. Consider the following case:

Ex. 3:

You should fill your car's tires with nitrogen instead of plain air for two reasons. First, nitrogen will diffuse through the tire walls much more slowly than plain air, since nitrogen molecules are bigger than molecules of oxygen. Second, filling your tires with nitrogen keeps water vapor from getting inside your tires.

Uncovering the structure of this argument means isolating all of the parts and determining their relationships. We will begin by using indicator words as our guide. On a first pass, we should be struck by the presence of two indicators: 'two reasons' and 'since'. The author is claiming that there are two reasons for thinking that "You should fill your car's tires with nitrogen instead of plain air." The author has numbered these premises using the term 'first' and 'second'. Thus, these two propositions are premises that support the conclusion that you should put nitrogen in your tires. However, there is one other indicator word here—'since'. This suggests that there is a second justification present. Indeed, the fact that "nitrogen molecules are bigger than molecules of oxygen" is given as a reason to believe that "nitrogen will diffuse through the tire walls more slowly than plain air will." This means the author's claim that "nitrogen will diffuse..." is being used as both a premise and a conclusion. On the one hand, it is a premise because it supports the

proposition that you should fill your tires with nitrogen. On the other, it is a conclusion because it is supported by the proposition that nitrogen molecules are bigger than oxygen molecules.



"Low Pressure Car Tyres" by David Blaikie CC BY 2.0

Our analysis is complete; we have uncovered all the parts of the argument and we know how they are related. In order to represent this argument's structure let's standardize it. As we learned above, we should assign a number to each relevant proposition reserving the highest number for the conclusion. Wait. There are *two* conclusions in this argument. Which one should we assign the highest number? We will reserve the highest number for the ultimate or main conclusion, which in this case is that "You should fill your car's tires with nitrogen instead of plain air." Can we simply assign numbers to the other propositions, and standardize the argument like this:

Standardization for Ex. 3?

1. Nitrogen will diffuse through the tire walls much more slowly than plain air.
2. Nitrogen molecules are bigger than molecules of oxygen.
3. Filling your tires with nitrogen keeps water vapor from getting inside your tires.
4. So, you should fill your car's tires with nitrogen instead of plain air.

No. While this standardization captures all the relevant propositions, it misses an important part of the argument's structure. Although 1-3 all ultimately support the conclusion, they do not do so in the same way. The proposition that nitrogen molecules are bigger than molecules of air only supports the conclusion because it gives evidence for the proposition that nitrogen will diffuse more slowly than plain air. In other words, 2 only supports 4 *through* its support of 1; let's say that 2 offers only **indirect support** of 4 whereas 1 and 3 offer **direct support**, and our representation of the argument needs to reflect this. This means that we'll need to supplement our basic standardization process.

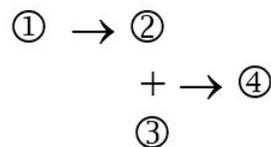
First, let's follow the convention that conclusions always come after their premises, so we'll want to assign the proposition 'nitrogen molecules are bigger than molecules of oxygen' a higher number than the proposition 'nitrogen will diffuse through the tire walls much more slowly than plain air'. Second, since the proposition 'nitrogen will diffuse...' is a conclusion, we should make it clear by adding an indicator word to our standardization. Last, when there are multiple inferences in an argument we need to know for sure what premises lead to what conclusion. To mark this, let us agree that after every conclusion we will note the premises from which the proposition is drawn. Following these additional rules gives us the following standardization:

Standardization for Ex. 3:

1. Nitrogen molecules are bigger than molecules of oxygen.
2. **So**, nitrogen will diffuse through the tire walls much more slowly than plain air. **(from 1)**
3. Filling your tires with nitrogen keeps water vapor from getting inside your tires.
4. **So**, you should fill your car's tires with nitrogen instead of plain air. **(from 2 and 3)**

What does a diagram of Ex. 3 look like? We begin our diagrams with the conclusion. Following the number system we used in the correct standardization of Ex. 3, our main conclusion is 4, so we should start by drawing the number four with a circle around it. Evidence is offered on behalf of this conclusion, so we should draw an arrow to the left of our circled number pointing to it. What evidence is offered on behalf of 4? The propositions numbered 2) and 3) above are the direct evidence for 4), and we should connect these pieces of evidence using a plus since it is clear that the author intends these pieces of evidence to be taken together. Last, as we've seen, the author offers 1) as evidence for 2), so we should draw an arrow to the left of 2) pointing to it, and 1) to the left of that. This gives us the following diagram:

Diagram for Ex. 3:



The diagram of this argument shows very clearly that this complex argument is built out of two simple arguments. Working backwards from the main conclusion, there is the simple argument with 2 and 3 as premises and 4 as the conclusion. In addition, because the author gives a reason for 2, we have another simple argument from 1 to 2. Let us turn to some exercises.

EXERCISES

Exercise Set 3A:

Directions: For each of the following determine whether the passage contains an argument. If it does not, write "no argument". If it does, then standardize and diagram the argument.

#1:

It seems likely that this year will be Morita's first as a professional without a major win on account of continuing problems with her short game.

#2:

Your health care provider will not cover this test on the grounds that it is neither medically necessary nor an expense covered by your policy.

#3:

Disney was the first studio to release a truly massive film originally set for theaters onto a streaming platform. To watch their latest version of *Mulan*, viewers needed to pay close to \$30 on top of their Disney+ subscription.

#4:

Judging from the astonishing range of daily life and human endeavor reflected in his poems and plays, we can only infer that Shakespeare was a keen observer.

#5:

You are not eligible for an upgrade, since you haven't signed up for our newsletter, and signing up is necessary for eligibility.

#6:

Advisory boards are limited in authority, and consequently in legal responsibility, to those powers granted by the local government.

Exercise Set 3B:

Directions: For each of the following determine whether the passage contains an argument. If it does not, write "no argument". If it does, then standardize and diagram the argument.

#1:

We can be sure that the murder was committed by the judge, given that it had to be either the butler or the judge, and we know it wasn't the butler since he was passionately in love with the victim.

#2:

Since goat's milk contains smaller fat globules than cow's milk, it is easier to digest than cow's milk. Consequently, goat's milk may be a viable alternative for children who have a difficulty digesting cow's milk.

#3:

To insert genes into a cell, scientists often prick it with a tiny glass pipette and inject a solution with the new DNA. The extra liquid and the pipette itself, however, can destroy it. In place of a pipette, scientists at BYU have developed a silicon lance. They apply a positive charge to the lance so that the negative charged DNA sticks. When the device enters a cell, the charge is reversed and the DNA is set free. In a recent study using this method, 72 percent of nearly 3,000 mouse eggs cells survived.¹

#4:

The proposed ban on high-capacity magazines doesn't make any sense. Think about it: a ban on high-capacity magazines wouldn't necessarily prevent any of these mass killings, since with practice a person can learn to swap out a depleted 7 round magazine in a couple of seconds or less.

#5:

Because attention is a limited resource—we can attend to only 110 bits of information per second, or 173 billion bits in an average lifetime—our moment-by-moment choice of attentional targets determines, in a very real sense, the shape of our lives.²

#6:

The chief reason painting is superior to sculpture is that painting as a medium affords the artist many more possibilities than sculpture does. After all, how can you sculpt mist or clouds, or the appearance of reflective surfaces? Likewise, in painting the artist can represent impossible objects, and this is not an option for the sculptor who is bound by the laws of space and time.

Exercise Set 3C:

#1:

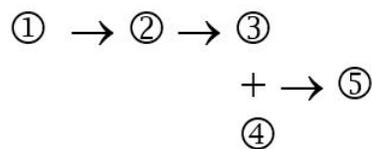
Make up a complex argument and write it out. Once you've done this, standardize and diagram your argument.

#2: Below is a standardized argument without any content. Draw the diagram that corresponds to this standardization.

- 1) xxxxxx
- 2) xxxxxx
- 3) So, xxxxxx (from 1 and 2)
- 4) So, xxxxxx (from 3)
- 5) So, xxxxxxxx (from 4)

#3:

Below is an argument diagram. Create the standardization that corresponds to this diagram (don't worry about content, just follow the xxxxx pattern from above).



#4:

What is a "rhetorical question"? Give an example. Can a rhetorical question be part of an argument?

Notes

1. Adapted from Giller, G. (2014, July 1). Hold Still. *Scientific American*, 311 (1), 19.
2. Adapted from Anderson, S. (2009, May 25). In Defense of Distraction. *New York*. 43 (18), 28-33, 100-101.

Complications to Argument Analysis

SECTION 1: INTRODUCTION

As we have seen, we are not in a position to fairly evaluate an argument until we have accurately determined its structure. That is, argument analysis is a prerequisite for argument evaluation, and in the last chapter we introduced a number of terms and techniques for analyzing and representing arguments. In this chapter, we take a look at a number of complications to the process of argument analysis. The most important of these complications has to do with the absence of indicator words. Authors and speakers often express arguments without using indicator words, and this can make the process of analyzing their arguments particularly difficult. The chapter closes with a general method for looking behind an author's words to see their intended argument. Overall, the chapter emphasizes the importance of paying close attention to the language and goals of argument.

SECTION 2: A TRICKY INDICATOR WORD: 'BECAUSE'

Consider the following example:

Ex. 1:

Dara and Chatri are sitting in Dara's car. They have spent the last 15 minutes trying (and failing) to start the car, when Chatri says: "I bet the car won't start because it is out of gas."

It is tempting to read 'because' as indicating a premise, and to interpret Chatri's comment as expressing the following argument:

1. The car is out of gas.
2. So, I bet the car won't start.

That seems straightforward, but something isn't quite right here. Dara and Chatri are sitting in the car together; they have been trying and failing to start the car for 15 minutes. Wouldn't it be a little weird for Chatri to offer evidence that the car won't start? Chatri and Dara already know that the car won't start! What they are looking for is an explanation for why it won't start.

The problem here is that the word 'because' can be used to express two closely related, but distinct, relationships. On the one hand, it can be used to indicate that a premise is being offered as evidence in support of a conclusion. Let us call this the *Evidential Use* of the term 'because'. The term is also commonly used to

indicate that an explanation is being offered for some state of affairs. This is how the term is being used in the example above. Let us call this the *Explanatory Use* of the term. Let's look at another example to help clarify this difference.

Suppose that there was an extended power outage at a local school yesterday. The next day you are discussing the outage with some other people in your community, and somebody says:

Ex. 2:

The school lost power because a construction crew dug through an electrical cable.

Is this an evidential or an explanatory use of the term 'because'? If it is an evidential use of 'because', the person is giving evidence that the school lost power. However, this wouldn't really make sense—everybody already knows the power went out at the school. The interesting question in this context is what caused the power outage in the first place, and that is what the speaker is offering—an explanation. We can compare this to an evidential use in the same conversation.

Ex. 3:

The school probably won't experience an extended power outage like that again because they are installing a backup generator.

In this example, the person speaking is talking about the possibility of a future power outage, and this is not the sort of thing that it makes sense to explain. After all, nothing has happened yet! Instead the speaker is giving us a reason for thinking that this prediction is accurate.



"Kodiak power outage" by jkbrooks85 CC BY 2.0

In general, when you see a 'because' you will want to keep the following questions in mind: What is the sentence about? Is it talking about something that is widely unknown or debatable? If so, it would make sense to give evidence for it, and the word 'because' is probably introducing an argument. Alternatively, is it talking about some commonly known or obvious fact? If so, the situation doesn't call for evidence, and the word 'because' is probably introducing an explanation.

It is important to raise one further issue when it comes to arguments and explanations, namely that we can argue on behalf of explanations. Put otherwise, although arguments and explanations are different, we can bring the two together by giving reasons for thinking that some specific explanation is right. To illustrate, let us return to Ex. 2. There are many possible explanations for why the school lost power, and it would be fair to ask the speaker "how do you know it was because a construction crew dug through an electrical cable?"

In reply, to this question about evidence suppose the speaker says: “Because it was in the newspaper this morning.”

In this case, the speaker is using the term ‘because’ in both evidential and explanatory senses. We can standardize the speaker’s as follows:

Ex. 4:

1. The newspaper said the school lost power because a construction crew dug through an electrical cable.
2. So, the school lost power because a construction crew dug through an electrical cable.

Here the speaker is offering evidence on behalf of the proposed explanation. That is, the speaker is arguing that a specific explanation is correct. Arguments for explanations are extremely common, and we will take a closer look at them in Unit #5. For now, however, we will turn to a different complication to argument analysis.

SECTION 3: SEEING BEHIND AN AUTHOR’S WORDS

Indicator words are a valuable aid in argument analysis, but in most cases indicator words alone will not reveal an argument’s structure. The chief reason is that authors and speakers do not always mark every inference they make using indicator words. In fact, sometimes an author or speaker will present an argument without using *any* indicator words. Why wouldn’t someone use indicator words? Sometimes, people do so because they take it as obvious that they are making an inference, other times people do so for rhetorical reasons.



“Hide and seek” by nicola.albertini CC BY-NC-ND 2.0

Whatever the reasons, most arguments do not wear their structure on their sleeves, so to speak. This makes the task of argument analysis significantly more complicated. Since the author or speaker has not made the structure of their argument explicit, we are forced to look behind their words to see their intended argument—and this can be tough. How are we supposed to know what an author’s intended argument is, if they do not explicitly tell us? Unfortunately, there is no easy answer to this question; however, there are some strategies that are fairly reliable means to finding an author’s intent.

The first, and most important, step is to find the main conclusion. This is the focal point of the argument to which all the relevant parts of the argument are attached. We find the main conclusion by asking: “what

is the point here?” and “What does the author want me to take from this piece?” Once you’ve found the main conclusion, you are in a position to work backwards to uncover the argument’s structure. Thus, the next task is to find the main reasons the author uses to directly support the main conclusion. We look for these reasons by asking: “what reasons does the author give for thinking the main conclusion is true?” Once we’ve uncovered these reasons, we are not done. After all, the argument may have multiple layers; it may be a complex argument in which the author has also argued for the main reasons. This means we need to look to see if the author has given any second-level reasons as evidence for the main reasons. We do this by asking for each main reason: “what reasons does the author give for thinking that the main reason is true?”. Once we have uncovered any second-level reasons, we can proceed to look for third-level, fourth-level, and fifth-level reasons (and so forth). You will want to continue this process until you’ve captured the whole extended argument. To see how this works, let’s take a look at an example.

Ex. 5:

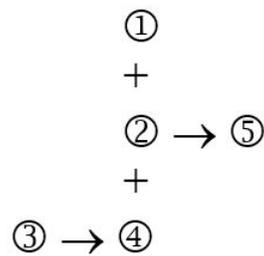
Mega-Mart executives often say that Mega-Mart stores ultimately benefit the community. This just isn’t right. When are people going to wake-up? The presence of a Mega-Mart store in a community actually drives down wages, and the money spent at Mega-Mart does not stay in the community. Further, its presence drives local businesses to close, since small family-run operations cannot compete with the buying power of a gigantic corporation like Mega-Mart.

The first step is to identify the main conclusion. To that end, we’ll ask: what is the point of this piece? In this passage, the author is vehemently disagreeing with the claim that Mega-Mart stores ultimately benefit the community, and he backs this up. This tells us that the point—and main conclusion—of this passage is that “Mega-Mart stores do not ultimately benefit the community.” Notice that the author does not say this explicitly, writing instead that “This just isn’t right” where “this” refers back to the claim that Mega-Mart stores ultimately benefit the community. Once we know the main conclusion, we can work backwards by asking, “what reasons does the author give for thinking the main conclusion is true?” In this case, it is not too difficult to identify these reasons, since the author immediately turns to three ways Mega-Mart stores harm communities: Mega-Marts drive down wages, the money spent at Mega-marts does not stay in the community, and that the presence of a Mega-Mart store drives local businesses to close. Consequently, these are the author’s main reasons.

Our next question is whether the author has second-level reasons. In other words, we need to figure out whether the author offers any evidence on behalf of the main reasons above. Helpfully, the author has given us an indicator word, claiming that “its presence drives local businesses to close, *since* small family-run operations cannot compete with the buying power of a gigantic corporation like Mega-Mart.” So, this is a second-level reason. That covers everything the author says in this passage. We’ve captured all the author’s reasons, and so we can go ahead and put together our final standardization and diagram as follows:

Standardization and Diagram of Ex. 5:

1. The presence of a Mega-Mart in a community actually drives down wages.
2. The money spent at Mega-Marts does not stay in the community.
3. Small family-run operations cannot compete with the buying power of a gigantic corporation like Mega-Mart.
4. So, the presence of a Mega-Mart drives local businesses to close. (from 3)
5. So, Mega-marts do not ultimately benefit the community. (from 1, 2, and 4)



SECTION 4: ANALYZING AMBIGUOUS TEXTS

It wasn't too difficult to identify the author's main reasons in Ex. 5, but it is not always so easy. Sometimes it will not be clear at all how an author is arguing, and this raises an important question: How do we find the author's reasons if they aren't obvious? Here is an example to help us think through the problem.

Ex. 6:

There is no room in the hall of fame for Jack Martinez. His batting statistics weren't earned since he took performance enhancing drugs. Not only that, but he was a divisive figure who undercut the success of his teams.

The main point of this passage is expressed in the first sentence. The next step is to look for the author's main reasons. The sentence that immediately follows begins by telling us that his batting statistics weren't earned. The author hasn't used an indicator word, but this is one of the author's main reasons since it supports the main conclusion. The author continues: "since he took performance enhancing drugs." This is evidence for the main reason, so let's put it aside until we get to a consideration of second-level reasons. The last thing the author says is that Martinez was a divisive figure who undercut the success of his teams. The author thinks it is relevant to the argument (saying "not only that"), but where does it fit into the overall argument? It is hard to tell. What should we do? Well, there are a number of ways it might fit into our analysis. Let's look at each one, and see if anything sticks out.

Option #1: The author thinks that the fact that Martinez was a divisive figure who undercut the success of his teams is evidence that he took performance enhancing drugs.

It seems unlikely that this is the inference the author is making. What does being a divisive figure have to do with taking drugs?

Option #2: The author thinks that the fact that Martinez was a divisive figure who undercut the success of his teams is evidence that his batting statistics weren't earned.

Again, it seems unlikely that this is what the author intends. What does being divisive have to do with whether a person's batting statistics were earned or not?

Option #3: The author thinks that the fact that Martinez was a divisive figure who undercut the success of his teams is evidence that he doesn't belong in the hall of fame.

Of the three options, #3 is the most plausible. It is not difficult to see how a person might think these characteristics of Martinez are relevant to whether he belongs in the hall of fame. As a result, it seems most likely that the author is using the fact that Martinez was a divisive figure as one of the main reasons for the main conclusion.



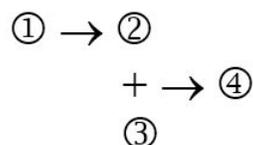
*"Baseball at Dodgers Stadium" by Rafael Amado Deras
CC BY 2.0*

Notice the procedure here. In the face of an ambiguous argument, we identified three different ways of eliminating the ambiguity. Of the three options, only one stands out as plausible. As a result, we attributed this inference to the author, and analyzed the argument accordingly. It is important to note that in saying that Option #3 stands out as the most plausible, we are not thereby saying that we agree with the author or think the argument is sound. Rather, we are engaging cooperatively with the text (see Chapter 2). We are assuming that the author is a reasonable person, and we are trying to interpret the argument in a way that is fair to them. For all this, we might still be skeptical of, or disagree with, the argument.

It is worth noting that we do this sort of thing all the time. The way we deal with typos is a good example. To illustrate, imagine you get a text from you mom that says: 'can you get some nilk at the store'. You will automatically assume that she meant 'milk'. Why? Well, you've never heard of a product called 'nilk', milk, on the other hand, is something you buy regularly, and it would be easy to hit an 'n' instead of a 'm'. Thus, like the case above, in this instance you've used your knowledge and experience to look behind what somebody has said and identify their likely intention.

Let us return to the analysis. We have resolved the ambiguity, and we have the main reasons for the conclusion. The next step is to ask about second-level reasons, and we saw above that the author gives one for thinking that Martinez's statistics weren't earned (namely, that he used performance-enhancing drugs). We have captured the whole argument, and we can standardize and diagram Ex. 6 as follows:

1. Martinez took performance enhancing drugs
2. So, Martinez's batting statistics weren't earned. (from 1)
3. Martinez was a divisive figure who undercut the success of his teams.
4. So, there is no room in the hall of fame for Jack Martinez. (from 2 and 3)



Again, what is important about this example is the cooperative approach to ambiguity. We were not sure where to put the last claim, and in order to figure it out we considered the possibilities, and discovered that

one, in particular, would make sense given that the author is a reasonable person who intends to give a strong argument. Put in slightly different terms, we proposed and tested several different possible analyses of the argument to ultimately arrive at a final analysis.

Before turning to the exercises, let's briefly review the method we've been using. We'll simply call it *The Method*.

The Method:

Step 1: Find the main conclusion by asking "what is the point of the piece?"

Step 2: Find the author's main reasons by asking "what reasons does the author give for the main conclusion?" Resolve ambiguity by considering competing analyses.

Step 3: Find secondary reasons by asking: "does the author offer any evidence for the main reasons?" Resolve ambiguity by considering competing analyses.

Step 4: Find tertiary reasons by asking: "does the author give any evidence for the secondary reasons?" Resolve ambiguity by considering competing analyses.

Step 5: Continue this process until the answer is 'no'.

As a final note, The Method is a tool for comprehensively analyzing an argument. That is, consistently applying this technique to an argument will allow you to see every (stated) element of an argument. It is important to point out, however, that we do not always want or need to go to such depth. Particularly when it comes to longer texts. Depending on our interest or the situation at hand, we may only need to identify the conclusion and the author's main reasons. Alternatively, we may be primarily interested in one of the sub-arguments and only need to analyze a part of the argument. In general, different interests or concerns call for different degrees of analysis. We will say more about this in Chapter 6. For now, however, we will practice fully applying the method and comprehensively analyzing arguments.

EXERCISES

Exercise Set 4A:

Directions: Determine whether 'because' is being used in an evidential sense or an explanatory sense.

#1:

You and your roommate walk into your dorm room; you flip the light switch and the light does not come on. Your roommate says: "the light didn't come on because the bulb is burned out."

#2:

The Broncos are likely to be in the championship hunt this year because of their offseason talent acquisitions.

#3:

A bank manager and her employees watch as the new vault door fails to close. At which point the repair-person says, "The door to the bank vault doesn't work correctly because it was not built with hinges strong enough to carry the weight of the door itself."

#4:

Fewer people will be flying this summer because of the recent spike in ticket prices driven by a surge in the price of jet fuel.

#5:

You limp into the doctor's office, and tell the nurse: "I hurt my ankle because I was wearing the wrong shoes".

Exercise Set 4B:

*Directions: For each of the following determine whether the passage contains an argument. If it does not, write "no argument". If it does, then standardize **and** diagram the argument.*

#1:

You don't have to worry about giving me the chicken pox. I've already had it.

#2:

The fact that the candidate only paid an effective tax rate of 15% in the last couple of years is distressing, but is it relevant to his candidacy for president? No way. It is not like he set the tax rate; he paid what he owed just like the rest of us.

#3:

Nearly two years after the catastrophic Equifax data breach of 2017 was announced, it looks like the company is readying to cough up damages to any of the 147 million people whose personal information was exposed in the breach.

#4:

The claim that the U.S. cannot afford to put a cap and price on greenhouse gases is mistaken. It costs less than many other government programs. Moreover, the cost to households, according to the Congressional Budget Office, would be small. Last, a good program would create more jobs than it cost.

#5:

Unemployment continued to fall in the second quarter largely because of a construction boom in the Sun Belt.

#6:

Headline: "Sculpture bears little resemblance to Kings"

The Jan. 19 issue of the newspaper had a picture of the new Martin Luther King Jr. and Coretta Scott King statue. Perhaps the city should have second thoughts about the purchase. The sculpture looks nothing like them. No likeness at all! It is not a fitting monument to them.

#7:

This governor has got to go! He doesn't have a clue—he actually thinks that the way to drive students to achieve is to buy new computers for our schools. You only have to look at the statistics to see that poverty is the obvious cause of educational under-achievement. Eighty percent of those who ultimately drop out come from homes whose income is at least 50 percent below the average.

#8:

Because virtually all states award all their electoral votes to the winner of the popular vote in the state, and because the Electoral College weighs the less populous states more heavily, it is entirely possible for the winner of the electoral vote to lose the popular vote.

#9:

An attorney says to a potential client: "I would not recommend going forward with your suit against the mayor for defamation for two reasons. First, there is a good case to be made that the mayor's remarks about you were privileged. Given this, they cannot be the basis for a defamation suit. Second, you admit that the mayor's remarks about you are, strictly speaking, true. This means that, legally speaking, you have not been defamed, since defamation only applies when statements about a person are false."

#10:

According to the U.N. Food and Agriculture Organization, more than 1 billion people in the world are hungry and about 25,000 people die every day from hunger or related causes. The United Nations estimates the cost of ending world hunger at \$195 billion a year. That sum — \$195 billion — is a lot of money. However, the world military expenditure was \$1.46 trillion in 2008, according to the Stockholm International Peace Research Institute. The Center for Arms Control and Non-Proliferation credits the U.S. military having spent \$660 billion or 43 percent of the total world expenditure for military purposes in 2007. How have our priorities become so disordered?

Missing Premises

SECTION 1: INTRODUCTION

As we have seen, when people do not use indicator words it can be difficult to determine their intended argument. There is a further fact of life that complicates argument analysis: not only do people sometimes leave out indicator words, but often they leave out *whole premises*. While this might be surprising, we will see that in many cases there are good practical and rhetorical reasons for leaving assumptions implicit. Moreover, being able to spot argumentative gaps and identify missing premises can be crucial for understanding arguments, and seeing their problems. After all, people do not tend to formulate or express arguments in ways that make the argument look obviously bad. Rather, when arguments are poor, often the fault lies in what has been merely assumed or left unsaid.

SECTION 2: ARGUMENTATIVE GAPS AND MISSING PREMISES

Let us begin with a relatively simple example. Suppose a student says:

Ex. 1:

This course is a waste of my time since it doesn't contribute to my major.

Given what we have learned so far, we should standardize and diagram the student's argument as follows:

1. This course doesn't contribute to my major.
2. So, this course is a waste of my time.



Moving forward, we will call this form of analysis, **Surface-Level Analysis**. This is an example of surface-level analysis because it represents only what the author has explicitly stated—nothing more. Given this, we can distinguish surface-level standardizations and surface-level diagrams accordingly. Overall, surface-level analyses can be very useful, and we will continue to use them throughout this book. However, sometimes we need to look behind the explicit statement of an argument, and this calls for a different form of analysis.

With this in mind, the important thing to notice about the argument in Ex. 1 is that it has a gap. The author makes one claim about the course in the premise, namely that it doesn't contribute to the speaker's major, but then leaps to a different claim about the course in the conclusion, namely that it is a waste of the

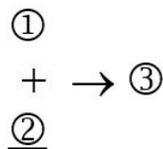
speaker's time. What does the premise have to do with the conclusion? After all, these are distinct characteristics of the course. Why does the fact that the course doesn't contribute to her major mean that it is a waste of her time? This is just one example of a gap, but we can define argumentative gaps more broadly along similar lines. We will say that an argument has a **gap** when (and only when) it does not explicitly connect the premises to the conclusion.

Now although the speaker has not explicitly spelled out what the premise has to do with the conclusion, she clearly thinks they are related. So what is the connection here? There are a number of ways the speaker might be thinking about this connection, but for now we will formulate it as follows:

1. This course doesn't contribute to my major.
2. Courses that don't contribute to my major are a waste of my time. (MP)
3. So, this course is a waste of my time.

We've bridged the gap in this argument by going beyond the stated argument and adding a premise that explicitly connects the claim in the premise that the course doesn't contribute to her major to the claim in the conclusion that the course is a waste of time. We will call any analysis that adds a connective premise, a **Deep Analysis**, and we can distinguish deep-level standardizations and deep-level diagrams, accordingly. It is important when using deep analysis to distinguish an author's explicitly stated premises from the premises we attribute to the author. Consequently, we will adopt the following convention: after an added premise in a standardization we will write **MP** (for missing premise) in parentheses as above.

How do we include missing premises on a diagram? We add missing premises to an analysis to show how the stated premise connects to the conclusion. As a result, in our diagrams missing premises are always added (+) to the stated premise. Thus, in this case we will add the missing premise (2) to the stated premise (1). In addition, as in standardizations, we will want to mark that we have added a missing premise. We will do so by underlining the number of the missing premise. Thus, the deep diagram of Ex. 1 will look like this:



Let us consider another example:

Ex. 2:

Jenna is not eligible to run for the U.S. Senate given that she is only 25 years old.

A surface-level standardization and diagram of this argument looks like this:

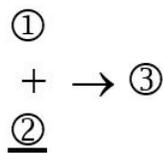
1. Jenna is only 25 years old.
2. So, Jenna is not eligible to run for the U.S. Senate.



As stated, this argument contains a gap; the author is leaping from a claim about Jenna's age to a claim about her ineligibility to run for the U.S. Senate. Again, it is fair to ask: what does one have to do with the

other? Clearly, the author sees the premise as evidence for the conclusion, but how so? The author seems to be assuming that there is some bridge between being only 25 years old and being ineligible to run for Senate. This unstated assumption taken together with the stated premise would connect to the conclusion. Again, there are a number of ways that the author might be thinking about this connection, but for now we will say that the missing premise is 'Anyone who is only 25 years old is not eligible to run for the U.S. Senate,' and we will give a deep analysis of the argument as follows:

1. Jenna is only 25 years old.
2. Anyone who is only 25 years old is not eligible to run for the U.S. Senate. (MP)
3. So, Jenna is not eligible to run for the U.S. Senate.



Again, note that we have marked the addition of the missing premise in the standardization by writing 'MP' after the missing premise, and have marked it in the diagram by underlining the number.

It is important to note that adding missing premises can sometimes feel redundant or repetitive. This feeling, in part, grows out of the fact that a missing premise will normally use the same terms found in the premise and conclusion. Nevertheless, a missing premise says something different (see the examples above). That is, a missing premise makes a claim that is distinct from either the premise or conclusion, and so is not redundant. Indeed, it is precisely because it is making a claim that differs from both the premise and the conclusion that it can serve to bridge or connect the two.



"Hoover Dam Bypass Bridge Construction 3" by squeaks2569 CC BY-SA 2.0

SECTION 3: THE GOLDEN RULE OF ARGUMENT INTERPRETATION

Let us briefly review what we did above. In both Ex. 1 and Ex. 2 we noticed an argumentative gap: namely that the premises didn't explicitly connect to the conclusion. In light of this, we might have inferred that these are simply poor arguments whose premises are irrelevant to the conclusion. But we didn't. In each

case, we gave the author the benefit of the doubt, and assumed they must have left something out. That is, we assumed that the author had drawn on some unstated connection between the premise and the conclusion. Why give the benefit of the doubt in these cases? That is, why not just conclude that these are poor arguments, and move on?

Well, if you think about it, you'll realize that we almost never make all our premises explicit. Often it would take too long to communicate every detail of our thinking, and in many cases, there is no need to make our assumptions explicit since our audience shares them. To return to our original question, the reason to give the benefit of the doubt is that *you* would want to be given the benefit of the doubt if your roles were reversed. After all, most of the time you don't make every premise of your thinking explicit, and you wouldn't want your arguments rejected on these grounds.

This brings us back to the *Golden Rule of Cooperative Dialogue*. As we've seen, the first rule of cooperative dialogue is to treat your interlocutor as you would like to be treated, and this is a specific application of the general rule. The rule tells us that we should try to interpret other people's arguments in a fair and plausible way based on the assumption that they are, like you, a reasonable person. This application of the *Golden Rule of Cooperative Dialogue* is so important that we'll give it its own name, and call it the *Golden Rule of Argument Interpretation*.

The Golden Rule of Argument Interpretation: When engaging cooperatively, interpret other's arguments in a fair and plausible way.

Although we are making this principle explicit here, this is not the first time we've drawn on it. Recall the example from the last chapter about whether Jack Martinez should be in the hall of fame. In that case we were faced with an ambiguous argument, and we ended up attributing the most plausible interpretation of the argument to the author, and we did so in an effort to understand the argument in a way that was fair to the author. Indeed, understanding is the ultimate goal here. Remember, cooperative dialogue is a way of engaging with others to increase understanding and improve decision making, and by following the *Golden Rule of Argument Interpretation* we put ourselves in a position to better understand another's argument.

A final point to emphasize here is that there is a big difference between understanding an argument and agreeing with it. To give an author or speaker the benefit of the doubt, and to interpret their argument in a fair and plausible way is to try to understand what another reasonable person thinks. However, reasonable people make mistakes, and reasonable people can disagree. To understand somebody else's position is simply to get their view straight, and once we've done so, we are free to evaluate it and form our own views about its truth or accuracy.

SECTION 4: UNCOVERING AN AUTHOR'S UNSTATED REASONS

Every time we draw an inference we leave a great deal implicit. The author of the argument in Ex. 2, for instance, is making a wide variety of assumptions including (among others) that:

- There is a person named Jenna.
- There is a place called U.S.
- The U.S. has a Senate.
- The Senate is something people run for.

- Running for the Senate is something that a person can be ineligible for.

Thus, in even simple cases it isn't practical (or interesting!) to try to identify all a person's assumptions. So, what assumptions should we include in a deep analysis of an argument?

When it comes to argument analysis, we are interested primarily in how the author's stated premises connect to the conclusion, and these are the only assumptions relevant to deep analysis. Put differently, in analysis we are only interested in the piece of information that would bridge the gap between the stated premise(s) and the conclusion. Given this, we will adopt the following rule with regard to an author's assumptions:

Rule for Missing Premises: *In the analysis of an argument, include as missing premises only those assumptions that are required to connect the stated premise (or premises) with the conclusion.*

The Rule for Missing Premises tells us that we should only include missing premises that are required to bridge an argumentative gap. Assumed information that is not required to connect the stated premises to the conclusion should be left out of a deep analysis of the argument. Given this, the next question we need to take up concerns how we identify the proposition that bridges the argumentative gap?

As we have seen, the missing premise needs to relate the stated premise to the conclusion, and in order to fill this role the missing premise(s) will need to say something about the subject matter of the premise *and* something about the subject matter of the conclusion. Let's go back and think about Ex. 1. Recall that we began with the Surface-Level standardization:

1. This course **doesn't contribute to my major**.
2. So, this course **is a waste of my time**.

We proposed the following missing premise:

(MP): Courses that **don't contribute to my major** are **a waste of my time**.

In the missing premise the subject matter of the premise is in red, and the subject matter of the conclusion is in blue. The proposed missing premise is one way of bringing these two subjects together to show how the truth of 1) supposedly bears on 2). Specifically, the fact that the course doesn't contribute to this student's major purportedly shows that the course is a waste of time because the student is assuming that courses that don't contribute to her major are a waste of time. In general, when trying to identify and formulate missing premises, it is useful to take the comparison of a missing premise to a bridge seriously. Just as a bridge over a river needs to physically connect both sides, so too, a missing premise needs to connect, in terms of its content, the stated premise(s) to the conclusion.

Let's put this to work in another example.

Ex. 3

Chris: Is *Avengers: Endgame* a good movie?

Sadie: Obviously! I mean, it is the highest grossing movie in history.

A surface-level analysis of Sadie's argument looks like this:

1. *Avengers: Endgame* is **the highest grossing movie in history**.

2. So, *Avengers: Endgame* is a good movie.



“So we have an opening in the Avengers...” by The mofoJT CC BY-NC 2.0

Again, we have an argumentative gap. The stated premise does not explicitly connect to the conclusion. Given the *Golden Rule for Argument Interpretation* we will take for granted that the author is drawing on some connection between the two. In order to bridge the gap, the missing premise will have to start from the content of the premise, and extend to the content of the conclusion. That is, it will need to relate being **the highest grossing movie of all time** to being **a good movie**. Moreover, in doing so, we will need to stick close to the author’s original language, since we don’t want to inadvertently transform the argument into something different. How do we do this? At this point, adding missing premises can get difficult. As it turns out, there are many ways of connecting these two characteristics into a single premise while sticking close to the author’s language. We might, for example, connect them in a direct and literal way:

(MP1): **The highest grossing movie of all time** is a **good movie**

While this bridges the gap, it seems overly specific. As an alternative, perhaps Sadie is making an assumption about *all* high-grossing movies, not just the highest grossing:

(MP2): All **movies that gross a lot of money** are **good movies**.

Again this bridges the gap, but this proposed missing premise seems too broad. Perhaps Sadie is assuming something a bit more tentative, but which nevertheless connects the relevant characteristics:

(MP3): **Movies that gross a lot of money** tend to be **good movies**.

There are many other ways we might bridge the argumentative gap here. So, how do we know which one Sophie is assuming? In short, without talking to Sophie, we don’t know. The best we can do is to choose the premise that strikes us as the most plausible, and MP3 is the most plausible among these options. This means that different people may attribute slightly different missing premises to the same speaker or author. That said, there will be a limited range of acceptable missing premises—namely those missing premises that stick closely to the author’s original language, connect the content of the stated premise to the conclusion in a way that shows why the truth of the stated premise supports the truth of the conclusion, and does so in a plausible way.

SECTION 5: DEEP ANALYSIS IN MORE COMPLICATED CASES

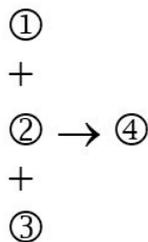
So far we've looked at simple arguments with only one premise. As we've seen however, there are all kinds of arguments. In this section, we will take a look at how deep analysis works in more complicated cases. Let us start with a simple argument with multiple premises. Consider the following:

Ex. 4:

There is no point to student government. They only have the power to make requests to the administration, these requests are rarely fulfilled, and the election process is mostly a popularity contest.

The first sentence of this example is the conclusion. The author then gives three reasons on behalf of the conclusion (without using indicator words). The surface-level standardization and diagram for this argument is:

1. Student government only has the power to make requests to the administration.
2. Requests from the student government to the administration are rarely fulfilled.
3. The election process for student government is mostly a popularity contest.
4. So, there is no point to student government.



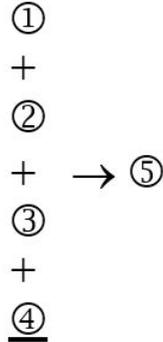
There is a gap in this argument. The first two premises seem to go together because they note that student government has limited power to effect change, and the third is a proposition about the election process. In contrast, the conclusion says there is no point to student government. What do the premises have to do with the conclusion? Given the *Golden Rule for Argument Interpretation*, we'll take for granted that the author is assuming that the premises support the conclusion, and so the task is to figure out how the content of the premises could support the content of the conclusion. Here is one way:

(MP): Student government doesn't have a point if it can't effect change and the election is about popularity.

The color coding shows how the content of the premises and the conclusion is contained within the missing premise. Green captures the idea that student government doesn't have much power to act or create change; red captures the idea that the election process is about popularity, and blue captures the idea that there is no point to student government. There are other ways of formulating the missing premise, but, again, any acceptable missing premise will have to connect the content of the premises to the content of the conclusion. Given this missing premise, the deep standardization and diagram look like this:

1. Student government only has the power to make requests to the administration.
2. Requests from the student government to the administration are rarely fulfilled.
3. The election process for student government is mostly a popularity contest.

4. Student government doesn't have a point if it can't effect change and the election is about popularity. (MP)
5. So, there is no point to student government.



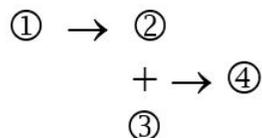
Let us turn to a different example, namely the case of a complex argument. Consider the following:

Ex. 5:

The chemical BPA has been shown to disrupt endocrine function in mice, and this suggests that humans should limit their exposure to this chemical. Since BPA is found commonly in plastics, we should try to limit our exposure to such wrappings.

A surface-level standardization and diagram of this argument looks like this:

1. The chemical BPA has been shown to disrupt endocrine function in mice.
2. So, humans should limit their exposure to BPA. (from 1)
3. BPA is found commonly in plastic food wrapping.
4. So, we should try to limit our exposure to such wrappings. (from 2 and 3)

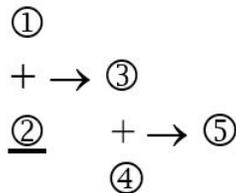


This argument illustrates two complexities. First, as a complex argument there are two inferences here: from 1 to 2 and from 2 and 3 to 4. This means that there are two possible argumentative gaps (one for each inference). A look at the first inference reveals a gap: the premise tells us about endocrine disruption in mice, whereas the conclusion tells us humans should limit their exposure to BPA. What, exactly, is the connection between these two facts? Connecting the content of 1 (red) and 2 (blue) will lead us to something like this:

(MP): Humans should limit their exposure to chemicals that disrupt endocrine function in mice.

What about the other inference? This brings us to the second notable element of this example. There is no gap in the argument from 2 and 3 to 4. The content of 2 and 3 directly relates to the content of 4, and it is clear how the truth of 2 and 3 establish the truth of 4. The takeaway here is that while most arguments contain gaps, not all do. Ultimately, our deep standardization and diagram of Ex. 5 will look like this:

1. The chemical BPA has been shown to disrupt endocrine function in mice.
2. Humans should limit their exposure to chemicals that disrupt endocrine function in mice. (MP)
3. So, humans should limit their exposure to BPA. (from 1 and 2)
4. BPA is found commonly in plastic food wrapping.
5. So, we should try to limit our exposure to such wrappings. (from 3 and 4)



We do not always need to take the time to identify missing premises. It can be quite useful, however, and is not too difficult to learn with some practice. Before we turn to the exercises, however, we'll take up one final issue: our awareness of our own reasons.

SECTION 6: AWARENESS OF ONE'S REASONS

Go back to Ex. 3. We might look at the missing premise and feel like something fishy is going on. After all, did Sadie really think to herself: 'Movies that gross a lot of money tend to be good movies'? The answer is probably 'no'. It seems much more plausible that she inferred the conclusion from the premise without explicitly identifying this proposition in her thinking. Yet, if she didn't think it explicitly, can we really say that it was part of her intended argument? That is, can a person intend an argument if they don't explicitly identify all of the argument's premises in their thinking? Sure.

We each have a vast and complex system of belief, though at any given time we are explicitly aware of only a few beliefs. For example, you believe that you are a particular age, that your birthday is on a particular date, and that your first pet was named so-and-so. However, most of the time you do not pay attention to these beliefs. In fact, you probably weren't thinking about any of these beliefs until just now. In addition, these dormant or implicit beliefs can inform our thinking and acting. Here is an example: suppose you wake up and walk to the bathroom to brush your teeth. You open the cabinet to get your toothbrush. In doing so, you probably did not think to yourself, 'the toothbrush is in the cabinet' and 'I need to open the cabinet door to get it'. Although you were not explicitly thinking these beliefs as you acted, it seems clear that you nonetheless have these beliefs and that they informed your action.



"Assumption, Minnesota" by afileer CC BY-SA 2.0

Furthermore, as we saw back in Chapter 1 there are unconscious and semi-conscious belief-forming processes constantly at work within us. These processes are a way of quickly updating our system of belief in light of new information from our environment. As we navigate the world we engage with it at different levels, and can soak up patterns, norms, and values from our surroundings. A notable consequence is that it is possible for us to have beliefs that we do not even know about. This means that our viewpoints and decisions may be a consequence of beliefs we are not aware we have—beliefs we maybe wouldn't endorse if we did think about them explicitly.

Uncovering the hidden assumptions of an argument is one of the most valuable consequences of the process of argument analysis. The reason is that when people's reasoning is poor, it is often poor *because* these hidden, presupposed, or suppressed assumptions are false. It is easy for people to hold false or implausible beliefs when these beliefs remain below the level of conscious awareness. Consider the argument in Ex. 1. A deep analysis of this argument shows that it assumes that all or most courses that don't contribute to a person's major are a waste of time. But is that really true? When we explicitly identify this assumption, this argument seems a lot less persuasive than it might have initially. Furthermore, reflecting on this assumption might ultimately lead the student to rethink her attitude towards course work.

EXERCISES

Exercise Set 5A:

Directions: For each of the following determine whether the passage contains an argument. If it does not, write "no argument". If it does, (i) give a deep standardization of the argument, making sure missing premises are added in accordance with the Rule for Missing Premises, and (ii) diagram it on the basis of this standardization.

#1:

Of course you can get direct international flights from Windhoek to New York—Windhoek is Namibia's capital city, after all.

#2:

I deserve at least a B on the assignment on the grounds that I worked really hard.

#3:

I don't have any symptoms, so I can't spread the virus.

#4:

Kyle is related to the Johanssons? I bet he is going to be trouble when he is older.

#5:

The PTA's suggestion that in order to protect children from sunburn, a rule should be instituted requiring all children at our school to wear a sun hat when they are outside after 11 a.m. is unacceptable. For clearly such a rule would be an infringement upon the freedom of the individual.

#6:

Judge Zito's policy of taking away the driving licenses of people charged with DUI who are awaiting trial is unconstitutional because the policy does not treat people as if they are innocent until proven guilty.

#7:

A: Illegal immigrants are stealing our jobs!

B: No they aren't. How can they steal jobs that nobody wants?

#8:

If I had to say, I'd say that Kelsey is probably a very good swimmer. Why? Well, she is on the swim team—or at least I think she is on the swim team. I mean I've seen her walking around wearing swim team gear.

#9:

The proposed ban on high-capacity magazines doesn't make any sense. Think about it: a ban on high-capacity magazines wouldn't necessarily prevent any of these mass killings, since with practice a person can learn to swap out a depleted 7 round magazine in a couple of seconds or less.

#10:

There is no need for me to wear a mask, since I'll be fine if I get the virus. After all, I'm young and don't have any health problems.

#11:

I am in favor of returning America to the gold standard. Not only would it limit inflation, but it would force greater fiscal constraints on governments because they couldn't simply print money to pay their debts or bail out bankers. Moreover, it would bring the kind of stability to the monetary system that it had a hundred years ago.

#12:

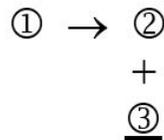
We cannot hold the Scandinavian countries' health care systems as a model to copy in the U.S., for their populations are almost completely homogeneous. The respective populations of Norway, Sweden, Denmark,

and Finland are almost wholly white, Protestant, economically affluent, and highly educated. Furthermore, we cannot even hope to model these health care systems, given that in these countries the ratio of resource wealth to population is much higher than our own. Thus, we will have to be creative and come up with our own solutions to the current health care crisis.

Exercise Set 5B:

#1:

What is wrong with the following diagram? Explain.



#2:

Make up an argument with a gap. Then identify the missing premise, and give a deep analysis of it.

#3:

In the discussion of Ex. 3, the text says that MP3 is the most plausible of the potential missing premises. Why is that do you think? Explain.

#4:

Consider the following argument:

█ I think she is an English major, since she reads a lot.

What is wrong with proposing the following missing premise?

█ (MP): People who are English majors usually read a lot.

Argument Analysis—Some Final Notes

SECTION 1: INTRODUCTION

We have looked at some of the most significant complications to the process of argument analysis. In this chapter, we will consider a number of other important considerations, and bring all of our skills together. We will start by distinguishing independent from conjoined premises, noting the difference between arguments and reports of arguments, and discussing peripheral information. We will close this chapter by putting our skills to work on a much longer, but more realistic, text.

SECTION 2: CONJOINED VERSES INDEPENDENT PREMISES

Authors will sometimes offer multiple distinct arguments on behalf of a single conclusion. Consider the following example:

Ex. 1:

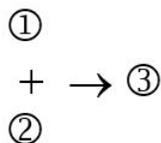
You shouldn't go this weekend because you can't afford it. But even if you could afford it, you shouldn't go because Laila will be there.

We might be tempted to give a surface-level analysis of this example as follows:

Standardization of Ex. 1?

1. You can't afford to go.
2. Laila will be there.
3. So, you shouldn't go.

Diagram of Ex. 1?



However, neither of these representations accurately reflect the argument's structure. The reason is that the author is not claiming that 1 *conjoined with* 2 shows or establishes 3. The author's claim is that 1 *all by*

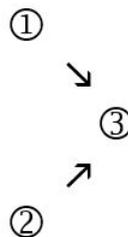
itself establishes 3; that is, the fact that you can't afford to go, means that you shouldn't go. Similarly, the author is claiming that 2 *all by itself* establishes 3. We know this because the author explicitly tells us that even if you could afford to go (that is, even if 1 were false), it would still follow that you shouldn't go since Laila will be there.

Given that the author takes these reasons to independently show the conclusion, we have two distinct arguments. This needs to be reflected in our analysis of the argument, and we will do so in the standardization by explicitly saying this after the conclusion:

1. You can't afford to go.
2. Laila will be there.
3. So, you shouldn't go. (*from 1 and 2 independently*).

In addition, let's agree to represent this kind of argument using a diagram as follows:

Diagram of Ex. 1:



Authors are not always explicit about how the premises of an argument are related to the conclusion. Imagine the person in Ex. 1 has said something slightly different, call it Ex. 1*.

Ex. 1*:

You shouldn't go this weekend because you can't afford it. Also you shouldn't go because Laila will be there.

Put in this way, the argument's structure is not obvious. Does the author mean for the premises to be taken together or independently? More generally, how should we represent an author's argument when it is ambiguous in this way? In order to deal with this situation, we will adopt the following rule:

Rule for Independent Premises: When multiple reasons are offered on behalf of a single conclusion we will assume that the author intends all the premises to be taken together, *unless* the author explicitly says otherwise.

Why adopt this rule? On the one hand, this rule makes sense because arguments with independent premises are much less common than arguments with conjoined premises. In addition, there is no harm in reading ambiguous arguments as having conjoined premises, since doing so captures the author's main point in offering the argument in the first place, namely that we should believe the conclusion on the basis of the premises. Interpreting ambiguous arguments as having independent premises, on the other hand, opens to the door to misinterpretation, since it treats each premise, all by itself, as sufficient for the conclusion. While the author may mean this, they may not. For these reasons, then, we will always interpret ambiguous arguments like this as having conjoined premises *unless* the author explicitly tells us they mean otherwise.

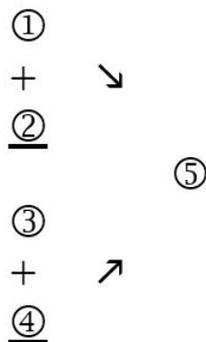
Finally, we should briefly consider what deep analysis looks like in cases of independent premises. Since there are two inferences here, there is the possibility of two argumentative gaps, and indeed this is the case in Ex. 1. First, there is a gap between the fact that the person can't afford to go, and the conclusion that she shouldn't go. Second, there is a gap between the fact that Laila will be there, and the conclusion that she shouldn't go. How are these distinct propositions supposedly related? In order to identify the missing premises here, we need to identify propositions that plausibly bridge the gap by including the contents of both the premise and the conclusion. Thus, we might end up with the following:

- (MP): You shouldn't do things that you can't afford to do.
- (MP): You shouldn't go to places that Laila will be.

Remember, missing premises are added to existing premises, and we can include these missing premises in our standardization by assigning them a number, and modifying the justifications as follows:

1. You can't afford to go.
2. You shouldn't do things that you can't afford to do. (MP)
3. Laila will be there.
4. You shouldn't go to places that Laila will be. (MP)
5. So, you shouldn't go. (*from 1 and 2, and 3 and 4, independently*).

Here is the diagram for this example:



Again, we have underlined 2 and 4 to indicate on our diagram that these are not stated premises, but are being attributed to the author.

SECTION 3: PERIPHERAL INFORMATION

So far in this unit we have practiced our argument analysis skills on short texts. However, there is a sense in which this isn't very realistic. Almost everything we read or listen to in everyday life is much longer, and comes embedded within a larger context. As such, it is important to be able to distinguish the arguments from the context. We will use the term **peripheral information** to identify information that is presented within the context of an argument, but which is not part of the author's main argument. To be clear, saying that some piece of information is peripheral to the argument is to say only that it does not figure into the argument as a premise or a conclusion, and so should not be included in an analysis of the main argument.

A peripheral claim may, nevertheless, be very important for understanding the significance of an argument, its history, its motivation, or any of a number of other background features. Let us look at an example.

Ex. #2:

For decades the city has funded three public golf courses. Roosevelt Park Golf course opened in 1925, while Kas-ten Park opened in 1959, and Westfield in 1991. Unlike private country clubs, anyone is eligible to play on these courses, and there are reduced rates for young people, senior citizens, and military veterans. In recent weeks, the city council has opened up the future of these courses to public discussion, and I'd like to be the first to propose closing the Westfield course. Westfield is the least popular course and it will cost the city at least 2 million dollars to make repairs after last summer's flooding. Moreover, the land this golf course occupies could better serve the residents of the city as a general use park.

We do not get to the author's argument until about halfway through this paragraph. The author's point in this passage is that the city should close Westfield golf course, and she gives a number of reasons for it. The rest of the information in this paragraph is peripheral. The names of each of the public courses and the dates they opened may be important for appreciating the argument, but this information is not being used as evidence for the main conclusion that Westfield course should be closed.



"Edmonton's Victoria Golf Course – Canada's oldest municipal golf course" by Edmonton Economic Development Corporation CC BY-NC-SA 2.0

In general, distinguishing the argument from peripheral information is not too difficult if you are using The Method. Recall that according to this method, we should start with the main conclusion and methodically work backwards to identify the all argument's premises. Once we've completed that process, we will have uncovered the main argument, and consequently, everything that is left will be peripheral. Nonetheless, there are two kinds of peripheral information that can be tricky: **qualifications** and **replies to objections**. Let's start with qualifications.

Often authors will stop to qualify or clarify their claims. Indeed, qualification and clarification are signs of a well thought-out argument. Imagine, for example, that the author of Ex. 2 above continues:

Ex. #3:

I realize that there are residents that enjoy playing regularly at Westfield, and I am not saying that the course is of no service to the community. I am saying only that there are better ways for the city to use this land.

This is an important part of the argument. The author is telling us that she recognizes that the course is a valuable public service. He has taken this into account, and should not be understood as saying Westfield is of no service. Because qualifications like this can tell us something important about the argument, it can be tempting to include them in the standardization. Nevertheless, we should not. Although qualifications and clarifications can be helpful for understanding an argument, they are not themselves a premise or conclusion, and so should be left out of an analysis.

A related situation occurs when an author pauses in the middle of their argument to answer or reply to an objection. In cases like this, the author foresees an objection and wants to cut it off or eliminate it in the audience's mind before proceeding with the rest of the argument. In the following example, the Asst. Manager of a shoe store is talking to the Manager.

Ex. #4:

I think the store should once again carry Campx brand boots. We have had several requests about them over the last few weeks, and they offer a distinctive product. It's true that we stopped carrying them in the past due to problems with product quality, but the company is under new management and claims to have addressed these issues. One other thing: the company is about to kick off a multi-million dollar ad campaign which will surely drive greater interest in these boots.

The assistant manager is arguing that the store should carry Campx brand boots. However, he foresees that the manager might object on the grounds that they've had problems with Campx products in the past. The author raises this objection and then offers a counter-argument. What can be tricky about this kind of case is that though the author is offering an argument, it is not part of the main argument. The reply is a distinct argument, and so is peripheral to the main argument.

SECTION 4: DISTINGUISHING ARGUMENTS FROM REPORTS OF ARGUMENTS

Argument reports are common in everyday communication. To **report** an argument is to describe somebody else's argument. It is important not to confuse a case in which an author reports an argument with a case in which an author puts forth an argument they endorse. Certainly a person can report an argument that they endorse; nevertheless, this need not be the case. Here is an example:

Ex. #5:

Kylie says she isn't going to register for a Film Studies course this semester because she is already enrolled in too many courses.

This is not an argument. In this case, the author is not making any claims about Kylie's argument, but is simply reporting Kylie's thought process. Although a report is not, in itself, an argument, reports can be parts of arguments. In fact, any time we argue against somebody else's argument we report it (after all, we have to say what we are objecting to). Consider this case from a newspaper editorial:

Ex. #6:

I read Wes Pauling's Jan. 13 column, "The death penalty can be justified," hoping for insights, but his argument lacked proof. Pauling says that the absence of the death penalty makes violent criminals more prone to excesses because they know they will not face death. While this sounds reasonable, few criminals weigh the penalties they face if convicted. Most are unaware of legal consequences, do not care, or feel they will never be caught.

In this example, the author’s conclusion is that Wes Pauling’s argument lacks proof. As part of his evidence for this proposition the author offers a report of Pauling’s argument (“Pauling says...”) and his reason for thinking this argument is poor.

SECTION 5: ANALYZING A LONGER PIECE—STAGE 1 AND 2 OF THE METHOD

In this last section, we will put our skills to work on a longer, and more realistic, example. The process of analyzing the main argument in a longer text using The Method is the same as a shorter case. Thus, the first step is to read through the passage and identify the main conclusion. Nevertheless, in a longer case, there is more information to keep track of, and so we will need to be very organized. Let’s turn to the example.

Ex. #7: “The Age of Twitter” by Brian L. Ott

Twitter is a microblogging platform—a form of blogging in which entries typically consist of short content such as phrases, quick comments, images, or links to videos. Users send and receive “tweets,” messages consisting of no more than 140 characters. Since its launch in March 2006, Twitter has grown rapidly in popularity, and by 2014, it had more than 500 million users.

Different communication technologies shape how users process information and make sense of the world, and Twitter is no different. However, Twitter is different in that it has an especially toxic effect on public discourse. I’m not suggesting, of course, that all content on Twitter is harmful. Much of the Twittersphere is relatively innocuous. The danger comes when issues of social, cultural, and political import are filtered through the lens of Twitter.

Twitter negatively influences the public discourse in three ways I will point out. First, because of its 140 character limitation, Twitter structurally disallows the communication of detailed and sophisticated messages. This isn’t to say that a Tweet can’t be clever or witty—it can, but overall Twitter is not a medium for communicating complex thoughts. Second, tweeting is a highly impulsive activity that one can do even if one has nothing considered or important to say. Using this platform requires almost no effort at all—one can tweet from virtually anywhere at any time with the push of a button. The last point is that Twitter encourages uncivil discourse. Twitter’s lack of concern with proper grammar and style undermines norms that tend to enforce civility. Further, Twitter “depersonalizes interactions” creating a context in which people do not consider how their interactions will affect others.

In our society we have seen the rise and mainstreaming of divisive and incendiary public discourse, and a growing intolerance for cultural and political difference over the last 10-15 years. It is difficult to say how much of this is a consequence of Twitter, but one thing is certain: the continued use of Twitter as a means of communicating important ideas will not help matters.¹

What is the main point of this piece? Overall, the author is critical of Twitter, but he seems to draw two conclusions about it. First, he tells us in the second paragraph that Twitter “has an especially toxic effect on public discourse,” and then proceeds to give three reasons on behalf of this proposition. Second, he brings the piece to a close by concluding: “the continued use of Twitter...will not help” to fix the growing intolerance for cultural and political difference in public discourse. Which one is the main conclusion? The propositions seem to be related, and this suggests one is being given as a reason for the other. This presents us with two options to compare:

- Option #1: The author is claiming that the fact that *Twitter has an especially toxic effect on public discourse* is a good reason to believe that *the continued use of Twitter will not help fix the problems in public discourse*.

OR

- Option #2: The author is claiming that the fact that *the continued use of Twitter will not help fix the problems in public discourse* is a good reason to believe that *Twitter has an especially toxic effect on public discourse*.

When set out this way, Option #1 is clearly preferable. While we might not agree with the argument expressed in Option #1, it is an argument one could reasonably make. Option #2, in contrast, is hard to understand, and is consequently less plausible to attribute to a reasonable interlocutor than Option #1. Given this, we've identified the main conclusion, and at least one of the author's main reasons.

In order to illustrate the process of analyzing this longer piece, we will forgo numbers until the end. Instead we will follow the terminology outlined in The Method of 'main reasons,' 'second-level reasons,' 'third-level reasons,' and so forth. Accordingly, we will call the proposition *Twitter has an especially toxic effect on public discourse*, Main Reason #1, or MR1 for short. It doesn't look like the author offers any other main reasons, so we will proceed to the next step of analysis.

- MR1: Twitter has an especially toxic effect on public discourse
- MC: The continued use of Twitter as a means of communicating important ideas will not help matters.

SECTION 6: ANALYZING A LONGER PIECE—LATER STEPS OF THE METHOD

Before continuing it is worth reminding ourselves that we do not always need to comprehensively analyze an argument. As we noted at the end of Chapter 4, different degrees of analysis are appropriate for different situations. In some cases, we may only need or be concerned with the conclusion and the author's main reasons for thinking it is true. In such a case we need not go through the work to identify secondary or tertiary sets of reasons. Alternatively, we might only be interested in a specific sub-argument at work in the piece, and so not need to reconstruct the whole thing.

Indeed, in everyday life the need or interest to comprehensively analyze lengthy arguments like this is rare. For example, in the "Age of Twitter" the author makes the provocative claim that Twitter has an especially toxic effect on public discourse, and we might only (or primarily) be interested in his main reasons for thinking this is true. Nevertheless, we will go ahead with a comprehensive analysis in this case in order to illustrate larger argumentative structures as well as to illustrate some of the interpretive situations and techniques discussed in this Unit.

Ok, so we have identified the main conclusion and the author's main reason for thinking it is true (MR1). The next step is to ask whether the author provides any evidence for MR1. In this passage, the author is very direct about this. He says: "Twitter negatively influences the public discourse in three ways I will point out," and proceeds to give three reasons labeling them 'first,' 'second,' and 'third,'. Thus, we know the second-level reasons, or SRs, on behalf of MR1 are as follows:

- SR1: Twitter structurally disallows the communication of detailed and sophisticated messages.
- SR2: Tweeting is, then, a highly impulsive activity that one can do even if one has nothing considered or important to say.
- SR3: Twitter encourages uncivil discourse.

We now need to investigate each second-level reason. Does the author offer any evidence on their behalf? Let's begin with SR1. The author gives us a reason for SR1 when he says: "because of its 140-character lim-

itation, Twitter structurally disallows the communication of detailed and sophisticated messages.” We thus have a third-level reason, call it, TR1, for SR1.

TR1: Twitter has a 140-character limit.

The author continues, saying “This isn’t to say that a Tweet can’t be clever or witty—it can, but overall Twitter is not a medium for communicating complex thoughts.” Here the author is being careful to note that the 140-character limit is not all bad, and then reiterates SR1 in slightly different terms. Since he is repeating himself, we don’t need to include this in our analysis.

Let’s turn to SR2: “Tweeting is a highly impulsive activity that one can do even if one has nothing considered or important to say.” Does the author give any evidence on behalf of this proposition? He follows SR2 by saying:

“Using this platform requires almost no effort at all—one can tweet from virtually anywhere at any time with the push of a button.”

He does not use an indicator word, but the proposition “Using [Twitter] requires almost no effort at all” does support the claim in SR2 that it is an impulsive activity. Given this, and that the proposition immediately follows SR2, the author is likely offering evidence for SR2 here. Furthermore, though he doesn’t use an indicator word, he offers evidence that using the platform is easy in saying that “one can tweet from virtually anywhere at any time with the push of a button.” Thus, we are given a third-level reason for SR2, call it TR2, and then given a fourth-level reason for TR2, call it FR1.

FR1: One can tweet from virtually anywhere at any time with the push of a button.

TR2: Using this platform requires almost no effort at all.

We have been working backwards through the layers of this argument, and have considered the reasons the author offers on behalf of two of his second-level reasons. One remains: SR3. Does the author give us any reason to believe SR3? Yes, he gives two reasons, saying:

Twitter’s lack of concern with proper grammar and style undermines norms that tend to enforce civility. Further, Twitter “depersonalizes interactions” creating a context in which people do not consider how their interactions will affect others. We’ll call these two reasons third-level reasons 3 and 4 (or TR3 and TR4).

TR3: Twitter’s lack of concern with proper grammar and style undermines norms that tend to enforce civility.

TR4: Twitter “depersonalizes interactions” creating a context in which people do not consider how their interactions will affect others.

Ok, next step: does the author give us any reason to believe TR1-4? We’ve already seen that he gives one reason for TR2, which we called FR1, but when we ask whether the author has given us any reason to believe that any of these other third-order reasons are true, the answer is ‘no’. This means that we are done with our Surface-Level Analysis. Let’s go ahead and reorder them.

TR1: Twitter has a 140-character limit.

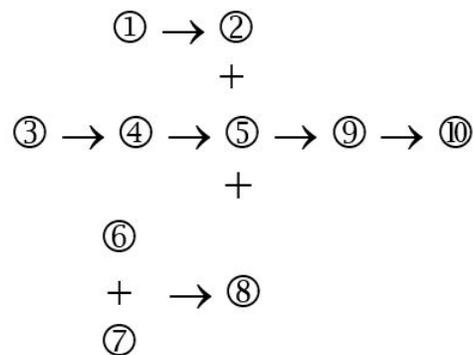
SR1: So, Twitter structurally disallows the communication of detailed and sophisticated messages. (from TR1)

- FR1: One can tweet from virtually anywhere at any time with the push of a button.
- TR2: So, using this platform requires almost no effort at all. (from QR1)
- SR2: So, tweeting is a highly impulsive activity that one can do even if one has nothing considered or important to say. (from TR2)
- TR3: Twitter's lack of concern with proper grammar and style undermines norms that tend to enforce civility.
- TR4: Twitter "depersonalizes interactions" creating a context in which people do not consider how their interactions will affect others.
- SR3: So, Twitter encourages uncivil discourse. (from TR3 and 4)
- MR1: So, Twitter has an especially toxic effect on public discourse. (from SR1-3)
- MC: So, the continued use of Twitter as a means of communicating important ideas will not help matters. (from MR1)

Ordering the propositions in this way shows how we have followed The Method in analyzing this argument. Now that we have everything, however, we can add numbers for our final Surface-Level Analysis of Ex. 7:

1. Twitter has a 140-character limit.
2. So, Twitter structurally disallows the communication of detailed and sophisticated messages. (from 2)
3. One can tweet from virtually anywhere at any time with the push of a button.
4. So, using this platform requires almost no effort at all. (from 3)
5. So, tweeting is a highly impulsive activity that one can do even if one has nothing considered or important to say. (from 4)
6. Twitter's lack of concern with proper grammar and style undermines norms that tend to enforce civility.
7. Twitter "depersonalizes interactions" creating a context in which people do not consider how their interactions will affect others.
8. So, Twitter encourages uncivil discourse. (from 6 and 7)
9. So, Twitter has an especially toxic effect on public discourse. (from 2, 5, and 8)
10. So, the continued use of Twitter as a means of communicating important ideas will not help matters. (from 9)

Surface-Level Diagram of Ex. 7:



A few final notes about this example are in order. First, this article contains a lot of peripheral information. As it turns out, the whole first paragraph is peripheral to the main argument. Second, the author raises and replies to an objection when he says “I’m not suggesting, of course, that all content on Twitter is harmful. Much of the Twittersphere is relatively innocuous.” However, this is not part of the main argument and so was not included in the analysis. The last point is that we have not added missing premises here. That is, we have not pursued a deep-analysis of this example (though we could have).

This brings us to the end of the unit on argument analysis. We have discussed how to identify arguments, and learned some techniques for representing them. Moreover, we have focused on spotting argumentative gaps, and have learned some strategies for identifying hidden or missing premises. This can be hard work, and it is important to emphasize that the recommendation here is not to standardize and diagram every argument that we run across. Doing so is not realistic. Rather, the hope is that as a result of having worked through this unit you are able to spot and pull apart arguments in a much more explicit, effective, and efficient way. More generally speaking, hopefully working through these exercises has shown you the argumentative structures that underlie a lot of our everyday communication. We will continue to use our argument analysis skills as we proceed throughout the rest of the text, and in the next unit we will see more clearly the role analysis plays in the process of argument evaluation.

EXERCISES

Exercise Set 6A

*Directions: For each of the following determine whether the passage contains an argument. If it does not, write “no argument”. If it does, then standardize **and** diagram the argument making sure to add missing premises. Also make a note of any reports or objections/responses the texts contain.*

#1:

Did you know the music department has a course on Avant Garde music? I am definitely taking it next time it is offered. Given that you want to take the course too, you’ll have to sign up for Music Theory III next semester, since Music Theory III is a pre-requisite for taking the course.

#2:

According to Dept. of Justice officials, it isn't realistic to think that Guantanamo prison will be closed this year since there just isn't anywhere else to keep its current inmates.

#3:

A: I am sure I remember that the Mona Lisa is by Rembrandt.

B: No, you're wrong.

A: I am pretty sure.

B: Look, the Mona Lisa was done in the early 1500s and Rembrandt lived during the 1600s. But suppose I am totally wrong about that. Even still, you are wrong since there is a picture of the Mona Lisa in this book about the works of Leonardo Da Vinci—see!

#4:

Smart growth — a politically convenient euphemism for increasing the density of residential neighborhoods and putting them nearer urban areas — is not an adequate solution. Most Americans clearly want more living space and countryside — not less, but even putting this aside, this solution fails since population density brings traffic—again, something nobody wants.

#5:

Parent to child: You know that you shouldn't go to the concert. You promised me that if we didn't give you a curfew any more, you would make sure that all of your homework was completed and submitted on time. And there is no way you can finish your history essay if you go to that concert. I know you are thinking that you can finish it, but c'mon you couldn't finish the essay and have it be any good.

#6:

It is no wonder that Fox is pulling "Lone Star," its acclaimed drama about a con man leading a double life. The show's first episode was a Texas-sized bomb—only 4.1 million viewers showed up for the "Lone Star" premiere.

#7:

I will be voting in favor of the proposed bike-sharing program. There is significant community interest, it is not expensive to install, and has the potential to reduce traffic gridlock. Critics of this proposal claim that the municipality's road and sidewalk system is not built for significant bike traffic, and they are right. However, the proposal introduces a limited number of bicycles, and according to officials from the police dept. this should not be a problem. One last benefit of this proposal is that it extends the reach of public transportation. A bike provides easy access to parts of the city not directly served by bus.

#8:

A British rabbi is encouraging couples to consider being buried on top of each other in single graves to save money and space in cemeteries. Rabbi Ian Morris, of Sinai Reform Synagogue in Leeds, suggested the idea in his newspaper. His argument was that we need to save land and double-depth burials will do that.

#9:

The idea that Iran is a threat to the U.S. itself is not credible. The Iranians do not have the capacity, technological or otherwise, to strike the U.S. in any meaningful way. But even putting that aside, it makes no sense for them to actually attack the U.S., since a U.S. response would be so devastating. Not only is there a carrier fleet just off their coast, but the whole country is virtually encircled by U.S. military installations in Iraq, Turkey, Afghanistan and elsewhere.

Exercise Set 6B: More Difficult Analyses

*Directions: Standardize **and** diagram the following arguments making sure to add missing premises.*

#10:

Chicago's subway system, the "L", is one of America's oldest and busiest. The oldest portions of the system date to the 1890's. Despite this history, the system in Chicago is atrocious. The "L" extremely loud, dirty, and inefficient (it takes forever to get from downtown to either airport). Not only that, but it doesn't provide service to some of Chicago's most important sites. Backers of the L argue that it can't be that bad since the number of people using the system has remained stable over the last 10 years, but the fact the people are using the system doesn't mean it is a good one! If we want Chicago to become a world-class city, we'll need a world-class subway system, this is why we should support an overhaul of the current system.

#11:

Southwest, the nation's largest low-fare carrier, announced Monday it had agreed to purchase smaller discount carrier AirTran for \$1.4 billion, to get a footprint in Atlanta, the world's busiest airport, and expand its presence in airports in New York, Washington, Boston, and Baltimore. AirTran will be absorbed by Southwest and adopt Southwest's policies, according to details of the \$1.4 billion deal. The question that remains, however, is whether this merger—if approved—will benefit consumers. The answer is yes and no. On the one hand customers will likely benefit from expanded offerings, since Airtran serves many smaller cities that Southwest currently does not. On the other hand, it is likely that the merger will ultimately result in higher prices. Mergers create less competition in the market. Consequently, there will be less pressure on airlines to keep their fares down. While the proposed deal's effects remain to be seen, it seems likely that this merger is a mixed bag for consumers.

#12:

Headline: "Proposed Change for Pennsylvania Electoral College Bad Idea"

The Republican leadership in the Pennsylvania Legislature has shown interest in scrapping the current Electoral College system—where the presidential candidate who wins the most votes in Pennsylvania gets all of Pennsylvania's electoral votes—and replacing it with a system where there is competition for single electoral votes in each of the state's 18 congressional districts. This is a bad idea.

Perhaps the deepest problem with this proposal lies in the fact that it builds the system on which electoral votes will be divided up in Pennsylvania on a rotten foundation. The proposal seeks to make Congressional districts the ground upon which presidential elections in Pennsylvania will be built, but Congressional districts throughout the state and the nation have been purposely designed, or gerrymandered, to kill off real competition between the political parties. As long as the core remains rotten, this plan to reform the Electoral College remains rotten too.

Moreover, if Pennsylvania turns to a system where Electoral College votes are chosen district by district, you will ensure that most of the state's electoral votes will be determined before the race even begins. Consequently,

of course, presidential candidates will spend neither time nor money campaigning where the outcome is preordained; it makes no sense to use precious resources on done deals.

Finally, the argument for this change offered by Republicans in the Legislature just doesn't make any sense. They claim that the system needs to change because presidential elections in Pennsylvania just aren't competitive, but the evidence just doesn't support the claim. While Barack Obama won here by about 10 percentage points in 2008, Al Gore won by less than 5 percent in 2000 and John Kerry less than 2 percent in 2004. Hardly evidence of a playing field on which Republicans can't compete.

To be sure, the current version of the Electoral College has many liabilities. The winner-take-all system does not ensure that the candidate receiving the most popular votes will win the presidency. Al Gore can attest to that fact. It's reasonable to explore options in which we as a people can have our voices heard in an equal and meaningful way. Unfortunately, this is not such an option.²

Notes

1. Adapted from the article by Brian L. Ott (2017). "The age of Twitter: Donald J. Trump and the politics of debasement," *Critical Studies in Media Communication*, 34:1, 59-68.
2. This example is significantly adapted from Borick, Christopher. (2011, Sept. 21). Proposed Change for Pa. Electoral College Rotten to the Core. *Morning Call*. <https://www.mcall.com/opinion/mc-xpm-2011-09-21-mc-electoral-college-pa-borick-yv-0922-20110921-story.html>.

Unit #2 Summary

Skill in argument analysis is a crucial prerequisite for both cooperative dialogue and argument evaluation, and in this unit we focused extensively on developing this skill. In order to spot arguments and identify their structure we will often have to actively interrogate the text by working backwards from the conclusion. In addition, we have learned two ways of representing arguments schematically: standardization and diagramming. Argument analysis is complicated because language is complicated, and in this unit we looked at a number of common problems and identified some strategies for dealing with them. One thing that can make analysis especially difficult is that fact that people often leave important parts out. In light of this, we have learned how to see behind an author's words and identify missing or suppressed premises. Last, we noted that arguments are almost always found within a broader context, and we learned how to distinguish arguments from peripheral material and reports or discussions of other arguments. We are now ready to turn, in the next unit, to the main event: argument evaluation.

KEY TERMS:

- Argument Analysis
- Standardizing an Argument
- Diagramming an Argument
- Simple vs. Complex Args
- Direct versus Indirect support
- 'Because'
- The Method
- Surface-Level Analysis
- Argumentative Gaps
- Deep Analysis
- Rule of Interpretation
- Rule for Missing Premises
- Peripheral Material
- Qualification
- Reply to an Argument
- Report of an Argument
- Conjoined Premises
- Independent Premises
- Independent Premises Rule

FURTHER READING

For a solid alternative approach to argument diagramming see Trudy Govier's *A Practical Study of Argument*. For a broader look at the history of argument diagramming and its place in critical thinking pedagogy see "Using Argument Mapping to Improve Critical Thinking Skills" by Tim van Gelder in the *Palgrave Handbook of Critical Thinking in Higher Education* (2015).

UNIT III

AN INTRODUCTION TO EVALUATION

CHAPTER 7

Evaluation and Logical Strength

SECTION 1: INTRODUCTION

Now that we have practiced and refined our argument analysis skills, we can turn to our ultimate goal: argument evaluation. As we've noted, you are already a skilled argument evaluator. You have a good intuitive sense for the difference between good arguments and bad ones, and our goal is to sharpen these existing skills by taking a closer look at this process. In the rest of this book we will pursue this aim in three ways. First, we will develop a vocabulary for talking about argument evaluation. Second, we will consider a variety of common mistakes in argument evaluation, and thereby put ourselves in a position to avoid making them. Third, we will identify some useful techniques for evaluating arguments in specific contexts. As we will see, accurate evaluation often means asking the right questions, and over the next few units we will talk specifically about what kinds of question to ask, and when to ask them. In this chapter, we take up some important preliminaries to evaluation. We begin by briefly reviewing the criteria we use to evaluate arguments, and the questions we need to ask in the process. We will then take a closer look at the notion of logical strength, and lay out some important vocabulary for moving forward.

SECTION 2: EVALUATING ARGUMENTS

You've found an argument, and isolated its premises and conclusion. Now it is time to evaluate: is it a good argument or a poor one? Normally, we evaluate arguments in fairly general terms. We might think "yep, that argument makes sense" or "nope, I don't buy it." However, let's walk through the process of making this judgment in more detail. First, recall that we said we'd call good arguments, *sound*, and poor arguments, *unsound*. In a sound argument the premises are true and strongly support the truth of the conclusion. Consequently, a person can rationally believe the conclusion of a sound argument on the basis of the premises. In an unsound argument, on the other hand, we cannot rationally believe the conclusion on the basis of the premises.

Given these definitions, there are two ways that an argument can fall short. First, an argument goes wrong if one (or more) of its premises are false. After all, a false claim can't serve as evidence for anything. Thus, a sound argument cannot have any false premises. We have a term for this: *factual correctness*. As we've noted, an argument is factually correct when all its premises are true, and factually incorrect otherwise. The second way an argument can fall short is when the premises do not offer enough support for the truth of the conclusion. An irrelevant claim, for example, even if true, is no reason to believe the conclusion. In a sound argument the premises have to offer strong support for the conclusion. Again, we have a special term

for this characteristic: *logical strength*. An argument is logically strong when the premises—if true—provide strong support for the truth of the conclusion, and it is logically weak otherwise.

Because arguments can fall short in terms of factual correctness and logical strength, we need to specifically look for each of these features when we evaluate an argument. When it comes to factual correctness, we need to look at each premise and ask: is this likely to be true? When it comes to logical strength, we need to ask: is the conclusion probable assuming the truth of the premises? If your answer to all these questions is ‘yes,’ then you should deem the argument sound. If your answer to one is ‘no’ then unsound, and if you are not in a good position to say one way or the other, then you can only step back, admit you don’t know, and wait for more information.

The Two Key Evaluative Questions:

Are the premises likely to be true?

Is the conclusion probable given the truth of the premises?

With this terminology in hand, let’s walk through a few cases to illustrate the evaluative process in action. Ava has just met Greg. She mentions this to a friend who says:

Ex. 1:

“I don’t know him, but I think he is on the tennis team; I stopped by his dorm room once to borrow something from his roommate, and noticed that Greg has a tennis racket.”

Ava’s friend has just given the following argument:

- 1) Greg has a tennis racket.
- 2) So, Greg is on the tennis team.

Presented with this argument, Ava turns to evaluation. She thinks to herself: *I don’t buy it. Yeah, I have seen that he has a tennis racket too, so the premise is true; but I don’t see how that makes it probable that he is on the tennis team. Tennis is a popular pastime, and many people play tennis without being on a team. So, although I agree that the argument is factually correct, I ultimately think the argument is unsound because it is logically weak.*



“kids tennis racket” by 5thLuna CC BY-NC 2.0

This example provides an opportunity to emphasize that in saying an argument is unsound, we are not thereby saying the conclusion is false. The fact that Greg has a tennis racket doesn't give a good reason for thinking that he is on the tennis team. That said, neither does it rule it out. In judging that the argument in Example 1 is logically weak Ava is saying only that the evidence provided isn't good enough for believing Greg is on the tennis team, and barring additional information, we can't say one way or the other. Let's consider a different example. Ava and a friend are looking through the course offerings for next semester, and her friend says:

Ex. 2:

"Hey, we could take neuroscience of gender, that sounds cool. Wait. I take it back; we can't sign up for the course since it has BIO 2 as a prerequisite, and we haven't taken that yet."

Ava's friend has just given the following argument:

- 1) The course has BIO 2 as a prerequisite.
- 2) We haven't taken BIO 2.
- 3) So, we can't sign up for the course.

Turning to evaluation, Ava thinks to herself: *that's not right. I agree that if BIO 2 were a prerequisite, we wouldn't be able to sign up for the course, since neither of us have taken BIO 2. The problem is that BIO 2 is not actually a prerequisite (Ava got an email from her advisor earlier in the day noting a number of mistakes in the course listings, including this one). So, while the argument is logically strong, it is unsound because it is factually incorrect.*

In both of these cases, Ava has been presented with an opportunity to update her system of belief with new information. Moreover, in each case Ava used the *Key Questions* to methodically think through each argument. In both cases she judged the evidence to be lacking: in the first case because it was logically weak and did not support the conclusion strongly enough, and in the second case because it was not factually correct. As a result of these judgments, she rejected the proposed additions to her system of belief.

Let us consider one more case. In the following example Ava's friend is talking about a proposed state law that would require all motorcyclists to wear helmets. Her friend disagrees with the proposal, and says:

Ex. 3:

"This is a clear case of government overreach since it infringes on an individual's right to accept the risks of riding without their helmet."

Ava's friend has just given the following argument:

- 1) The proposed law infringes on an individual's right to accept the risks of riding without their helmet.
- 2) So, the proposed law is a case of government overreach.

Ava thinks to herself: *that doesn't seem right. On the one hand, the premise is true—the law does infringe on motorcyclists' right to choose whether or not to wear a helmet. But all by itself this doesn't mean the law is a case of government overreach. It seems like there are other relevant factors here. A helmet is not only a fashion choice— it concerns matters of life and death. Wearing a helmet can save a person's life in a crash and prevent irreparable brain damage. I mean most states have seatbelt laws, and these laws infringe on an individual's right*

without counting as government overreach, right? As a result, I think this argument is logically weak—the mere fact that a law infringes an individual's right does not mean it is a case of government overreach.

This example is different from the previous ones, and raises some interesting questions. Although we might not know whether Greg is on the tennis team or not, or whether BIO 2 is really a prerequisite, there are clear and commonly accepted ways to settle these issues by determining the facts of the matter (e.g. checking with the tennis coach or team members, asking the instructor or registrar's office). But this case is not so clear. How do we determine whether this is "government overreach" for example? There is no easy and commonly accepted procedure to answer this question. Indeed, although Ava thinks this argument is unsound, it is not difficult to imagine others arriving at the opposite conclusion. Perhaps, for example, they would say that seatbelt laws are cases of government overreach as well. How can Ava be sure that her evaluation is right?

In thinking about cases like this it is important to begin with an observation: the world is complex and in a state of constant change. This can make it difficult to understand the world and its contents, since doing so often requires integrating disparate pieces of information—information that itself can be complex and changing. Put differently, some subjects are difficult and complicated, and accordingly there may be no simple, easy, or obvious way to evaluate arguments about them. Indeed, we are called upon to evaluate ideas, arguments, and proposals in these circumstances commonly.

Moreover, when issues are complicated in this way, we should expect disagreement or controversy. Think about it: when we evaluate arguments we do so from our perspective and in light of our system of beliefs. The more complex an issue is, the more likely that two or more people will not share all the same beliefs about it. In these cases, it is common for one person to know something the other person does not. This also shows that disagreement, all by itself, does not mean you are wrong, or that there is no fact of matter. After all, this disagreement may grow out of misunderstanding, ignorance, or mistaken reasoning (we will talk more about disagreement in Chapter 18).

To return to Ex. 3, Ava is evaluating this argument in light of what she takes herself to know about the world, and from her perspective the argument seems logically weak. Nevertheless, she might be wrong. Maybe her friend knows something she doesn't. Perhaps there are good reasons to think that the rights of motorcyclists outweigh concerns about their safety. But in the absence of any positive information to this effect, it makes sense for Ava to evaluate the argument as unsound. That said, her evaluation of the argument is not fixed. Because Ava recognizes that the world is complicated and she may have missed something, she keeps an open mind, and is willing to adjust her views in light of new information or arguments. Indeed, this would be a good opportunity for cooperative dialogue with her friend as a way to learn more about why her friend disagrees as well as to share her perspective.

As a final point, let us consider a comparison that may be useful for conceptualizing argument evaluation: a baseball umpire. An umpire has standards for calling strikes, fair and foul balls, and whether a runner is safe or out (among others things). These standards tell an umpire what to be looking out for in certain situations. Moreover, some events may be more difficult to call than others, and it takes practice to learn how to consistently apply these standards. Nevertheless, even the calls of experienced umpires are subject to error and controversy from time to time. In some cases, the situation is ambiguous or difficult to decide. In other cases, an umpire might not be in a good position to see something relevant, and sometimes umpires simply make a mistake. Evaluating arguments is similar. In evaluating an argument our job is to put the standards of factual correctness and logical strength to work. These standards tell us what we should be

looking for, and the questions to ask. Some arguments are more difficult than others to evaluate, and it takes practice to learn how to apply them well. Further, like an umpire, sometimes we are faced with ambiguity or disagreement, are missing some crucial piece of information, or simply make a mistake. There is an important difference, however. Unlike the umpire we do not have to make a call. That is, we don't have to say 'factually correct' or 'factually incorrect', or 'sound' or 'unsound'. In most cases, we can simply admit that we don't know, and need more information before we can say one way or the other. We will have more to say about the process of evaluation as we proceed through the coming chapters. For now, let us take a closer look at one of the key concepts in argument evaluation, logical strength.

SECTION 3: DEGREES OF LOGICAL STRENGTH

As we've seen, an argument is logically strong when the premises—if true—give strong support for the truth of the conclusion, and logically weak otherwise. Logical strength is not all or nothing; it is a variable characteristic of arguments that comes in degrees. That is, one set of premises might offer more or less support for a conclusion than another. As we will see, this is important because it allows us to distinguish different kinds of arguments, and argumentative standards.

Let's take a closer look: imagine degrees of logical strength on a scale that runs from 0 on the left to 100 on the right. Let 0 represent the absence of logical strength, and 100 represent maximal logical strength. In an argument with no logical strength, the truth of the premises is wholly unrelated to the truth of the conclusion. We don't see arguments like this too often, but here is an example:

Ex. 4—*Logical Strength 0*:

Since Harrisburg is the capital of Pennsylvania, the Denver Nuggets will probably make the NBA playoffs this year.

This is a total *non sequitur* (Latin for "does not follow"): which city is the capital of Pennsylvania tells us nothing about the NBA playoff picture. Moving up the scale from 0, we have arguments whose logical strength lies somewhere between 1 and 50. In these arguments the truth of the premises would give some reason to believe the conclusion, but would not make it probable. For example, suppose that a painting has been stolen, and Lucy's fingerprints have been found at the scene of the crime. The fact that Lucy's fingerprints were found at the scene of the crime is relevant and gives some reason to think that she committed the crime, but does not make her guilt probable all by itself, since there may be other reasons her fingerprints were found there. That is, it would be logically weak to argue:

Ex. 5—*Logical Strength 1-50*:

Since Lucy's fingerprints were found at the scene of the crime, she is probably guilty.

Because the fingerprint is only suggestive, we'd need more evidence to conclude that she probably committed the crime. Continuing up the scale, let us imagine that 50 is the tipping point or threshold for strong arguments. In arguments with a logical strength of 51-99 the truth of the premises makes it probable that the conclusion is true. These are logically strong arguments, and most of the logically strong arguments we run across fall within this range. To give an everyday example, you've got a really sore throat and go to your trusted family doctor. The doctor does a quick test and tells you that you've got strep throat, and prescribes some antibiotics. In this case, it would be logically strong to reason:

Ex. 6—*Logical Strength 51-99*:

Since the doctor says I have strep throat, I probably have strep throat.

In this example, your doctor is a trusted expert whose job involves accurately diagnosing and communicating with patients, and this gives you good reason to conclude you've got strep. This brings us, last, to maximal support at 100. In these arguments, the truth of the premises guarantees the truth of the conclusion. Put differently, in an argument with maximal logical strength the conclusion cannot be false if the premises are true. To illustrate:

Ex. 7—*Logical Strength 100*:

We can be sure that all whales have livers, since all whales are mammals and all mammals have livers.

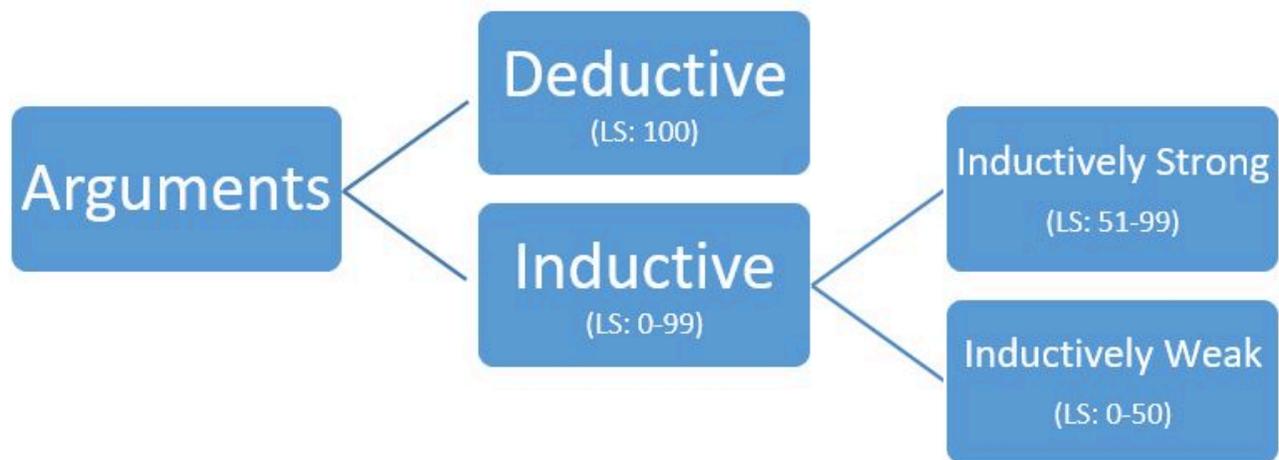
In this case, if it is true that all mammals have livers and that all whales are mammals, it cannot be false that whales have livers. Though arguments like this are not particularly common in everyday life, they are very important. As it turns out, arguments with maximal logical strength have a number of interesting and useful characteristics, not the least of which is that their conclusions can be known with certainty (if their premises are true). In fact, there is a whole discipline called 'Formal Logic' dedicated to the study of this kind of argument. Because of their distinctiveness, arguments with maximal logical strength are used as a basis for classifying arguments more broadly. Thus, arguments with maximal logical strength are called *deductive arguments*, and arguments with less than maximal logical strength are called *inductive arguments*. Let us say that:

An argument is **deductive** when (and only when) it is maximally logically strong insofar as the truth of the premises guarantees the truth of the conclusion.

An argument is **inductive** when (and only when) it is less than maximally logical strong insofar as the truth of the premises does *not* guarantee the truth of the conclusion.

It is important to draw one more set of distinctions when it comes to inductive arguments. After all, just because the truth of an argument's premises does not guarantee the truth of its conclusion, doesn't mean it is a weak argument. We regularly and effectively reason using arguments in which the truth of the premises makes the truth of the conclusion probable without guaranteeing its truth. Thus, we can distinguish between those inductive arguments that are logically strong, and those that are logically weak. We will say that the former are **inductively strong**, and the latter are **inductively weak**.

Given these distinctions, we can classify all arguments in terms of their logical strength. In the following diagram 'LS' stands for logical strength.



“Classifying Arguments by Logical Strength” CC

Why does this matter? Different levels of logical strength are appropriate for different circumstances, and we seek to create and evaluate arguments of differing strength accordingly. Although we can find deductive arguments in many contexts, they are most common in mathematics. For example, think back to high school geometry class. Suppose that you’ve been assigned the task of *proving* that every angle has a unique bisector. Success, in this case, is a matter of coming up with a deductive argument for this conclusion, and that is what you’ll try to offer. However, most contexts do not call for such a high degree of logical strength, and most of the arguments we use and run across are intended to be merely inductively strong. Nevertheless, this distinction means that we need to keep in mind that some arguments are crafted to be deductive, and should therefore be evaluated by the standard. Similarly, some arguments are crafted to be only inductively strong, and should therefore be evaluated by that standard. To illustrate the worry, go back to Ex. 6. Imagine a critic objecting: “this is a poor argument, after all, it is possible that you don’t have Strep throat even though the doctor said you do.” To object in this way, is to apply a deductive standard to a case in which an inductive standard is appropriate.

How do we figure out what standard is appropriate? Context will help, of course, but the most straightforward way is in terms of indicator words. When a speaker or author intends an inductively strong argument, they will commonly qualify their conclusions by saying that they are ‘probable’ or ‘likely’ or that the conclusion ‘seems to follow’. Similarly, when a speaker or author intends a deductive argument, they will commonly indicate this by saying that the conclusion *must* be true, that it necessarily follows, that it is proven, etc. Of course, there are many different ways to express that a conclusion is likely or necessary, and we cannot look at every possible synonym. It will be enough for our purposes to note that i) authors and speakers often make this qualification and ii) that this qualification gives us a guide for thinking about the standards by which to evaluate the argument. If an argument is offered as a proof, then it must have perfect logical strength—anything less qualifies it as a deductive failure. If an argument is offered with a recognition that the argument is not conclusive, then the standard of inductive strength is appropriate. There is a lot more to say about both deductive and inductive arguments, and we will discuss these topics in detail later.

EXERCISES

Exercise Set 7A:

#1:

Create a standardized argument that is logically strong, but factually incorrect.

#2:

Create a standardized argument that is logically weak, but factually correct.

#3:

Create a standardized argument that is logically strong, but has a false conclusion.

Exercise Set 7B:

Directions: Decide whether the following arguments are intended to be deductive or inductive, and then try evaluating the argument by that standard of logical strength.

#1:

I doubt Callista will be at Spin class, since she broke her ankle last week.

#2:

We know that either the Butler or the Chauffer committed the murder, but it can't have been the Butler since he was 20 miles away at the time of the murder. So, it must have been the Chauffer.

#3:

Both Alex and Hakim live in Philadelphia; I bet they go to the same dentist.

#4:

Vaughn is an employee; so he must have access to office, since everyone who has access to the office is an employee.

Exercise Set 7C:

Directions: Evaluate the following arguments for soundness. In making your judgment, please explicitly comment on the argument's factual correctness and logical strength.

#1:

Since Ottawa is the capital of Canada, it is probably Canada's most populous city.

#2:

The multiples of 3 are all odd, since 3 is an odd number and every multiple of an odd number is itself odd.

#3:

An inductively weak argument is one whose premises, if true, do not make the conclusion probable. If the premises, even if true, do not make the conclusion probable, then the conclusion is false. So, inductively weak arguments always have a false conclusion.

#4:

Many books have already been digitized, and there is a digital version of almost every new book that is published. Because of this, brick and mortar libraries will likely become obsolete in coming years.

#5:

Smoking and vaping should be legally banned in all their forms, since both smoking and vaping are bad for your health.

#6:

Christopher Columbus initiated the colonization and subjugation of indigenous peoples of the Americas, and these are not the actions of a man who should be honored or commemorated in our public spaces. Consequently, we have a duty to remove monuments of Columbus from public spaces.

Exercise Set 7D:

#1:

As you were evaluating the arguments in exercise set 7C, what questions or problems did you have? If you didn't have any questions or problems, what are some questions other people might have in completing these exercises?

#2:

In order to find a defendant guilty in a criminal trial, jurors need to believe the evidence shows guilt "beyond a reasonable doubt." Understood in terms of logical strength, what degree of logical strength is "beyond a reasonable doubt?" Explain.

CHAPTER 8

Factual Correctness

SECTION 1: INTRODUCTION

At the end of the last chapter we focused on logical strength, and some of the complications involved in deciding whether an argument is, or is not, logically strong. In this chapter, we will look at factual correctness. We will get a better sense for what factual correctness means, look at the process we must go through in deciding whether an argument passes this test, and think about some of the obstacles to making these judgments. In addition, we will take a closer look at conditional claims (first discussed in Chapter 1), as well as a particularly common failure of factual correctness—the straw man argument.

SECTION 2: MAKING JUDGMENTS ABOUT FACTUAL CORRECTNESS

As we've seen, to evaluate an argument for factual correctness is to ask of each premise: *is it likely to be true?* Although this is a straightforward question, and we make judgments of truth all the time, it will be useful to slow down and clarify a few elements of this process. Here is an example we can work with:

Ex. 1:

Binge-watching TV shows is a superficial way of experiencing a story, since it doesn't allow you to appreciate each episode.

Focusing on factual correctness we will ask: is it likely to be true that binge-watching doesn't allow you to appreciate each episode?



“*Watching*TV*” by Annelogue CC BY-NC 2.0

One of the first things you'll notice when you start being more intentional about argument evaluation is that it involves paying close attention to language and the meanings of words. What, for example, is meant by 'binge-watching' in this premise? Moreover, who, exactly, is the premise talking about? Is this talking about everybody, or most people, or what? The point in asking these questions is not to be hyper-skeptical or to cast doubt on the premise. After all, we often express ourselves incompletely or in ambiguous ways. Rather, the point is to clarify the premise, so that we can decide whether it is likely to be true. Often, the context will make the author's intended meaning clear, but sometimes we'll have to fall back on the *Golden Rule of Argument Interpretation*—what would make sense for the author to mean here? In this case, the author probably means something like this: for most people, watching many episodes of a show in a row makes it difficult to appreciate each episode.

This process of clarifying premises highlights the dynamic between analysis and evaluation. It has been useful to this point in the text to characterize argument analysis and evaluation as if they are wholly distinct processes. In reality, however, we often need to go back and forth to effectively think through an argument. In giving an author or speaker the benefit of the doubt, we implicitly bring our evaluative standards to bear. At the same time, the process of evaluation can force us to return to analysis to clarify the issue at hand.

A second thing you'll notice as you reflect on judgments of factual correctness is that we can be more or less confident that a premise is true. There are some things that we are certain about. You are probably certain that $1 + 1 = 2$, for example. We believe plenty of things that we aren't certain about, however. You might be very confident, but not certain, that it won't snow (or be 90 degrees) today. Furthermore, you might be less sure, yet nonetheless confident that traffic will be bad this afternoon or that the library is open until 9 pm. Last, some things we think are true, but wouldn't be all that surprised to be wrong about. In these cases, we think only that the proposition in question is more likely true than false. Now, the fact that we can be more or less confident about the truth of a premise raises a question for evaluation, namely: how confident do we have to be about a premise to judge that it is likely true? Do we have to be certain about it? No. Certainty is too high a standard. Rather, in claiming that a premise is likely true, we are saying only that in our view it is more likely true than false (although, of course, we can be more confident as well).

A third point to emphasize is that in judging whether a premise is likely true or not, we've got three options: '*yes—I think the premise is likely to be true*', '*no—I think the premise is not likely to be true*' (either because it is likely to be false or because it is as likely to be true as it is false), and '*I don't know—I can't say whether it is more likely to be true than false or not.*' Because there are three options when it comes to individual premises, there are also three options when it comes to making judgments about the factual correctness of an argument. When you think all the premises are likely to be true, then you think the argument is *factually correct*. When you think at least one of the premises is not likely to be true, then you think the argument is *not factually correct*, and when you can't say one way or the other whether both premises are likely to be true, then the argument is *undetermined*. Again, soundness requires factual correctness, so if you think the argument is not factually correct or it is undetermined, then we cannot claim the argument is sound.

A final point has to do with the fact that there are different senses of the terms 'true' and 'truth' at work in everyday conversation, and this can lead to confusion when it comes to judging factual correctness. As we will understand the term, a proposition is true when (and only when) the world matches what the proposition says. This is really straightforward; for example, if you believe your friend Manan's birthday is on Aug. 12th, and his birthday is on this date, then this proposition and your belief are true. If not, then they are false. While we often use the term 'true' in this way, we do not always do so. In disagreements people sometimes say things like: "that's your truth, but this is mine." Similarly, sometimes people advise us

to “live your truth.” In both of these cases, the word ‘truth’ is being used to mean ‘belief’: “that’s your belief, but this is mine,” “live your belief”. It isn’t surprising that people sometimes use ‘truth’ to mean ‘belief’ since to believe something is to think that *it is true*. Nevertheless, in this text we will carefully distinguish between ‘belief’ and ‘truth’ because we will need to clearly distinguish between true beliefs and false beliefs (and you can’t do this if ‘truth’ just means ‘belief’).

SECTION 3: A SPECIAL CASE—CONDITIONALS

In many cases it is not too difficult to assess the truth of an author’s stated premises. One kind of claim, however, can present some questions—conditionals. We briefly introduced conditionals in Chapter 1, and in the meantime we’ve seen that they commonly appear as premises. Recall the following examples:

Ex. 2:

You are legally eligible to purchase and consume alcoholic beverages in this state only if you are 21 years of age or older.

Ex. 3:

If the candle is lit, then oxygen is present.

Ex. 4:

You are not permitted to play poker at the Platinum table if you are not willing to bet at least \$20 per hand.

These examples each express a relationship between two things; more specifically, that one thing is dependent or conditional on another. Conditional claims can be expressed in a variety of ways. However, we will call the ‘if....then...’ form seen in the Ex. 3 above, the **Standard Form for Conditionals**. As we will see, each of the other examples can be expressed in Standard Form as well. Standard-form conditionals have two parts: the part immediately following the ‘if’ which is called the **antecedent**, and the part immediately following the ‘then’ which is called the **consequent**.

If (.....antecedent.....), then (.....consequent.....)

As we’ve seen, premises can be ambiguous and conditionals are no different. Sometimes people use conditionals to express a strict conditional relation between two things. When a person says, for example, ‘if a shape is a square, then it is a rectangle’ they are using a conditional to say that every square is a rectangle—no exceptions! Let’s call conditionals that allow for no exceptions **Strict Conditionals**. Not all conditionals are strict, however. Sometimes people use these claims to express a strong conditional relation between two things, but one that nevertheless admits of exceptions. Suppose somebody says: ‘if you fall out of a third story window, then you will get hurt’. In making this claim, they are not ruling out the possibility you won’t get hurt, they are just saying that it is very likely. We will call conditionals that describe a regular pattern, but which nonetheless allow for exceptions, **Strong Conditionals**.

This distinction is important because determining whether a conditional is true or not depends on what kind of conditional it is. If an author has proposed a strict conditional as a premise, then we need to ask whether this relationship always holds without exception. If there is even one exception, then the strict conditional will be false. Consider the following dialogue:

Ex. 5:

Kelly: If a person lies, then they have done something that is wrong.

Elena: I disagree; think about white lies. If somebody asks if you like their outfit, and although you don't think so, you say 'yes', I don't think you have done anything wrong. Do you?

Kelly: Well, I guess not.

Kelly has put forth a strict conditional that says lying is always wrong. Elena is suggesting that Kelly's strict conditional is false, since it seems like there are exceptions—there are cases in which the antecedent is true (a person might lie), but the consequent false (it isn't wrong). Kelly agrees with Elena's claim, and consequently agrees that the original strict conditional is false. In general, a strict conditional will be false when there is even a single exception. More specifically,

A **Strict Conditional** is false when (and only when) there is a case in which the antecedent is true and the consequent false, i.e. an exception.

In everyday reasoning, strict conditionals are relatively uncommon. Most of the time we use strong conditionals; that is, most of the time we use conditionals to say only that the truth of the antecedent is usually accompanied by the truth of the consequent. For example, this is what we tend to mean when we say things like:

Ex. 6:

If you push the 'off' button, the computer will turn off.

Pushing the 'off' button usually turns the computer off, but there are exceptions. We know that it is possible—though unlikely—that the button is broken somehow. So when we make this claim, we are not saying that the truth of the antecedent is *always* accompanied by the truth of the consequent, only that most of time it is.

Given this, strong conditionals are false in different circumstances than strict conditionals. Such conditionals do not, after all, claim that there are no exceptions. Consequently, the fact that there is some circumstance in which the antecedent is true and the consequent false does not show that the conditional as a whole is false. It follows that:

A **Strong Conditional** is false when (and only when) the truth of the antecedent is not usually accompanied by the truth of the consequent.

Let's look at an example of a false strong conditional: suppose that a friend of yours will be doing some email correspondence with a person from Denmark, and worries that this person probably doesn't speak English on the basis of the following strong conditional:

Ex. 7:

If a person is from Denmark, then they probably don't speak English.

This strong conditional is false. The truth of the antecedent (a person's being from Denmark) is *not* usually accompanied by the truth of the consequent (not speaking English). In fact, quite the opposite: over 80% of Danes speak English.

In most cases it will be clear that an author intends one kind of conditional or the other. In ambiguous cases, however, we should assume the author intends a strong conditional. The reasons for this is grounded in

the *Golden Rule for Argument Interpretation*. An argument with a strict conditional is, all things being equal, more likely to be factually incorrect (and hence unsound), since all it takes is one exception for the premise to be false. So, if we are giving our interlocutor the benefit of the doubt, we should interpret their conditional claim as a strong one.

SECTION 4: A SPECIAL CLASS UNSOUND ARGUMENT—THE STRAW MAN

When discussing factual correctness, it is important to note a special class of arguments called “straw man arguments”. Perhaps the best way to start is with an example. At the first meeting of the year, Jasmin and Lauren are talking about issues the student council should take up this year:

Ex. 8:

Jasmin: The dorms are too loud at night. Some of us have early classes or practice, but on most nights there are people talking loudly and playing music until 1 or 2 in the morning. I'd like the council to think about some steps that we could take to limit noise at night.

Lauren: I disagree; college students are college students, and this idea that we are going to completely eliminate late-night noise is simply not realistic.

In response to Jasmin's concern Lauren has given a straw man argument. First, Lauren has misrepresented Jasmin's concern. Jasmin has asked the student council to consider some steps to limit noise, not to completely eliminate it. Second, Lauren has criticized the misrepresented concern (on the grounds that it is unrealistic). Thus, Lauren's straw man creates the impression that Jasmin's concern can be justly set aside, when really it hasn't been discussed at all. Hopefully, Jasmin will recognize this, and jump in to correct Lauren's misrepresentation.

In general, a **straw man argument** is an argument that draws a negative conclusion about a claim, viewpoint, or organization on the basis of a misrepresentation of it. A straw man looks like a real person from a distance, and is a lot easier to knock down. So too, it is often easier to reject a misrepresented position than the real one.



“Scarecrow” by Marxchivist CC BY 2.0

Importantly, straw man arguments can be unintentional—people sometimes misrepresent their opponents views because of because of misunderstanding or confusion. Whether intentional or not, straw man argu-

ments are common. They are especially common in heated public debates in which cooperative dialogue has broken down. Perhaps the most notorious straw man argument in recent public discourse comes from former Alaska Governor, Sarah Palin. In opposition to the Affordable Care Act, Palin posted the following on her Facebook page:

“The America I know and love is not one in which my parents or my baby with Down Syndrome will have to stand in front of Obama’s ‘death panel’ so his bureaucrats can decide, based on a subjective judgment of their ‘level of productivity in society,’ whether they are worthy of health care. Such a system is downright evil.” (August 7th, 2009)

Former Gov. Palin is completely and obviously right that such a system would be downright evil. We can all agree about that. The problem is that her argument is not factually correct; the argument grossly misrepresents the health care proposal it was criticizing. To be clear, to call Palin’s argument (or any other) a “straw man” is not to make a claim about the truth of the conclusion. It is only to say that this argument misrepresents the position it is criticizing, and so does not succeed in really criticizing that at which it is ostensibly aiming.

A different kind of straw man misrepresents a claim, view, or organization by quoting their own words out of context. A sentence can seem to have a very different meaning if viewed in isolation or out of its original setting, and people sometimes make use of this fact to criticize their opponents. Consequently, we will call these examples **Contextual Straw Man Arguments**. Most contextual straw man arguments are intentional misrepresentations. Here are a couple of cases from political history.

In 2011 presidential candidate Mitt Romney’s campaign released an advertisement which quoted President Obama as saying “If we keep talking about the economy, we’re going to lose.” The ad suggested that Obama himself recognized that in the 2012 election he would be vulnerable when it came to the economy. On its face, these were Obama’s own words. How could he have meant anything different? As it turns out, this quotation was taken from a speech during the 2008 campaign and Obama was actually quoting someone within his challenger’s campaign. Not only wasn’t Obama talking about the 2012 election, he was quoting his opponent!

During the same campaign the Democratic National Committee released a video which repeatedly showed Mitt Romney saying “I like being able to fire people” interspersed with video clips from movies and television of other people saying “you’re fired.” The video suggested that Romney was a callous jerk out of touch with average people. Again—these were Romney’s own words—what else could he have meant? As it turns out, Romney was not referring to people, but rather to insurance providers. He said:

“I want individuals to have their own insurance...That means the insurance company will have an incentive to keep you healthy. It also means if you don’t like what they do, you can fire them. I like being able to fire people who provide services to me. You know, if someone doesn’t give me a good service that I need, I want to say I’m going to go get someone else to provide that service to me.” (Jan. 9th, 2012)

Romney’s point was that choice in the marketplace is a good thing. His claim was that being able to switch from one insurance provider to another gives your provider an incentive to deliver high-quality service. One can disagree with Romney’s claim, but he certainly was not saying that he likes to fire people!

SECTION 5: RECOGNIZING AND AVOIDING STRAW MEN

A straw man's "natural habitat" is criticism, and we want to be on the lookout for them whenever a claim, viewpoint, or organization is criticized. Now, to spot a straw man argument is to spot a misrepresentation, and so the more one knows about a person or organization's view, the easier it is to spot misrepresentations of that view. Unfortunately, we do not always have this knowledge. In this case, there are a few rules of thumb to keep in mind.

First, if you don't know much about the idea or organization that is being criticized—look it up! To be sure, we do not have the time, interest, or energy to look up the subject of every criticism. However, in cases where it matters, we should not endorse an objection to a view without having good reason to believe that the representation on which it is based is accurate.

Second, ask yourself: "how plausible is the view the critic attributes to his or her opponent?" The less plausible the attributed view, the more suspicious you should be that the critic has misrepresented the view. To give just a few examples, if a person or group is represented as thinking that: criminals should be allowed to run free, that our national parks should be strip-mined, or that we should return the economy to the barter system, you should suspect a straw man. It is not that people or organizations never have extreme or implausible views, it is just that, again, we should be confident that they have the view before we criticize them for it.

Third, who is the critic? Is the critic a neutral third-party, or is the critic part of an organization or entity that is opposed to the idea or organization in question? As we will see, we are subject to a variety of biases (conscious and unconscious) that can lead us to misrepresent ideas, and this is especially the case if there is some incentive or benefit for doing so. Thus, the less neutral the critic, the more suspicious we should be that we are looking at a straw man argument.

Turning to straw man arguments in our own reasoning, it might seem easy, given the examples above, to avoid using them. We might think that all we need to do is to abstain from purposefully misrepresenting other's ideas. It is not so simple, however. The problem is that although some cases of straw man argumentation are intentional, many are not. As we will see, it can be very easy to accidentally and unintentionally misrepresent others' views. Sometimes this happens because the views in question are unfamiliar or complicated, other times this happens because we are unknowingly biased in one way or another. As such, we need to be careful before criticizing others' viewpoints that we have accurately represented them. One strategy for avoiding accidental straw man arguments is to ask yourself prior to developing your criticism:

■ *Would a proponent of this view accept my characterization of it?*

Asking this question also puts you in a position to develop a more nuanced criticism. After all, since you already have a proponent of this view in mind, it is easier to imagine how they might reply to your objection, and easier to avoid biased reasoning (as we will see in the next chapter).

EXERCISES

Exercise Set 8A:

#1:

List in descending order of confidence three propositions you believe are true.

Exercise Set 8B:

Directions: Drawing on your basic knowledge, determine whether the following conditionals are likely being used to express strict or strong conditionals. Write 'Strict' or 'Strong'.

#1:

If I am late, then my mom will be mad.

#2:

If a person has murdered another person, then they have killed another person.

#3:

If the Sears tower is the tallest building in the world, then it is taller than the Empire State Building.

#4:

If (s)he graduates from college with high honors (and is trying to get a job), then (s)he will get a good job right after graduation.

#5:

If you ever saw the Beatles play live, then you are older than 45.

#6:

If you eat expired yogurt, then you are going to get sick.

Exercise Set 8C:

Directions: Identify the misrepresentation in each problem. Then say how you would reply.

#1:

A: I don't think that the U.S should have invaded Iraq in 2003.

B: Oh, I see, so you don't support the troops. I, for one, think that it is our obligation as citizens to support the troops during wartime.

#2:

A: I am not in favor of legalizing recreational marijuana use. I just don't see many benefits, but a lot of costs.

B: Are you serious? What do you think...that everybody is going to walk around high all the time?

#3:

A: The new statewide standardized tests don't do a good job measuring what is really important. Moreover, it leads teachers to spend undue time teaching test-taking strategies.

B: Critics of the new statewide standardized tests think that schools do not need to be held accountable. But public schools are receiving taxpayer funds and the public needs to know what they are paying for.

#4:

The TSA recently announced that it will be using full-body scanners in many of America's busiest airports. These scanners are excellent tools for uncovering contraband without an intrusive personal search. Predictably, all kinds of people have objected to the use of these scanners, claiming that they violate their privacy. To the critics I say: I am sorry, that you are uncomfortable with your body, but public safety outweighs your body issues. I suggest that folks stop complaining, and start thanking TSA employees for their important work.

#5:

A: I won't be voting for Bill 891 because there are cheaper and more effective ways to monitor who and what is coming across the border.

B: Unlike A, I am interested in the security of the American people, and will do everything within my power to protect them from external dangers to their health and prosperity.

#6:

Proponents of gun control are totally unrealistic. They want to ban the use and ownership of all firearms in the United States so that not even the police will be able to carry guns, and hunters and target shooters will have to give up their sports.

#7:

A couple car shopping:

A: I like this one; it fits our budget and is pretty reliable.

B: True, but I am not convinced. It is just such a boring car in terms of style.

A: I'm sorry, I didn't know you were looking for a Ferrari—c'mon you know we can't afford that!

Exercise Set 8D:

#1:

What is the problem with Ben's reasoning in the following case?

Hansa: A train is coming, get away from the edge of the platform!

Ben: Why?

Hansa: Uh...because if you fall off the edge, you'll get hit by the train and die!

Ben: That is a poor argument Hansa, I mean, I heard that one time somebody fell off a platform, got hit by a train, and lived!

#2:

Do you recall having seen a Straw Man Argument? If so, explain. If not, create your own straw man argument. Try to make it something that a person might realistically offer.

Bias and Motivated Reasoning

SECTION 1: INTRODUCTION

We know how to unpack and analyze arguments and we know the relevant features for evaluating them. Now all we need to do is put our knowledge and skill to work, right? Well...no. Unfortunately, things are not so simple, since we are subject to a variety of biases and cognitive illusions that can make accurate and objective evaluation of arguments difficult. Although these illusions and biases are pervasive and often unconscious, we can adopt a number of strategies to limit their effects over us. In this chapter we will start with a consideration of one of the most pervasive forms of bias, *myside bias*, and see that simply having a preference for an idea can bias our thinking about it. We will then look at some of the factors that influence our preferences, and conclude with some techniques to mitigate the effects of bias.

SECTION 2: MYSIDE BIAS

We have all seen cases of biased thinking. Take for instance the serious sports fan that unfairly dismisses the possibility that her team might lose, or the political partisan who can't see past his own views, or the mother who simply won't believe that her child could have done such a thing. Although we see it in others, we tend to think that our judgments are, on the whole, fair and impartial. Other people are biased, but we are not. Unfortunately, we are often wrong about this. At least sometimes, we are not the neutral detached evaluators we think we are, and are guilty of the same biased thinking we see in other people.

In general, a **bias** is a preference that inhibits impartial evaluation. In this chapter we will focus on a broad category of bias, namely *myside bias*. **Myside bias** is the propensity to let our impressions, beliefs, and interests preferentially influence our evaluation of evidence. In general, we have a preference for views that conform to our existing beliefs and interests, and it can be difficult to separate our preferences from the evidence. What is especially troubling about *myside bias* is that biased reasoning can feel the same "from the inside" as unbiased reasoning! That is, the fact that we *feel* like we are being fair and neutral in our evaluation of the evidence is no guarantee that we are actually *being* fair and neutral.

The effects of *myside bias* can vary. In general, however, the more personally invested we are in the truth or falsity of some claim, the more difficult it is to objectively assess the evidence for or against it. When we are personally invested in the truth or falsity of some claim, we tend to focus selectively on information that supports our preferred view. Moreover, we tend to over-estimate the evidential value of this supportive information, and to under-estimate the evidential value of information that challenges our preferred view. For example, a person who has politically liberal beliefs will tend to pay closer attention to information that

supports his political views, and less to information that raises doubt about them. Furthermore, such a person will tend to see the evidence for his political beliefs as being stronger than it really is, and the evidence against his view as being weaker than it really is.¹

In this vein consider the old saying that, “a man who is his own lawyer has a fool for a client.” Unlike many proverbs, this one offers good advice and myside bias explains why; when your guilt or innocence is on the line it can be very difficult to objectively consider the merits and faults of the case against you, as well as the strength of your own defense. Thus it is best to step aside and let someone who has less personal involvement take over.

However, we need not have a strong personal investment in some claim in order for the effects of myside bias to kick in. The impartiality of our evaluations can be undermined by our own perspective in a variety of ways. A particularly infamous case of this is the justification for the U.S. led invasion of Iraq in 2003. In February of 2003 the U. S. Secretary of State, Colin Powell, stood before the United Nations and made a case for the invasion of Iraq. Central to this case was his assertion that Iraq was in possession of weapons of mass destruction. Soon thereafter Iraq was invaded, but no evidence of such weapons was ever found. A subsequent review of how the intelligence community could have made such a mistake came to the conclusion that what we have called myside bias played a big role. The review noted that:

The Intelligence Community (IC) suffered from a collective presumption that Iraq had an active and growing weapons of mass destruction (WMD) program. This...dynamic led Intelligence Community analysts, collectors and managers to both interpret ambiguous evidence as conclusively indicative of a WMD program as well as ignore or minimize evidence that Iraq did not have active and expanding weapons of mass destruction programs.²

According to this report, the intelligence community collectively expected Iraq to have such a program, and this shared expectation led to an unintentionally biased evaluation of the evidence, and ultimately contributed to a decision to take military action.

Consider another case. Scientists who study the efficacy of new drugs or treatments know that peoples' wishes, desires, and expectations can bias their studies. Obviously, a patient will want the new drug or treatment to work, and this can affect how they evaluate their own state. In this way, a patient might feel as if the drug or treatment is working even when it is not, for example. In part to limit these kinds of effects, researchers normally split subjects into two groups: a group that gets the experimental treatment and a group that does not. The group that does not get the experimental treatment gets a placebo instead (a treatment with no medicinal value, e.g. a sugar pill). Neither group knows whether they've received the experimental treatment or a placebo—they are “blind” to this factor as researchers put it. Since the test subjects do not know what treatment they've received, researchers can separate biased effects that are the result of peoples' preferences from effects that are the result of the treatment in question. (Note: we will talk more about the reasoning behind this kind of experimental design in the unit on Scientific Reasoning.)

Subjects are not the only possible source of bias in this kind of context. The researchers' own preferences for one result or another can lead to an unconsciously impartial evaluation of the evidence, as well. After all, researchers want these experimental drugs or treatments to work too. In order to prevent these biases from coloring the results, researchers can effectively blind themselves so that—like the subjects—they do not know who has received the placebo and who has not (until the end of the study). Studies conducted in this way are called **double-blind studies** since neither patients nor researchers know who is getting the treatment and who is merely getting a placebo.

Moreover, we know that double-blinding makes a difference. Partly to emphasize the importance of double-blinding, a group of researchers studying treatments for multiple sclerosis (a debilitating and currently incurable disease that attacks the nervous system) decided to conduct two versions of the same study: a blinded version and an unblinded one.³ The medical aim of the study was to discover whether a promising new treatment was really more effective than a placebo. Here is how they set up the two versions of the study: first, researchers divided patients into different categories: some patients received actual treatment while others received a placebo. Second, they split the researchers into two groups. One group knew which patients had received the placebo and which had not—that is, they conducted the study *unblinded*. The other group of researchers were *blinded*—that is, they were prevented from knowing who had received the treatment and who hadn't. Both blinded and unblinded researchers were asked to examine the test subjects over a period of 2 years at 6 month intervals.



"testing tube" by wader CC BY-NC-SA 2.0

Unfortunately, the treatments turned out to be no more effective than a placebo. However, on the whole, the unblinded researchers determined quite the opposite—they took the treatments to be effective! The unblinded researchers surely did not set out to make biased judgments. Presumably, they wanted to know whether the treatment was effective as much as anybody else and sought to be as objective as possible. Nonetheless, their assessment was biased, and this illustrates the deceptiveness of myside bias. Again, the problem is not just that our judgments can be biased by our own preferences, but that they can be biased *despite our sense that they are not!*

SECTION 3: SOCIAL INFLUENCES

Given that even slight preferences can, unbeknownst to us, influence our evaluations, it is important to have a sense for the forces that influence our preferences. Of course, we live in a complex world, and there are all kinds of factors that influence us. Nevertheless, perhaps the most important influence on our preferences is social. The crucial observation here is that beliefs can have **social value**. We first raised this in Chapter 2, where we noted that our desire to be seen by others in a particular way can undermine the goals of cooperative dialogue. In addition, the social value of an idea can create preferences that undermine fair evaluation. Let's see how this works.

The social value of a belief has to do with what other people will think about *you* upon finding out that you have it. More specifically, when somebody else thinks better of you because you have a particular belief, that belief may have positive social value for you, and when they think worse of you in virtue of it, the belief

may have negative social value for you. Whether a specific belief or idea has positive or negative social value for you depends on a number of factors that it is worth making explicit.

First, it is important to understand that not every idea has social value—in fact, many do not. For example:

Ex. 1

Noura: I think the grocery store should get new grocery carts.

Here Noura is expressing her opinion, and while people might agree or disagree with her about this, it is unlikely that anybody is going to think better or worse of Noura herself in virtue of this opinion.

Second, the social value of an idea depends on who is doing the judging. Suppose, for example that Noura says:

Ex. 2

I think that people who go to church on Sunday are pretty much wasting their time.

We can easily imagine that people's response to Noura's claim might vary. Suppose that upon hearing this Noura's friend Maddy responds with disappointment, saying "Geez, I didn't realize you had become so arrogant." Whereas another friend of Noura, Lori, responds in a different way saying, "At last, somebody brave enough to tell it like it is." In this case, Noura's belief will have positive social value for Noura *with respect to Lori*, but negative social value *with respect to Maddy*.

Third, the social value of an idea depends on whether, and to what extent, we *care* what the person who is doing the judging thinks. In general, we want to be well-regarded by other people, but this doesn't mean that we care what every single one of our acquaintances thinks of us. Moreover, we care to varying degrees. What your mom or best friend thinks probably matters a lot more to you than what your neighbor or dentist thinks about you. To illustrate: suppose that Noura has recently fallen out with Maddy, no longer regards her as a friend, and as a result doesn't particularly care what Maddy thinks. In this case, Noura's belief may have little to no social value with respect to Maddy—even though Maddy thinks less of Noura because of her belief.

How does the fact that ideas or beliefs can have social value relate back to the question of bias? The answer is that our desire to be well-regarded by people who matter to us can give us a *preference for* ideas and beliefs that we think might have positive social value for us, and a *preference against* ideas that we think might have negative social value for us. As we have seen, simply having a preference for an idea can unknowingly lead us to biased and unfair thinking. More specifically, this means that we will tend not to fully investigate or fairly evaluate ideas that have negative social value for us, and not be adequately critical of beliefs that have positive social value for us. This is something we need to keep in mind, so that, at least when it really matters, we stop to think about the extent to which our thinking is being influenced by our desire to be liked (or not disliked) by other people.

SECTION 4: PEOPLE AND IDEAS

Ideas can have social value for us because we care about what our friends and family think. However, we can turn this around and think about it from the opposite perspective as well. After all, just as you care what your friends and family think about you, so too your friends and family care about what you think about

them. As such, you contribute to the social value of ideas for your friends and family. That is, if you'll think less of one of your friends for believing X and they care about you think, then believing X may have negative social value for your friend. The fact that we can have this effect raises a question about what kinds of attitudes we should take towards people on the basis of their ideas and beliefs. How should we think about people when they disagree with us, propose dubious ideas, or ask questions we don't like? The short answer to this question is that we should only think worse of a person on this basis if we are justified in doing so. Of course, it is a much tougher question to settle when we are justified in doing so, and like many questions we will take up, there is no one-size-fits-all answer. However, there are a couple of observations it is worth keeping in mind.

First, recall from Chapter 2 that we have a tendency to unfairly vilify people who disagree with us, or who disagree with our in-group (a group of people we identify with as a member). That is, we have a tendency to unjustifiably think less of people in these circumstances. There is no doubt that there are people who are ignorant and unethical in ways that justify thinking less of them as people. However, it is important to emphasize that simply because somebody disagrees with us doesn't, all by itself, give much reason to think less of them as a person. After all, there are many explanations for why someone might disagree with you outside of some moral or intellectual flaw: it may be that you have relevant information or experiences they do not have. It may be that *they* have relevant information or experiences that you don't; alternatively, it may be that you have the same information, but draw on different but reasonable principles or values to evaluate it. Given this, in many cases, leaping to the conclusion that there is something wrong with a person or thinking less of them solely because they disagree with you is not only poor reasoning, but is also unfair to the other person.⁴

This brings us to the second point. We should be careful about thinking worse of people in cases of disagreement for purely self-interested reasons—namely because doing so can undermine our own ability to think clearly and fairly. After all, judgments about a person's character or intelligence often themselves embody and generate preferences for and against individual people. As we have seen, once we have even a slight preference for or against a particular person, the engine of myside bias can kick in with respect to *other things* that person says. Once a person disagrees with us, for example, we might begin to think less of them as a person. If we think less of them, then we may be more inclined to unfairly evaluate things they subsequently tell us. This can hurt us; after all, when we think less of others and thereby prematurely discount what they say, we can unnecessarily lose out on relevant information or points of view that might otherwise inform and improve our own thinking.

In sum, the fact that ideas can have social value raises a question about what we should infer about a person when they disagree with us. This is complicated, but given that (i) we have a natural tendency to unfairly think worse of those who disagree with us, and (ii) doing so can lead to us to think in biased ways, together suggest that we should be very careful about thinking less of a person on the basis of disagreement.

SECTION 5: PREFERENCES FOR PEOPLE—OTHER INFLUENCES

We have just seen that a preference for or against a person can influence how we think about what they say. As it turns out, this is very common, and happens in all kinds of ways and circumstances. People who have worked in sales, for example, know that if customers are favorably inclined toward them personally, they are more likely to accept what they say (all things considered); and if they are less inclined toward them, their customers will be less likely to accept what they say (again, all things considered).

In his book *Influence: Science and Practice* Robert Cialdini summarizes a number of the elements which can influence our feelings about other people.⁵ He begins by pointing out that we tend to be favorably inclined toward people we find physically attractive in some way. Cialdini notes that a great deal of psychological research has shown that we have a propensity to unconsciously attribute a wide-array of favorable traits to people we find physically attractive. We tend to see people we find as attractive in some way as being more intelligent, more honest, and more talented than we would otherwise.⁶ There is evidence that attractive job applicants are more likely to get hired than less attractive, *but equally qualified*, job applicants, and that attractive defendants (in criminal cases) tend to get lighter sentences than unattractive people in the same position. Moreover, there are a number of studies that suggest that physical appearance is an element in our decisions about who to vote for. In a recent study, Alexander Todorov and his colleagues presented people with pictures of candidates for the U.S. Senate and House of Representatives and asked them, on the basis of the picture alone, to rate each candidate's competence. He then compared the perceived competence of the candidates to election outcomes and found that the more competent looking candidate for Senate won 71.6% of the time, while the more competent looking candidate for the House of Representatives won 66.8% of the time.⁷ These results suggest something surprising: namely that the mere physical appearance of a candidate is contributing to our voting choices.

A second factor is similarity. We tend to be favorably inclined toward people who are similar to us in some notable way, e.g. dress, background, beliefs, hobbies, etc. Thus, we are likely to be favorably inclined toward people who are from the same area of the country as we are, or like the same music we do, or have the same political beliefs. On one level, this is probably not all that surprising. What may be surprising is the extent to which this influences our behavior. For example, Cialdini draws attention to a study which suggests that people are twice as likely to complete and return a survey if the survey is sent by a person with a similar name!

Third, a great deal of research in social psychology has shown that we tend to be favorably inclined toward people with whom we are working on a shared task. Thus, for example, players on a team tend to be favorably inclined toward their teammates. Again, this is probably not too surprising. What is important to note about this is that people can take advantage of this phenomenon by creating a shared task or goal for us. Cialdini notes that a common technique among car salespeople during price negotiations is to act as if they have taken your side against the sales manager. In doing so they have created a shared task: together you and the salesperson are working against the sales manager. As such, you will tend to be more favorably inclined toward the salesperson than you might have been otherwise, and consequently more likely to be less skeptical of his or her claims. Another example is the police tactic of playing Good Cop/Bad Cop during interrogations of suspects. In these cases, one officer plays the role of "Bad Cop" by adopting an actively hostile and suspicious stance towards the suspect. The other officer plays the role of the "Good Cop" by adopting a friendly and helpful stance toward the suspect. In this situation it is easy for the suspect to see the officer playing the Good Cop as an ally against the Bad Cop, and to consequently be favorably inclined and less guarded toward them.

As these examples show, we have a propensity to assume that when a person has one positive feature, they probably have other positive features as well. Psychologists refer to these kinds of associations as **halo effects**. The idea is that positive qualities radiate a halo that makes other features of a person look positive too (note that this works in reverse too—negative features can radiate a halo that makes other features look negative as well).



"22° moon halo tonight. With Jupiter, palm trees, and mountains." by slworking2 CC BY-NC-SA 2.0

It is important to emphasize that in these cases nobody is thinking to themselves: "Since he is handsome, I bet what he says is true too!" Put so explicitly we all recognize that this is a poor argument. Nor are most people consciously deciding who to hire, convict, or vote for on the basis of their appearance. Most of us recognize that making decisions on such a basis would be manifestly unfair. These halo effects, like the biases we have looked at, occur largely below the level of conscious awareness. Further, these factors do not dictate or determine our choices. The effects are much more limited. Nevertheless, they are still worrying. After all, factors like appearance and similarity are not relevant in most cases, and so should not influence our judgments and decisions.

In light of this, we need to keep in mind that when we walk away from a personal interaction with vague fondness for or aversion to a person, we may well be experiencing a halo effect based on their similarity to us, their attractiveness, their agreement with our ideas, etc. This is not to say that our interactions with other people cannot give us good evidence that they are honest or trustworthy or competent—they surely can. The problem is that our impressions are not always based on relevant features of people, and moreover that often we don't have any idea what features or cues our impressions have been based on in the first place. Ultimately, then, when it comes to important decisions involving other people (e.g. hiring) we should take the extra time to explicitly identify relevant factors as a way of limiting the impact of inappropriate halo effects.

SECTION 6: COUNTERING THE EFFECTS OF BIAS

As we have seen in our discussion of biases, we are often unaware of our preferences and unaware that they are biasing our reasoning. There is something disturbing about this; we think that we are in control of what we believe and decide, but these biases suggest otherwise. They suggest that our thinking and decisions are often influenced by external forces outside of our conscious awareness. When it comes to biases, then, the question is: how can we take control of our thinking and avoid biased thinking and decision-making?

The most important thing we can do is be aware that we are subject to these biases, and pay attention to cases where they are especially likely to be at work. Simply knowing that you have a real personal stake in a conclusion, and that as a result, you are likely to be biased towards it, can have dramatic effects. After all, this realization puts you in a position to monitor your own thinking, and to ask yourself whether you've

over-valued the evidence in favor of your preferred view, ignored relevant information, or under-valued the evidence against it.

This might sound like obvious advice, and it is—but it is advice that we tend not to follow! As we saw in Chapter 1, in day-to-day life we naturally respond to our environment in a variety of ways. Automatic and semi-automatic reasoning processes generate intuitive reactions to, and impressions of, people, places, circumstances, and ideas, and these impressions inform our conscious reasoning processes. Halo effects are good examples: it is a largely automatic reasoning process that gives us a favorable impression of a person that we find appealing or similar to us in some way. Many of the intuitions and impressions formed through these processes rise to the level of consciousness (though not all), at which point they become available for use in conscious reasoning processes. Think back to the ball and bat example from Chapter 1. When you first looked at the question, the answer ‘\$1.00’ presumably came immediately to mind. At the conscious level we can either accept the result of automatic processes, that is, we can take for granted our intuitions, impressions, and ways of thinking or we can treat them skeptically, as we should in the case of the bat and the ball. As it turns out, however, people rarely question their impressions, intuitions, and ways of thinking. We normally take our impressions for granted and use them as the starting points for our thinking. We do this partly because overriding our impressions and normal ways of thinking is hard work, and in general we tend to avoid expending mental energy if we can. The Nobel laureate Daniel Kahneman explains:

A general “law of least effort” applies to cognitive as well as physical exertion. The law asserts that if there are several ways of achieving the same goal, people will eventually gravitate to the least demanding course of action. In the economy of action, effort is a cost, and the acquisition of a skill is driven by the balance of benefits and cost. Laziness is built deep into our nature.⁸

Just as we tend to complete physical tasks using as little energy as possible, so too do we tend to avoid mental exertion if we can. Kahneman puts this in terms of laziness, but other psychologists have used different terms. Keith Stanovich, for example, makes this point by saying that we are **cognitive misers** who are stingy with our energy. Stanovich makes a particularly effective case for the importance of learning to distinguish and evaluate our automatic impressions and ways of thinking. He writes:

Humans are cognitive misers because their basic tendency is to default to [automatic] processing mechanisms of low computational expense...Nevertheless, this strong bias to default to the simplest cognitive mechanism—to be a cognitive miser—means that humans are often less than rational. Increasingly in the modern world, we are presented with decisions and problems that require more accurate responses than those generated by [automatic] processing. [These] processes often provide a quick solution that is a first approximation to an optimal response. But modern life often requires more precise thought than this. Modern technological societies are in fact hostile environments for people reliant on only the most easily computed automatic response. Think of the multimillion-dollar advertising industry that has been designed to exploit just this tendency...When we are over-reliant on [automatic] processing we lose personal autonomy. We give up our thinking to those who manipulate our environments, and we let our actions be determined by those who can create the stimuli that best trigger our shallow automatic processing tendencies.⁹

Stanovich’s point is that it is important that we learn to question our intuitive responses and ways of thinking—i.e. to learn that we cannot always trust the way things seem to us. Not only are our impressions sometimes inaccurate, but people can take advantage of these largely automatic processes to manipulate us.

The obvious solution to this problem is to treat our impressions more skeptically. Of course, we cannot do this all of the time. We have too much to do, our environment changes too fast, and we do not have enough energy to stop and think about all our impressions. The simple fact is that we will have to trust most of our intuitive responses and ways of thinking. However, when something important is at stake we should slow down and carefully identify and evaluate the impressions that drive our thought. We do so by track-

ing down and identifying the sources of our impressions: “What about this person or product or situation is striking me as good or bad?” “Are those features really indicators of goodness or badness?”. We will not always be able to isolate these features, but in taking the time to briefly ask ourselves these questions we thereby exercise more control over our thinking and decision-making than we would have otherwise.

In addition to simply being aware of the possibility of bias, what other steps can we take to mitigate bias? The most important step you can take is to separate yourself from the inquiry or idea. In some cases, you can literally do this, as in the case of the lawyer who trusts his defense to someone else, or the researcher who blinds herself to the experimental identity of the subjects. This strategy, however, is not realistic for most everyday situations. A second strategy for mitigating bias is to discuss the issue with other people—that is, to seek cooperative dialogue. A person without the same set of preferences may be able to see weaknesses (or strengths) that you would have had a hard time seeing on your own.

Unfortunately, cooperative dialogue is not always an option either. In this event, the best we can do is try to separate ourselves from the idea or claim—in our imagination. The idea is to try to envision the claim in question from a critic’s perspective.¹⁰ How would the critic object? Put otherwise, the idea is to play **devil’s advocate** with your own views. There are a number of ways to do this. One way to do this is to imagine that you have to defend your view in front of a critical audience. What objections would such an audience raise? How would they argue against your position? What evidence would they draw upon? An alternative is to pretend that you have the opposite view to the one you actually have. What would you say on behalf of “your” view? How would you criticize your opponents? Alternatively, you might imagine a future in which it turns out your view was mistaken or your decision the wrong one. You can then think about the information you would have to gather in order to understand why you made the mistake. Once you have that information in hand, you can take it into account in coming to your view or making your decision.

In sum, taking a different perspective in these ways can reveal unseen strengths and weaknesses in our preferred view. There is no doubt that playing devil’s advocate with your own views is hard cognitive work, but when we want to know the truth about important matters it can be well worth the effort.



“Contrarian” by drodesign CC BY-NC-ND 2.0

SECTION 7: BUBBLES AND ECHO-CHAMBERS

In closing this chapter, it is important to point out some of related situations that are wider in scope. On this score, philosopher C. Thi Nguyen makes a useful distinction between “epistemic bubbles” and “echo-chambers”.¹¹ Let us start with epistemic bubbles. The term ‘epistemic’ means having to do with knowledge, and

an **epistemic bubble** refers to an information source or network that omits, overlooks, or filters out relevant facts, arguments, or perspectives.

In order to illustrate this point, consider the fact that our social media accounts tend to be connected with friends or people we like, respect, or are positively inclined toward in some other way. In itself, this is not problematic. The problem is that many people get their news, analysis, and commentary from social media as well. Because people we like tend to be similar to us, the information, analysis, and commentary we get from social media tends to be skewed toward our existing beliefs. That is, our social media feeds can act as a filter that screens out information and perspectives that do not fit with our existing beliefs or the beliefs of people who are similar to us. As Nguyen points out, friends make for good parties, but not necessarily good information networks. Further, we may not be aware that our information has been filtered, since we don't see the information or arguments we are missing. Beyond social media, the proliferation of news channels and the 24-hour news cycle make it easy to pursue social, political, and economic, news and commentary that coheres with our existing beliefs and values (almost whatever they are!). This need not be intentional. As we saw in our discussion of myside bias, an inclination toward an idea or perspective can start up the engine of myside bias and lead us to look for and pursue information that fits our existing system of beliefs.

To be in an epistemic bubble is to lack certain information, arguments, or perspectives. An **echo chamber**, on the other hand, is a community that actively undermines dissenting voices. Within such a community, dissenting voices are unfairly and pre-emptively dismissed as insincere, corrupt, or otherwise untrustworthy. A person in an echo chamber may well be exposed to dissenting perspectives, arguments, and data; but they will set these opposing views aside since, according to trusted voices within the echo chamber, these views come from unreliable sources. Given this description, we can see that echo chambers are simply an extension of people's predisposition to vilify those who disagree with them. However, in an echo-chamber this tendency is amplified, and has an isolating effect, since it insulates members of the community from questions and information that challenge the echo-chamber's views.

Importantly, there is nothing wrong with setting aside a person's claims because they are being insincere or because they are an unreliable source on the issue. The problem with an echo-chamber is that it does so unfairly. Dissenting voices are undermined and dismissed simply because they question or dissent from the echo-chamber's views. That is, in an echo-chamber voices are set aside or disparaged when there is no good prior reason to doubt them. In addition, we can think of both epistemic bubbles and echo-chambers as coming in degrees. When it comes to echo chambers, a community can go to greater or lesser lengths to discredit opposition, and similarly a person within an echo-chamber can be more or less insulated from inconsistent information or dissenting voices. As a final point, echo-chambers can grow up around all kinds of phenomena. Nguyen gives as examples communities centered around, political positions, specific diets and exercise programs, activism, child-rearing techniques, and marketing programs.

Avoiding epistemic bubbles and echo-chambers is not complicated if we are paying attention. In the case of epistemic bubbles, we need to make sure to get our information from varied sources, and to pay special attention to other people's sources of information. This makes sense particularly when people disagree with us. This is an opportunity to ask: where is this opposing information coming from? Is it a good source? When it comes to echo chambers, we need to pay attention to how a community handles dissent. Does it consider new information or perspectives or does it seek to vilify those who raise questions? Does the community fairly represent criticisms, or do they misrepresent them (straw man) or mock them?

EXERCISES

Exercise Set 9A:

Directions: In order to practice combating bias, briefly lay out the strongest case a critic might make against each of the following. Note: this might be uncomfortable.

#1:

The legal drinking age should be lowered to 18.

#2:

Each person has only one true soulmate.

#3:

Corporations only care about profit.

#4:

Firefighters are heroes.

#5:

Banning books always amounts to unacceptable censorship.

#6:

Standardized testing is a waste of everyone's time.

Exercise Set 9B:

#1:

Some of the claims in exercise set 9A are provocative. Why do you think you were asked to consider those claims in particular?

#2:

Evaluate the following argument:

1. It seems to me that I have impartially evaluated the evidence in this case.
2. So, I probably have.

What do you make of it?

#3:

Outside of sports and political contexts, where have you seen biased reasoning? Give at least one example.

#4:

As we've seen, ideas, questions, and arguments can have social value. How can this contribute to the generation of echo-chambers in particular?

#5:

What sources of news or information do you trust? Are there sources you don't trust, but that others do, or vice versa? In general, what features or characteristics of a source give you reason to trust it, or not trust it?

#6:

Have you ever seen, experienced, or heard of anything like an echo chamber as defined above? Explain.

Notes

1. Lord, C. G., Ross, L., & Lepper, M. R. (1979). "Biased assimilation and attitude polarization: The effects of prior theories on subsequently considered evidence." *Journal of Personality and Social Psychology*, 37(11), 2098–2109.
2. U.S. Congress, Senate Select Committee on Intelligence, *Report on the Select Committee on Intelligence on the U.S. Intelligence Community's Prewar Intelligence Assessment on Iraq*, 18.
3. Noseworthy, J. H. et al. (1994). "The impact of blinding on the results of a randomized, placebo-controlled multiple sclerosis clinical trial." *Neurology*: 44, 16-20.
4. Garcia and King call this kind of inference the "attitude-to-agent fallacy". See Garcia, Robert K. and King, Nathan L. (2016), "Towards Intellectually Virtuous Discourse: Two Vicious Fallacies and the Virtues that Inhibit Them" in *Intellectual Virtues and Education*. Ed. Jason Baehr. New York: Routledge, 202-220.
5. Cialdini, R. B. (2008). *Influence: Science and practice, 5th ed.* New York: HarperCollins College Publishers.
6. This is a general tendency. Indeed, people who are typically regarded as very attractive in some way may be more likely to be regarded unfavorably in some contexts.
7. Todorov, Alexander T., Mandisodza, Anesu, M., Goren, Amir, and Hall, Crystal C. (2005). "Inference of Competence from Faces Predict Election Outcomes." in *Science* 308 (5728), 1623-1626.
8. Kahneman, Daniel. (2011). *Thinking, Fast and Slow*. New York: Farrar, Strauss, and Giroux, 35.
9. Stanovich, Keith E. (2009). *What Intelligence Tests Miss: The Psychology of Rational Thought*. New Haven: Yale University Press, 29-30. Stanovich distinguishes between Type 1 and Type 2 systems of reasoning. Since I have not adopted this same terminology, I have substituted the term 'automatic' in parentheses.
10. Lord, C. G., Lepper, M. R., and Preston, E. (1984). "Considering the Opposite: A Corrective Strategy for Social Judgment." *Journal of Personality and Social Psychology*, 47 (6), 1231-1243.
11. Nguyen, C. Thi. (2020). "Echo Chambers and Epistemic Bubbles" in *Episteme*, 17, 2: 141-161.

Unit #3 Summary

In this unit we took a closer look at the process of evaluating arguments. We began by discussing how to apply the standards of factual correctness and logical strength. We then characterized different levels of logical strength, and introduced the related terminology of 'deductive' and 'inductive'. We then turned to factual correctness, focusing on two potential complications. First we looked at conditional claims, and discussed when claims like this are true and false. Second, we looked at a common kind of fallacious argument—the Straw Man. Last we looked at how bias can undermine our reasoning, and identified some techniques for mitigating its effects.

KEY TERMS

- Soundness
- Factual Correctness
- Logical Strength
- Undetermined Arguments
- Deductive Arguments
- Inductive Arguments
- Inductively Strong
- Inductively Weak
- Not Factually Correct
- Undetermined
- True
- Standard Form for Conditionals
- Antecedent
- Consequent
- Strict Conditionals
- Strong Conditionals
- Ambiguous Conditionals
- Straw Man Argument
- Contextual Straw Man
- Bias
- Myside Bias
- Double-Blind Studies
- Social Value
- Halo Effects
- Cognitive Miser
- Devil's Advocate
- Epistemic Bubble
- Echo-Chamber

FURTHER READING

For a broad overview of the psychological literature on motivated reasoning see: "Motivated Cognition in Self and Social Thought" by David Dunning in *APA Handbook of Personality and Social Psychology*. For more on associations and some interesting examples see Chapter 1 of *Mindware: Tools for Smart Thinking* by Richard E. Nisbett. If you are interested in learning about how advertisers use halo effects and other techniques to persuade see *Sold on Language: How Advertisers Talk to You and What this Says about You* by Julie Sedivy and Greg Carlson. Finally, for an accessible discussion of epistemic bubbles and echo chambers see Nguyen's contribution "Escape the Echo Chamber" in the online magazine *Aeon*.

UNIT IV

AN INTRODUCTION TO DEDUCTIVE ARGUMENTS

Deduction and Argument Form

SECTION 1: INTRODUCTION

To this point we have focused on learning how to identify and represent arguments, and have introduced the basics of argument evaluation. However, the practice of putting these evaluative principles to work can vary from one kind of argument to another. The remainder of this book is largely devoted to learning how to evaluate some of the most common and important kinds of arguments. As we saw in Chapter 7, one way we can classify arguments is on the basis of their logical strength. Recall that deductive arguments have maximal logical strength: the premises, if true, guarantee the truth of the conclusion. In contrast, inductive arguments have less than perfect logical strength. In Unit #5, we will focus in on inductive arguments. In this unit, on the other hand, we will take a closer look at deduction. Although deductive arguments are not particularly common in everyday life, they have a number of distinctive features. Clarifying these features will not only help us understand deduction, but will give us a deeper understanding of argument evaluation more broadly. In this chapter, we will start by identifying one distinctive feature of deductive arguments in particular, indefeasibility, and we will discuss its relation to argumentative form. We will then identify some common deductive argument forms, and end the chapter with a look at some argumentative forms that people commonly confuse with deductive ones.

SECTION 2: DEFEASIBLE AND INDEFEASIBLE ARGUMENTS

Most arguments are defeasible. To say an argument is **defeasible** is to say that its logical strength is sensitive to new information. Consider the following example.

Ex. 1:

Dominic: How are you going to get to the airport tonight?

Zoe: The interstate seems like the best option, since normally traffic is pretty light at that time of the evening.

Dominic: But the baseball game will be finishing right about the time you'll be on the road, and all those people will be trying to get home.

Zoe: Oh right; I'll guess I'll take the back way.

Let us set aside numbers for a moment and represent Zoe's initial inference like this:

Premise 1: Typically, traffic on the interstate is pretty light at that time of evening.

Conclusion: So, taking the interstate is the best option.

Though we don't know all the relevant details, Zoe's initial inference seems like a reasonably good argument. How does adding the new information about the baseball game affect Zoe's reasoning?

The new information Dominic adds into the mix does not show that the premise in Zoe's argument is false, and so doesn't challenge the factual correctness of the argument. Rather the new information undercuts the logical strength of Zoe's initial inference. The fact that the baseball game will be finishing at that time means traffic will probably be pretty heavy (though normally it is light at that time). As a consequence, taking the interstate is *not* the best option. Zoe's final argument looks like this:

Premise 1: Typically, traffic on the interstate is pretty light at that time of evening.

Premise 2: The baseball game will likely be finishing about the time you'll be on the road, and all those people will be trying to get home.

Conclusion: So, taking the interstate is *not* the best option (and the back way is).

Zoe's initial inference was defeasible: although considered by itself P1 was good evidence for C1, this inference was sensitive to new information. Indeed, in this case the introduction of P2 undermined the evidential support P1 gave to C1, so that when taken together P1 and P2 do not give us good reason to believe C1 (as Zoe's reasoning shows).



"Dodger Stadium-Downtown L.A." by kla4067 CC BY 2.0

Here is another example. Suppose a detective is investigating a murder. In the course of the investigation she discovers that the murder weapon was owned by the victim's sister, the sister's fingerprints were found on the gun, and the sister had a strong motive. These facts strongly suggest that the sister is the murderer. But again, this argument is sensitive to new information. Say the detective subsequently discovers, for example, security camera footage showing she was 25 miles away at the time of the murder. This piece of information completely undermines the previous argument. The addition of this new evidence doesn't undermine the argument by showing that one of the premises is false; rather the premises no longer support the conclusion given this addition.

We are constantly updating our beliefs in light of new information, and anytime we do so we are thinking in terms of defeasible arguments. What does all of this have to do with deductive arguments? Introducing the concept of defeasibility is useful for highlighting a distinctive feature of deductive arguments, namely that

they are not defeasible. That is, *deductive arguments are indefeasible*. An argument is **indefeasible** when (and only when) its logical strength is not sensitive to new information. Add as much relevant information as you want to a deductive argument; the information will never change the logical support the original premises lent to the conclusion. To be clear, new information *can* challenge or contradict the truth of a premise in an indefeasible argument; that is, we can certainly discover that a premise we initially thought was true, is actually false. Thus, while new information can lead us to change our mind about the factual correctness of an indefeasible argument, it can never lead us to change our mind about its logical strength. This should strike you as odd. How can an argument be insensitive in this way?

The answer falls right out of the definition of a deductive argument. If the truth of the premises genuinely *guarantees* the truth of the conclusion, then nothing anyone might learn can change the logical support the premises give to the conclusion. After all, that is what ‘guarantee’ means in this context: in a deductive argument it is impossible for the premises to be true and the conclusion false.

SECTION 3: DEDUCTION AND ARGUMENT FORM

As we just saw, deductive arguments are indefeasible on account of the fact that in a deductive argument the truth of the premises guarantees the truth of the conclusion. This raises a question: namely, how do they do this? That is, how do deductive arguments guarantee the truth of the conclusion (if the premises are true)? The answer is that in a deductive argument the conclusion is already contained within the premises. By granting the premises you thereby automatically grant the conclusion. Another way of putting this is that deductive arguments are indefeasible because of the way information is arranged within them—that is, through their *form*. Let’s look at an example of a deductive argument.

1. If Alpha Centauri is more than a light-year away from Earth, then it is further away from Earth than Jupiter.
2. Alpha Centauri is more than a light-year away from Earth.
3. Alpha Centauri is further away from Earth than Jupiter.

To see this argument’s form, we need to simplify its content. To this end, we will replace each independent clause with a variable (we will use capital letters, A, B, C, etc.). The first premise is a conditional claim that contains two independent clauses. The first is in the antecedent and is prefaced by the word ‘if’: *Alpha Centauri is more than a light-year away from Earth*. Call this A. The second independent clause is in the consequent of the conditional, and is prefaced with the word ‘then’: *it [Alpha Centauri] is further away from Earth than Jupiter*. Call this B. By substituting the variables into the argument, we can see its form:

1. If A, then B.
2. A.
3. B.

It is because the argument exhibits this basic structure that it is deductive. It may be surprising, but no matter what information you substitute for the variables ‘A’ and ‘B’, the resulting argument will be deductive if the information is presented in this pattern. This particular argumentative form is very common and is called **Modus Ponens**.

Before looking at some other deductive forms, it is important to make two notes. First, for deductive forms the order of the premises doesn't matter. A standardized argument like the one above claims only that, when taken together, 1 and 2 imply 3. It does not say that a person has to discover or think 1 first and 2 second (or vice versa). Thus, the following is also an instance of the form *Modus Ponens*:

1. A.
2. If A, then B.
3. So, B.

Second, many common deductive forms have conditional premises. This raises a question about what kind of conditional is at work here: strict or strong? Recall that in Chapter 8, we said that strict conditionals claim that the truth of the antecedent is always accompanied by the truth of consequent, no exceptions! (e.g. if a shape is a square, then it is a rectangle). Whereas strong conditionals claim only that the truth of the antecedent is usually or mostly accompanied by the truth of the consequent (e.g. if you jump out a three-story window, then you will get hurt). Conditionals in deductive arguments are always strict. They have to be: think about the case of *Modus Ponens* above. If the conditional were only strong and at least sometimes A is true without B, then knowing A is true wouldn't guarantee that B is true as *Modus Ponens* asserts.

SECTION 4: COMMON DEDUCTIVE FORMS

While there are many deductive argument forms, we will isolate here only the most common ones. We have already seen the form known as *Modus Ponens* (MP). A similar deductive form is called **Modus Tollens** (MT). Arguments of this kind have the following form (note: the symbol '~' is shorthand for 'not').

1. If A, then B.
2. ~B.
3. ~A.

Arguments of this form are deductive because the strict conditional tells us that if the antecedent of the conditional is true, so is the consequent. As a result, if the consequent is not true, then neither is the antecedent. To see this, consider the following example:

1. If Kendall has been convicted of murder, then Kendall is a felon.
2. It is not true that Kendall is a felon.
3. So, it is not true that Kendall has been convicted of murder.

A similar argumentative form makes use of two strict conditionals. These arguments are called **Hypothetical Syllogisms** (HS) and have the following form:

1. If A, then B.
2. If B, then C.
3. So, if A, then C.

The term 'hypothetical' is referring to the fact that both premises are conditionals and tell you will happen in a hypothetical situation. Here is a hypothetical syllogism with content:

1. If Kendall has been convicted of murder, then Kendall is a felon.
2. If Kendall is a felon, then she cannot vote.
3. So, if Kendall has been convicted of murder, then Kendall cannot vote.

One last form is called a **Disjunctive Syllogism** (DS). A disjunctive syllogism is perhaps the most obviously deductive of the argument forms discussed here. These arguments have the following form:

- 1) Either A or B.
- 2) $\sim A$.
- 3) So, B.

'Disjunction' is the grammatical term of an 'or' statement (and the terms of disjunction are called "disjuncts"). We have seen examples like this in previous homework. Consider the following:

1. Either the butler committed the murder or the judge did it.
2. It was not the butler.
3. So, the judge committed the murder.

The first premise is a disjunction, and its disjuncts are 'the butler committed the murder' and 'the judge committed the murder'. It is important to be clear that there are many deductive argument forms besides these. Moreover, these argument forms can be strung together into more complicated deductive arguments. For example, imagine the conclusion of a Hypothetical Syllogism is If C, then D, and the conclusion of a Disjunctive Conclusion is C. Bringing these two pieces of information together you could deduce, by Modus Ponens, that D is true.

SECTION 5: ILLUSORY FORMS

There are a number of argumentative patterns that resemble genuinely deductive forms and are commonly (but mistakenly) seen to be deductive forms themselves. The first of these is called **Denying the Antecedent** (DA). Arguments of this kind have the following form:

1. If A, then B.
2. $\sim A$.
3. $\sim B$.

Instances of denying the antecedent are not deductive arguments because the premises do not guarantee the truth of the conclusion. Strict conditionals say that if the antecedent is true, then so is the consequent. It does not tell us what is true when the antecedent is false. For example:

1. If Will is enrolled in Biology 101, then Will is enrolled in a science course.
2. Will is not enrolled in Biology 101.
3. So, Will is not enrolled in a science course.

If the premises are true, they do not guarantee the conclusion because there are many different science courses Will might be enrolled in besides BIO 101 (e.g. Chemistry 101, Astronomy 101).



"Illusory House II" by Frankenstein CC BY-NC 2.0

The second illusory argument form we will discuss is called **Affirming the Consequent (AC)**. Arguments of this kind have the following form:

1. If A, then B.
2. B.
3. So, A.

Instances of affirming the consequent are not deductive forms because the strict conditional only tells us that if A is true, then B is too. It doesn't tell us what is the case when B is true. Let us turn to Will again:

1. If Will is enrolled in Biology 101, then Will is enrolled in a science course.
2. Will is enrolled in a science course.
3. So, Will is enrolled in Biology 101.

Again, the premises, if true, do not guarantee the truth of the conclusion for the same reason. There are many different science courses Will might be enrolled in, and the fact that he is enrolled in a science course doesn't tell us that it must be Biology 101. Importantly, there are many other non-deductive argument forms besides these two. We have focused on Denying the Antecedent and Affirming the Consequent because they are particularly deceptive. That is, many people treat arguments with these forms as if they were deductive, even though they are not.

EXERCISES

Exercise Set 10A:

Directions: Determine whether each of the following is true or false, and explain your answer.

#1:

A deductive argument can have a false premise.

#2:

A deductive argument can have a false conclusion.

#3:

A deductive argument can have all false premises and a false conclusion.

#4:

A deductive argument can have all false premises and a true conclusion.

#5:

A deductive argument can have all true premises and a false conclusion.

Exercise Set 10B:

Directions: For each of the following, first, determine whether the argument is an instance of one of the forms identified above, and second whether the argument is deductive or not. Assume that the conditionals are all strict. Note: some arguments may not be of any of the kinds we have discussed so far.

#1:

If Zala is a senior, then she can register for this class, but since she is not a senior, I guess she cannot register for the class.

#2:

If an infectious disease has an R_0 value less than 10, then it is less contagious than chickenpox. This virus has an R_0 value less than 10, and so the virus is less contagious than chickenpox.

#3:

If the proposed explanation is correct, then the specimen will dilate. Therefore, since the specimen did not dilate, we know the proposed explanation is incorrect.

#4:

Your problem has got to be a bad fuel pump, since it is either a bad fuel pump or bad filter and I know it isn't the filter.

#5:

If the state adopts water restriction policies, then you won't be allowed to water your lawn, and if you aren't allowed to water your lawn, then it will die. So, if the state adopts water restriction policies, then your lawn will die.

#6:

I strongly suspect the head of the accounting office has been cheating the company somehow. I know how much she gets paid, and she has been living beyond her means for years.

#7:

If proposed explanation is correct, then the specimen will dilate. Since the specimen did dilate, we can conclude that the proposed explanation is correct.

#8:

If your business is certified by the SBA, then it is eligible for tax credits on the installation of new solar panels. Since your business is not certified by the SBA, then it is not eligible for these tax credits.

#9:

If public high school education in this state were truly equitable, then all high-school students would have equal access to AP courses in History, English Language and Composition, and Calculus. But less than $\frac{1}{2}$ the high schools in the state offer students AP courses in these three subjects. It follows that public high school education in this state is not equitable.

Exercise Set 10C:

#1:

Given what we've learned in this chapter, what do you make of the following inference?

If Japan makes its airports more convenient for the Japanese public, Japanese tourism in Europe will increase. If Japanese tourism in Europe increases, then either European facilities in key centers will be enlarged or crowding in key centers will occur. Japan is making its airports more convenient for the Japanese public, but European centers are not expanding their tourist facilities. So we can expect crowding in main European tourist centers.

#2:

Suppose that your team is going to a tournament, and you need to find somebody to drive the college van. The driver has to be at least 22 and have an unrestricted driver's license. A friend of a friend, suggests Sam. You talk to her briefly and she claims she meets both conditions, and you reason as follows:

1. Sam has an unrestricted driver's license and is at least 22 years old.
2. If Sam has an unrestricted driver's license and is at least 22 years old, then she is eligible to drive the college van.
3. So, Sam is eligible to drive the college van.

First, what form is this deductive argument? Suppose that you subsequently learn from Sam's roommate that she actually doesn't have an unrestricted license (despite what she said). This will surely lead you to revise your earlier thinking, but doesn't this show that that not all deductive arguments are defeasible? After all, a new piece of information led you to change your mind. Is this really an example of a defeasible deductive argument? Why or why not?

The Language of Deduction

SECTION 1: INTRODUCTION

In the last chapter, we identified some distinctive features of deductive arguments, and saw that the unique logical strength of deductive arguments is a result of their form. We then looked at a number of common deductive forms, and closed the chapter by identifying two argument forms that are often mistakenly taken to be deductive. In this chapter, we will focus on the language of deduction, specifically on the language of conditionals. As we saw in the last chapter, strict conditionals play an important role in a number of common deductive argument forms. However, we can express conditional relationships in a variety of ways, and this means that we can express common deductive arguments in a variety of ways. As a result, we will be in a better position to spot deductive argument patterns when we run across them. Moreover, some of the lessons we will take from our considerations here will be relevant when we turn to a consideration of the language of inductive arguments.

SECTION 2: WHAT STRICT CONDITIONAL STATEMENTS SAY (AND WHAT THEY DON'T)

As we have seen, conditionals are common in everyday reasoning. We can use conditionals in all kinds of ways, e.g. to make promises, threats, or offers, to talk about causal relationships, or to talk about the way things might have been (e.g. “if the Broncos had won Super Bowl XLVIII ...”). Given this, let us distinguish between conditionals which are used to describe existing relations—call them **conditional statements**—and conditionals that are used for other purposes (e.g. conditional promises, conditionals threats, conditional offers, counter-factual conditionals). For the sake of thinking about deductive arguments, we will set aside these other kinds and focus solely on conditional statements. Thus, we will focus on conditionals like the following:

- If sugar is put in the water, then it will dissolve.
- If you are driving without a license, then you are breaking the law.
- If the bird is a mallard, then it is a member of the genus *Anas*.

And not on conditionals like:

- If you clean the garage, then I'll lend you the car.
- If you don't hand over your wallet, then I'll take it from you.

To this point we have discussed conditional statements as if they were always expressed in the following way: ‘if...then...’. We have called this the Standard Form for conditionals, and have identified the ‘if’ part as the antecedent, and the ‘then’ part as the consequent. However, everyday language is a rich vehicle for expression, and gives us many ways to convey or articulate ideas. Conditional claims are no exception. Because of their importance to deductive arguments, we will identify and discuss some of the most common alternative ways of expressing conditionals.

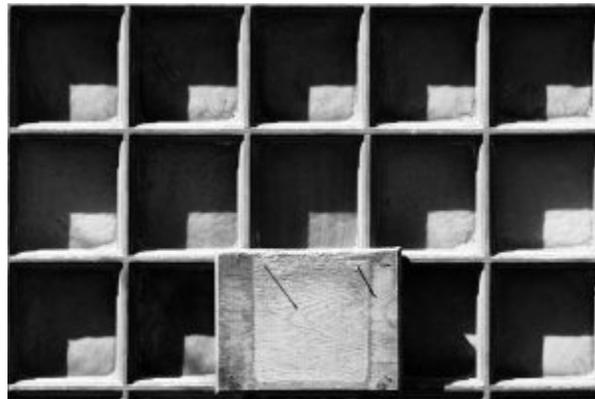
Towards this end, let us begin by thinking about the relationship(s) that a strict conditional statement expresses. Let’s start with an example we’ve seen before:

Ex. 1:

If a shape is a square, then it is a rectangle.

This strict conditional is claiming that anytime the antecedent is true, and there is a shape that is a square, the consequent is also true: the shape is a rectangle as well. This generalizes to all strict conditionals; they claim that the truth of the antecedent is always accompanied by the truth of the consequent—no exceptions.

Let us say that a strict conditional claims that the antecedent is **sufficient** for the consequent because the conditional is saying that knowing the antecedent is true is sufficient for knowing that the consequent is true too, since the one is always accompanied by the other. If the relationship between the antecedent and the consequent in a strict conditional is one of sufficiency, what is the relationship between the consequent and the antecedent? It can be tempting to think that the relationship from the consequent to the antecedent is the same as the antecedent to the consequent, but the example shows us that this is not the case. If a shape is a square, then it is a rectangle, but just because it is a rectangle doesn’t mean it is a square. After all, a square is only one kind of rectangle.



“squares and rectangles” by tommpouce CC BY-NC 2.0

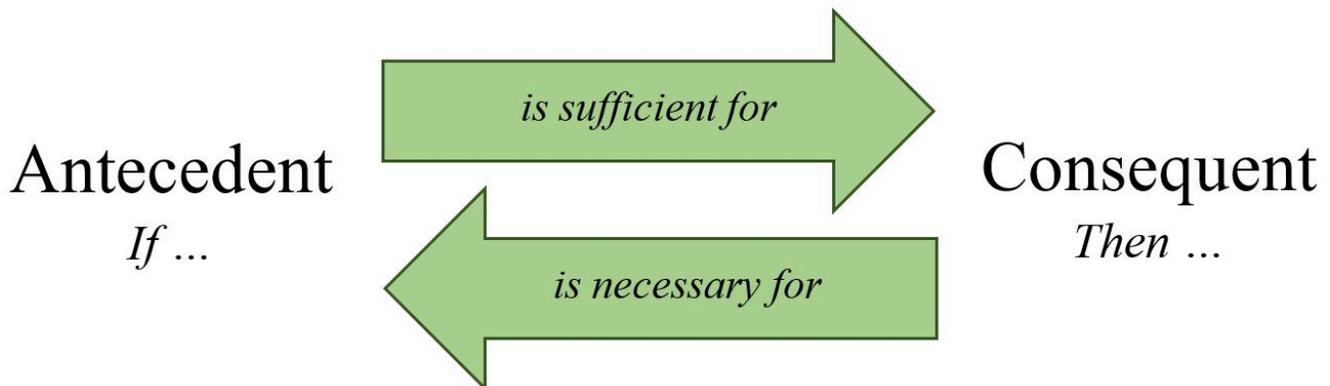
So what is the relationship that the consequent has to the antecedent? Go back to the conditional ‘If a shape is a square, then it is a rectangle.’ This tells us that squares are rectangles, but it also tells us that shapes that are not rectangles also are not squares. After all, since a square is a kind of rectangle, if you knew that shape was not a rectangle you would automatically know that it couldn’t be a square. In other words, it tells us that being a rectangle is crucial to being a square. We can say, then, that in this case, the truth of the consequent is required, or is necessary for, the truth of the antecedent. This conclusion generalizes to all strict conditionals. In light of this, we will say that in a strict conditional the consequent is **necessary** for the antecedent.

The fact that the relationship the antecedent has to the consequent is not the same as the relationship the consequent has to the antecedent it is worth emphasizing. According to psychologists and linguists who study reasoning, we have a tendency to misread conditional claims. The problem is that in certain circumstances we read conditionals as if the consequent were sufficient for the antecedent. We tend to read 'If A, then B' as saying also that 'If B, then A'. Of course nobody reads the conditional above about squares and rectangles in this way, because we know that would be false. It is in cases where we aren't so sure about the antecedent and the consequent that we become prone to making this mistake.¹

That said, sufficiency and necessity are, in a sense, two sides of the same coin. From the perspective of the antecedent the relation is one of sufficiency for the consequent, while from the perspective of the consequent the relation is one of necessity to the antecedent. More specifically:

A strict conditional of the form *if A, then B* simultaneously expresses two relations. It says that A is sufficient for B AND that B is necessary for A.

The idea that a conditional expresses two relationships strikes some people as counter-intuitive at first. However, we deal with this situation all the time in everyday life. Consider a different relation: 'to the left of'. When I say that I am standing to the left of you, I am simultaneously saying that you are standing to my right. So too, to say that A is the grandparent of B is at the same time to say that B is the grandchild of A. The same goes for conditionals—to say that A is sufficient for B is to say simultaneously that B is necessary for A.



Two Sides of the Same Coin: Necessary and Sufficient Conditions

SECTION 3: NECESSARY AND SUFFICIENT CONDITIONS

We began Section 2 by noting that there are many ways to express conditionals other than the standard (if...then...) form. Our discussion of the meaning of conditionals allows us to easily see two alternatives. To say that one thing is sufficient for another, or that one is necessary for the other, is to express a conditional relationship. Consider:

Ex. 2:

A person's being legally eligible to buy alcohol in this state is sufficient for that person's being at least 21 years of age.

Because this sentence asserts a sufficiency relation we can express it with a conditional. We can put it into standard form as follows:

■ If a person is legally eligible to buy alcohol in this state, then they are at least 21 years of age.

There is nothing unique to this example. Using 'A' and 'B' as variables, we can generalize on this and say that any sentence with the form 'A is a sufficient condition for B' can be expressed using the standard-form conditional: 'If A, then B' or:

■ *A is sufficient for B = If A, then B*

The same goes for necessary conditions. Consider the following example:

Ex. 3:

A person's being at least 21 years of age is a necessary condition for a person's being legally eligible to buy alcohol in this state.

Because this sentence asserts a necessity relation we can express it with a conditional. Since 'a person's being at least 21 years of age' is the necessary condition, then it will serve as the consequent in the standard-form conditional, as follows

■ If a person is legally eligible to buy alcohol in this state, then they are at least 21 years of age.

Again, there is nothing unique to this example. We can generalize on this and say that for a person to say that A is a necessary condition for B is just to express the standard-form conditional: If B, then A or:

■ *A is a necessary condition for B = If B, then A*

The conditional above tells us that everybody who is eligible to legally buy alcohol in this state is at least 21—no exceptions. Given this, we can formulate an equivalent variant of this conditional as follows. Consider somebody who is not at least 21 years of age, say they are 19. What do we know about them given the truth of the conditional above? We know that they are not eligible to legally buy alcohol, and we can express this with the following conditional:

■ If a person is **not** at least 21 years of age, then they are **not** legally eligible to buy alcohol in this state.

Being at least 21 years of age is a necessary condition for being legally eligible, so if this necessary condition is not met, it will also be true that a person will not be legally eligible. Again, there is nothing distinctive about this conditional expressed above, and we can say in general that:

■ *If A, then B = If not B, then not A*

There is a special name for this equivalent variant of 'If A, then B'; it is called the **contrapositive**.

SECTION 4: ALTERNATIVE WAYS OF EXPRESSING CONDITIONALS

There are ways of expressing conditionals beyond the language of necessary and sufficient conditions. Perhaps the most obvious alternative to the standard form occurs when we use the term 'if' out of order. For example, a person might say: The bird is a member of the genus *Anas*, if it is a mallard. The presence of the 'if' here is a big clue that we've got a conditional, and we can reformulate it into standard form as follows: 'If the bird is a mallard, then it is a member of the genus *Anas*'. We can generalize on this:

B, if A = if A, then B

The term ‘only if’ is another way that we sometimes express conditionals. The use of ‘only if’ to express conditionals is especially common in legal language. The word ‘only’ can be tricky, however, and it is important to see that the addition of ‘only’ to an ‘if’ fundamentally changes the meaning of the sentence. Consider the following example:

- (i) A person is breaking the law *if* they are speeding.
- (ii) A person is breaking the law *only if* they are speeding.

These two sentences say very different things. The first sentence says that if a person is speeding, then they are breaking the law. That’s true. The second sentence is telling us that the *only way* a person might break the law is by speeding. But this is certainly false, since we can break the law by littering, blowing through a stop sign, etc. Adding an ‘only’ to the ‘if’ significantly changes the meaning of the sentence, and so we cannot transform ‘only if’ sentences into standard-form conditionals in the same way we would an ‘if’ sentence. So how do we do it? Let’s look at a new example:

Ex. 4:

A person is considered an “active member” of the club *only if* they have paid their yearly dues.

This sentence tells us that paying yearly dues is a crucial part of being qualified as an “active member”. Put otherwise if a person hasn’t paid their yearly dues, they are not considered an “active member”. This gives us a standard-form conditional, and we can generalize on this result.

A only if B = If not B, then not A

Now that we also are aware of the contrapositive, we can see that ‘A only if B’ is equivalent to ‘if A, then B’ as follows:

A only if B = if not B, then not A = if A, then B

Perhaps surprisingly ‘all’ and ‘every’ claims like ‘all dogs are mammals’ or ‘every student at this school has been issued an ID card’ express conditional claims as well. In order to see this, think about what they say. To say that every student has been issued an ID card is to say that every single student has an ID card, *no exceptions*, and so being a student is sufficient for having an ID card. Putting this into the standard form for conditionals requires modifying the sentence a bit (since in standard form conditionals the antecedent and consequent are independent clauses), and this can be a little awkward. We can transform the sentence ‘every student at this school has been issued an ID card’ as follows:

If somebody is a person at this school, then they are a person who has an ID card.

This pattern generalizes and we can say that:

All As are Bs = If something is an A, then something is a B

Claims like ‘all As are Bs’ tell us about membership, namely that things that are As are also members of B. Similarly, claims about exclusion that use ‘no’ or ‘none’ can also express conditionals. For example, to say that ‘no reptiles are warm-blooded’ is to say that all reptiles are not warm-blooded—*no exceptions*. Put otherwise, being a reptile is sufficient for not being warm-blooded. Like ‘all’ claims, translating a ‘no’ claim into

a standard-form conditional requires modifying the sentence structure just a bit. The standard-form conditional in this case is:

■ If something is a reptile, then it is not warm-blooded.

The general pattern here is:

■ **No As are Bs = If something is an A, then it is not a B**

There are other ways of expressing conditionals, but this list captures the most common ways. In the next section, we will see how these alternatives figure into the deductive argument forms we've learned.

SECTION 5: RECOGNIZING DEDUCTIVE FORMS

Being able to recognize conditionals will allow us to recognize instances of the argument forms discussed in Chapter 10. Consider the following argument:

■ Ex. 5:
Don't worry, you will graduate. I know that the registrar cleared you, and this is sufficient for graduating.

The claim that being cleared by the registrar is sufficient for graduating expresses a conditional. Transforming this into a standard-form conditional we can see that this argument is an instance of the form *Modus Ponens*.

1. The registrar cleared you.
2. If the registrar clears you, then you will graduate.
3. So, you will graduate.

Here is another example:

■ Ex. 6:
Look, I told you—the battery doesn't have a charge, and the car will start only if the battery has a charge. So, I am sorry to say, the car won't start.

We know that the claim: 'the car will start only if the battery has a charge' expresses the standard-form conditional: 'if the battery does not have a charge, then the car won't start'. Given this, we can see that this argument is actually an instance of *Modus Ponens*, and is consequently a deductive argument.

1. If the battery does not have a charge, then the car won't start.
2. The battery does not have a charge.
3. So, the car won't start.

Wait. We have seen that a strict conditional is equivalent to its contrapositive. Thus, in this case although we transformed the claim 'the car will start only if the battery has a charge' into the standard-form conditional 'if the battery does not have a charge, then the car won't start' we could have transformed it into a different (but equivalent) standard-form conditional. We might have transformed it to say: 'If the car starts, then the battery has a charge.' Had we transformed it this way, we would have standardized the argument as follows:

1. If the car starts, then the battery has a charge.
2. The battery does not have a charge.
3. So, the car won't start.

But this is a case of *Modus Tollens*. So which is it? Is this argument an instance of *Modus Ponens* or *Modus Tollens*?

The answer, in short, is that it depends on how you transform the conditional claim into standard form. This might seem problematic, but there is no need to worry. How we transform the conditional into standard form will dictate what form the argument takes, but we will never be able to accurately represent a deductive argument as not-deductive if we transform our conditional claims correctly.

EXERCISES

Exercise Set 11A:

#1:

Identify at least one necessary condition for being a citizen of the United States. You might need to do a little bit of research.

#2:

Identify a necessary condition for getting an A in this class.

#3:

The terms 'consequent' and 'conclusion' sound very similar. How are they different?

Exercise Set 11B:

Directions: Transform each of the following conditional claims into the correct standard form conditional.

#1:

The chemical will dissolve, if the theory is correct.

#2:

Earning 93% of all available points in this course is sufficient for earning an 'A' for the course.

#3:

You may enter only if you are exercising your first amendment rights.

#4:

All submitted photos are the property of Insta.com.

#5:

No current Naval Commanders are convicted felons.

#6:

Having a social security number is a necessary condition for obtaining a driver's license.

Exercise Set 11C:

Directions: For each of the following arguments standardize it using standard form conditionals and if the argument is one of the kinds we have studied, identify it as such (e.g. Modus Tollens, Affirming the Consequent, etc.)

#1:

Understanding is impossible if words refer only to private sensations in the minds of speakers. We clearly do understand each other. Therefore, words do not refer merely to private sensations.

#2:

Getting a 4 on your AP exam is sufficient for getting college credit. Since I got a 4 on the exam, I will get college credit.

#3:

Future presidents will be allowed to serve a third term only if the Twenty-second Amendment is repealed. The Twenty-second Amendment will not be repealed. Therefore, future presidents will not be allowed to serve a third term.

#4:

Every time Louis is tired, he's edgy. He's edgy today, so he must be tired today.

#5:

You can do well in math classes only if you keep up with the assignments. You keep up with the assignments, so you do well in math classes.

#6:

Most senior citizens go to bed before 11 p.m. so my great-grandma probably goes to bed before 11 p.m.

#7:

The alternator is not working properly if the ammeter shows a negative reading. The current reading of the ammeter is negative. So, the alternator is not working properly.

#8:

Manuel will play only if the situation is hopeless. But the situation is hopeless. So Manuel will play.

#9:

Candidate Flores will lose the election if he doesn't win both Jackson and Northampton County, but he won't win either of these counties, so Flores will lose the election.

#10:

I passed 'Go' and passing 'Go' is a sufficient condition for collecting \$200, so hand over \$200!

#11:

Elephants have been known to bury their dead. Elephants would bury their dead only if they have a concept of death. So, elephants have a concept of death.

#12:

Having a good technical education is a necessary condition for being a good engineer. Given the quality of Charity's technical education, she will surely make a good engineer.

#13:

We cannot worship a god or gods if we don't have a capacity to form a concept of the divine. And we cannot have a capacity to form a concept of the divine if we don't have a capacity to form concepts transcending sense perception. So we cannot worship god or gods if we don't we have a capacity to form concepts transcending sense perception.

#14:

We know that Philip is a nonresident. Since no citizens are nonresidents, it follows that Philip is not a citizen.

Notes

1. Wason, P. C., & Johnson-Laird, P. N. (1972). *Psychology of Reasoning: Structure and content*. Harvard U. Press, 61.

Unit #4 Summary

In this unit we focused on deductive arguments. As we have seen, deductive arguments are indefeasible and have maximal logical strength due to their form. Although there are many argumentative forms, we specifically identified four especially common ones. We also looked at two forms people commonly confuse for deductive forms. Because strict conditionals are such an important part of many deductive arguments, we talked through what a strict conditional is really saying, and identified a variety of different ways we express conditionals. This led, finally, to a consideration of the many ways that deductive arguments, themselves, can be expressed.

KEY TERMS

- Deductive Arguments
- Defeasible Arguments
- Indefeasible Arguments
- Argument Form
- Strict Conditionals
- *Modus Ponens (MP)*
- *Modus Tollens (MT)*
- *Hypothetical Syllogism (HS)*
- *Disjunctive Syllogism (DS)*
- *Denying the Antecedent (DA)*
- *Affirming the Consequent (AC)*
- Conditional Statement
- Sufficiency Relation
- Necessity Relation
- Contrapositive

FURTHER READING

People have been studying deductive arguments for over 2000 years going back at least to the work of Aristotle in the 4th century BCE. The study of deductive arguments is known as deductive or formal logic and there are many solid introductory books on the topic. Two in particular that stand out are *A Concise Introduction to Logic* by Patrick Hurley and *The Power of Logic* by Frances and Daniel Howard-Snyder.

UNIT V

COMMON INDUCTIVE ARGUMENTS

Arguments from Analogy

SECTION 1: INTRODUCTION

In the last unit we focused primarily on deductive arguments. In this unit, we will turn our attention to inductive arguments. Recall that in an inductive argument the truth of the premises does not guarantee the truth of the conclusion. Nevertheless, there are logically strong inductive arguments, and our primary goal will be to learn how to distinguish these arguments from the logically weak ones. We will start with a brief look at inductive arguments more generally. As it turns out, inductive arguments have two distinctive and important characteristics. First, inductive arguments are defeasible. Second, unlike deductive arguments, the logical strength of an inductive argument is not due solely to its form. As a result of these differences, the process of argument evaluation differs in important ways. In this chapter, we will discuss these distinctive features, show how they impact evaluation, and then take a close look at a particular argument form: *Argument from Analogy*. In the three chapters that follow we will look at other common forms. For each form, we will identify what makes it distinctive, learn the difference between logically strong and logically weak instances of that form, discuss some common pitfalls, and end with some questions that will help us to evaluate for logical strength.

SECTION 2: DEFEASIBILITY AND THE RULE OF TOTAL EVIDENCE

Unlike deductive arguments, all inductive arguments are defeasible. That is, the logical strength of inductive arguments is sensitive to the introduction of new information. Recall the detective example from Chapter 10: a detective is investigating a murder, and has uncovered a number of important pieces of information. She has discovered that the murder weapon was owned by the victim's sister, the sister's fingerprints were found on the gun, and the sister had a strong motive. While these three facts strongly point to the sister, they do so in a defeasible way since the support they give to the conclusion might be undermined (or enhanced) by the addition of new information. Indeed, the detective did subsequently find new information, namely security camera footage showing the sister was 25 miles away at the time of the murder. This new information rightly changed the detective's thinking because the body of evidence available to her no longer supported the sister's guilt.

Let's take this example in a different direction. Imagine the detective ignores the fact that the sister was 25 miles away at the time of the murder, and argues instead that the sister committed the crime. We can all agree that the detective has made a big mistake. But what, exactly, is the mistake? It is not a problem with factual correctness or logical weakness: the evidence the detective is drawing on (gun ownership, fin-

gerprints, motive) is accurate and supports the conclusion that the sister is guilty. That is, the problem here is not internal to the argument as formulated. Instead, the problem is with *the way* the argument has been formulated in the first place. The detective has left something crucial out of the argument, so the conclusion is based on only part of the total available evidence. The same problem can crop up when we evaluate other people's arguments. After all, if we fail to bring all the relevant facts to the table in deciding whether an argument is sound or not, we are guilty of precisely the same mistake. This kind of mistake is called a violation of *The Rule of Total Evidence*. This is a rule for evaluating and formulating arguments which tells us that we can't leave out information. More specifically,

According to *The Rule of Total Evidence*, when evaluating and formulating an inductive argument, all available relevant evidence must be taken into account.

This rule tells us that in order to adequately evaluate inductive arguments we have to look outside or beyond the argument as stated. We have to think about the claim the conclusion is making, and consider whether there is any available information that would undermine (or enhance) the support the premises lend to the conclusion. Again, the issue here is that if there is available information, and you don't take this into account, then any conclusion you draw about the argument's soundness will be premature. Moreover, as premature, the conclusion will be more likely to be mistaken since it isn't based on the total available evidence.

It is important to note a number of qualifications to *The Rule*. First, we almost never have "all" the evidence, and *The Rule* doesn't say we have to. What *The Rule* says is that we shouldn't set aside or ignore evidence that is relevant and available.¹ Second, by 'available' we mean information that we know or have easy access to in some sense. Third, it can be difficult to spot violations of *The Rule*. In order to spot a violation of *The Rule*, you have to see that something is missing. That is, you have to see what is not in the argument, and this requires relevant background knowledge and the effort to sort through it, as we will see. Fourth, *The Rule* probably sounds like an obvious requirement—and it is. Nevertheless, people regularly violate this rule—both intentionally and unintentionally. People can unintentionally violate this rule for all kinds of reasons. Recall, for example, that myside bias can lead us to unintentionally leave out relevant evidence. Other times, we simply don't think to check, or don't know the right questions to ask. That said, people sometimes purposefully leave out information in order to make a conclusion look more plausible or appealing than it really is.

To intentionally violate *The Rule* in this way is called **cherry-picking**. Cherry picking is especially common in persuasive contexts. After all, a proposal, policy, or product can seem especially appealing when we can only see their positives. Here are a couple of recent examples to illustrate the point. In his State of the Union address in January of 2018, President Trump took credit for rising wages. Trump was correct in claiming that wages had risen during his presidency. What he didn't say, and what undermines his claim, is that wages had been steadily increasing since at least 2014—2 years prior to his becoming president.² In short, Trump cherry-picked statistics to make it sound as if his administration had been the chief cause of these gains in wages. Indeed, cherry picking is common in political contexts.

As another example, near the end of the 2020 Democratic Presidential primaries, Joe Biden tried to paint his chief opponent, Bernie Sanders, as against "common sense" gun-control measures like background checks for those purchasing firearms. In support of his contention, Biden pointed out that Sanders had voted against the Brady Bill 5 times. The Brady bill, which was passed into law in 1993 did, among other things, include a background check provision. Moreover, it is true that Sanders voted in this against this bill as it worked its way through Congress in the early 1990s. However, Biden was cherry-picking this statistic. In the

intervening 25 years Sanders has voted in favor of a variety of gun control measures that have come before Congress. Taking all of these votes into consideration makes Biden's claim that Sanders is against common sense gun-control measures more difficult to maintain.



"Cherry picking" by barnimages.com CC BY 2.0

The lesson from all of this is that in evaluating inductive arguments we need to look outside the argument itself, and we will incorporate this insight into our basic evaluative procedure. As we have seen, in evaluating an argument we need to check for factual correctness by asking: *are the premises likely to be true?* We also need to check for logical strength by asking: *is the conclusion probable given the truth of the premises?* Finally, in light of the Rule of Total Evidence, let us ask: *is there any other relevant information available?* This gives us...

Three Key Evaluative Questions:

- *Are the premises likely to be true?* (check for Factual Correctness)
- *Would the truth of the premises make the truth of the conclusion probable?* (internal check for inductive strength)
- *Is there any other relevant information available?* (external check for inductive strength)

This third question reminds us to look beyond the argument and to check its logical strength against what we know, or have access to, more broadly. For this reason, it is helpful to think of this third question as an external check on (inductive) logical strength, and the more straightforward consideration of the support lent by the premises themselves as an internal check.

SECTION 3: DETERMINING LOGICAL STRENGTH FOR INDUCTIVE ARGUMENT FORMS

In the last unit, we identified a number of deductive argument forms (e.g. *Modus Ponens*, *Hypothetical Syllogism*, etc). Identifying arguments by form is a useful way of classifying arguments, and we will continue to do so in this unit by identifying common inductive forms. However, there is an important difference between deductive and inductive argument forms, namely that while deductive forms guarantee logical strength, inductive forms do not. To illustrate, there are no logically weak arguments with the form of *Modus Ponens*. Any argument with this form will be logically strong because the form itself guarantees logical strength. In contrast, consider a common inductive form: *Arguments from Analogy*. While there are logically strong arguments with this form, there are logically weak *Arguments from Analogy* as well, and this is because inductive forms do not guarantee logical strength. The upshot here is that when it comes to inductive forms, simply

knowing an argument's form won't tell you whether it is logically strong or not. There are other relevant factors at work, and we'll need to take these into account in determining logical strength.

What factors? There is no one-size-fits-all answer to this. Nevertheless, we will see that the relevant factors vary from one inductive form to another. Thus, the factors relevant to determining whether an *Argument from Analogy* is logically strong differ from the factors relevant to determining whether an *Inductive Generalization* is logically strong. Consequently, even though knowing an argument's inductive form won't automatically tell you whether it is logically strong or not, it *will* tell you what factors are relevant, and what questions to ask to figure it out. It is important to emphasize that in saying that evaluation differs from argument form to argument form, it is the *process* of evaluation that differs, not the *criteria*. We evaluate all arguments by the criteria of factual correctness and logical strength. The difference between evaluating an *Argument from Analogy* and an *Inductive Generalization* has to do with the appropriate steps to take for determining logical strength. So, what is an *Argument from Analogy* anyway, and how do we evaluate them? Let's take a look.

SECTION 4: IDENTIFYING ARGUMENTS FROM ANALOGY

Observations of similarity and difference play a number of roles in everyday contexts. In many cases they are used as tools for explaining difficult concepts or relationships. For example, a science teacher might explain how electricity flows by comparing it to water. Moreover, judgments of similarity and difference underlie many other reasoning processes. To give just one example, think about categorization. To decide that two things belong in the same category is to decide they are similar in some important way. Our primary concern, however, will be arguments in which one or more similarities are used to infer the presence of a further similarity, and we will refer to arguments with this form as *Arguments from Analogy*.

Let us begin with some examples.

Ex. 1:

Sofia is a very good dancer, so I bet she has the potential to be very good at fencing, since fencing, like dance, requires excellent balance and fine motor skills.

Ex. 2:

You think it is fine to eat pork, and biologically speaking there are only minor differences between pigs and dogs, so you should be willing to eat dog too.

Let us take a closer look at these examples by standardizing them.

Surface-Level Standardization of Ex. 1:

1. Sofia is a very good dancer.
2. Fencing, like dance, requires excellent balance and fine motor skills.
3. So, I bet Sofia has the potential to be a very good fencer.

Surface-Level Standardization of Ex. 2:

1. You think it is fine to eat pork.
2. Biologically, pigs and dogs are extremely similar.
3. So, you should be willing to eat dog too.

What these arguments have in common is that each draws its conclusion on the basis of a comparison between two objects or situations (note the underlined terms), and indeed this is what is distinctive about arguments from analogy. It is important to remember that just because an argument contains an analogy or comparison does not mean it is an argument from analogy. Again, arguments from analogy are inferences in which the conclusion is drawn on the basis of a comparison.



“Fencing duel” by uwdigitalcollections CC BY 2.0

We will refer to the objects or situations that are compared in an argument from analogy as the **analogues**. Thus, the analogues in Ex. 1 are fencing and dancing, and the analogues in Ex. 2 are pigs and dogs. People are creative and draw comparisons between all kinds of things; in some cases, people offer arguments from analogy in which the analogues are themselves arguments! Consider the following:

Ex. 3:

Sometimes people argue that we should not take prescription drugs because they contain unnatural substances and unnatural substances may be harmful to the body. But this is a poor argument since by similar reasoning we could conclude that you should not consume ice cream (since it is an unnatural substance and may be harmful to the body). Since this is a bad argument, so too is the first one.

This argument draws a comparison between arguments. Let us call the first argument the *Prescription Drug Argument* and standardize it as follows:

1. Prescription drugs contain unnatural substances.
2. Unnatural substances may be harmful to the body.
3. So, we should not take prescription drugs.

Let us call the second argument the *Ice Cream Argument* and standardize it as follows:

1. Ice cream contains unnatural substances.
2. Unnatural substances may be harmful to the body.
3. So, we should not (eat) ice cream.

In short, the argument in Ex. 3 is claiming that the *Prescription Drug Argument* is like the *Ice Cream Argument* in that both draw their conclusions on the basis of the fact that their subjects contain unnatural substances, and that unnatural substances may be harmful to the body. Given these similarities and that the *Ice Cream*

Argument is a poor argument (the author is assuming that you will agree that it is not wrong to eat ice cream), the author concludes that the *Prescription Drug Argument* is a poor one as well.

SECTION 5: EVALUATING ARGUMENTS FROM ANALOGY

What is the difference between logically strong *Arguments from Analogy* and logically weak ones? Let us begin with a weak *Argument from Analogy*. We will modify the argument in Ex. 1 to make it uncontroversially weak.

Ex. 4:

Sofia is a very good fencer, so I bet she has the potential to be very good at throwing the shot put, since fencing, like throwing the shot put, is an Olympic sport.

We can agree that this is a weak argument, but what makes it weak? The problem is that the noted similarity is *irrelevant* to the similarity that is inferred in the conclusion. The fact that both fencing and shot put are Olympic events has nothing to do with whether Sofia would be a good shot putter. There are many different Olympic events, and the skill and expertise required for one may not translate to another. Since all *Arguments from Analogy* rely on a comparison, this is a point we can generalize upon. We can say that in a logically strong *Argument from Analogy*, the noted similarity or similarities are relevant to the inferred similarity.

Two brief notes about this requirement. First, a noted similarity is relevant to the inferred similarity when (and only when) the presence of the noted similarity makes the presence of the inferred similarity more probable. Second, this is only one condition for logical strength. Even if the noted similarities are relevant, it does not follow that the argument is logically strong. After all, arguments from analogy are defeasible, and so in accordance with *The Rule of Total Evidence*, we have to consider whether there is relevant available information that bolsters or undermines the support the noted similarities give to the inferred similarities. To illustrate, let's modify the example above again.

Ex. 5:

Sofia is a very good fencer, so I bet she has the potential to be very good at throwing the shotput, since fencing, like throwing the shot put, requires quickness and agility.

This argument does not have the same problem as Ex. 4. In this case, the author is noting that both fencing and the shot put require quickness and agility, and these characteristics are relevant to the inferred similarity. Knowing that somebody is quick and agile would give you some positive reason to think that they would be good at throwing the shot put. Nevertheless, this argument is problematic, since even a passing familiarity with the two sports reveals fundamental differences. For example, body mass and power are crucial for success in throwing the shot put, but are less so in fencing. Similarly, fine hand-eye coordination and physical endurance are crucial for success in fencing, but less so for shot put. Thus, although the noted similarities are relevant to the inferred similarity, there are *relevant differences* between the two that undermine the logical strength of the existing argument. Again, since all arguments from analogy depend on a comparison, we can generalize on this example and say that the presence of relevant differences between the analogues weakens the strength of the analogy—sometimes decisively.

One last note: it is important to recognize that not all differences are relevant differences. What makes a difference relevant is that it makes the inferred similarity less probable, but there will always be differences

that do not tell us anything about the probability of the inferred similarity. For example, in fencing you wear a mask, and in the shot put you do not. This difference doesn't really tell us anything about whether Sofia would be a good at the shot put or not. Again, there will always be many differences between analogues, and the task is not to identify differences, but relevant differences.

There are a number of other features relevant to the evaluation of *Arguments from Analogy*. For example, the more relevant similarities there are between the analogues, the stronger the argument. Furthermore, relevant similarity and difference are a matter of degree; that is, one similarity (or difference) might be more relevant than another. Nevertheless, the relevance of the noted similarity to the inferred similarity, and whether there are relevant differences are by far the most important issues when it comes to evaluating *Arguments from Analogy*. Thus, after having identified an argument as an *Argument from Analogy*, and having identified the noted similarity and the inferred similarity, you will always want to ask yourself two questions.

Two Questions to Ask of Any Argument from Analogy:

- Are the similarities relevant?
- Are there relevant differences?

The first question asks about the argument as stated. The second question asks about whether there is other available information that is relevant to the conclusion. While we could ask the question more broadly to look for similarities as well, looking for differences is especially important. After all, the person giving the argument has already identified similarities that purportedly support the conclusion. If they've left something important out, it is likely to be in the form of a difference, and so it makes sense to focus your thinking on them.

In general, it is worth observing that people tend to find it much easier to detect weak *Arguments from Analogy* when the noted similarity is not relevant to the inferred similarity, than it is to detect arguments that are weak due to relevant differences. One plausible explanation for this is that taking up the first question above is relatively easy—the similarities in question are right there in front of you, and all you need to do is determine whether one is relevant to the other. Taking up the second question above requires more effort, since whether there are relevant differences is *not* right there in front of you. If there are relevant differences, they are missing; to find them we have to look beyond the argument itself, thinking about what else we know of the analogues. Thus, it takes extra work to identify differences, and this is worth keeping in mind when evaluating *Arguments from Analogy*.

EXERCISES

Exercise Set 12A:

#1:

Have you ever noticed or suspected that somebody was cherry-picking their evidence? When? If not, make up a realistic example of cherry-picking.

#2:

Invent an argument from analogy in which the noted similarity is not relevant to the inferred similarity.

#3:

Think about the last time you felt you were treated unfairly. How might your judgment that you were treated unfairly rely on an argument from analogy?

#4:

In law the concept of a legal precedent is very important. What is a legal precedent and how might its use in legal proceedings be related our discussion of arguments from analogy?

Exercise Set 12B:

Directions: Each of the following is an argument from analogy. Assume each argument is factually correct, and evaluate each argument's logical strength. Explain your evaluation.

#1:

You wouldn't think eating a person is okay, would you? A cow is a living breathing being just like a person is, so why do you think that eating a cow is fine?

#2:

Last week I sold my 2012 Toyota Camry for \$5000. The body was in good shape and so was the engine, and the mileage on the car was pretty low. Your 2013 Camry is in about the same condition, and the mileage is about the same, so I'd guess you could get at least \$5000 for it.

#3:

Blaming guns for the mass shooting at Columbine is like blaming spoons for America's obesity epidemic!

#4:

Having children is just like buying a house. Like buying a house, raising a child properly is an expensive prospect, and just as you shouldn't take on a mortgage you can't afford, so too you shouldn't have children unless you can afford to do so.

#5:

Covid-19 is an easily spreadable virus just like the flu. But we don't lock down the country or take extreme measures to prevent the spread of the flu, so we shouldn't do so for Covid-19 either.

#6:

I am completely on board with the proposed law to require voters to present valid ID before they can vote. I mean you have to present ID to buy alcohol or cigarettes, to board a plane, legally drive a car, and to get into an R-rated movie, so why should voting be any different?

#7:

You wouldn't commit to buying a car without test driving it first, so why would you commit to marrying someone without living together first?

#8:

Copy machine repair person: Your copy machine is pretty good—I mean it only breaks down on average once a month.

Copy machine owner: You think that is good? I have a dishwasher and it has never broken down!

#9:

Comedian George Carlin:

I don't understand this notion of ethnic pride. "Proud to be Irish," "Puerto Rican pride," "Black pride."...Being Irish isn't a skill; it's genetic. You wouldn't say. "I'm proud to have brown hair," or "I'm proud to be short and stocky." So why...would you say you're proud to be Irish? I'm Irish, but I'm not particularly proud of it. Just glad!...[G]lad to be Irish!

#10:

Hydraulic fracking produces, as a waste product, a lot of salty chemical-laden water. The EPA prefers that when the water cannot be recycled that it be put into injection wells which deposit the waste far underground. However, in Ohio this process may be responsible for a recent spate of earthquakes. This had led to a moratorium on permits on new injection wells. In response the owner of a company that runs inspection wells says:

- "If you've got a car that has a problem with it, does that mean we should take all the cars off the road and go back to horse and buggy?"³

Notes

1. See Fetzer, James (Ed.). (2001). *The Philosophy of Carl Hempel*. Oxford: Oxford University Press, 114.
2. Jacobson, Louis. (2018, Feb. 5). "The Age of Cherry Picking." *Politifact*. <https://www.politifact.com/truth-o-meter/article/2018/feb/05/age-cherry-picking/>
3. Source: Detrow, Scott. (2012, Apr. 11). "From Pa. Waste To Ohio Quakes" *NPR*. <https://www.npr.org/2012/04/11/150445837/from-pa-waste-to-ohio-quakes>

Inference to the Best Explanation

SECTION 1: INTRODUCTION

In the last chapter we looked at *Arguments from Analogy*. In this chapter we will turn to another common non-deductive argument form: *Inference to the Best Explanation*. We are constantly seeking explanations: why did my computer stop working? Why did the housing market crash? Why does my foot hurt so badly this morning? This is especially common in our dealings with other people: why hasn't she called me back? Why is the baby crying? Why does that guy always sit by himself in the cafeteria? In general, when we reason to an explanation, we start with an observation and work backward to identify its likely source or cause. In many cases, it is not difficult to come up with an accurate explanation, and overall we tend to do a reasonably good job at it. Nonetheless, we do make mistakes. In order to understand and recognize these mistakes, we will need to slow down and learn what makes *Inferences to the Best Explanation* distinctive, the difference between logically strong and logically weak instances of this form, and some questions to help us distinguish between the two. We will close by considering a bias we are subject to when it comes to explaining other people's behavior in particular.

SECTION 2: INFERENCE TO THE BEST EXPLANATION

In most cases, to offer an explanation for some event or state of affairs is to reason to a conclusion. To illustrate this, let us begin by looking at some examples of explanatory reasoning.

Ex. 1:

Andre: Oh no! Somebody stole my bike!

Bobby: What do you mean?

Andre: I left it right here 10 minutes ago and now it is gone!

Andre returns to find his bike missing. This state of affairs calls for an explanation, and Andre concludes that his bike has been stolen. We can represent his reasoning as follows:

1. I left my bike right here 10 minutes ago and now it is gone.
2. So, somebody stole my bike.

Here is another example:

Ex. 2:

Jessie says: “Kayden didn’t return my text; I must not be very important to Kayden.”

Again the fact that Kayden didn’t return Jessie’s text is something that calls out for an explanation (at least to Jessie!). Jessie reasons as follows:

1. Kayden didn’t return my text.
2. So, I must not be very important to Kayden.

Notice that in both cases there is a significant gap between the premise and conclusion. In each case, the only stated premise is an observation, and the argument then jumps to the conclusion that this observed state of affairs has a specific explanation. On what grounds should we think that these arguments are logically strong? Let’s focus on Ex. 1—why does the fact that Andre’s bike is missing make it probable that it was stolen? Given the Rationality Assumption, the author is assuming a connection here, so let’s go deeper and try to identify an appropriate missing premise.



*“RESPECT: 3-Speed Flying Pigeon Racer w. Wald Basket”
by ubrayj02 CC BY 2.0*

Andre expects that his bike will be right where he left it, and when he finds that it is missing he immediately wonders what has happened. Although he quickly concludes that it has been stolen, it is important to recognize that there are many other explanations for why his bike might be missing *other than* that it was stolen. After all, a friend might have borrowed it, Andre might be misremembering where he left it, his friends might be playing a joke on him and have camouflaged it in some way, and so forth. Thus, in claiming that the bike was probably stolen, Andre is identifying one explanation for why his bike is gone as the most likely one. Thus, the missing premise says something like this:

1. I left my bike right here 10 minutes ago and now it is gone.
2. The most likely explanation for why my bike is gone is that somebody stole my bike. (MP)
3. So, somebody stole my bike.

The missing premise in this case shows us how the stated premise likely connects to the conclusion, and we can generalize on this. After all, for any given state of affairs there are many possible explanations, and to pick one of these possible causes or explanations out is to say that it is the most likely or best explanation.

This is why such arguments are called inferences to the best explanation. That is, an *Inference to the Best Explanation* argues that some specific explanation is probable given that it is the best or most likely explanation among those available for an observed state of affairs.

We should add a few brief notes about *Inferences to the Best Explanation*. First, it is important to recognize that by the ‘best explanation’ we are not speaking objectively here. We have limited knowledge and in most cases cannot realistically expect to have the actual best explanation of some event within our grasp. What we *can* do is to determine the best explanation given the information we have or can access. Second, *Inferences to the Best Explanation* are defeasible, and what explanation counts as the “best” can change given the acquisition of new information. A third note is that in everyday thinking and conversation we rarely make the “best explanation” premise explicit. That is, most cases are like Ex. 1 and Ex. 2 in that we simply jump from an observation to an explanation.

It is worth drawing special attention to a couple of contexts in which *Inferences to the Best Explanation* play a particularly important role. First off, consider court cases. Suppose Jack is charged with the murder of Xavier. The prosecutor will make a case that the best explanation for what is known about the murder of Xavier is that Jack did it. The defense, of course, will maintain just the opposite. They will try to undercut the prosecutor’s evidence and to show that there is reasonable doubt that this is the best explanation (perhaps by presenting an alternative explanation for the facts surrounding Xavier’s murder). A second prominent context is the diagnosis of illness. You don’t feel well and go to the doctor. The doctor will take account of your symptoms, and run some tests. The doctor will consider various explanations for your illness in light of your symptoms and test results, and the best explanation will be the diagnosis of your illness.

SECTION 3: EVALUATING FOR LOGICAL STRENGTH: STEP ONE

When it comes to evaluating *Inferences to the Best Explanation* for logical strength, the first and most obvious step is to determine whether the proposed explanation is plausible or not. This first step amounts to a check on the internal logical strength of the argument. To illustrate, go back up to Ex. 2. In this case, we need to ask ourselves whether the proposal that Jessie isn’t important to Kayden is a plausible explanation for why Kayden didn’t return Jessie’s text. Although we don’t know all the specifics of the case, in general terms, the answer seems to be, yes, it is a plausible explanation. After all, we often don’t return phone calls or texts promptly when the person involved is not important to us, and so this proposal would plausibly explain why Kayden didn’t return Jessie’s text. That said, we are only part way through our evaluation, and this brings us to Step Two.

SECTION 4: EVALUATING FOR LOGICAL STRENGTH: STEP TWO

The second step in evaluating *Inferences to the Best Explanation* is to do an external check for logical strength by looking outside of the proposed explanation for plausible alternatives. This is a crucial step since the fact that the proposed explanation is plausible does not mean that it is the best or most likely one. There may be other more likely explanations available, and we won’t know unless we look. Despite its importance to evaluating *Inferences to the Best Explanation*, this step can be less obvious than the first one. In fact, we often fail to consider whether there are alternative explanations, and simply take it for granted that the first explanation that we are given, or that occurs to us, is the best one.

Before saying more about Step Two, let us briefly pause to note that going with the first explanation that occurs to us is not always a bad thing. Indeed, sometimes the first explanation that pops into our heads does so precisely because it makes the most sense in light of our knowledge and experience. For example, suppose that class starts and you notice that your friend Ellis is absent. There are many possible explanations for this: she dropped the class, she had to take care of a family emergency, she broke her leg on the way to class, etc. Let's say, however, that over the last few years Ellis has been late or missed morning classes a lot, and in every case it has been because she overslept. Given this background information it would probably naturally and automatically occur to you that she overslept again, and this really would be the best and most likely explanation given your knowledge and experience.

Importantly, however, not all cases are like this one. A particular explanation can occur to us for a variety of reasons—some of which have nothing to do with the explanation's being the best or most likely one. Return to Ex. 2: Jessie automatically infers from the fact that Kayden hasn't returned her text that she is not important to Kayden. Why is this the first explanation that occurs to Jessie? It is difficult to say, of course, but it may be that she has recently been wondering about the status of her relationship with Kayden. But the fact that she has been wondering about this, does not necessarily mean that this explanation is the most likely one. It follows that—at least in cases where it is important—we need to take the time to consider whether there are plausible alternative explanations. So, are there plausible alternatives for why Kayden hasn't returned her text? Yes: maybe Kayden hasn't noticed the text, maybe the phone is turned off, maybe the battery is dead, and so forth. Ultimately, then, the second step in evaluating *Inferences to the Best Explanation* is to ask yourself whether there are plausible alternative explanations, and in doing so prevent rash or premature explanations. Let us call a failure with respect to this second step, a *hasty explanation*. More specifically, we will say that:

A person gives a **hasty explanation** when they mistakenly conclude that some explanation is the most likely because they did not adequately consider whether there were plausible alternatives.

Again, a person can't legitimately argue that some explanation is the most likely one, if they haven't considered whether there any others that might be better. Thus, if Jessie concludes that she isn't very important to Kayden without having considered alternative explanations, she risks a premature conclusion, that is, a hasty explanation.

SECTION 5: EVALUATING FOR LOGICAL STRENGTH: STEP THREE

Suppose that you've considered the plausibility of the proposed explanation, and looked outside for plausible alternatives. Of course, if there are no plausible alternatives, your evaluation is complete. This is uncommon, however. In most cases, there will be plausible alternatives, and this brings us to the third step in evaluating *Inferences to the Best Explanation* for logical strength—identifying the best or most likely one. This is another external check for logical strength, and the task is to compare the plausible alternatives to identify the best or most likely among them. Returning to Ex. 2, which plausible explanation is best or most likely? Is it that Kayden hasn't noticed the text, that his phone is off, that his battery is dead, or that Jessie isn't important to him? In order to answer this question, we need to think about what else we know about Kayden. Thus we might ask: does he constantly check his phone? Does he regularly turn his phone off? Does he regularly let his phone's battery go dead? Has he been ignoring Jessie in other ways? Answering these questions can help identify the most likely explanation all things considered.

Nevertheless, comparing explanations can be difficult. It can be easy to get bogged down in the information that supports or detracts from the plausibility of multiple explanatory alternatives.

Fortunately, there is a short cut, or heuristic, for identifying the best or most likely explanation. The heuristic is:

When comparing alternative explanations, look for the one that would be **least surprising**.

In order to see why this technique is effective, think about what it means to say that an explanation is plausible. When we are considering the plausibility of an explanation, we are trying to figure out how likely that explanation is in light of our experience and in light of what we know about how the world works. Moreover, when things happen in accordance with our knowledge and experience we barely notice them. For example, when you turn on the tap and water comes out, you don't think about it at all. This is exactly what you'd expect given your experience. However, if you turn on the tap and nothing comes out, you'd be surprised. This is unexpected given your knowledge and experience of the world. More generally, how surprising you find some situation tells you how unlikely its occurrence was, given your knowledge and experience. Consequently, we can use surprise to compare competing explanations.



"Card Tricks" by canopic CC BY-NC-ND 2.0

To illustrate, imagine that you are walking down the street and come across a street magician. He asks you to pick a card randomly out of the deck. He turns around and you draw the 8 of diamonds. He turns back around to face you and says he is going to identify the card you chose by reading your mind. He closes his eyes and is silent for a moment, and then says: "The card you chose was the 8 of diamonds." You are surprised! How could he have known? You ask him how he did it, and he replies: "Just like I said...I read your mind." As you walk away, you consider some explanations. He could have used a mirror, he could have gotten a sign from somebody behind you, the card is marked, all the cards in the deck could have been the 8 of diamonds, or... he could have really read your mind. As you think about these, you note that you didn't see any mirrors or anybody standing behind you who could have seen your card, and he showed you beforehand that the deck had different cards in it. This seems to leave only mind-reading. At this point, however, you might want to use the "least surprising" heuristic by asking yourself: *which would be less surprising, that he really can read minds or that you just failed to see how he did it?* If you are like most people, you would be very surprised if the magician can really read minds. We do not have any experience of mind-reading, and probably don't know of any cases of it. In short, it would be surprising precisely because it is unlikely given your experience and knowledge. Thus, even though you might not know exactly how the magician identified your card, when we look for the least surprising explanation, we know that the best explanation is that the magician somehow identified your card by conventional means, *not* that he read your mind.

In sum, the third step in evaluating *Inferences to the Best Explanation* is to compare the proposed explanation with the alternatives you've identified in order to determine which explanation is most plausible. As we've seen, one way of determining this is to ask which explanation is least surprising. Let us call a failure to accurately compare and select from your alternatives, a *poor explanation*. More specifically, we will say that:

A person gives a **poor explanation** when they endorse an explanation that is not the best or most likely given their information.

To say that a person has given a poor explanation is not to say that the explanation isn't plausible, it is to say that it is not the most plausible given the information available. As an example, a person who concludes that the magician probably did read their mind will have drawn a poor explanation since this won't be more likely than other available explanations (e.g. mirrors, somebody behind you, marked cards).

Where does this leave us? We have identified three steps to evaluating *Inferences to the Best Explanation*. First, we determine whether the proposed explanation is plausible. Second, we consider whether there are plausible alternative explanations. Third, we determine which of the plausible alternatives is best or most likely, and we can do this by asking which alternative would be least surprising.

Three Questions to Ask of *Inferences to the Best Explanation*:

- How likely is the proposed explanation?
- Are there other plausible explanations? (Failure: Hasty Explanation)
- Would the truth of the proposed explanation be less surprising than the truth of any competitor? (Failure: Poor Explanation)

SECTION 6: EXPLAINING PEOPLE'S BEHAVIOR

Although we put a lot of effort into explaining why the world is as we find it, we tend to focus in particular on other people's behavior. That is, we want to know why people say the things they say, and do the things they do. It is important to briefly emphasize this, in part, because we face distinctive challenges when it comes to other people. Explaining other people's behavior can be difficult. We are complicated creatures, and we have different knowledge and experience, different personalities and values, and pursue different kinds of goals. Thus, part of the difficulty in explaining other people's behavior comes out of the fact that other people can be so different from us. There is a further, and related, complication—namely our thinking about other people is often biased in a variety of ways. We will briefly discuss one bias in particular that social psychologists call *Fundamental Attribution Error*. In order to understand the bias, let's start by identifying two broad categories of factors that influence our behavior: personal factors and situational factors. We have goals, interests, and preferences, and of course these personal factors are often the primary factors driving our behavior. On the other hand, sometimes our behavior is primarily a response to a situation or circumstance that we are in. This distinction is not an exclusive one, and often our behavior is a consequence of both personal and situational factors. The problem is that we have trouble appropriately explaining other's behavior in terms of these categories. More specifically,

Fundamental Attribution Error is our tendency to overestimate the influence of personal factors, and underestimate the influence of situational factors, when explaining other people's behavior.¹

To illustrate this, suppose that somebody is driving very fast and cuts Ben off in traffic. Ben immediately thinks, "what a jerk!" and wonders what is wrong with that guy. The first thing to notice about this example is that it is an *Inference to the Best Explanation*. Ben has made an observation: a man is driving very fast and

has cut him off, and he's jumped to an explanation of this behavior in terms of a personal factor: namely because he is a jerk. Of course there are other plausible explanations—many of them situational. The driver might, for example, be having a medical emergency; alternatively, maybe he is late for his sister's wedding. However, Ben didn't think of these, and this illustrates *Fundamental Attribution Error* in action. That is, Ben risks making a hasty explanation. Ben's first and natural impulse was to explain the man's behavior in terms of a personal as opposed to a situational factor. Noting this bias is important because it tells us to double-check our explanations of other people's behavior by taking a moment to consider whether there might be unappreciated situational factors at work.

EXERCISES

Exercise Set 13A:

#1:

Give three examples of inferences to the best explanation from your everyday life.

#2:

Inferences to the best explanation are defeasible. For each of the examples presented in the text (Ex. 1-3) specify one fact that you might learn that would likely change a person's judgment of the best explanation.

#3:

You have agreed to meet up with your friend for coffee, but she cancels at the last minute. What does Fundamental Attribution Error suggest we are prone to think, and how might you combat this?

Exercise Set 13B:

Directions: Consider each of the following situations. First give an explanation that might immediately occur to you. Second, identify a distinct explanation. Third, choose between the two. Which is more plausible and why? The goal of this exercise is to practice identifying and comparing different explanations.

#1:

You were just singing a song in your head, and then you hear it on the radio.

#2:

Japan has the highest life expectancy of any country in the world.

#3:

You look up into the night sky and see a tiny light moving slowly across the sky.

#4:

During the economic downturn following the coronavirus pandemic women lost their jobs at a higher rate than men.

#5:

You are at a restaurant, and although the group at the adjoining table ordered after you did, they get their food before you do.

#6:

More people come down with colds and the flu in the winter than any other season.

#7:

The state of Utah has the lowest per capita rate of beer consumption in the U.S.

#8:

Briteshine is the best-selling brand of flashlight.

#9:

Students who regularly eat breakfast have a higher GPA than those who do not.

#10:

State A had more than double the number of COVID-19 cases as compared to State B.

Notes

1. Taylor, Shelley E., Peplau, Letitia Anne, Sears, David O. (1997) *Social Psychology 9th ed.* Upper Saddle River, NJ: Pearson/Prentice Hall, Chapter 3.

Inductive Generalizations

SECTION 1: INTRODUCTION

The ability to reason using generalizations is one of our most basic rational functions. We generalize all the time, and once we believe a generalization we readily apply it to new cases. Put otherwise, we have a natural capacity to reason from experience to a generalization and once we believe some generalization is true, we are able to easily put this knowledge to use. We never suspect, for example, that the elderly woman walking behind us down the sidewalk intends to rob us. We never suspect this because we have the general belief that few elderly women are street thieves. Reasoning to and from generalizations is largely an inductive process, and in this chapter we will focus on the practice of reasoning to a generalization from particular instances. We will discuss what makes these arguments distinctive, identify some crucial questions for determining logical strength, and will end by considering a common bias that influences the process of generalization.

SECTION 2: GENERALIZATIONS

Generalizations are common in our reasoning. In order to think about them clearly, let's get some examples on the table (starting with one of the examples from the introduction):

- Ex. 1: Few elderly women are street thieves.
- Ex. 2: Most science textbooks cost more than \$150.
- Ex. 3: 77% of felony criminals are repeat offenders.
- Ex. 4: All birds are warm-blooded.

You can probably see a pattern here. In each of these examples, we are being told something about the members of a group. Among the group of elderly women there are few street thieves, and among the group of science textbooks most cost more than \$150, and so on. Accordingly, we will say that a **generalization** is a claim about how many members of a particular group have some property, feature, or characteristic.

There are three notable parts to any generalization. First there is a quantity term ('few', 'most', '77%', and 'all'). As you can see from the examples above, the quantity term can be more or less precise. Second, there is the subject term. In a generalization, the subject term refers to the group we are talking about a quantity of ('elderly women', 'science textbooks', 'felony criminals', and 'birds' respectively). This group—the group

referred to by the subject term—is called the **subject class**. Third, there is the predicate term. The predicate term refers to the property, feature, or characteristic had by the members referred to by the subject term ('street thieves', 'expensive', repeat offenders', and 'warm-blooded' respectively). We will refer to the group of all things that have the property, feature or characteristic identified by the predicate as the **predicate class**. Using these terms, we can understand Ex. 4 as saying that every individual that is a member of the subject class 'bird', is also a member of the predicate class 'warm-blooded things'. This will probably strike you as a somewhat stilted and unnatural way to read this simple sentence. It is, but putting things this way will help us understand more clearly what is going on in generalizations moving forward.

Generalizations that make a claim about *all* or *none* of the members of the subject class are called **universal generalizations** (see Ex. 4 above). For the time being we will put aside these generalizations to focus on generalizations that claim something about most or few of the subject class. Such generalizations are called **statistical generalizations**. We are not always clear about whether we intend a universal or a statistical generalization, and in many cases people leave the quantity term unstated. Consider the following:

Ex. 5: Dogs like to chew on bones.

The quantity term is unstated, and so this generalization is ambiguous. While it isn't always important to identify precisely what an author or speaker intends, sometimes it is. In these cases, normally it is charitable to interpret ambiguous generalizations as expressing statistical generalizations. The reason for this is that it only takes one counter-example to show that a universal generalization is false. That is, if we interpret Ex. 5 as a universal generalization all we need to do is find one dog that doesn't like bones. A statistical generalization claims only that most dogs like to chew on bones and so will not be falsified by a single counter-example.

How do we use generalizations in our reasoning? Perhaps the most important way we use generalizations is to make predictions about particular people, objects, or events. To return to a previous example, given the generalization that few elderly women are street thieves as a premise, we can apply this to a particular case and conclude that the elderly woman walking down the street behind us is probably not going to try to rob us. We not only use generalizations as premises, however; often we reason to a generalization in a conclusion. Put otherwise, sometimes a generalization is *the result* of our reasoning, not a starting point for it. In what follows we will separate our discussion of generalizations into these two contexts—reasoning *from* a generalization (Chapter 15) and reasoning *to* a generalization (this chapter).

SECTION 3: INDUCTIVE GENERALIZATIONS

Let's begin by looking at a couple of examples of reasoning *to* a generalization. Here is example that might appear in everyday conversation:

Ex. 6:

Imani says, "The textbook prices for science classes are out of control! Most science textbooks cost more than \$150! I mean over the last two years I've bought 6 science textbooks, and they have all cost more than \$150."

Imani is drawing two conclusions here. First, on the basis of her experience buying science textbooks, she is concluding the most science textbooks cost more than \$150. Second, she is using the generalization about textbook prices to argue that textbook prices are out of control. Let's zero in on the first argument because this is where she reasons to a generalization.

She argues:

1. I've bought 6 science textbooks and each one cost more than \$150.
2. So, most college textbooks cost more than \$150.

Here is another example:

Ex. 7:

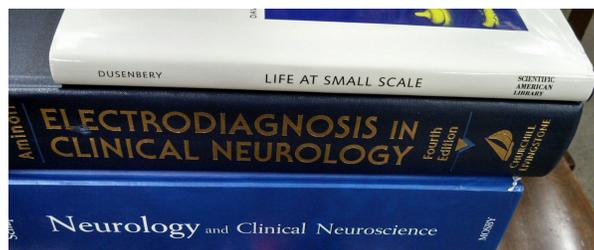
Carly: I'm thinking about trading in my laptop for a new Copperbook.

Jason: Why would you do that? I had Copperbook before, and I was always having problems with it. If you ask me Copperbooks are junk!

Jason is reasoning to the conclusion that Copperbooks are junk, on the basis of his experience owning a Copperbook. Standardized, Jason is arguing:

1. I had a Copperbook before, and was always having problems with it.
2. So, Copperbooks are junk!

When these examples are standardized, it is not difficult to see that both arguments have a significant gap. Both arguments start with an experiential statement: Imani's experiences buying textbooks and Jason's experience with his Copperbook. Both arguments then jump to a much broader generalized conclusion. Imani draws a conclusion about most science textbooks, and Jason concludes that most Copperbooks are junk. Given the *Golden Rule of Argument Interpretation*, we will take it for granted that each speaker is assuming a connection between the stated premise and conclusion. What is it likely to be?



"Biology and Neuroscience books" by brewbooks CC BY-SA 2.0

Let's focus for a moment on Ex. 6. Despite the fact that Imani is drawing a conclusion about most science textbooks, Imani has not bought, or even seen, most science textbooks; after all, there are hundreds of science textbooks on the market, and she has experience shopping for only a few of these. In this argument, she is applying what she knows about the textbooks she's bought, to textbooks outside of her experience, and this only makes sense if she is assuming that the science textbooks she *doesn't* know about are just like the ones she *does* know about. We can capture this idea using the term 'representative':

1. I've bought 6 science textbooks and each one cost more than \$150.
2. The books I've bought are representative of most college textbooks with regard to price. (MP)
3. So, most college textbooks cost more than \$150.

The deep analysis of this argument shows us how the premise connects to the conclusion, namely because Imani takes her experiences to be like or representative of science textbooks more generally. More broadly speaking, in arguments of this form a conclusion is drawn about a group on the basis of observations or experiences of only a part of that group. Arguments with this form are known as *inductive generalizations*.

In thinking about *inductive generalizations*, it will be helpful to add two more terms to our vocabulary: sample and population. In an *inductive generalization*, the **sample** is the group of experiences or observations that serve as premises, whereas the **population** is the group the conclusion generalizes about. In different terms, the population is the subject class of the concluding generalization. Thus, Jason's sample is the one Copperbook within his experience, and the population is Copperbooks in general. As a brief aside, the term 'population' is often used to refer only to groups of people, as in "the population of Cincinnati is about 300,000." However, we are using this term much more broadly to cover anything that is the subject of a generalization. Given these terms, we can define *inductive generalizations* more specifically:

An *inductive generalization* concludes that the population has some characteristic because the sample has that characteristic.

Inductive generalizations are a particularly well-known and well-studied inductive form, and this is, in part, due the importance they play in public discourse. After all, polls and surveys are *inductive generalizations*. Consider the following:

Ex. 8:

According to a *Muhlenberg College/Morning Call* poll taken in early fall of 2016, Hilary Clinton was leading Donald Trump in the presidential race 40% to 32% among likely Pennsylvania voters.¹

This makes a claim about a very large population of people—namely likely Pennsylvania voters. How did the pollsters behind the *Muhlenberg College/Morning Call* poll figure this out, that is, what was their evidence for this conclusion? They did not go out and ask every likely Pennsylvania voter (doing so would cost too much and take too long); rather they took a poll of a small group of likely Pennsylvania voters (as it turns out—405 people) and generalized upon those results to make a claim about Pennsylvania voters as a whole. We can represent the reasoning here as follows:

1. In early fall of 2016, 40% of the 405 polled likely Pennsylvania voters planned on voting for Clinton while 32% planned on voting for Trump.
2. The sample is representative of the population with respect to voting plans. (MP)
3. So, in early fall of 2016, close to 40% of all likely Pennsylvania voters planned on voting for Clinton while close to 32% planned on voting for Trump.

Ultimately, in an *inductive generalization* we extend or extrapolate upon our experience to draw a conclusion about a whole class of unexperienced objects or events. In doing so, we claim that our experiences of some class match or are representative of that class more broadly. Our goal, when it comes to *inductive generalizations*, is to distinguish the logically strong ones from the logically weak ones. Let's take up this task.

SECTION 4: REPRESENTATIVE SAMPLES AND LOGICAL STRENGTH

Our deep analysis above shows that the core element of *inductive generalizations* is the claim that the sample matches, reflects, or is otherwise representative of the population. This means that the most important

part of evaluating an *inductive generalization* for logical strength is to determine whether the sample really is representative. Before we walk through the key steps in this process, we need to take note of the fact that a sample can be more or less representative of the population with respect to some characteristic. A sample can be perfectly representative in which case the sample is exactly like the population, or it can be less than perfectly representative in which case it approximates the population to a certain degree. In a logically strong *inductive generalization*, the sample approximates the population to a significant degree. Put this way, this is a pretty vague requirement, but in most everyday contexts we won't be able to be any more specific. When it comes to professional surveys and polls, however, a number of techniques allow pollsters to be very precise. The *Muhlenberg College/Morning Call* poll above, for example, gives a margin of error of 5.5 percent. The **margin of error** in a survey or poll is a measure of how representative the sample is of the population with respect to the characteristic in question. In giving the margin of error in the poll above, the authors are saying that in early fall of 2016, somewhere between 34.5% and 45.5% of all likely Pennsylvania voters planned on voting for Clinton (40% plus or minus 5.5%), and somewhere between 26.5 and 37.5% planned on voting for Trump (32% plus or minus 5.5%). Overall, when we are evaluating *inductive generalizations* for logical strength, we are trying to determine whether the sample approximates the population to a significant degree.

SECTION 5: IS THE SAMPLE REPRESENTATIVE?

So how can we know that a sample approximates the population to a significant degree? The answer is twofold. First, the sample must be sufficiently **large**. In order for a sample to be representative it must be large enough to capture the diversity of the population. To illustrate, return to Ex. 7 above. Jason is drawing a conclusion about most Copperbooks, namely that they are junk, on the basis of one Copperbook. Clearly this is a weak argument. The fact that one Copperbook laptop is unreliable is no more a guide to the reliability of Copperbooks overall than the fact that one human is 8 feet tall is a guide to the height of humans more generally. Why? Just as one finds a lot of variability in the height of humans, so too one would expect to find a lot of variability in the reliability of a particular brand of laptop, and a sample that is not large enough to reflect this variability cannot be a representative sample. To generalize on a sample that is too small is to draw a premature conclusion that is not justified by the evidence. Let us call a failure with respect to sample size, a *hasty generalization*. More specifically, we will say that:

A person gives a **hasty generalization** when they generalize on the basis of a sample that is too small to capture the diversity of the population.

Thus, Jason has made a *hasty generalization* in drawing his conclusion that Copperbooks are junk. Again, his sample size is way too small to capture the diversity of Copperbooks with respect to reliability.

Sample size itself does not guarantee that a sample is representative. After all, a sample might be large enough to capture the diversity of the population without actually doing so. This brings us to the second key issue when it comes to evaluating *inductive generalizations*: is the sample **sufficiently diverse**? A sample is not sufficiently diverse if it disproportionately includes observations or experiences that will likely bias the sample. If we wanted to know, for example, what percentage of American baseball fans favor for the Yankees, it would be a bad idea to get our sample by polling people at the entrance to a Yankees home game! After all, we would get a disproportionate number of Yankee fans in our sample. Rather, you would want to collect a sample of baseball fans from all over the country. An *inductive generalization* whose sample is not sufficiently diverse is called a *biased generalization*. More specifically, we will say that:

A person gives a **biased generalization** when they generalize on the basis of a sample that does not adequately reflect the diversity of the population.

As is probably clear, sample size and diversity are related. A sample cannot be diverse if it is too small. Again, just because a sample is large does not mean that the sample is sufficiently diverse. One of the most famous mistakes in polling history illustrates just this fact. The candidates in the presidential election of 1936 were Franklin Roosevelt and Alf Landon. As part of an effort to predict the outcome of the election a magazine called *Literary Digest* sent out 10 million questionnaires and received 2.5 million back. The returned questionnaires indicated that Landon would win 56% to 44%. In reality, the results of the election were reversed! Roosevelt beat Landon 62% to 38%. The sample size in this case was sufficiently large—in fact it was massive. The problem is that the addresses to which the questionnaires were sent were taken from phone books and various club membership lists, and in 1936 many people did not have phones—particularly the poor. Similarly, these same people tended to not show up on club membership lists. The sample did not reflect the views of this sub-group of the American population and most importantly, this sub-group voted almost exclusively for Roosevelt. Thus, although the sample was large enough, because its diversity was compromised it was a biased generalization.

Consider another illustration of the relationship between sample size and diversity: a population with little to no variability or diversity with regard to some characteristic. Here is an example:

Ex. 9:

A chef is making soup as part of the evening's menu. As she finishes, she takes a quick taste of the soup, and decides the soup needs a little salt.

It might be surprising, but this is an *inductive generalization*. Is it a good one? On its face, the answer seems to be no. The chef took one taste and drew a conclusion about the soup as a whole, and so we might think this argument is as weak as Jason's in Ex. 7. But remember, we insist on large sample sizes primarily to ensure diversity. Soup is something that is largely homogenous—one spoonful of soup tends to be just the same as any other. There just isn't much difference between one spoonful of soup and another with regard to taste. Because there isn't much diversity among spoonfuls, a very small sample size—maybe even one—can suffice. There are many examples of populations that are mostly or entirely homogeneous with respect to some characteristic or another: most fish have gills, most cars have radios, and so forth. Overall, then, when we have reason to believe that some characteristic is not likely to vary much across a population, then a small sample sizes can be sufficient to ensure logical strength.



"coconut pumpkin soup" by stu_spivack CC BY-SA 2.0

As a final note, when it comes to scientific polls and surveys, it is important to be aware that there are sophisticated mathematical and statistical methods pollsters can use to find out how large a sample must be. Similarly, there are a variety of sampling techniques one can call upon to try to ensure a diverse sample. A discussion of these methods and techniques is outside the scope of this introduction. Nevertheless, there are a few things a non-specialist should keep in mind when it comes to polls and surveys in particular. First, finding out how big a sample was and how it was collected will often give you enough information to identify suspect polls and surveys. Second, statisticians have found that a sample of around 1000 people—if *properly selected*—can accurately represent the view of millions of people. For example, according to the National Council on Public Polls, the average result of the 19 major nationwide polls collected right before the 2008 presidential election varied less than 1% from the actual results of the election.² While this is an amazing level of precision, it is important to point out that national polls are not always this accurate (as the U.S. presidential election of 2016 showed). Third, whether a sample is biased or not is largely a matter of how the sample was collected. The single best way to prevent bias from creeping into your sample is to take a **random sample**. A sample is random when every member of the population has an equal chance of being chosen as a member of the sample. That a sample is random does not guarantee that it is unbiased, but it greatly limits its likelihood. The *Muhlenberg College/Morning Call* poll, for example was drawn from a random sample of telephone numbers (conventional and cellular).

Where does this leave us? The crucial question to answer for any *inductive generalization* is whether the sample is representative of the population. We have identified two steps to answering this question. First, we determine whether the observations or experiences in question constitute a large enough sample. Second, we determine whether the observations or experiences in question constitute a diverse enough sample.

Two Questions to Ask of *Inductive Generalizations*:

- Is the sample large enough? (Failure: Hasty Generalization)
- Is the sample diverse enough? (Failure: Biased Generalization)

Section 6: The Psychology of Generalization

In everyday contexts we rarely have access to carefully and rigorously collected information. More commonly when we generalize we merely consult our experience. Think about the example at the beginning of the chapter: few elderly women are street thieves. Why do you believe this generalization? Presumably it is

not because you have examined crime statistics; rather you likely believe this because you can't recall ever having heard of a mugging perpetrated by an elderly woman. That is, you have likely never experienced this nor even heard of a case. Because this is an unheard-of event, you conclude that this is probably pretty unusual overall.

In general, using this procedure to judge the frequency of events is very common; in many everyday contexts we use remembered experiences as a guide for drawing and assessing generalizations. Of course we might be wrong. After all, as we have seen, it may be that our experiences are simply not representative. But there is a deeper potential problem—it may be that our *memory* of our own experiences is not representative, even if our actual experiences are!

In the 1970's, psychologists Amos Tversky and Daniel Kahneman used a number of simple experiments to show that when we consult our experiences in this way, we unreflectively tend toward one strategy in particular. When we appeal to this strategy, which they called the **Availability Heuristic**, we use the *ease* with which instances of some event come to mind, or are *available*, as an indicator of the frequency of the event. On one hand this might not sound like much of a discovery—after all, it is probably difficult to think of a mugging perpetrated by an elderly woman because this simply doesn't happen very often! Tversky and Kahneman were well aware of this, and observed that "Availability is a useful clue for assessing frequency or probability, because instances of large classes are usually recalled better and faster than instances of less frequent classes."³ Despite its usefulness, psychologists have found that this heuristic lies at the root of a number of different kinds of mistakes.

In general, the problem with the availability heuristic is that some events come to mind more easily than others, *but not because we have experienced them more often*. As it turns out, there are a number of factors which tend to make events stick out in our minds more easily than others. Events that are especially vivid, interesting, or surprising are easier to recall; so too recent events are easier to recall than events that occurred long ago, as well as those that have some kind of personal significance. Because of these kinds of factors, our judgments of the frequency of particular kinds of events will be skewed, and so too therefore, will the generalizations we make from them.

For example, consider the prevalence of heavy drinking among college students. Students who are themselves heavy drinkers tend to overestimate the number of other students at their institution who are heavy drinkers, and the availability heuristic explains why. Because drinking tends to be a social event, it wouldn't be surprising to find out that it is much easier for a heavy drinker to bring to mind other heavy drinkers, than it is to bring to mind non-drinkers. Another interesting illustration of the availability heuristic can be found in people's (mis)judgments of risk. People tend to think that dying in an accident is much more likely than dying by stroke, when, in fact, the opposite is the case. So too people mistakenly tend to think that dying by tornado is more likely than dying by asthma attack.⁴ Given the availability heuristic it is not difficult to understand these kinds of mistakes. Accidents and tornados are highly significant events to which people (and the media!) tend to pay a lot of attention. The same cannot be said for strokes or asthma attacks.

One final example of the availability heuristic is found in the judgments that people living together make about each other's contributions to household chores and other shared responsibilities. In general, people tend to overestimate their own contribution and underestimate the contribution of others. This is actually an example of a much broader phenomenon. Overestimation of one's contribution is common in cases where two or more people are working together towards a shared goal.⁵ Again, the availability heuristic explains why. We tend to remember our own contributions to the group's work much more easily than the

contributions of others. Given the availability heuristic, it seems to us as if our contributions were more frequent than others'.

Again, this is not to say that we should not make use of this strategy. In fact, we can't help but make use of it, and it gives us accurate frequency judgments at least some of the time. Rather, the lesson to take from these observations about the psychology of generalization is that we should pause to second-guess important generalizations that are based on intuitive frequency judgments.

EXERCISES

Exercise Set 14A:

Directions: For each of the following first determine whether it is a generalization or not. If it is a generalization, say whether it is a universal or a statistical generalization, and then identify the quantity term or likely quantity term (if the claim is ambiguous), the subject class, and the predicate class.

#1:

Most new video games cost more than \$50.

#2:

No prisoner has ever succeeded in breaking out of this prison.

#3:

The diamond was stolen either by the student or by the lawyer.

#4:

For 10 percent of the time U.S. drivers are behind the wheel, their eyes are off the road due to eating, reaching for the phone or texting.⁶

#5:

Cattle are frightened by shiny objects.

#6:

Almost everybody who started the race finished it.

Exercise Set 14B:

Directions: Determine whether each of the following is an inductive generalization or not. If so, identify the sample and the population and then comment on the argument's logical strength. Is the argument likely to be logically strong or not, why?

#1:

Almost everybody on the soccer team is closely following the World Cup tournament, so I suspect that most Americans are following it too.

#2:

Your friends keep telling you about this hilarious show and they have watched several seasons of it. So, you sit down and watch an episode. You don't think it was very funny, and conclude "I don't know what my friends were thinking, that show isn't very funny."

#3:

You dip your foot into the river, and conclude that the water is too cold for swimming.

#4:

Most species of tomato love hot weather, so these cherry tomatoes are probably doing pretty well.

#5:

A nationwide poll of a random sample of thousands of college graduates revealed that 75 percent of them are in favor of the proposed debt forgiveness program for student loans. Therefore, roughly 75 percent of the adult population is in favor of this program.

#6:

I don't think that true, selfless love really exists, since I've never experienced it or seen it in other people.

#7:

It is unfair that you won't let me dance in the recital because of my injury. You let Peter dance last year, and my situation is no different from his.

#8:

A major internet news website asks users to contribute to a daily poll. Today's question is: when you buy gas do you normally buy the lowest octane fuel available at the station or do you buy higher octane gas? About 2500 people responded, and according to the results, only 20 percent of these people normally buy the lowest octane fuel. Commenting on this poll, a writer for the website writes a piece entitled: "Why about 80% of Americans buy higher octane fuel".

Exercise Set 14C:

#1:

We generalize on or extrapolate from our experience all the time. Give three examples of generalizations you have made or heard.

#2:

Pick one of the examples you just gave and evaluate it for logical strength.

#3:

Other than those mentioned above, what kinds of events do you think people overestimate due to the availability heuristic?

Notes

1. Olson, L. (2016, Sept. 17). New Morning Call/Muhlenberg College poll shows Clinton ahead in Pennsylvania. *Morning Call*. <https://www.mcall.com/news/local/mc-pa-trump-clinton-poll-20160917-story.html>.
2. National Council on Public Polls. (2008, Dec. 18). NCPP Analysis of Final Presidential Pre-Election Polls, 2008. http://www.ncpp.org/files/NCPP_2008_analysis_of_election_polls_121808%20pdf_0.pdf.
3. Tversky, A. and Kahneman, D. "Judgment under Uncertainty: Heuristics and Biases." *Science* 185 (4157), 1127.
4. Lichtenstein, S., Slovic, P., Fischhoff, B., Layman, M., & Combs, B. (1978). "Judged frequency of lethal events" *Journal of Experimental Psychology: Human Learning and Memory*, 4, 551-578.
5. Ross, Michael and Sicoly, Fiore. (1979). "Egocentric Biases in Availability and Attribution." *Journal of Personality and Social Psychology* 37: 322-336.
6. US Drivers Take Eyes Off the Road 10 Percent of the Time. (2014, Jan. 1); UPI.com;http://www.upi.com/Health_News/2014/01/01/US-drivers-take-eyes-off-the-road-10-percent-of-the-time/UPI-23811388634974/)

Inductive Applications

SECTION 1: INTRODUCTION

Statistical generalizations tend to be used in two ways in our reasoning: as premises and as conclusions. As we have seen, in an *inductive generalization* we start with premises about individuals or small groups and move to a generalization about a whole population. However, we can also do the reverse, and reason *from* a conclusion. That is, we can start with a generalization about a whole population and move to a conclusion about an individual or small group. In this chapter we will take a close look at the process of applying generalizations to specific cases. As we will see, these arguments bear a similarity to some of the deductive argument forms discussed in Chapters 10 and 11.

SECTION 2: REASONING FROM A STATISTICAL GENERALIZATION

We will refer to an argument which draws a conclusion about an individual or small group on the basis of a statistical generalization an *Inductive Application* (these arguments are also sometimes called statistical applications or statistical syllogisms). Let's take a look at an example.

Ex. 1:

Chandler: I doubt that Tony is a Cubs fan.

Alex: Why?

Chandler: I know Tony loves baseball, but he grew up on the south side of Chicago, and most baseball fans who grew up on the south side of Chicago are White Sox—and not Cubs—fans.

Chandler is reasoning *from* a generalization. She starts with a statistical generalization about people who grew up on the south side of Chicago, and uses it to draw a conclusion about an individual—Tony. We can standardize Chandler's argument as follows:

1. Most baseball fans who grew up on the south side of Chicago are White Sox—and not Cubs—fans.
2. Tony is a baseball fan who grew up on the south side of Chicago.
3. So, Tony is probably a White Sox—and not a Cubs—fan.

Inductive applications follow a pattern. Although they may be stated with more or less specificity, all *inductive applications* include a statistical generalization as a premise. The statistical claim tells us that many, most, or few members of the group that constitute the subject of the sentence are members of the group that

constitute the predicate of the sentence. Put in terms we introduced in the last chapter, the statistical claim in the argument above claims that most members of the subject class, ‘baseball fans who grew up on the south side of Chicago’, are members of the predicate class, ‘White Sox—and not Cubs—fans’.

Inductive applications also rely on a premise which tells us that an individual is or is not a member of one of the classes. In this case, we are told that Tony is a member of the subject class—that is, Tony is a baseball fan who grew up on the south side of Chicago. Given that Tony is a member of the subject class, and that most members of the subject class are also members of the predicate class, Chandler concludes that Tony is also probably a member of the predicate class—that is, that he is probably a White Sox—and not a Cubs—fan.

Alternatively, we could have applied the generalization to draw a different conclusion. Suppose that we only knew that Tony was a Cubs—and not a White Sox—fan. We could then reason as follows:

1. Most baseball fans who grew up on the south side of Chicago are White Sox—and not Cubs—fans.
2. Tony is a Cubs and not a White Sox fan.
3. So, Tony probably didn’t grow up on the south side of Chicago.

This is a different way of formulating an inductive application, since although it appeals to the same statistical generalization in premise 1, it adds that Tony is *not* a member of the predicate class, and concludes he is *not* a member of the subject class.



“Chicago Skyline reflection in the BEAN” by mikeyexists
CC BY-NC-ND 2.0

Perhaps you have already noticed that inductive applications have a lot in common with some other arguments we have looked at—namely *Modus Ponens* and *Modus Tollens*. When we know that the individual in question is a member of the subject class, and thereby infer that the individual is a member of the predicate class, we are drawing an inference that is structurally similar to *modus ponens*. Similarly, when we know that the individual in question is not a member of the predicate class, and thereby infer that the individual is not a member of the subject class, we are drawing an inference that is structurally similar to *Modus Tollens*. Despite these similarities, there is one fundamental difference, of course, namely that *inductive applications* are just that—inductive.

Like other kinds of argument, inductive applications are not always explicit. As noted above, sometimes the quantity in question is ambiguous. In other cases, the statistical generalization is left implicit. Consider the following example:

Ex. 2:

Sydney: I am not sure what I am going to do when I get out of college, but I am pretty sure I don't want to be a teacher.

Gina: You don't want to be a teacher? I thought that you were an English major!

In this case, Gina is surprised that Sydney doesn't want to be a teacher, because she is assuming a connection between students who are English majors and students who are pursuing teaching as a career. It is hard to know exactly what the assumed connection is in this case without asking Gina, but it is likely that she is assuming something like this: 'most English majors want to be teachers'. Given this, the unstated assumption is a statistical generalization, and Gina is reasoning using an *inductive application*. This kind of ambiguously stated inductive application is relatively common in our everyday thinking. We naturally and intuitively generalize on our experiences, and once these generalizations are part of our overall system of belief we can use them to draw conclusions about specific cases. As we saw in our discussion of missing premises, we often leave out premises that we take to be obvious and generalizations are no exception. Thus, inductive applications are more common in our reasoning than we might expect.

SECTION 3: EVALUATION: ILLUSORY FORMS

Most errors in reasoning with generalizations are found in the process of generalizing, not in the process of applying the general claim to a particular instance. This is not, perhaps, surprising since the standards for logically strong inductive applications would seem to be clear: in general, the higher the likelihood that members of the subject class are members of the predicate class, the logically stronger the inductive application. Put otherwise, you have better reason to believe that Sydney wants to be a teacher if 90% of English majors want to be teachers, than you do if only 75% do. It is important to note, however, that there are two forms of inductive applications. We will call them (a) and (b):

There are, however, two kinds of mistake we are especially prone to make when dealing with inductive applications. First, recall from our discussion of deductive arguments that there are illusory argument forms. As we saw in Chapter 10, arguments with the form *Affirming the Consequent* and *Denying the Antecedent* often strike people as deductive forms, even though they are not. Similarly, we need to be clear that argument forms we will call *Affirming the Predicate Class* and *Denying the Subject Class* can seem logically strong, but tend to be logically weak. We will add 'probably' to the conclusions below to remind ourselves that these are inductive arguments.

Affirming the Predicate Class (AP—watch out!)

1. Most As are Bs.
2. x is an B.
3. So, probably x is a A.

Denying the Subject Class (DS—watch out!)

1. Most As are Bs.
2. x is not a A.

3. So, probably x is not an B .

Instances of *Affirming the Predicate Class* and *Denying the Subject Class* are often logically weak because in statistical generalizations the predicate class is often much larger than the subject class. Thus, in most cases of *Affirming the Predicate Class* knowing that an individual is a member of the predicate class doesn't tell you anything about whether it is also a member of the subject class. Similarly, in most cases of *Denying the Subject Class*, knowing that an individual is not a member of the subject class doesn't tell you that the individual is not a member of the predicate class. To illustrate:

Ex. 3: *Affirming the Predicate Class* (AP)

1. Most sodas have a lot of sugar.
2. Sweet tea has a lot of sugar.
3. So, probably sweet tea is a kind of soda.

Ex. 4: *Denying the Subject Class* (DS)

1. Most sodas have a lot of sugar.
2. Sweet tea is not a kind of soda.
3. So sweet tea does not have a lot of sugar.

While we might want to clarify what counts as 'a lot' of sugar, these two examples nevertheless illustrate the problem with *Affirming the Predicate Class* and *Denying the Subject Class*. These are both logically weak arguments because, like most statistical generalizations, the predicate class is much larger than the subject class. In this case, there are many drinks with a lot of sugar (predicate class) that are not sodas (subject class). That is, sodas are only one of many kinds of drinks with a lot of sugar (energy drinks, juices, lemonades, iced teas, etc.). For this reason, knowing that sweet tea has a lot of a sugar doesn't allow us to infer, as in Ex. 3, that sweet tea is a soda. The same goes for Ex. 4. Since there are many drinks with a lot of sugar besides sodas, simply knowing that sweet tea is not a kind of soda does not allow us to infer that it doesn't have a lot of sugar.

We need to be clear, however, that there are logically strong instances of AP and DS (*unlike Affirming the Consequent and Denying the Antecedent*). This is unusual, however, and occurs only when the subject class and the predicate class overlap almost completely. Here is an example. Suppose that the employees of a small jewelry store all know the code to the security system that allows them into the store. The store vigilantly updates its code every three months, and has instructed the employees not to share the information with anybody else. Assuming that the employees are trustworthy in this respect, the class of 'people who know the security code' will almost perfectly overlap with the class of 'people who work at the store'. Most, if not all, of the people who know the security code will be employees of the store, and most, if not all, of the employees of the store will know the security code. In this case, the subject class and predicate class almost perfectly overlap, and we could formulate a logically strong instance of *Affirming the Predicate Class* as follows:

Ex.5:

1. Most of the people who know the security code are employees of the store.
2. Jan is an employee of the store.
3. So, Jan probably knows the security code.

Again, what makes this kind of case unusual is that the subject and predicate classes almost perfectly overlap. In the vast majority of our everyday arguments, when we express a statistical generalization we are talking about subject classes that are much smaller than predicate classes.

The fact that instances of AP and DS are commonly weak arguments means that we need to be on the lookout for them. More specifically, we need to be on the lookout for arguments with premises telling us that an individual is a member of the predicate class or that an individual is *not* a member of the subject class. Again, it is not that all such arguments are logically weak, but they often are, and we should take a closer look if we find one. We will call a logically weak instance of AP or DS a *misapplication*.

More specifically, a person gives a **misapplication** when they mistakenly conclude that (i) an individual is a member of the subject class because they are member of the predicate class or (ii) an individual is not a member of the predicate class because they are a member of the subject class.

SECTION 4: RELEVANT SUBJECT CLASSES

When it comes to evaluating inductive applications, the second thing to keep in mind arises from the fact that inductive applications are, like other inductive arguments, defeasible. That is, the discovery of new information can weaken the logical strength of the argument. The chief concern, in this context, arises from the fact that individual objects and people are always members of many subject classes. A person, for example, might be a student, a resident of Pennsylvania, a dog owner, a Steelers fan, a Toyota owner, and so forth, and we need to take this into account in our evaluation of *inductive applications*. Consider the following example. Suppose that we know that Jocelyn is from Nashville, Tennessee. We might draw the following inference.

Ex. 6:

1. Most people from Nashville, Tennessee have been to the Country Music Hall of Fame.
2. Jocelyn is from Nashville, Tennessee.
3. Jocelyn has probably been to the Country Music Hall of Fame.

However, since Jocelyn is a member of many different subject classes she may be a member of a subject class that would undermine the support the premises lend to the conclusion in Ex. 6. That is, Jocelyn may be a member of another subject class that is relevant to the logical strength of the argument. Let us call this a **relevant subject class**. Suppose, for example, that she strongly dislikes country music. Knowing this would significantly undermine our confidence that she's probably been to the Hall of Fame.



"Country Music Hall of Fame & Museum – Nashville, Tennessee" by Timothy Wildey CC BY-NC 2.0

This shows that in evaluating an *inductive application* we have to look outside or beyond the argument as stated to think about what else we know, or can find out about, the individual in question. Again, doing so is in accordance with the *The Rule of Total Evidence* (see Chap 12), and helps us avoid drawing premature conclusions based on partial evidence. We will call an *inductive application* that violates this rule a *hasty application*. More specifically:

A person gives a **hasty application** when they mistakenly conclude that an individual is, or is not, a member of some class because they did not adequately consider other relevant subject classes.

It is important to see that a concern about hasty applications is no merely academic matter—it can be absolutely crucial that we avoid the use of hasty applications. Here are a couple of examples. For most women, oral contraceptives (birth control pills) are perfectly safe means of family planning; consequently, an individual woman might reasonably conclude on the basis of an *inductive application* that using oral contraceptives will be perfectly safe. However, birth control pills are not a safe means of family planning for women with one of a number of different blood clotting disorders. In these cases, the use of oral contraceptives can raise the chances of pulmonary embolism and other dangerous blood clots. Thus, for individuals in the subject class ‘women with blood clotting disorders’ taking oral contraceptives is not a perfectly safe means of family planning.

Here is another example: for most people drinking grapefruit juice is perfectly fine and has no adverse health consequences (other than its taste!). In fact, you might never have heard of anyone for whom grapefruit juice is potentially harmful. Thus you might reason as follows: for most people, grapefruit juice is perfectly fine, so it is probably perfectly fine for me. However, if you are taking statins (a class of drug commonly used to lower cholesterol), drinking grapefruit juice can be dangerous. As it turns out grapefruit juice prevents the body from breaking down some kinds of statins, and this has potentially dangerous consequences. Thus, if you fall into the subject class ‘person taking statins’ then the consumption of grapefruit juice is not perfectly safe.

It is not difficult to see how important it can be for a physician to have all the relevant available information before him or her prior to making even routine suggestions to patients. Of course many of our own *inductive applications* are less important than these; nevertheless, the lesson is the same—in making *inductive applications* we cannot ignore an individual’s membership in other relevant classes.

Two Questions to Ask of Inductive Applications:

- Is the individual in question a member of the subject class *or* not a member of the predicate class? (Failure: Misapplication—AP or DS)
- Is the individual in question a member of other relevant classes? (Failure: Hasty Application)

SECTION 5: EVALUATING INDUCTIVE ARGUMENTS—A SUMMARY

This brings us to the end of Unit #5. We have taken a close look at four of the most common forms of inductive argument, and we have identified the most important questions to ask of these arguments. Of course in the real world, arguments do not come neatly packaged and categorized, and they are often put together to form long chains of argument. We have to do the work of breaking down, identifying, and evaluating for ourselves, and the hope is that you are now in a better position to do so. That is, hopefully you are able to better spot and distinguish *Arguments from Analogy*, *Inferences to the Best Explanation*, *Inductive Generalizations*, and *Inductive Applications*. Moreover, you are now in a position to ask the right questions as

you evaluate these kinds of argument. Of course, these questions will not always settle the matter. Evaluating arguments can be hard, and we don't always have the information we need. Nevertheless, knowing the right questions to ask can push conversations and inquiries forward, and can help direct us toward further information that bears on the logical strength of these arguments.

EXERCISES

Exercise Set 15A:

Directions: Determine whether the following are inductive generalizations or inductive applications. For inductive applications identify the subject class and predicate class of the appropriate generalization. For inductive generalizations identify the sample and population.

#1:

There is no way I am getting one of those little dogs! Little dogs bark all the time.

#2:

She loves going to the movie theater. I've been with her 3 or 4 times and each time she really enjoyed it.

#3:

I don't think I am going to call him this afternoon; he just finished paying bills and that typically puts him in a bad mood.

#4:

Math teachers have weird personalities. I've never had a math teacher I would call normal.

#5:

What do you mean you aren't pre-med? You are part of the campus emergency medical service (EMS) and everyone there is pre-med!

#6:

Jill is an excellent baker, judging from these macaroons.

#7:

I think Anne will probably enjoy the cake, I mean most people like chocolate.

Exercise Set 15B:

Directions: For each of the following inductive applications identify the subject class and the predicate class. Then evaluate the argument for logical strength. If there is a problem with the argument, identify it.

#1:

Casey is probably in a motorcycle club; I mean he has tattoos and most guys in motorcycle clubs have tattoos.

#2:

Most people from Japan do not speak English, so the newly appointed Japanese ambassador to the United States probably does not speak English.

#3:

I doubt that Casey is in a motorcycle club; I mean he doesn't have any tattoos, and most guys in motorcycle clubs have tattoos.

#4:

Most college football players never play professionally after college, so the Heisman trophy winner will probably not play professionally after he is out of college.

#5:

Most obese people either suffer from diabetes or are at high risk for developing diabetes. But since Jack isn't obese, it is likely that he neither suffers from diabetes nor is at risk for developing it.

#6:

Typically, jobs in retail do not come with health benefits. So, I doubt these benefits will be part of your new job.

#7:

I am worried I am suffering from heavy metal poisoning. After all, most victims of heavy metal poisoning suffer from fatigue, and I have been fatigued lately.

Unit #5 Summary

In this unit we looked at four of the most common types of inductive arguments: Arguments from Analogy, Inference to the Best Explanation, Inductive Generalization, and Inductive Application. We use each one of these types of argument multiple times every day. In each case, we learned how to identify these arguments by type, learned what makes them logically strong, and isolated key questions to ask in evaluating for logical strength.

KEY QUESTIONS FOR SPECIFIC INDUCTIVE ARGUMENT TYPES

Two Questions to Ask of Arguments from Analogy:

- Is the noted similarity relevant to the inferred similarity?
- Are there differences that are relevant?

Three Questions to Ask of Inferences to the Best Explanation:

- How likely is the proposed explanation?
- Are there other plausible explanations?
- Would the truth of the proposed explanation be less surprising than the truth of any competitor?

Two Questions to Ask of Inductive Generalizations:

- Is the sample large enough?
- Is the sample diverse enough?

Two Questions to Ask of Inductive Applications:

- Is the individual in question a member of the subject class *or* not a member of the predicate class?
- Is the individual in question a member of other relevant classes?

KEY TERMS

Inference to the Best Explanation

Poor Explanation

Hasty Explanation

Generalization

Subject Class

Predicate Class

Universal Generalization

Statistical Generalization

Inductive Generalization

Sample

Population

Margin of Error

Sample Size	Denying the Subject Class
Sample Diversity	Hasty Application
Hasty Generalization	Misapplication
Biased Generalization	Arguments from Analogy
Random Sample	Analogues
Availability Heuristic	Relevant Similarity
Inductive Application	Relevant Differences
Affirming the Predicate Class	Fundamental Attribution Error

FURTHER READING

For a deeper engagement with many of the issues raised in this chapter see the *Stanford Encyclopedia of Philosophy's* entries on "Analogy and Analogical Reasoning," "Inductive Logic," and "Abduction." See also *Choice and Chance: An Introduction to Inductive Logic* by Brian Skyrms. For more about inference to the best explanation see Peter Lipton's aptly titled book *Inference to the Best Explanation*.

UNIT VI

SOCIAL ARGUMENTS

Testimony and Trust

SECTION 1: INTRODUCTION

Earlier we noted that other people are probably the single most important influence on our thinking. We learn a lot about the world and how it works from others, and in this unit we will focus on this process. More specifically, we will focus on arguments that draw a conclusion about what to believe on the basis of what other people (either individually or in groups) say or are like. As we will see, when it comes to learning from others perhaps the most important step involves judgments of credibility, and we will look at how we make these judgments and what makes a judgment of credibility logically strong. In addition, we often check with others about our beliefs, and use their agreement or disagreement with us as evidence for or against the accuracy of those beliefs. Finally, we will look at what groups can teach us, and under what conditions. After all, other people are rational and make good choices, and when they are doing so, we can use their actions and beliefs as a guide for our own. Of course, this is not always the case, and in Chapter 18 we will distinguish the two. However, in this chapter we will begin with learning from what people tell us, that is, with testimony.

SECTION 2: TESTIMONY AND TRUST

If you stop and think about it, a lot of what you believe, you believe because somebody has said so. For example, you might believe that your mom had a cavity last time she went to the dentist because that is what she told you; you might believe that x-rays were discovered in 1895 because that is what the textbook says; you might believe that *Slumdog Millionaire* won the Oscar for Best Picture in 2009 because that's what came up when you did an internet search. These are all cases of believing something on the basis of testimony. As we will use the terms,

A person believes something on the **basis of testimony** when (and only when) they believe it primarily because somebody else claims it is true.

The term 'testimony' is probably most familiar to us from legal settings. In a trial, for example, a witness might be asked to give her testimony to the jury. In this case, the jury is being asked to believe that what the witness says is true. Testimony, as we will understand it, need not take place in a courtroom or legal context, however. The definition above is wide and includes *any* case a person comes to believe something because somebody else claims it is so. This includes believing things that other people have written down, as well as things they say. To believe that your roommate will be back at 5 pm because that is what the note

says or that Congress passed a bill because that is what the newspaper reports are both cases of believing on the basis of testimony.

As the examples above suggest, we rely on the testimony of other people all the time—and for good reason. After all, the knowledge and experience of a single individual is quite limited. Think about how you know about something as basic as the circumstances of your own birth, for example. You can't remember it, and although there might be pictures or videos of you as a baby, you know many of the details only because your parents or relatives have told you about them. Testimony allows us to easily extend our knowledge far beyond the limits of our own inquiry and experience. In this way, you do not have to go with your mom to the dentist to know about her cavity or do the research to find out the history of x-rays—you can simply take others' word for it. Indeed, it is difficult to overstate the importance of testimony, and hopefully it is clear that our understanding of ourselves and the world is due, in large part, to our wider community.

Importantly, to accept somebody's testimony is to trust that they are a reliable source. When we trust others, we open ourselves to new information, but we also open ourselves up to deception and exploitation. After all, we know that sometimes people lie to us or otherwise try to deceive us. For example, imagine your phone rings and the person on the other side identifies himself as a representative from Microsoft's technical support. He goes on to tell you that your computer is infected with one or more viruses or pieces of malware, and that he will need remote access to your computer to fix it. Put in these terms, this is obviously a bad idea. It is important to remember, however, that smart people fall for these kinds of scams all the time. Presumably they do so for all kinds of reasons, but anybody who succumbs to this scam does so, at least in part, because they have trusted the stranger on the other end of the phone call—they have accepted this person's claim that they are, in fact, who they say they are.



"Consumers Energy warns of phone scams" by Consumers Energy is licensed under CC BY-NC-ND 2.0

Ideally, then, we would trust only reliable sources, and would be able to distinguish reliable from unreliable ones. As it turns out, our normal approach towards the testimony of other people aims toward this goal. Cognitive scientist Dan Sperber and others call our natural approach to the testimony of other people, *vigilant trust*: we tend to accept what other people say as true while at the same time keeping an eye out for signs that they are not trustworthy.¹ Thus, we seek to learn as much as we can from the testimony of other people without being gullible or naïve. Moreover, this is an effective strategy, and in many cases people are both sincere and accurate in what they say. There are plenty of exceptions, of course. Everybody has lied to

or misled others, and sometimes even our sincere beliefs turn out to be false. Thus, an attitude of vigilant trust roughly corresponds to the world as we find it, and is an expedient approach to navigating our social situation. Vigilance takes different forms. On the one hand, much of our vigilance is automatic. As people talk, we implicitly monitor their tone, body language, and so on for cues of insincerity or incompetence. On the other hand, we can be vigilant much more explicitly by actively looking for relevant characteristics of the speaker or source. Before we take a closer look at vigilant trust, it is important to make a couple of brief notes about the evidential value of testimony.

Testimony is normally a quick and easy way to add to our system of beliefs. Nonetheless, it is important to be aware of the limits of this kind of evidence. First, testimony is not particularly weighty evidence. Although, it can certainly give us good enough reason to believe in some cases, in many cases it does not—even when the source is credible. Imagine, for example, that your friend tells you that your professor has moved the final exam to a different time. You may trust your friend, but it is always possible they have made a mistake—and if so, you'll miss the final, and miss your chance for a good grade as well. Since the consequences of being wrong are substantial, you'll want to check on this for yourself instead of simply taking your friend's word for it. More broadly, in cases where the claim in question really matters or has significant potential costs, you will want to look for additional evidence rather than relying solely or primarily on testimony.

Second, compare the following two cases: (a): you work a math problem and come to the conclusion that the answer is '42'; (b): you don't work the math problem, but come to believe that the answer to it is '42' because your friend tells you the answer. Let's say that the answer really is '42', and so you'll have a true belief either way. What is the difference? One big difference is that in (a) you will *understand* why the answer is '42' and in (b) you won't. We can generalize on this point to say that, at least in many cases, testimony does not help us understand why something is true, only *that* it is true. Of course, in many cases we do not need to have a deeper understanding; knowing that it is true is good enough. Moreover, testimony can provide information that will allow us to explain or understand broader issues. What a witness tells the jury can help it figure out who committed the crime, for example. Nevertheless, understanding something for yourself, and knowing because somebody told you are very different ways of knowing. In at least some cases, it is important to understand *why* something is true, and in these cases testimony will not suffice.

SECTION 3: TESTIMONY AND *THE CREDIBILITY ASSUMPTION*

As we saw in previous chapters, there are largely automatic and implicit reasoning processes that can inform our conscious judgments, and trust is no different. As we navigate interactions with others, we are constantly and implicitly monitoring them for signs of dishonesty or incompetence. Thus features of the situation or source can leave us with a vague sense that they are trustworthy or suspicious or something in between. While it is important to be aware of these automatic processes, we will focus in this chapter on conscious and explicit inferences. What is involved in choosing to be vigilant? Let's start with a simple case. Say your friend Sierra tells you that your English class is cancelled for today, and on this basis you do not go to class. We can represent your reasoning this way:

1. Sierra says class is cancelled for today.
2. So, class is cancelled for today.

Formulated in this way, we can see there is a gap in this argument. What connects the premise that Sierra says it with the conclusion that it is true? In general terms, we might say—well...Sierra wouldn't tell me that

it is cancelled if she wasn't sure about it. That is, because Sierra is a trustworthy, reliable, or credible source on this issue. More specifically, we will say that a source is credible on some issue only if:

- (i) They are saying what they believe to be true, and
- (ii) They are in a good position to have an accurate belief.

In light of this, we will call the assumption that connects a person's claim that something is true, to the conclusion that it probably is, the **Credibility Assumption**.

Because (i) and (ii) are requirements for credibility, if we have reason to suspect that either (i) or (ii) is not the case, then we have good reason to doubt the source's credibility. That is, we are justified in doubting a person's credibility on some issue if we have reason to think they are not saying what they believe to be true, or if we have reason to believe that they are not in a position to have an accurate belief on the issue in question. To be clear, to doubt a person's credibility on some issue is not to conclude that they are lying or saying something false. Rather, to doubt a person's credibility on some issue implies only that their word is not, by itself, good enough reason for you to believe what they are claiming; you will need additional evidence before you will accept what they've claimed.



"Court Gavel – Judge's Gavel – Courtroom" by weiss_paarz CC BY-SA 2.0

So, what might give us reason to doubt *The Credibility Assumption*? The most natural place to start is with the source itself. Indeed, there are a number of characteristics a source might have which would undermine its credibility. The most obvious example are people who have a *history* of regularly lying, deceiving, or exaggerating. The known perjurer, for example, is not somebody to whom we should give the benefit of the doubt. There are a variety of other characteristics that give us good reason to doubt a source's testimony. In order to get these on the table, consider the following (fictional) example:

Max is on trial for the murder of Ethan. While on the witness stand, Isabel claims that she saw Max leaving Ethan's house around the time of the murder.

Let us suppose that Isabel does not have a history of regularly lying, deceiving, or exaggerating. What other kinds of information might undermine Isabel's credibility? Imagine that we discover the following facts about Isabel and Max's relationship.

Isabel is Max's ex-wife, and the split was not amicable. While the two were married Max gambled away all of the couple's wealth, including a substantial inheritance Isabel had received upon the death of her parents. More-

over, Max repeatedly cheated on Isabel during their marriage, and Isabel is now aware of this fact. As a result of this, Isabel blames Max for ruining her life.

This sad story undermines the value of Isabel's testimony, because it suggests a possible motive for lying. After all, it is possible, given what we know about human nature, that Isabel might lie in order to get revenge on Max for ruining her life. Put otherwise, it is possible that Isabel has an interest in Max's being convicted of the crime—regardless of whether he actually committed the crime or not. Whether a person has *an interest* or something to gain is perhaps the single most important factor in assessing a person's credibility. In general, the more a person has to gain, the less credible they are, and correspondingly, the less they have to gain the more credible. There is a deeper point here, as well. In this example, the fact that Max has ruined Isabel's life gives her an interest or motive for lying, but as we saw in Chapter 9, interests can influence us in less explicit ways. Recall that simply have a preference for one claim or outcome over another can bias our evaluation of information. Thus, knowing that a person has an interest in some claim gives us some reason to think they might be biased in their thinking, and consequently, some reason to be suspicious of their testimony.

Of course, in spite of all these factors Isabel might be telling the truth—she might have seen Max leave the house. Nonetheless, in light of these facts, Isabel's testimony has much less weight than it would have otherwise; the jury cannot merely take her word for it and accept her testimony at face value once they know the background. Let us put aside the story that Max and Isabel were married, etc., and consider an alternative.

Suppose that Isabel was 50 feet away when she claims to have seen Max leaving the crime scene. Furthermore, it comes out during cross-examination that Isabel normally wears glasses because she is extremely near-sighted (she has an uncorrected visual acuity of only 20/200), and she was not wearing her glasses that day.

Again, this information does not necessarily offer evidence that Isabel is lying, but it does undercut her credibility. The fact that her uncorrected vision is so poor suggests that she was not *capable* of seeing accurately enough to identify Max as the person leaving the crime scene. Thus, even though she may believe what she is saying, she is making a claim that is beyond what she could reasonably know, and thereby gives us reason to doubt that her claim is correct.

There is one other related factor it is important to note. Consider the following:

Suppose that Isabel did not see the murderer's face, but did get a look at the handgun used in the crime. Isabel specifies that the gun used in the crime was a 9mm Beretta PX4 Sub-Compact—exactly the same kind of handgun owned by Max. It comes out during cross-examination, however, that Isabel is unable to distinguish between the Beretta Compact and Sub-compact models, and moreover is not even able to distinguish between different brands of 9mm handguns.

That Isabel cannot distinguish these three kinds of 9mm handgun strongly suggests that she does not have enough background knowledge about handguns to reasonably know that the gun she saw was a 9mm Beretta PX4 Sub-Compact. This is not something she could reasonably know because she does not have the knowledge base or *expertise* to make this distinction.

We have considered four especially common factors to be on the lookout for in considering a person's credibility on a specific issue. There are a variety of additional complications that we will take up in the next chapter, but for now we will turn to a second factor to keep in mind in deciding whether to believe somebody's testimony.

SECTION 4: CHECKING FOR PLAUSIBILITY

When we are being vigilant, we need to check for credibility, but we also need to check the claim against our general knowledge of the world (as per the *The Rule of Total Evidence*). That is, we need to ask—is the claim plausible given what we know? To illustrate, imagine that a friend tells you they saw former President Bill Clinton buying a toothbrush at the local pharmacy. Your friend seems sincere and there is no reason to think she is deceiving you. Moreover, she certainly knows what Bill Clinton looks like. In short, there is no reason not to trust her. The problem is that, as you think about it, it just seems so unlikely. You know that Bill Clinton doesn't live in the area, and what are the chances he would be buying a toothbrush here? Put in different terms, this claim seems implausible given your knowledge and experience. You are thus left with competing pieces of evidence, and you need to weigh them against one another. Which is more likely: that Bill Clinton really was at the local pharmacy buying a toothbrush *or* that your friend has misidentified somebody who looks a lot like Bill Clinton?

The main point here is that credibility is only one factor and can be outweighed by the implausibility of the claim. Overall, then, in being vigilant we'll want to ask two main questions.

Two Questions to Ask of Testimony

- Is the source credible? (history of lying? something to gain? capability/expertise?)
- Is the claim plausible given what you know about the world?

Again, to set aside what somebody has said because you find them less than fully credible or because you find their claim implausible is not to disagree with them or to conclude that they are wrong. To set their claim aside is simply to say that testimony is not enough in this case, and that more evidence is required before you'll be willing to add it to your system of belief.

SECTION 5: TESTIMONY AND THE WEB

We are constantly turning to the web for information. We can, for example, type a question into a search engine, and get a wide variety of answers almost immediately. Even when we are not actively searching for information, we are continually presented with news, commentary, opinions, humor, memes, factoids, and so on as we scroll through our friends and contacts on social media. This observation is relevant because it highlights that the internet is what philosopher Michael Lynch calls a *testimony machine*.² Think about it: the information on the internet was originally posted by somebody, and so when we come to believe something because a website, meme, or tweet says so, we come to believe it on the basis of testimony. Indeed, the web vastly increases our access to information in large part by vastly increasing our access to other people's knowledge and experience. At the same time, however, it also increases our exposure to false, misleading, and deceptive sources. For example, in recent years headlines shared through social media have proclaimed that the Pope endorsed Donald Trump for President, Minnesota made Arabic classes mandatory in high school, and the U.S. government issued warnings about bananas contaminated by the Zika virus. These are all false, but were widely shared on social media. As another example, according to a study performed in the Spring of 2021, close to 65% of all anti-Covid vaccination content shared on Facebook and Twitter could be traced back to just 12 people!³ Thus, while the web greatly expands our access to information, it simultaneously allows false and misleading claims to be quickly spread and amplified.



"A social network visualization" by brewbooks CC BY-SA 2.0

A big part of the reason that people are fooled when it comes to internet sources is that "web environments" make it difficult to be vigilant. As we've seen, an important part of being vigilant is looking to the credibility of the source, but the web presents us with many unfamiliar sources. Indeed, in many cases the source's identity is hidden. That is, online we are often not in a position to see whether there are factors undermining a source's credibility at all. Further, even when we are familiar with a source, it can be difficult to pick up on signs of dishonesty, insincerity, or incompetence on the web. Thus, in many cases online we are flying, at least partially, blind.

There are two online contexts it is worth emphasizing in particular on this issue: using a search engine and using social media. A search engine is a handy tool for navigating the sea of information available on the web, and often the first thing we do when we have a question is to "Google it". In thinking about web searches it is important to distinguish at least two kinds of question you might use a search engine to answer. On the one hand, you might have a question about a specific, straightforward, and uncontested fact. For example, we might need to know how many cups are in a gallon or the capital of Peru. We can trust search engines to quickly deliver accurate answers to these kinds of questions. Often, however, we are interested in questions that are more evaluative, interpretive, complex, or disputed. Indeed, most of the important questions we have fall into this second category, and when we use search engines to look for answers we need to be careful. Why?

When we type our question into a search engine, we get a ranked list of links. Some of these links may be to trusted and familiar sources. However, most of the links will be to sites we've never heard of. This means that we won't know anything about who is behind the information posted on these sites, why they posted it, and whether they are in a position to know what they are claiming (without extra work). In short, we won't be in a position to say anything about the credibility of most of these sites. This is not a criticism of search engines as much as it is reminder that internet searches quickly and easily take us to "strangers" whose credibility we may not be able to assess (at least immediately), and whose testimony we should therefore be cautious about.



"Hidden Fire." by Ardinnnn 😊 CC BY-NC-SA 2.0

But wait. The search engine's algorithms have identified this particular ranked list of sites in response to your question. Doesn't that make these results trustworthy? No. While it is true that a search engine brings potentially relevant sites to your attention, search engines do not work by identifying only trustworthy sources. Instead they work, in very general terms, by looking for key words, and finding sites with those key words that are popular (among other factors). Indeed, when you do a search, we end up with all kinds of results. You might get links directing you to online forums like Reddit or Quora, news sites, blogs, Wikipedia or other online encyclopedias, think tanks, trade magazines, videos, advocacy organizations, academic sites, and online businesses to name just a few. And these sources will not be equally reliable or trustworthy sources for answering your question. There is a deeper point here as well, namely that while it is true that a very reliable source might be popular as a result of its reliability, the reverse doesn't hold—the popularity of a site doesn't tell us a lot about its reliability, since there are many reasons a site might be popular besides it's being a reliable source of information as we will see in Chapter 18 (it is entertaining; it tells people what they want to hear, and so on). Consequently, we can't take the mere fact that a link came up in response to our question to mean that it is a trustworthy or reliable answer.

Turning to social media, we all know that you can't trust everything you see on Facebook (for example). Nevertheless, people are regularly fooled by headlines, stories, and pictures that are shared by their connections. What accounts for this? In some cases, we know the source of information that is shared, for example when a friend shares a news story from a well-known and trusted site. There are many cases, however, where the origins of the shared information are unknown. Again, in these situations we are not in a position to see whether the source is credible or not. This is made worse on social media by a number of other factors. First, think about your normal state of mind when you turn to social media. We go to social media for entertainment and to see what our friends and acquaintances are up to, and as we causally scroll through our feed we tend to be in a relaxed and passive state of mind. Second, the way social media feeds *look* can lend the claim of an anonymous, hidden, or unknown source an undeserved air of credibility. After all, the headline, story, or picture was posted by your friend, and is situated right under their name and picture. It can be easy, especially if you aren't being too careful, to extend the credibility of your friend to the claims they've shared. Third, sometimes people simply don't care whether what they share is true. A person might share a headline, story, blog post, or picture on social media primarily because they find it entertaining or interesting. Alternatively, they might do so primarily to express their identity or membership in some group. As we've seen, ideas can have social value, and just as this fact can lead us to have preferences for or against certain ideas, so too can it lead us to share things without a particular interest in their truth.

We've highlighted situations in which we are especially likely to run into unfamiliar sources of information. As we've seen, in these cases we simply can't say anything about the credibility the source, and we should be careful. This isn't to say that we can never trust a source we don't know much about. If an online source tells you something simple, objective, undisputed, and about which there is no obvious reason to mislead, then it might make sense to take its testimony (e.g. How many feet in a mile? What planet is closest to the sun? What is the official currency of Brazil?). However, most interesting claims on the web are not like this. If a claim is surprising, too good to be true, or the sort of thing a person might lie about, then you won't want to take a stranger's word for it. Again, this is not to conclude the claim in question is false or that the source is trying to mislead you, only that it is going to take more than the source's word to convince you the claim is true.

SECTION 6: SUMMARIZING TESTIMONY

There are three points to emphasize in summary. First, vigilant trust aims to steer between being naïve or overly trusting on the one hand, and being overly skeptical on the other. It is to strive for the informational benefits of reliable sources without being taken-in or otherwise misled by unreliable ones. Thus, we want to be open to learning from the knowledge and experience of others, while simultaneously keeping an eye on their credibility and the overall plausibility of their claim.

Second, as we've seen testimony is an important source of information, but it has limitations. In particular, although a person's testimony can give us good enough reason to believe it, testimony does not help us understand. That is, although testimony can give us good reason to think *that* something is true, it doesn't explain *why* it is true. If the claim in question becomes contested or has significant practical consequences, you'll want a more comprehensive understanding than testimony can normally provide, and you'll want to look for further evidence.

Finally, we have been talking about how to navigate people's claims on social media and elsewhere, but this also opens the door to thinking about our own behavior and the kind of reputation we'd like to cultivate. Thus, we can turn things around and ask: do other people see us as a credible source? Are we being sincere in what we say, post, and share? Do we share information we are in a position to know? Or do we suggest things are true, when we really aren't sure? More broadly, it puts us in a position to think about our responsibilities to others when it comes to testimony. This is a particularly pressing question when it comes to social media, given its ability to spread and amplify information. What do we owe to our friends, followers, and connections on social media when it comes to information?

EXERCISES

Exercise Set 16A:

#1:

List three things you believe on the basis of testimony.

#2:

What are the last three things you shared on social media?

Exercise Set 16B:

Directions: Let's practice being vigilant. In each case, can we take their word for it? Why or why not?

#1:

Mr. Ziegler has testified that his son Scott was home with him at the time when Scott is alleged to have shot the victim, so Scott can't have been the shooter.

#2:

"This place has the best Chinese food in Chicago" says your friend who just moved there from out of state last month.

#3:

CEO of Massive Dynamics (a company that makes vitamins among other things): There is no need to be concerned; our vitamin supplements are perfectly healthy.

#4:

Used car saleswoman: this car is a peach; the previous owner was a little old lady who only drove it to church on Sundays.

#5:

You are on a campus visit, and are looking for the admissions office. You walk up to a young person walking down the sidewalk, ask, and they say, "the admissions office is on the second floor of that building right there" as he points.

#6:

"There is no way introducing this little frog into the ecosystem will disrupt it," says the local fire chief.

#7:

The dentist says that the tooth is compromised and has to be removed. I guess I better make an appointment to get it extracted.

#8:

Watch out for Bigfoot on your camping trip this weekend. My friend says she saw one when she was hiking there last summer.

#9:

Dr. Johnson says that man-made global warming is a myth. He has a Ph.D. in geology—he should know.

Exercise Set 16C:

#1:

Think back to the discussion of cooperative dialogue and cooperative disagreement in Chapter 2. What role does trust play in these forms communication?

#2:

Suppose that a person claims to have been abducted by aliens, but is unable to offer any evidence other than their word. What would it take for you to legitimately take this person's word for it—that is, what would you need to know about this person to believe their testimony?

#3:

Look at your social media feed, and find something that strikes you as untrustworthy. What is it, and why does it strike you as untrustworthy. Be specific.

#4:

How can our actions online have positive or negative effects on others? Think of some specific cases. What does this say about our responsibilities to others online or on social media?

Notes

1. Sperber, D. et al. (2010) "Epistemic Vigilance." *Mind and Language* 25:4, 359-393.
2. Lynch, Michael Patrick. (2016). *The Internet of Us*. New York: Liveright, 24-25.
3. Center for Countering Digital Hate (2021, Mar. 24). "The Disinformation Dozen." <https://www.counterhate.com/disinformationdozen>.

Trust and the Ad Hominem

SECTION 1: INTRODUCTION

As we saw in the last chapter, our normal attitude toward the testimony of others is one of vigilant trust. That is, although we tend to accept what other people tell us, we are constantly on the lookout for signs of deception, insincerity, and incompetence. As we have seen, the fact that a person has something to gain or is somehow lacking in skill or expertise can give us a reason not to trust what they say. However, the fact that some features of a person *can* undermine the credibility of what they say, does not mean that they always *do*. This raises a number of questions: namely what features of a person undermine credibility, and under what circumstances? In this chapter we take up these questions. We will zero-in on arguments that appeal to some feature of a person to conclude that they are not credible. Arguments like this are called *ad hominem* arguments (Latin for “against the man”). Perhaps unsurprisingly, *ad hominem*s are extremely common in everyday thinking. Of course, like other inductive argument forms, there are logically strong *ad hominem*s and logically weak ones. Accordingly, in this chapter we will take a close look at how *ad hominem* arguments are used in everyday life, talk about what makes an *ad hominem* logically strong, and identify some strategies for accurate evaluation.

SECTION 2: TO TRUST OR NOT TO TRUST: AD HOMINEM ARGUMENTS

It is important to recognize that there are all kinds of conclusions we might draw about what a person has said based on what they are like, who they are, or what they have done. One common kind of inference appeals to what a person is like to draw a conclusion about their *right to speak* in some particular context. This is what we do when we say, for example: “since you aren’t part of the family, this is none of your business”. Similarly, when we say “people in glass houses shouldn’t throw stones” we are questioning the *fairness or appropriateness* of a comment. Neither of these are *ad hominem* arguments, however. What is distinctive about an *ad hominem* argument is that it is an inference from what a person is like to a conclusion about whether we can trust what they say in a particular case. As we have seen, our trust in a source is embodied in our assumption that they are both sincere and in a good position to have a true belief in this specific case. This is what we called *The Credibility Assumption*. Given this, we can define an *ad hominem* as follows:

An *ad hominem* argument draws a conclusion about *The Credibility Assumption* on the basis some feature of a source.

In principle, almost any feature of a person can serve as the foundation for an *ad hominem* argument; it is common for *ad hominem*s to draw upon a person's physical features, group membership (e.g. religion, class, political affiliation), achievements and failures, motives or interests, morality, and personal history (among many others). However, to say that these features can be used to challenge other people's credibility is not to say that they actually do so. Again, as we will see there are both logically strong and logically weak *ad hominem*s.

In addition, we can distinguish between positive and negative *ad hominem*s. A **positive ad hominem** draws on some feature of the source as support for *The Credibility Assumption*, and to thereby bolster the importance of their testimony. Here is an example:

Ex. 1:

If you are in the market for a new car, I'd hold off until December. My aunt worked in car dealerships for 20 years, and always says the best time to buy a new car is at the end of the calendar year.

In this case, the speaker is treating his aunt as credible on this issue because of some feature of her—namely her decades in the car sales industry. In contrast, a **negative ad hominem** draws on some feature of the source to raise doubt about *The Credibility Assumption*, and to thereby dismiss or set aside their testimony. For example, the cases from Section 4 of the last chapter are all examples of *ad hominem* arguments:

Ex. 2:

We cannot simply take Isabel's word for it that she saw Max leaving Ethan's around the time of the murder, since she has a motive to lie.

The author here is concluding that we cannot accept Isabel's testimony because of some feature of her—namely that she has a motive to lie. While we sometimes use positive *ad hominem*s, negative *ad hominem*s are much more common. Consequently, in the rest of this chapter we will focus on negative *ad hominem*s.



"Don't Trust Anyone Under 30" by gidsey_ CC BY-NC-ND 2.0

SECTION 3: NEGATIVE AD HOMINEMS

Let us begin our discussion of negative *ad hominem*s with some examples.

Ex. 3:

Taylor: Don't worry. The CEO promises that Company A will never sell customers' email addresses.

Cho: You can't take her word for it; she made the same promise when she was working for Company B and it turned out that Company B was selling customer's information as soon as it got it!

In this example, Cho gives an *ad hominem* that challenges the credibility of the CEO. He argues that we cannot take the CEO's claim for granted because of a feature of the CEO, namely that she has broken similar promises in the past. Let's take a look at another example.

Ex. 4:

Lara: I think pot should be legalized, I mean there is no good reason for legally distinguishing between alcohol and marijuana.

Ryan: You love smoking weed. Of course that's what you think.

This is also an *ad hominem*, but Ryan's argument is not nearly as clear as Cho's was above. Ryan is attacking Lara's credibility on this issue on the grounds that she loves smoking weed. He seems to be questioning Lara's objectivity: she can't think about the issue of legalization clearly because she has a clear preference. Moreover, he doesn't come right out and conclude that "we can't take her word for it." Instead he insinuates that we cannot trust her by dismissing her claim, and this is common in *ad hominems*.

In the effort to identify *ad hominem* arguments, it is important to remember that not every argument that appeals to a feature of a person is an *ad hominem*. Consider the following example:

Ex. 5:

Brandon: E. M. Waterhouse claims that we would all be better off if we lowered the tax rate on millionaires.

Omar: That's not true—I mean this is coming from a guy who is a millionaire himself!

Omar's argument in this example is not an *ad hominem*, although it certainly looks like one. Omar is appealing to a feature of E.M. Waterhouse, namely that he is a millionaire, in order to criticize his claim. However, he is not challenging his credibility. Instead she claims that what Waterhouse says is false (he says, "that's not true"). We will call this kind of argument, a **denier**, since they deny what the speaker has said.¹ That is, a denier uses a feature of a person to conclude that what they've said is false or incorrect. In contrast, *ad hominem* arguments only challenge a source's credibility—they do not conclude that what the speaker says is false or incorrect (thought it might be). It is important not to confuse *ad hominems* and deniers, since deniers are almost never logically strong. That is, a feature of a person rarely offers sufficient evidence for thinking that what a person has said is false.

Not only can negative *ad hominems* draw different kinds of conclusions, but they can draw conclusions about the credibility of different subjects. In general, people draw *ad hominem* conclusions about three kinds of communication. On the most basic level, an *ad hominem* can draw a conclusion about a person's credibility with respect to a specific **claim** (see Ex. 1, 2, and 3). Second, an *ad hominem* can draw a conclusion about a person's credibility with respect to an **argument** (as is the case in Ex. 4). Last, we can draw conclusions about a person's credibility when it comes to more **substantial productions**, e.g. a speech, article, book, or movie. For example:

Ex. 6:

I wouldn't trust anything you see in that so-called "documentary"; the director is a well-known conservative activist.

The speaker in this case is giving an *ad hominem* which does not target any specific claim or argument, but challenges the credibility of the documentary as a whole based on the director's political activity. Thus, negative *ad hominem*s come in many different forms. They are often stated in somewhat ambiguous terms, and can target different forms of communication. Moreover, it is important to distinguish *ad hominem*s from a similar form we've called Deniers. Suppose we've done that: we have spotted an *ad hominem*. What now? How do we evaluate it?

SECTION 4: EVALUATING NEGATIVE *AD HOMINEM* ARGUMENTS

As we have said, negative *ad hominem*s all point to some feature of the source to challenge *The Credibility Assumption*. Again, the credibility of a source is embodied in our trust in its honesty, and our trust that the source is in a position to have an accurate belief. Thus, any feature of a person that gives good reason to doubt either of these conditions, thereby gives reason to doubt the source's credibility. To evaluate an *ad hominem*, then, is simply to ask whether the feature in question gives us a reason to doubt either of these conditions. In Ex. 4 above, Cho gives a logically strong *ad hominem*: the fact that the CEO has failed to fulfill similar promises in the past calls into question her honesty in this case. In contrast, suppose Cho had responded this way:

Ex. 7:

Taylor: Don't worry. The CEO promises that Company A will never sell customers' email addresses.

Cho: You can't take her word for it; she has been married three times!

Having been married three times does not give us a reason to think that the CEO is not being honest, nor does it give evidence that the CEO is not in a position to have accurate beliefs about this issue. Consequently, Cho's *ad hominem* in this case would be logically weak. In general, to conclude that there is good reason to doubt a source's credibility is to conclude only that the source's testimony—all by itself—does not give us enough reason to believe what they have said. Again, to draw this conclusion is not to say they are wrong or mistaken. It is simply to refuse to take their word for it, and to agree to wait for more evidence before making a decision one way or another.

The fact that *ad hominem*s can have different targets introduces a complication to the process of evaluating for logical strength. In order to illustrate the issue, consider the following two arguments. One is logically strong, and the other is not. Before reading on, see if you can distinguish the two.

Ex. 8:

Mike: The salesman says that the previous owner of this car was a little old lady who only drove the car to church on Sundays.

Ali: This guy earns a commission if we buy the car, so I am not just going to take his word for it.

Ex. 9:

TV News analyst: Mega Petroleum argues on the basis of the 1990 Oil Pollution Act that they are only legally responsible for paying 75 million in damages for the oil spill. We should be skeptical, however, since of course they want to avoid paying for the clean-up.

The crucial difference between these two cases has to do with the target of the *ad hominem*. In Ex. 8 Ali is questioning the credibility of the car salesman with respect to his *claim* that the previous owner was a little old lady. In Ex. 9, on the other hand, the analyst is challenging the credibility of Mega Petroleum with respect to their *argument*. This matters because when a person simply makes a claim, they are asking you to take their word for it. In contrast, when a person gives an argument, they are not. Rather they are appealing to evidence that they take to justify the conclusion, and this evidence will either stand or fall independent of their credibility or reliability.



"NFL foul-ad hominem attack" by LittleRoamingChief CC BY 2.0

The practical consequence of this is that, in most cases, to use an *ad hominem* against an argument is to violate *The Rule of Total Evidence*. Recall that according to this rule we are required, in formulating a conclusion, to take all the available evidence into account. When a person uses an *ad hominem* to draw a conclusion about an argument they ignore relevant evidence. Consider Ex. 9 above. The analyst in this case is skeptical of Mega Petroleum's claim that they are only responsible for paying 75 million in damages. This makes sense given that Mega Petroleum clearly has an interest in this matter. However, Mega Petroleum has offered an argument for this claim. They are not asking us to trust them, but have presented independent evidence on behalf of this claim. Thus, given *The Rule of Total Evidence* we must take this argument into account if we want to draw a conclusion about the claim that they are only legally liable for 75 million in damages.

The upshot is that if we want to express skepticism about Mega Petroleum's conclusion, we have to suggest that there is something wrong with the argument by challenging either its factual correctness or logical strength. The analyst in Ex. 9 has completely ignored Mega Petroleum's argument by drawing the conclusion solely on the basis of the company's interest in paying only 75 million. As with most inductive arguments, there are exceptions. For example, if a person offers an argument the premises of which rely on the credibility of the person, one might have a logically strong *ad hominem* for an argument. Nonetheless, most *ad hominems* when applied to arguments violate *The Rule of Total Evidence*.

In closing this chapter it is important to point out how useful *ad hominem* arguments are. They can help us know when we ought to ask more questions, as well as guide us to more reliable sources (among other things). Similarly, if a person has a history of exaggerating, then it makes perfect sense to doubt her claim that she saw "like 100 whales" on her vacation. So too when the mechanic says you need 4 new tires instead of just 1, you could just take his word for it. But if he has an interest in your buying 4 new tires you should be skeptical on the basis of an *ad hominem* and ask for more information.

Ad hominems are also useful for sorting through information. Suppose, for example, that you are interested in learning about some controversial topic—say abortion. You want to get a solid sense of the relevant

issues so you can think about it for yourself. To this end there are all kinds of sources available, but you cannot read them all. So how are you going to decide what to read? To simplify, suppose you have narrowed the field to three books. You find out that one of the books is written by the director of a prominent Pro-Life organization, one of the books is written by a prominent member of a Pro-Choice organization, and one is written by a professional bioethicist. Presumably you will choose the book written by the bioethicist, and presumably you will do so on the basis of an *ad hominem*—there is some reason to think that the overall credibility of the two other books might be compromised (they may not offer a fair presentation of the issues, for example). To be clear: this is not to say that these authors cannot offer a balanced presentation of the issues (we are not appealing to a Denier after all), but only that we have some reason to be suspicious that this is the case. In this case we can justifiably use *ad hominem*s to decide among many possible sources.

Nevertheless, we have to be careful with negative *ad hominem*s. We do not want to unfairly or inaccurately challenge a person's credibility. This means that we have to be sure that the feature in question genuinely gives us good reason to suspect the other person's honesty or their ability to have an accurate belief on the issue in question. Often simply taking the time to ask this question will allow us to quickly distinguish between logically strong and logically weak *ad hominem*s. That is, once you've identified the feature in question ask:

Two Questions to Ask of Negative *Ad Hominem*s: *Does this feature...*

- Give us reason to think the source is not being honest or sincere in this case?
- Give us reason to think the source is not in a good position to know in this case?

EXERCISES

Exercise Set 17A:

Directions: For each of the following (i) decide if it is an ad hominem argument or not. If it is a Denier say so. (ii) Using the questions above, briefly comment on the arguments' logical strength. Be ready to share your answers.

#1:

Health Inspector to Restaurant Owner: Sorry, but I just can't take your word for it that everything is up to code. In recent years you have claimed to be up to code, but have actually had numerous health code violations.

#2:

Nixon is surely our worst president; after all, no other president has had to resign the office.

#3:

A: According to Senator X, without a boost to defense funding, the U. S. will be at increased risk for a terrorist attack.

B: That has got to be false; Senator X has been a staunch supporter of the defense industry, nobody in the Senate receives more campaign contributions from this industry than X.

#4:

A: The Company's accountant claims that she is not aware of any accounting irregularities.

B: Whatever, don't you know that she owns thousands of shares of the company's stock!

#5:

A: C thinks my relationship with D is not healthy and that we should go out on more dates instead of hanging around the dorm all the time.

B: I wouldn't exactly take relationship advice from C. She's never had a long-term relationship.

Exercise Set 17B:

Directions: Assume that all the following are ad hominem. For each case, evaluate the argument for logical strength using the questions above. Explain your answer.

#1:

A: I am totally on board with limiting the capacity of gun magazines. I think that the danger to society of large-capacity magazines outweighs individual's preferences for them.

B: Whatever, you've never even held a gun, much less fired one!

#2:

You can't take what he says seriously! I mean the guy's name is 'Cletus'!

#3:

A: The football team is underfunded. We have by far the lowest operating budget of any team in our conference. We can't even afford to have our jerseys washed after every game!

B: It is no surprise you'd say that since you're on the team.

#4:

You can't take the candidate's economic policy seriously, I mean the guy has an elevator—for his cars—and gets a \$70,000 tax break—on his horses!

#5:

Kid to mom: "How can you stand there and tell me I shouldn't smoke pot? I know that you did when you were in college!"

#6:

A: The American prison system needs to be reformed. Per capita the U.S. imprisons more of its citizens than any other country in the world.

B: *This*, coming from a convict. Yeah, right.

Exercise Set 17C:

#1:

This chapter is all about distinguishing logically strong *ad hominem*s from logically weak ones. But why does this matter? That is, why should we care whether our *ad hominem*s are good or not?

#2:

In this chapter we have focused primarily on interpersonal *ad hominem*s. But as we've seen, we use *ad hominem*-style reasoning to sort out trustworthy and untrustworthy sources more broadly. Are there any sources of information you tend to dismiss? On what grounds? Explain.

#3:

What is a conflict of interest? Give an example. How are conflicts of interest related to *ad hominem* arguments?

Notes

1. Use of the term 'denier' follows Sinnott-Armstrong, Walter and Fogelin, Robert. (2010). *Understanding Arguments 8th ed.* Belmont, CA: Wadsworth Cengage, 355-56.

Agreement, Disagreement, and Popularity

SECTION 1: INTRODUCTION

To this point we have highlighted the importance of testimony, and discussed the conditions under which we should accept what other people tell us. However, we can look to other people for other kinds of information as well. In some cases, the mere fact that other people agree or disagree with us can tell us something important about the accuracy of our beliefs, and give us reason to update them accordingly. Moreover, we can look to groups of people for information. In these kinds of cases we take the fact that something is popular as evidence of its merit. In the right circumstances, and when the conclusion is drawn carefully, this can make perfect sense. In this chapter, we will look at these ways of learning from others, identify key questions to ask, and think about some of the factors that influence what we infer from other people's beliefs.

SECTION 2: AGREEMENT AND DISAGREEMENT AS EVIDENCE

Recall our discussion of cooperative disagreement in Chapter 2. There we began by noting that disagreements naturally raise all kinds of good questions: *Why don't we agree? How could somebody see this differently? What are my reasons for thinking this in the first place? Am I missing anything?* The goal of the discussion in Chapter 2 was to show the value of engaging cooperatively with others to understand disagreements, and to highlight some of the social and personal obstacles to doing so. As we saw, a disagreement can indicate that we've missed something or made a mistake, and can thereby present us with an opportunity to correct ourselves. Similarly, as we will see, the fact that somebody else agrees with us can bolster our confidence in our belief. In this section, we will set aside the issue of cooperative dialogue as a social practice to focus on agreement and disagreement as evidence. That is, we will focus on clarifying when, and under what conditions, agreement and disagreement tell us something about the accuracy of our own beliefs.

In order to illustrate how agreement and disagreement can serve as evidence, let us start with a couple of examples.

Ex. 1:

Imagine that you and a friend individually work through the same complex math problem. You get the answer '42'. Your friend has roughly the same math skills you do, but she got '38.5'. Upon finding this out you are less confident your answer is correct, and go back to double-check it.

Ex. 2:

You are in the same situation, but when you check with your friend she says '42' as well. Hearing that she agrees with you makes you even more confident your answer is correct, and you move on to the next one.

In these examples, your friend's agreement/disagreement tells you something about your belief. Importantly, however, this is not always the case. Sometimes agreement and disagreement don't suggest anything about the accuracy of your own beliefs. To illustrate, say that in Ex. 1 you methodically worked through the problem, but your friend was distracted and did it quickly and carelessly. Upon finding this out, your friend's disagreement shouldn't make you less confident, since you are more likely to have gotten it right than they are. Alternatively, imagine in Ex. 2 you find out that your friend did not actually work the problem, but instead looked at your answer, and reported it as her own. Upon finding this out, your friend's agreement shouldn't make you any more confident—after all, she is just repeating back your own answer. Consequently, the fact that somebody agrees or disagrees with us is informative in some cases but not others. What is the difference?

Let's start with Ex. 1. In this case, it makes sense for you to be less confident about your answer, but why? On the one hand, you have diligently worked through the problem, and so you have good reason to believe the answer is '42'. On the other hand, in checking with your friend about the answer, you are treating her as a credible source for this problem. That is, you are taking for granted that she is being sincere and is in a good position to have figured out the problem correctly. Given that she is credible in this case, you have a reason to believe the answer is 38.5. But now you have conflicting pieces of evidence. Overall, then, the total evidence available to you now supports the answer '42' less strongly than it did prior to hearing from your friend, and you should update your system of beliefs accordingly. In contrast, if your friend was distracted and careless in working the problem, then she may not be credible in this specific case (even if you think that under normal circumstances she would be). But if she doesn't have credibility in this specific case, then the fact that she disagrees wouldn't give you a reason to doubt your view.

Let's say that your friend is credible in this particular case of disagreement. How should you update your system of beliefs? In general terms, this is going to depend on how you balance the conflicting pieces of evidence. First, how confident are you that you did the problem correctly in the first place? To put it differently, how confident are you that you didn't make a mistake when you worked the problem? Second, how credible is your friend in this case? Thus, in updating your beliefs you need to compare the weight of these two pieces of evidence. For example, it might be that you were very confident in your belief that '42' is the answer, and you persist in believing this is the correct answer even after you've taken into account your friend's credibility (although with a lesser degree of confidence). That is, your reasons for thinking the answer is '42' carry more weight than your reasons for thinking it is not. Alternatively, if your friend is very credible, and you weren't very confident in your answer to begin with, then it can make sense to abandon your belief that it is '42' for the moment, and say 'I don't know'. To push this example even further, imagine you find out that solving this kind of problem was part of your friend's job last summer, and she has solved this kind of problem many times at work. In this case your friend is highly credible and her conclusion carries a lot of weight; if, in addition, you were not very confident in your answer in the first place, then your friend's belief might fully outweigh your reasons for believing it is '42', and tentatively accepting 38.5 would make sense. Importantly, regardless of how you ultimately update your belief, each of these cases is a good opportunity to pursue a cooperative dialogue and figure out what, exactly, accounts for this discrepancy so that you'll both know better how to do the problem next time.

What about cases of agreement? As with disagreement, credibility is crucial. In Ex. 2, you have good reason to think the answer is '42' since you've worked through the problem yourself. When your friend is credible,

then the fact that she agrees with you gives you an independent reason to believe '42' is correct. Consequently, the total evidence available to you at this point supports the answer '42' more strongly than it did prior to hearing from your friend, and it makes sense to subsequently raise your confidence. If she is not credible, either because she doesn't know how to do the problem or didn't try very hard or was distracted, then there is no particular reason to think that she got it right. Consider a different case.

Ex. 3:

You are a big fan of NBA basketball. You watch a lot of games on TV, follow scores and statistics on a daily basis, collect basketball cards, and so on, and you've done so for years. One day at lunch, you say, "I think that San Antonio will be one of the best teams in the league this year". One of your friends agrees, saying "yeah, I think they are going to be really good." However, you know that your friend grew up in San Antonio. He doesn't really follow basketball, but is always optimistic about his home team—he thinks they are going to be good every year.

In this case, the fact that your friend agrees with you doesn't give you any additional reason to think San Antonio will be very good this year, since his belief is not based on a consideration of any of the relevant evidence for that claim (the relative strengths and weakness of the team, for example). It is important to make two points about this. First, notice that had your friend disagreed with you, it shouldn't lead you to change your mind either. Your friend's lack of credibility on this issue means that their view shouldn't lead you to update your beliefs in any way, since you know a lot more about the issue than they do. Second, unlike disagreement, in cases of credible agreement you do not need to balance competing weights, you need only add to your existing confidence.

Overall, that fact that somebody else agrees or disagrees with us can tell us something noteworthy about our own beliefs. The key issue to determine in these cases is whether the agreement/disagreement is coming from somebody who is credible on the claim at issue.

A Question to ask about Cases of Agreement or Disagreement:

- Is the other person credible about this specific claim?

When others agree or disagree with us, that agreement or disagreement matters only when it is credible. That is, we should update our beliefs in light of their agreement or disagreement only when the person is in a good position to have an accurate belief in this case. One important question remains: what if you don't know whether the other person is credible or not? In this case, you can't say whether they are likely to be correct or not, and so there is no reason to update your beliefs. Nevertheless, like the other cases, this can be an opportunity for open-minded evidence gathering and cooperative dialogue.

SECTION 3: APPEALS TO POPULARITY

A different way that we use other people to guide our thinking is called an Appeal to Popularity. An Appeal to Popularity is an argument that draws a conclusion about the truth or merit of some claim, behavior, or product on the basis of its popularity. Consider the following examples:

Ex. 4:

It's okay to photoshop pictures of myself and others and post them online, everybody is doing it.

Ex. 5:

I think the change would be great for the sport. After all, over 80% of players favor the change.

Appeals to popularity can be subtle; sometimes the conclusion is not explicitly stated but is merely suggested. Consider the following case:

Ex. 6:

Ford Trucks: the best-selling trucks in America.

On its face this is just a claim about Ford's sales figures, but Ford is not just telling you about their sales because they think you might be interested. Rather, they are hoping that you will take this fact as indicative of something; they hope you will reason as follows:

1. Ford has the best-selling trucks in America
2. So, Ford trucks are good in some important way (reliable, high quality, etc).

Indeed, Ford would probably like you to go even further and conclude that Ford has the best trucks in America.

In calling this kind of argument an *Appeal to Popularity*, it is important to note that 'popularity' is being used broadly to capture a variety of different notions. In Ex. 4, the author says that "everybody is doing it", but they probably don't mean that literally. What they probably mean is something like this: "many people I know of are doing it" or "it is popular among many of my friends". In contrast, Ex. 5 is more specific. In this case, the proposal is popular among players of the sport. Thus, the first point is that saying a belief, proposal, or product is popular is really saying that it is popular among a particular group of people. A second point is that a belief, proposal, or product may be popular to different degrees, and to say that something is popular is not necessarily to say that it is popular among all or most people in the group. Take Ex. 6: in pointing out that Ford has the best-selling truck, they are not thereby saying that all or most Americans have bought a Ford Truck. Rather they are making a relative claim—Ford sells more trucks than other companies. Overall, then, the term 'popularity' is not very precise in this contexts.

Keeping this in mind, let us lay out the basic structure of an *Appeal to Popularity* as follows:

1. x is popular among the members of some group.
2. So, x is probably good in some respect.

Put in these terms, we can see that there is a gap in this argument. What does the fact that something is popular have to do with its truth or merit? It can be easy to miss, but arguments like this are actually instances of *Inference to the Best Explanation*! An *Appeal to Popularity* begins with an observation that some subject is popular, and then jumps to a specific explanation for its popularity, namely that it is good. As a result, we can use what we know about *Inferences to the Best Explanation* as a guide for thinking about *Appeals to Popularity*. Thus, we can characterize the missing premises in *Appeals to Popularity* accordingly:

1. x is popular among the members of some group.
2. The most likely explanation for why x is popular among members of this group is that x is good in some respect. (MP)
3. So, x is good in some respect.

Setting out the argument in this way is especially helpful for understanding the difference between logically strong and logically weak Appeals to Popularity. However, before we turn to this issue we will briefly consider what motivates us to use this form of argument more generally.

SECTION 4: THE APPEAL OF APPEALS TO POPULARITY

As we will see, *Appeals to Popularity* are often logically weak. Nevertheless, they are common in everyday thinking, and so it is worth taking a moment to think about why we use *Appeals to Popularity*. Perhaps the most straightforward answer is that they are easy short-cuts. Go back to Ex. 6 above. Suppose that you draw the inference that Ford hopes, and come to the conclusion that Ford Trucks are probably pretty good trucks, since they are the best-selling ones. Think about your evidence here. In coming to this conclusion, you didn't look at any of the factors that might actually make a truck good in some way. You didn't look at its reliability, its towing capacity, its gas mileage, its acceleration, the size of the truck bed, and so forth. Instead, you've looked at what people believe about the truck (namely that it is worth buying). This is a lot easier—especially when Ford tells you about it—than doing the work for yourself. Many *Appeals to Popularity* are like this: we use what is popular as a shortcut for identifying what is probably true, or good, or meritorious in some way. In general, when popularity is genuinely an indicator of truth or merit, there is no problem with taking this shortcut, though again, this is not common. Moreover, even when *Appeals to Popularity* are sound, they do not help us understand *why* something has merit or is true, only that it does. Consequently, when it is important to understand something for ourselves, an *Appeals to Popularity* won't suffice.

Appeals to Popularity are easy, but there are other factors pushing us to look to what is popular as a guide for our belief and action, namely an interest in conformity. In a sense, this is probably not too surprising; after all, everybody has experienced peer-pressure in one way or another. What is, perhaps, more surprising is the power these forces can have over us. In a well-known series of experiments psychologist Solomon Asch sequentially presented a group of participants with two cards. On the first card was a single line. The second card had three lines on it.

The researchers asked each individual in the group to publicly judge which of the three lines was of the same length as the line of the first card. This is an easy task, the answer is obvious, and the first few times this task was completed everyone agreed about the obvious correct answer. What makes this experiment interesting is that all the participants—but one—were part of the experiment and were acting off of a script. They put this to work in the next trial: the first actor gave an obviously wrong answer, and all the other actors agreed. Researchers found that in about a third of the trials the last participant would conform to the rest of the group and give the obviously wrong answer. These kinds of experiments have been conducted in a variety of ways, but the result remains the same. In the face of a unanimous—but *obviously false* opinion—people will surprisingly often act in conformity with the group, as opposed to in conformity with the truth.¹ Given the motive to conform, it is not surprising that we often use the evidence of other's beliefs and choices as a guide for our own.



"conformity" by BitHead CC BY-NC-SA 2.0

SECTION 5: EVALUATING APPEALS TO POPULARITY

When is an *Appeal to Popularity* logically strong? Given that *Appeals to Popularity* are instances of *Inference to the Best Explanation*, we can follow the same guidelines. Thus, as we saw in Chapter 13 step one is to ask whether the proposal that the subject is good in some respect would plausibly explain its popularity. In general terms, yes. We prefer products and proposals that are good in some way, and we prefer beliefs that are true. Let's turn to step two. Step two will be to look for plausible alternative explanations for the popularity of the subject. As we know, there are many possible explanations for any state of affairs. Popularity is no different. Take a product for example. It might be the most popular because it is the best advertised, because it is the cheapest, or because it is carried by every store, and so on. As an illustration consider this case:

Ex. 7:

McDougal's sells more cheeseburgers than any other company in the U.S. So, McDougal's cheeseburgers must be the best.

The fact that McDougal's cheeseburgers are the most popular cheeseburger in the U.S is a consequence of a number of factors which surely include price and easy access to McDougal's stores. This brings us to step three. Now that we have some competing explanations, we have to compare them and see if the proposed explanation comes out on top. In the McDougal's example, it does not. That is, the facts that McDougal's cheeseburgers are inexpensive and are easily available to millions of people seem like a better explanation for the popularity of McDougal's cheeseburgers than the proposal that they are the best.



“two single cheeseburgers” by stu_spivack CC BY-SA 2.0

The argument in Ex. 7 is a logically weak *Appeal to Popularity*, but what would a logically strong one look like?

A logically strong *Appeal to Popularity* would be a case in which a belief, proposal, or product’s popularity really is best explained by its merit. So, when is this the case? For sake of ease, let us talk simply in terms of a product. A product’s popularity would be best explained by its merit if it is popular precisely *because* many people investigated the relative merits of the product, and on that basis, decided to purchase it. Thus, when we are looking for logically strong *Appeals to Popularity*, we need to look for cases where the best explanation for a subject’s popularity lies in many people’s thoughtful determination of the subject’s merit. As you can probably see, this is not common. Nevertheless, there are some logically strong *Appeals to Popularity* in our everyday experience. Let’s take a look at one.

SECTION 6: POPULARITY AMONG EXPERTS

There are some cases in which something is popular as a result of individual investigations of merit. Perhaps the most common case has to do with populations of experts. Consider the following example.

Ex. 5:

Delia: I don’t get why people are so concerned if there are fewer bees. I mean, they sting you. The less of them the better.

Sam: Biologists agree that bees are a crucial part of the ecosystem because of their role in pollination.

Sam is not just reporting what one Biologist says, but is using the fact that there is consensus among biologists on this issue. This is a logically strong *Appeal to Popularity* because biologists are experts who in their professional work seek to make judgments about what is true on the basis of the available evidence. More broadly, when there is consensus about some thesis within the community or population of experts, and that thesis is within the community of experts’ field of expertise, then a logically strong *Appeal to Popularity* can be drawn.

As another example, consider the fact that some of the current debate about global warming is actually a debate about whether a logically strong *Appeal to Popularity* can be drawn by the public. Most people are not in a position to personally evaluate the evidence for anthropogenic climate change (that is, the thesis that climate change is occurring and that it is caused by humans). After all, the vast majority of people do not have access to relevant data sets, they do not have the requisite statistical/mathematical expertise, they

do not have the background knowledge required to read and assess the literature, and last, they do not have the time. For these reasons, most people's beliefs about anthropogenic climate change are based on the limited evidence they can personally evaluate and on their knowledge of what the population of experts believe. In recent years a number of studies have found that well over 90% of climate scientists believe in anthropogenic climate change. Given the criteria we have discussed, it would seem that this fact offers good evidence for thinking that anthropogenic climate change is really occurring—after all, there seems to be consensus on the issue from experts in the field.

The idea of appealing to a population of experts is not limited to reasoning about science. Consider the well-known website rottentomatoes.com. The chief function of this site is to aggregate movie reviews. For each movie that is currently in theaters (and for some that are coming soon) the site offers a measure on the “tomatometer” scale. This is just a measure of how many critics have positively reviewed a movie. Thus, if 75% of the available reviews positively evaluate a movie, it has a rating of 75. This information is intended to help the movie-going public decide what movies to see. The underlying idea is that people who write movie reviews are (or tend to be) experts, and since popularity among experts is indicative in some sense of quality, the tomatometer can be a useful guide for making movie-going decisions.

SECTION 7: WHAT DOES POPULARITY TELL US?

Appeals to popularity are not often logically strong. There are many reasons a belief, behavior, or product might be popular other than that each member of the population investigated the belief, behavior, or product in question and determined it to be meritorious, and when compared to these alternatives merit is rarely the best explanation. Nevertheless, it seems like popularity often tells us *something*—especially when it comes to products. Consider the following situation: you are visiting a new city and have been staying in a hotel for a few days. Outside your hotel you have noticed that there are two similar restaurants. Both restaurants offer pub-style food and are located right across the street from one another. You have noticed that one of these restaurants is consistently busy, while the other is consistently dead. One evening you decide to eat near your hotel and these two restaurants are your only options. Which restaurant do you choose?

Most people choose the busy restaurant. But does this choice make sense? After all, this choice relies on an *Appeal to Popularity* and we have learned that there are many reasons a restaurant might be popular other than that the food is good—it might have been doing a lot of advertising, it might have really good drink specials, it might have attractive waiters/waitresses, it might be exceptionally cheap, etc. This choice does make sense, but not because it will likely have good food. The choice makes sense, rather, because it is likely that the food is not immediately and obviously poor (whereas you have no such evidence for the dead restaurant). At least when it comes to products, the fact that it is popular over time typically offers evidence that it isn't obviously awful or a clear failure, since products that are obviously awful or clearly fail to do what they claim are not popular over time. This is defeasible evidence, of course, and this is no guarantee, but it is evidence nonetheless. Consider Ex. 6 again. Although we cannot infer from the fact that Ford Trucks are the bestselling trucks that Ford's trucks are the best, we can reasonably infer that they are not immediately and obviously bad (of course that is not exactly the inference that Ford wants you to draw—*Ford Trucks: Not Awful*).

SECTION 8: THE PSYCHOLOGY OF POPULARITY JUDGMENTS

Using an appeal to popularity presupposes a belief that some idea, behavior, or product is popular; but how do we know whether something is popular? In some cases our belief about something's popularity among a specific group of people is based on solid statistical data, but most of the time it isn't. In fact, most of the time we simply rely on our intuitive sense that something is common or uncommon, frequent or infrequent. When making these kinds of judgments we need to be careful since psychologists have shown that we tend to overestimate the popularity of beliefs we share in, and underestimate when we do not. This phenomenon is called **false consensus**. As the social psychologist Ziva Kunda writes:

False consensus has been shown to color people's estimates of the prevalence of just about any choice, attitude, or behavior that has been examined, in domains as diverse as everyday habits and preferences, personality traits, and political opinions. For example, fans of white bread think more people would choose white bread over brown bread than do fans of brown bread, optimists believe optimism is more common than do pessimists, Americans who prefer the Republican presidential candidate believe that support for the Republican is more widespread than do those who prefer the Democratic candidate, and those who support cuts in public spending believe that more people support such cuts than do those who oppose them.²

To be clear, to say that we have this tendency is not to say that we always overestimate the relative frequency that others share our beliefs and attitudes. Nonetheless, this is a common phenomenon, and since appeals to popularity rely on a judgment that something is frequent or common, it is something we should be aware of in thinking about appeals to popularity.

Overall, *Appeals to Popularity* are common, and it is not hard to see why. They are an easy shortcut for determining subject's merit. Nevertheless, they are logically strong only in cases where merit is the best explanation of a subject's popularity. We find this in cases where experts agree, but in general, logically strong *Appeal to Popularity* are few and far between. Moreover, the phenomenon of false consensus shows us that we are subject to mistakes in judging popularity in the first place. The upshot here is that we should be extremely wary of *Appeals to Popularity* in our everyday thinking. In particular, we need to remember the following:

Two Questions to Ask of Appeals to Popularity

- Are there other plausible explanations for the subject's popularity?
- Would the truth of the proposed explanation be less surprising than the truth of any competitor?

EXERCISES

Exercise Set 18A:

Directions: Consider the following cases of agreement/disagreement. In each case comment on whether the agreement/disagreement in question implies that you should update your confidence or even change your belief. Explain your answers.

#1:

You think the midterm is in two weeks, since that is what it says on the syllabus. You overhear two people from class saying that the midterm is next week.

#2:

You are sitting in class, and it sounds like the teacher has just said the midterm was moved up. But somebody coughed at the same time the teacher said this, so you wonder whether perhaps you misheard. You turn to the student next to you and ask, “did she just say *moved up*?” The student says ‘yes’.

#3:

At lunch one day, the topic of CPR comes up and somebody asks which you should do first when you give CPR, rescue breaths or chest compressions. Because you had to complete a course in CPR and be proficient for your summer job as a lifeguard you know the answer, and say ‘chest compressions’. The person sitting next to you disagrees, saying, “no, in our school’s production of *Wit* we did rescue breaths first.”

#4:

You are reading a magazine and there is an interview with a chef from Santa Fe, New Mexico. In the article the chef mentions the challenge of cooking in Santa Fe where water boils at less than 200 F. This strikes you as wrong. You learned in school that water boils at 212 F and water is pretty much the same anywhere you go.

#5:

You are catching a ride with a friend to a city across the state for a wedding. You think route A is the fastest route because that’s what your phone’s mapping app says, but she thinks Route B would be quicker because that’s what her phone’s mapping app says.

Exercise Set 18B:

Directions: Determine whether each of the following is an Appeal to Popularity or not. If not, try to identify the argument’s form. Then comment on the argument’s logical strength.

#1:

Most small business owners agree that they pay too many taxes, so the legislature should take tax increases on small businesses off the table.

#2:

Evaluate B’s argument

A: I guess I just don’t think hunting is ok—we don’t need to do it, and I don’t think killing something for sport is acceptable

B: But c’mon. People all over the world enjoy hunting, and have been doing so for thousands of years.

#3:

You have recently moved to a new neighborhood and don’t know what day you should put out your garbage. You decide to wait and see what everybody else does. You see that on Wednesday night all your neighbors put out their garbage, so you decide to do so too.

#4:

A: I think we should just buy the new truck and call it a business expense so we can write it off on our taxes.

B: I don't know. That sounds like cheating to me. We wouldn't really use the truck very much in the business, you know.

A: Oh, don't worry about it. This kind of thing is done all the time.

#5:

9/10 podiatrists recommend against wearing the same pair of shoes every day, so I guess I should add some variety to my shoe wearing.

#6:

There are very good reasons for the death penalty. First it serves as a deterrent to those who would commit capital offenses. Second, it is just and fair punishment for the crime committed. Third, reliable opinion polls show that over 70 percent of all Americans favor it. If so many people favor it, it has to be right.

Exercise Set 18C:

#1:

In the Asch conformity experiment discussed in Section 4 above, one of the reasons that people gave for conforming in spite of the obvious evidence is that they started to doubt their own belief. They reasoned that everybody else must have known something they didn't. How does this connect to our discussion of disagreement? What question should people in this experiment have asked themselves, and how should they have subsequently evaluated the situation?

#2:

What about unpopularity? Is unpopularity ever a good reason to think that something is false or poor? When?

#3:

Suppose that you are an unscrupulous CEO of a corporation that makes product x, and there is consensus among scientists that your chief product is toxic. It seems like it is only a matter of time before your product is banned by the government. How might you try to manipulate public opinion about your product to prevent its being banned?

Notes

1. Asch, S. (1955). "Opinions and social pressure." *Scientific American* 193 (35).

2. Kunda, Ziva. (1999) *Social Cognition: Making Sense of People*. Cambridge, MA: MIT Press, 397.

Unit #6 Summary

Other people are a great source of information, and in this unit we've looked at some of the ways other people and their beliefs can serve as evidence. The largest single influence other people have on us is through what they tell us. As we've seen, we learn a great deal about the world through testimony in one form or another. Learning through testimony requires trust however, and other people are not always trustworthy. We looked at some of the factors that can undermine a person's credibility, and what we can (and cannot) legitimately infer about a person's testimony on the basis of other things we know about them. In addition, other people can serve as a good guide for whether our own beliefs are correct. The fact that somebody agrees with us can justly bolster our confidence in its truth, while disagreement can signal that we've missed something and should lower our confidence. Last, we looked at what we can infer from the fact that many people believe or like something. As we saw, popularity is not always an indication of truth or of merit.

Questions to ask of Testimony:

- Is the source credible? (history of lying? something to gain? capability/expertise?)
- Is the claim plausible given what you know about the world?

Questions to ask of Negative *Ad Hominem*:

- Does the feature give reason to suspect that the source is not being honest or sincere in this case?
- Does the feature give reason to suspect that the source is not in a good position to know in this case?

A Question to Ask about Cases of Agreement and Disagreement:

- Is the other person credible about this specific claim?

Questions to ask of *Appeals to Popularity*

- Are there other plausible explanations for the subject's popularity?
- Would the truth of the proposed explanation be less surprising than the truth of any competitor?

KEY TERMS/IDEAS

- Testimony
- Trust
- The Credibility Assumption
- History of Lying or Deception
- Interest
- Capability
- Expertise
- Internet as Testimony Machine
- *Ad hominem* Argument
- Positive/Negative *ad hominem*
- Deniers
- Cooperative Disagreement
- Agreement/Disagreement as Evidence
- *Appeal to Popularity*

- Conformity
- False Consensus

FURTHER READING

In recent years there has been a dramatic increase in research in psychology and philosophy on testimony. If you are interested in the most recent psychology literature a good place to start is with the overview in “Cognitive Foundations of Learning from Testimony” by Harris et al. in *Annual Review of Psychology* (2018). For the philosophical point of view see *Testimony: A Philosophical Introduction* by Joseph Shieber. For more on relationships between testimony, the web, and social media see *The Internet of Us* by Michael Lynch. Finally, for an excellent introduction to disagreement as evidence, see the aptly titled *Disagreement* by Bryan Frances.

UNIT VII

SCIENTIFIC REASONING

Dealing with Statistics

SECTION 1: INTRODUCTION

In this unit, we will take a look at some of the distinctive reasoning that goes on within the sciences. We will not learn how to engage in these complicated reasoning tasks ourselves (after all, these often involve specialized techniques for inquiry and problem solving). Instead, we want to get a better sense for some of the common techniques used across the sciences, and thereby put ourselves in a position to think a bit more critically about some of the scientific claims that we run across in the news and elsewhere. We will start with the use of statistics. We have already talked about statistics in our discussion of statistical generalizations back in Chapters 14 and 15. As we will see, statistical generalizations are just one kind of statistic—though an especially important one. The goal, in this chapter, is to get a better sense for what a statistic is, and to look at some common pitfalls in understanding statistics and their visual representations.

SECTION 2: STATISTICS

A **statistical claim** is traditionally defined as a statement expressing a numerical fact, and we will take this definition to include vague numerical statements that use terms like ‘most’ or ‘few’. Just as there are all kinds of numerical facts, there are all kinds of statistical claims. For example, a statistical claim might simply tell us how many things of some kind there are, as in:

- There are 3 eggs left in the fridge.
- The University of California System awarded 55,350 Bachelor’s degrees in 2017-18

While statistics like these are important, often we are concerned with statistics that tell us something about a whole group. Averages are statistical claims, and so are generalizations.

- The average score on the exam was an 83.
- 39% of high-school graduates in Greenville go on to enroll in college.

Similarly, we might be interested in statistics that are comparative. **Comparative statistics** give us a numerical measure of how much something has changed, or how two different groups are related. For example:

- Sales have decreased 8% from last quarter.
- Citations for driving under the influence (DUI) have gone up 90% under the new mayor.

- Teen pregnancy rates in Greenville are $\frac{1}{2}$ of what they are in Bluefield.

Precise statistical claims often carry a special authority, and this makes it particularly important to understand what they are saying. We have already looked at statistical generalizations, and we have a reasonably good sense of what counts as a good statistic of this kind. As it turns out, statistical claims come in many different forms, and we will not be able to discuss each kind individually. Nonetheless, we can talk about some general strategies for understanding statistics, as well as some questions to keep in mind when it comes to some of the more common kinds.

SECTION 3: A KEY QUESTION FOR UNDERSTANDING STATISTICS

When we look at a statistical claim, the most important thing to remember is that this information has been collected by one or more *people*. As Joel Best puts it in his book *More Damned Lies and Statistics*:

[W]e tend to assume that statistics are facts, little nuggets of truth that we uncover, much as rock collectors find stones. After all, we think, a statistic is a number, and numbers seem to be solid, factual proof that someone must have actually counted something. But that's the point: people count. For every number we encounter, some person had to do the counting. Instead of imagining that statistics are like rocks, we'd do better to think of them as jewels. Gemstones may be found in nature, but people have to create jewels. Jewels must be selected, cut, polished, and placed in settings to be viewed from particular angles. In much the same way, people create statistics: they choose what to count, how to go about counting, which of the resulting numbers they share with others, and which words they use to describe and interpret those figures.¹

Best's point is that people collect and "polish" statistics, and so although a statistic might accurately report some quantity, it will also reflect the interests and motives of the people who originally collected the information. After all, any statistic is the result of a person or group of people deciding to count something in particular, and deciding to pursue that information for the sake of some goal. One important consequence of this is that in order to really understand a statistic, we have to understand what it was so important to count. That is, we need to ask:

A Key Question for a Statistic:

- What, exactly, was counted?

While this probably sounds obvious, people regularly misinterpret statistics because they are not clear about the subject of the claim. To illustrate, take the generalization above that 39% of high-school graduates in Greenville go on to enroll in college. Many people would regard this number as unacceptably low, and this is the kind of statistic that might move a community to make substantive changes to its educational system. However, before taking action on this basis it is important to fully appreciate this statistic. So let's ask the question above: what, exactly, was counted in this case? Well...high school graduates who enrolled in college. But how is the term 'college' being used here? For example, what kind of college are we talking about—was this limited to 4-year institutions, or did they count community colleges or trade schools? What time frame were they working with—did they count only those graduates who enrolled immediately after graduation, or do these numbers include people who decided to take a year or two off? What does 'enroll' mean? Does it mean that they counted only those graduates who enrolled as full-time students, or did they count part-time students as well?

Getting answers to these questions is crucial to understanding the significance of this claim. If the author of this statistic was counting only those graduates who immediately enrolled as full-time students at 4-year colleges, then this statistic tells us a lot less than we might have thought. If, on the other hand, the author

was counting those graduates who enrolled as part-time or full-time students at 2 or 4 year institutions in the 3 years after graduation, then this statistic *is* quite telling.

Here is another example which shows the importance of understanding the relevant terms in a statistic. In March of 2014 the following headline appeared in *Time* magazine: "U.S. Autism Rates Jump 30% from 2012".² The story reports the findings of a recent study by the Center for Disease Control (CDC). What should we make of the statistical claim in the headline? It is easy to read this as saying that there has been a dramatic increase in the number of people with autism in just two years, and to worry, consequently, that we are in the middle of an epidemic. We need to be careful, and to think about what is being counted here, however. Whether a person has autism spectrum disorder (ASD) is something that must be diagnosed by a professional. Accordingly, in order to produce the statistic above researchers from the CDC did not count people with ASD, but counted people *diagnosed* with ASD. Consequently, what the statistical claim tells us is that there has been a dramatic increase in the number of people diagnosed with ASD in the last two years. However, this fact does not necessarily mean that there has been an increase in the number of people with autism. After all, a person can have ASD without having been diagnosed. The authors of this study attributed this increase at least partially to (i) changes in the medical definition of ASD which broadened its application and (ii) the fact that people with ASD are more likely to be diagnosed than in the past.

This particular statistical claim offers another lesson. The authors of the study were careful to warn against generalizing too far from their results. They wrote: "Because the [sampling technique does] not provide a representative sample of the entire United States, the combined prevalence estimates presented in this report cannot be generalized to all children aged 8 years in the United States population". Despite this warning, many news outlets reported this story as if the researchers had concluded that there had been a 30% increase in autism rates among **all** American children.

Here is one final example. Filtered water stations are marketed as way to reduce plastic bottle usage, and some include a digital counter that keeps track of how many disposable plastic bottles the station has "helped to eliminate." This is a statistic, so let's think about what is being counted here.



"Filtered Water Digital Counter" CC 0

Well, it is not actually counting the bottles this station has helped eliminate. In order to count those bottles the machine would have to know when a person has filled up at the station *instead of buying a new disposable plastic bottle of water*. But the machine does not know a person's reasons for filling up their bottle or what they would have done. All the machine "knows" is the volume of water it has dispensed, and this is

what the machine is actually counting. That is, the statistic the station can accurately report is that it has dispensed a volume of water equivalent to 199,963 disposable plastic bottles, something very different from what the machine actually claims. Has this machine helped eliminate waste from some plastic bottles? It almost certainly has, but ‘some’ is about all we can say in this case (and ‘some’ doesn’t look very good on a digital counter!). While this is a fairly minor inaccuracy, it is not difficult to see how it might lead to wider errors. Imagine, for example, somebody arguing that the filtered water system was worth the cost on the grounds that it prevented almost 200,000 plastic bottles from being used. Returning to the wider point, these examples show that we are not in a good position to understand a statistic if we don’t know what its terms mean, and consequently what, exactly, has been counted.

SECTION 4: THE AVERAGE IS NOT NECESSARILY TYPICAL

Averages are among the most common statistics we come across on a day-to-day level. Although we all know how to find the average (or mean) of a set of numbers, we do not always understand what an average is telling us. We tend to use averages as a quick way to characterize a whole body of information, and to interpret the average in terms of what is common or typical. Sometimes this works, and the average really does give you a good sense for what is common. Take height, for example. The average height of women in the U.S. older than 20 is 5 feet 4 inches, and this is a very common height. Most women in the U.S. cluster around 5’4 and the numbers fall off as you get significantly taller or shorter from this point.

However, this does not hold for all averages. An average is a simple way of summarizing numerical information, and does not capture the way the information *varies*. For example, say that you are moving to a new city and are looking to rent an apartment in a neighborhood you don’t know much about. As you look around online, you run across the following statistic: “in this area rent averages \$850/mo.” You might conclude on this basis that you’d have to pay roughly \$850 for rent. That is, you are thinking that rents in this neighborhood look something like Diagram 1 below:



Diagram 1



Diagram 2

Average rent may not be the typical rent

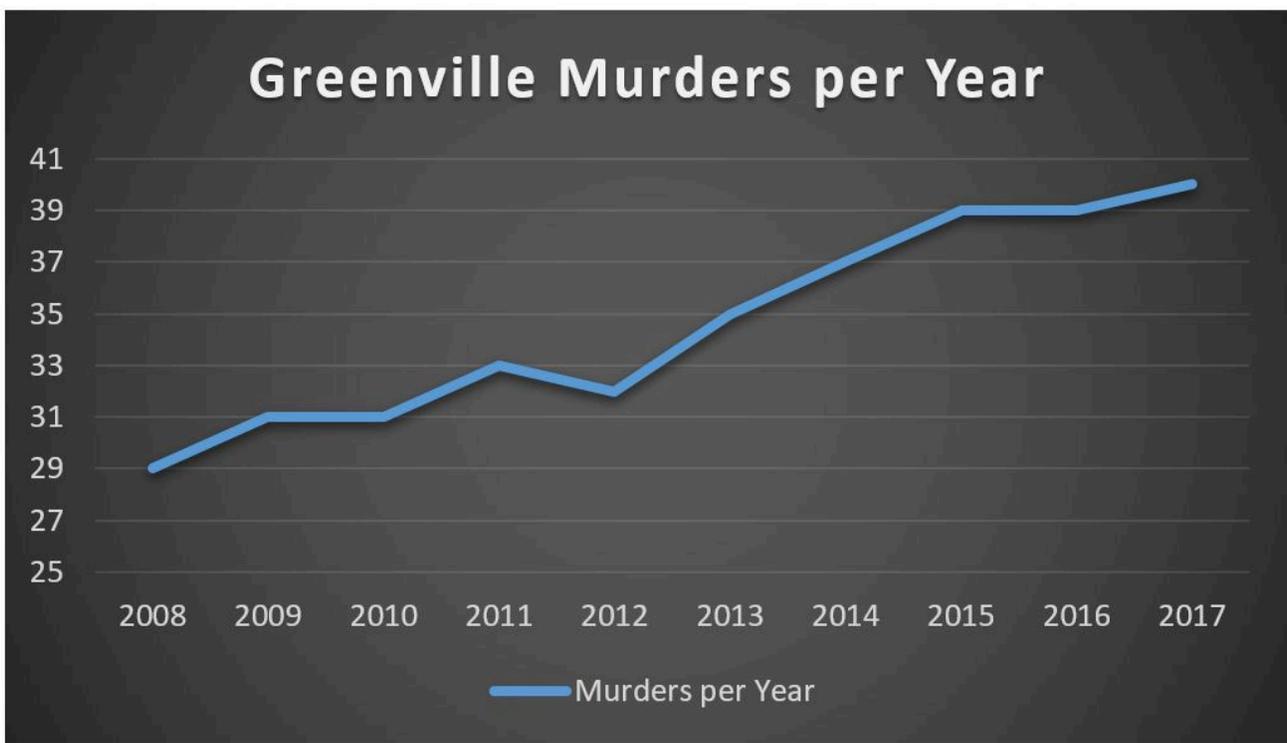
Maybe, but this fact alone certainly doesn’t guarantee this. In fact, it may be that there are few if any apartments that rent for around this price. Say the area is pretty clearly split between more and less desirable

areas. In the less desirable areas rent might run around \$600, but closer to \$1100 in the more desirable areas. In this case, your choice is likely to be between paying \$600 and paying \$1100 even if the average is \$850 as is represented in Diagram 2. This illustrates that the same average rental price is consistent with very different rental markets, and this holds more broadly. Knowing the average of a set of numbers alone doesn't tell you about how those numbers vary or are distributed, and this can make a big difference.

To give a final example, imagine you are interviewing for a new job. The boss tells you that employees are on track to get year-end bonuses averaging about \$1500. That sounds pretty good, and you start thinking about what you could do with the bonus if you took the job. But, again, it does not follow that you'd get anything close to \$1500. If, for example, company executives got giant bonuses, then the average might be \$1500 even though most employees got a much smaller bonus. In this case, a few very high bonuses will dramatically raise the average. The executive bonuses are **outliers**—they are numerical values that differ significantly from the bulk of the other values in the group. The presence of outliers can pull an average away from what is common or typical. The take-away from these examples is that averages all by themselves do not necessarily tell you about what is common or typical. Importantly, knowing this puts you in a position to ask effective questions. So, in the interview case, you might say: "that bonus sounds great, but what is the projected bonus for a person in the position I am applying for?"

SECTION 5: COMPARATIVE STATISTICS

Another common type of statistic compares one numerical value with another. To illustrate, imagine a headline: "Murders top 40 for the first time in the City's history". The story goes on to report that more murders were recorded this past year than any prior year, and is accompanied by the following graph.

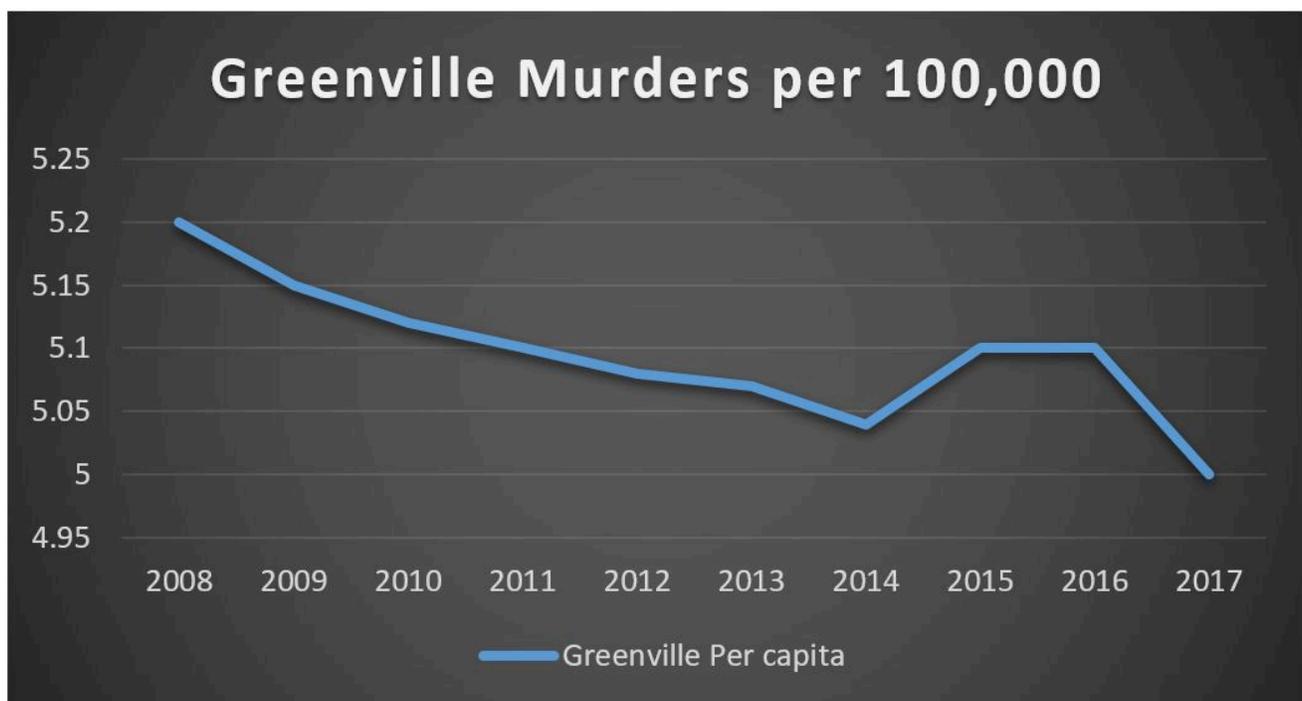


Greenville Murders by Year (2008-2017)

Like all graphs, this one compares statistics. Specifically, it compares the number of murders in Greenville on a year-by-year basis, and when we look at this graph we have to admit that this looks bad. A graph like this might lead you to think that crime in Greenville is getting worse, and that something needs to be done to address this growing problem. Maybe it does, but the comparison represented by this graph is meaningful only to the extent that it is comparing **relevantly similar statistics**.

In order to illustrate this point, let's add to the scenario that the population of Greenville has increased dramatically during this interval—from around 580,000 to 775,000. In this case, the graph is comparing the murder rates for different sized communities. For example, the number of murders in 2008 was 29 when Greenville had about 580,000 people. To straightforwardly compare this number to Greenville's 40 murders in 2017 when its population was 775,000 misses something important. The reason for this is that you would automatically expect more murders in a bigger city (all things being equal).

If you want to compare murder rates in Greenville from year to year, then you need to find a way to make the yearly murder statistics similar enough to compare them. One way to do this is to find the per capita murder rate. The per capita murder rate is the number of murders per person. In 2008 Greenville had a population of 580,000 people and 29 murders, and this gives us a per capita murder rate of .00005 (dividing 29 by 580,000). We can do this for each year and then compare the per capita murder rate. Doing so gives us the following graph:



Greenville Murders (2008-2017) per 100,000

This graph shows that despite the increase in the number of murders in Greenville, overall the murder rate in Greenville has gone down.³ Even though there were more murders in Greenville in 2017 than in 2008, overall there are fewer murders *per capita*.

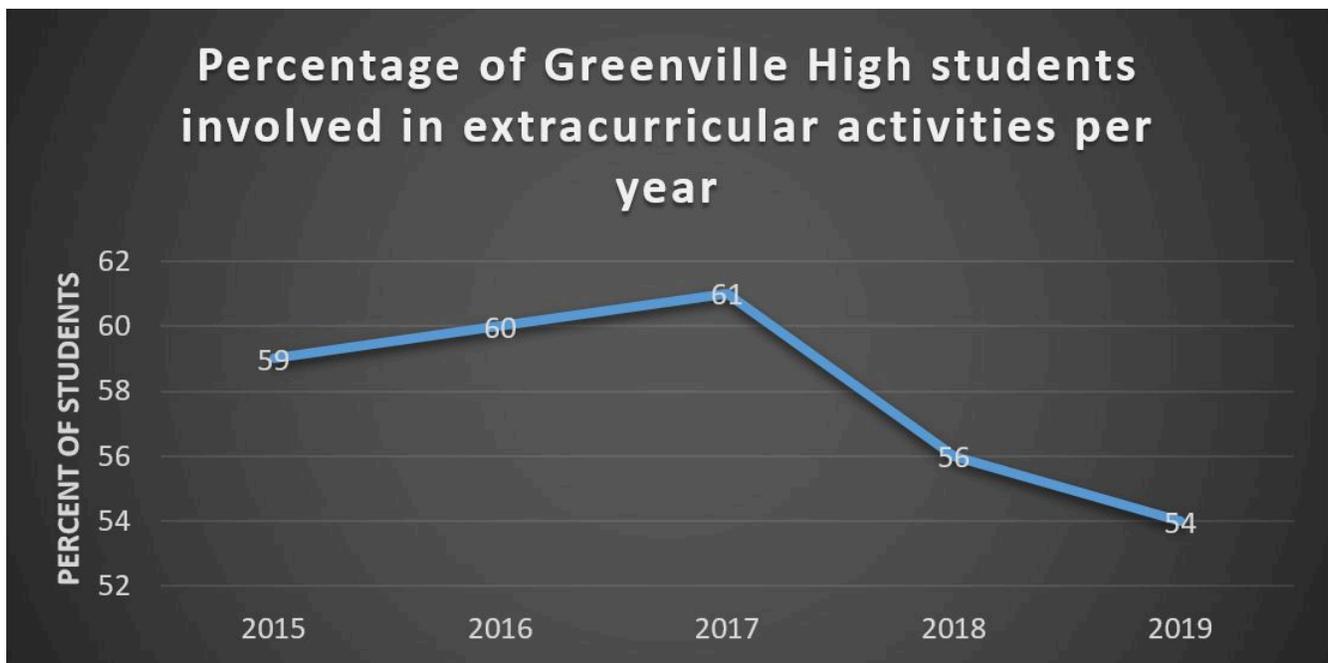
Given only the first graph, citizens of Greenville might be alarmed. However, when compared appropriately these numbers tell a very different story. The point of this example is to show that it is important to com-

pare apples to apples; comparative statistics are meaningful only to the extent that the comparison makes sense—and at least sometimes it does not.

One other point to make about comparative statistics has to do with comparing *percentages* in particular. For example, suppose you hear from a reliable source that regularly taking small doses of aspirin doubles your chance of developing a stomach tumor. This sounds terrible, and suggests that if you are a regular aspirin user, then you should stop taking it. However, in fact, by itself this information does not justify this conclusion—after all, you don't know the chances of getting a stomach tumor in the first place. Suppose that the probability of developing a stomach tumor in your lifetime is 1 in 200,000. If regular aspirin consumption doubles your chances, your probability of developing a tumor will be 1 in 100,000. This is still pretty unlikely; in fact, it matches the probability of dying in a parachuting accident. However, many people are willing jump out of airplanes fully knowing the risks involved, and so too a person might be willing to live with this level of risk of developing a stomach tumor for the sake of the benefits provided by aspirin.

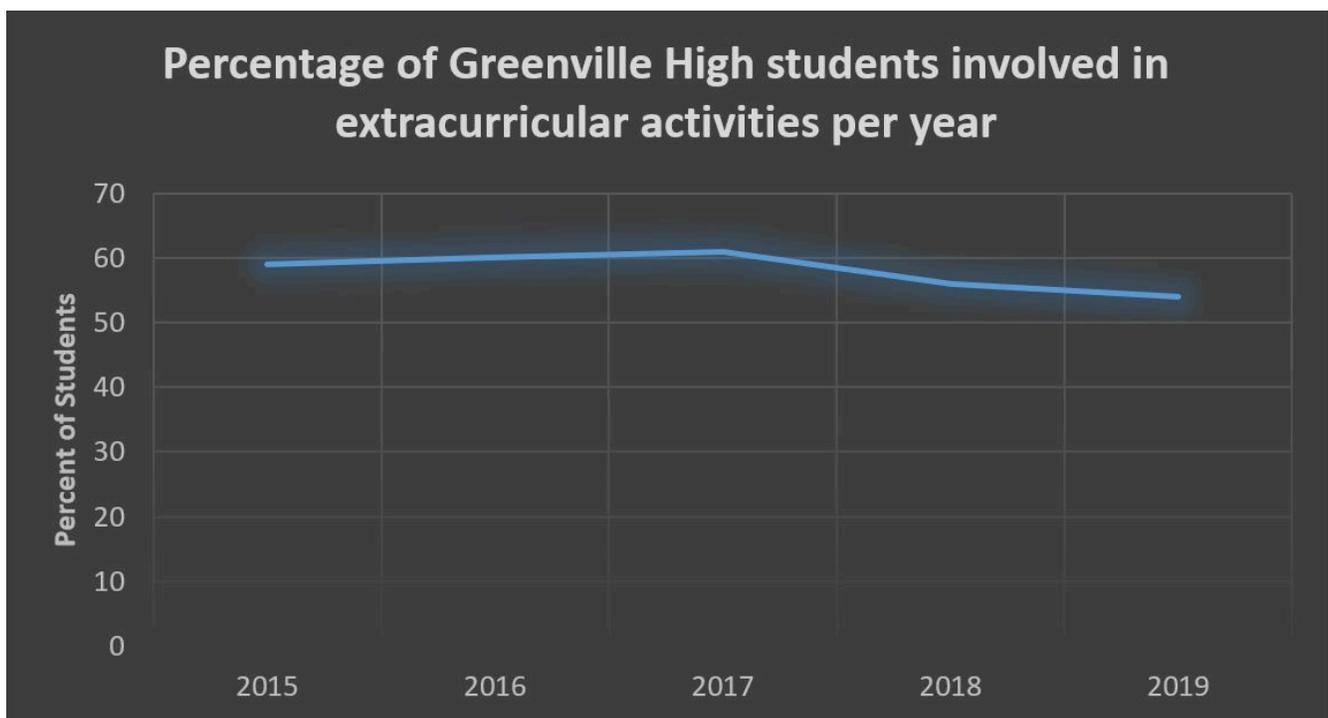
SECTION 6: SCALE

Another important characteristic to keep an eye on when it comes to understanding graphs is scale. On a graph the **scale** is a numerical interval or gap that is used to represent quantities on one of its axes. In the example above, the scale for the vertical axis (or y-axis) is murders per capita and the horizontal axis (or x-axis) is a time scale measured in years. The scale is an important, and easily overlooked, element of a graph and is crucial for understanding the significance (or insignificance) of the information provided. Moreover, it is important to remember that even when a graph objectively reports the facts, it is a representation of those facts chosen by one or more people. In creating graphs, people not only choose what information to present, but also *how* to present and frame the information. These choices, and in particular choices about scale, can make changes or differences seem more or less significant than they really are. Consider the following example. The local school board is interested in the percentage of high-school students who participate in extra-curricular activities because they think it reflects student attitudes about school more broadly. In response to their request for information, the school system's administration gives them the following graph:



Extracurricular Participation by Percentage 2015-2019

Presenting the information about participation using this graph makes it look as if the percentage of students who participate in extracurricular activities has fallen off significantly in the last couple of years. This graph suggests this by fitting the scale to the changes in participation. Consequently, the scale visually highlights the decline of the last two years. Consider this alternative, however:

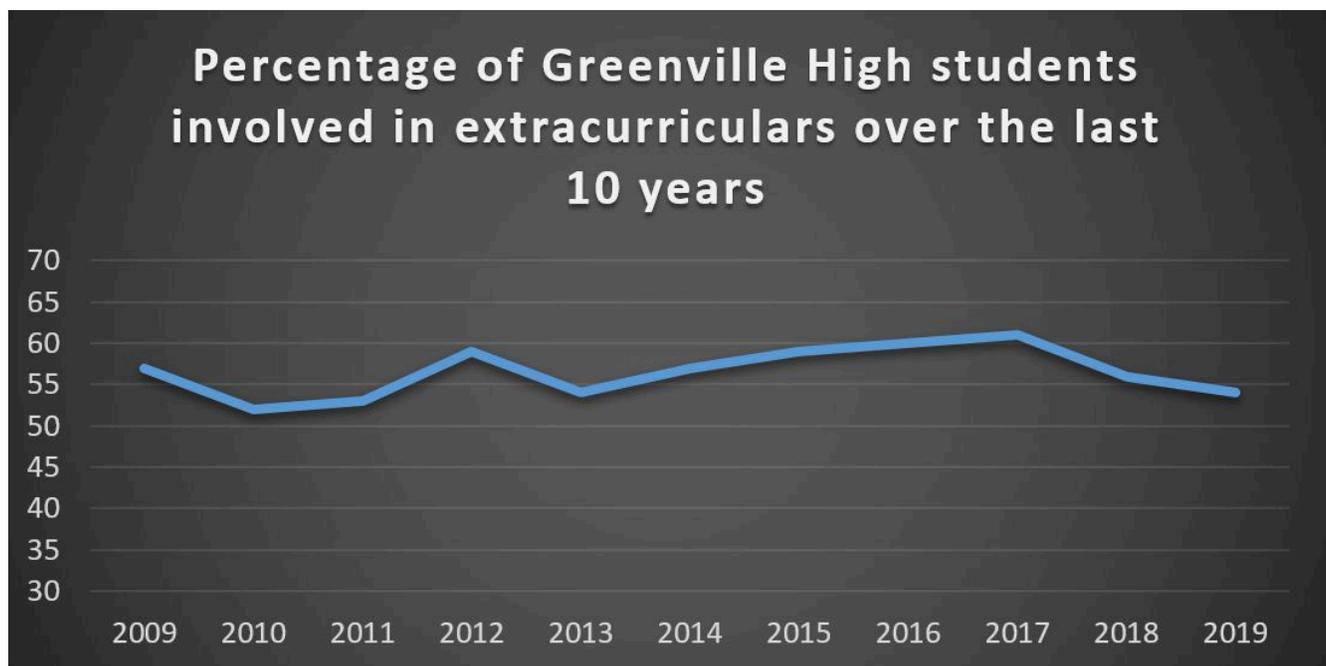


Same Information, Different Scale

Importantly, this graph represents exactly the same information. The only difference is that the vertical axis has been changed to go all the way to zero. Constructing the graph using this different scale highlights the relatively stable percentage of students involved in extracurricular activities at Greenville High and thereby minimizes the visual impact of the changes. It shows a decrease, but relative to the overall percentage of students involved, the change appears modest.

The fact that we can change the visual impact of a graph in this way automatically raises the question: what scale *should* the graph use? There is no one-size-fits-all answer here. What determines the appropriate scale will depend on what would count as a significant or important change or difference when it comes to the specific factors being compared. To illustrate, consider a very different case: levels of CO₂ in the atmosphere. A change as seemingly small as 100 parts per million (.01 percent) to the level of CO₂ in Earth's atmosphere is meaningful and can have significant consequences. As a result, a graph measuring global CO₂ levels should have a scale that shows changes of this degree to be visually significant as well (so, for example a scale with intervals of 10 or 20 parts per million). Alternatively, we might consider the United States annual defense budget. This budget is measured in the hundreds of billions of dollars, and so a scale with intervals of even 1 million dollars would be too fine grained to capture changes in this budget over time.

As a final point, often the significance of a change or difference becomes clearer given more information. To use the example of Greenville High again, the significance of the recent dip in participation in extracurricular activities might become clearer by looking at the last 10 years of data instead of only the last 5 as follows:



Extracurricular Participation by Percentage 2009-2019

This additional information shows that the participation rate has varied from about 50% to about 60% over the last ten years. The recent drop in participation to 54% is still well within what has been normal at the school. Indeed, this new information highlights that extracurricular participation was especially high in 2017 (as opposed to especially low in 2019). Notice the change in scale as well. Although there are surely other

variations of scale that are appropriate, running the interval from 30 to 70 in increments of 5 allows the reader to see the changes in participation without making these changes seem overly significant in light of the available information.

EXERCISES

Exercise Set 19A:

#1:

Give an example of an interesting or important non-comparative statistic.

#2:

Give an example of an interesting or important comparative statistic.

Exercise Set 19B:

Directions: For each of the following, say (i) what this statistic initially implies or suggests to you, and (ii) give at least one reason the statistic, as stated, might be misleading.

#1:

"Police officer deaths on duty have jumped nearly 20%."

#2:

"Illegal border crossings have fallen 81.5 percent since 2000."

#3:

"The number of homeless students in Seattle has increased by 55% since 2012."

#4:

"The state's unemployment rate is at its lowest level in 10 years."

#5:

"The average full-time working woman earns 78 cents for every dollar a man earns."

#6:

"Rates of violent crime in our state are twice that in a neighboring state."

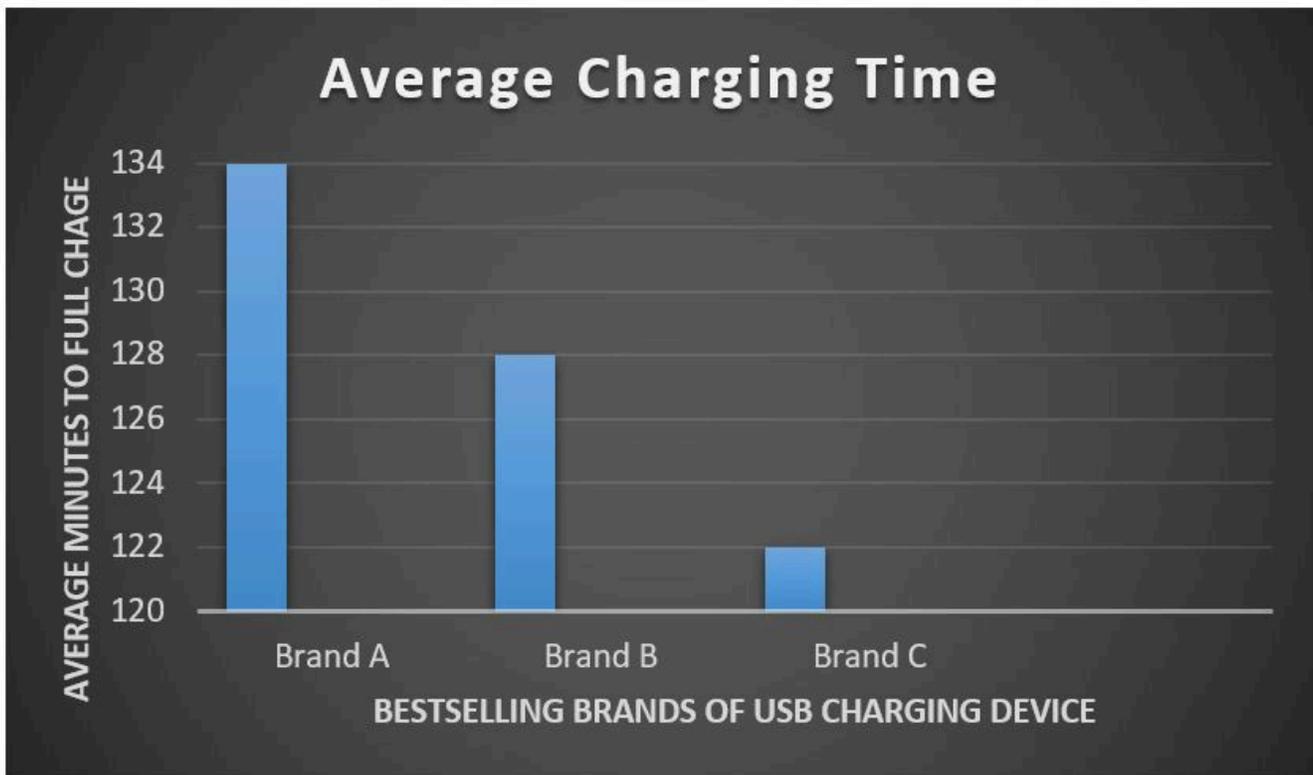
Exercise Set 19C:

#1:

Imagine that you've been asked to create a graph that represents changes in average price for unleaded gasoline over time in your city. What would be an appropriate scale to use? Why would it be appropriate?

#2:

Brand A makes USB charging cords for phones and other devices and features this bar graph in their advertisement. Evaluate this graph.



#3:

Find a graph in the news or on social media that is being used to support or illustrate a particular claim. Copy and paste it here, and then comment on its scale(s). Does the scale make sense? Why or why not? Be prepared to share.

Notes

1. Best, Joel. (2004). *More Damned Lies and Statistics*. Berkeley, CA: University of California Press, xii-xiii.
2. Park, Alice. (2014, Mar. 27). "U.S. Autism Rates Jump 30% from 2012." *Time*.
3. In order to avoid all the decimal places, the graph represents the murder rate per 100,000 people. Doing so allows us to present the rate in terms that are easier to read.

Correlation and Cause

SECTION 1: INTRODUCTION

The ability to identify causes is one of our most fundamental intellectual skills. It is because we can identify the likely causes of events that we can accurately predict how objects and people will behave, and this allows us to exercise control over our circumstances. More specifically, it allows us to pursue desired effects and to prevent undesirable ones. We will use the term **causal inference** to refer to a family of arguments all of which ultimately conclude that one thing causes or caused another. Most causal inferences are *Inferences to the Best Explanation*, and although we have already talked about such inferences it will be useful to take a closer look. In this chapter we will focus specifically on inferences that draw a causal conclusion on the basis of a correlation.

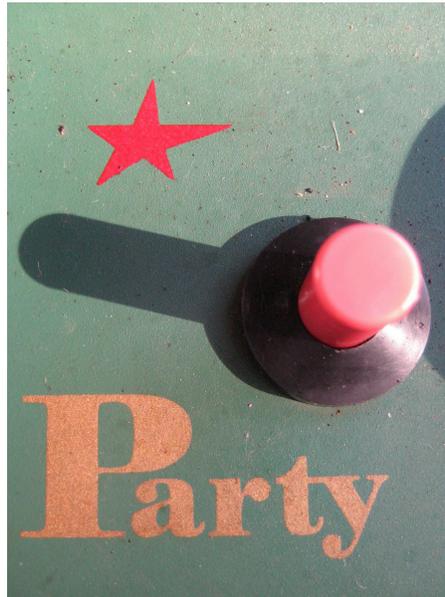
SECTION 2: CAUSAL LANGUAGE

We talk about causes constantly, though we do not always use the word 'cause'. The word 'cause' has many synonyms: for example, to 'make', 'create', 'generate', or 'produce', and many action terms express causation, for example: 'push', 'pull', 'hit', 'kick', and 'move'. Similarly, it is important to recognize that 'prevention' is a causal term, and so are its synonyms, 'hinder', 'hamper', 'impede', and so on. Before we look at causal reasoning itself, we need to clarify what we mean when we use causal terms. As it turns out, our everyday causal language tends to oversimplify the relations between events, and we often forget that virtually all events are the effects of a complicated series of prior events or circumstances.

Consider a very simple case. Suppose your computer is off and you want to turn it on so you can start working on your homework. What do you do? Presumably, you would simply push the 'on' button. Does pushing this button cause your computer to turn on? Well, yes, but this is a qualified 'yes'. After all, pushing the 'on' button does not, all by itself, turn on the computer. Pushing this button causes the computer to turn on only if it is also connected to a power source of some kind, the button functions correctly, the CPU is functional, and so forth. In an important sense, then, the cause of the computer's turning on is actually this whole set of conditions taken together. None of these circumstances alone are responsible for turning on the computer; rather it is the combination that brings about the desired effect. Let us call a group or set of conditions which together causes an effect, a **complete cause** of the effect, and let us call any individual circumstance or event that is part of a complete cause, a **partial cause** of the effect.

Despite the fact that pushing the 'on' button is only a partial cause of the computer's turning on, it is not surprising that in everyday language we tend to say that it is *the cause* of the computer's turning on. After

all, most of the time, you only need to push the button. In this way, pushing the button is a reliable way of bringing about the desired effect. Let us say in general that a **reliable cause** of some effect is an event the occurrence of which makes it very likely the effect will occur. Though pushing the button is only a partial cause of the computer's turning on, it is reliable because the other partial causes in the complete cause are largely fixed—they do not fluctuate much (the computer is normally plugged in, functional, etc.). Since these other conditions tend to hold, all you need to do is push the button.



"Party button" by Phillie Casablanca is licensed under CC BY 2.0

Knowing the reliable cause of some effect is incredibly useful, and when engineers create machines or other tools they normally design them so users can easily and reliably use the device. In other cases, however, we are not so lucky.

Often, many of the partial causes themselves vary or fluctuate. Take exposure to asbestos, for example. Asbestos is a mineral that, for many years, was mined and used as a building material. However, asbestos is now known to cause cancer. More specifically, however, asbestos is one among an incredibly complex array of bio-chemical factors which, taken together, can result in cancer. Further, many of the circumstances that figure into a complete cause of a person's developing cancer are not fixed. In other words, many of these factors are variable to some degree, and may not hold, or may not hold in the right way, to figure into the complete cause. Since a person develops cancer only if all the causal factors that figure into a complete cause of cancer hold, an exposure to asbestos alone will not reliably cause cancer. Indeed, many people who have been exposed to asbestos do not develop cancer. Of course, this is not to say that asbestos isn't dangerous—it is! Given this, it will be useful to introduce the notion of a probabilistic cause. Let us say that a **probabilistic cause** of some effect is a partial cause that raises the likelihood the effect will occur, but does not reliably bring about the effect. Indeed, it need not even raise the likelihood substantially—and an appropriately related event that raises the likelihood of the effect only 10% is still a probabilistic cause. Overall, probabilistic causes are common—especially when it comes to the kinds of complex systems found in organisms, ecosystems, or the atmosphere (to give a few examples).

In sum, our everyday causal language is remarkably imprecise. We often simply refer to one event as the cause of another, but this way of speaking glosses over different ways the events might be related. Indeed, in most cases what we actually mean when we express a causal claim is that one thing is the partial and reliable/probabilistic cause of another. Again, in both cases the cause makes the effect more likely, the difference between the two has to do with *how likely*. A reliable cause makes it highly likely the effect will occur, whereas a probabilistic one merely raises its chances to some degree.

SECTION 3: CORRELATIONS

We spend a lot of time trying to identify causes. This happens in an informal way all the time in our daily lives: “I wonder what caused my raised taste bud,” “why did I sleep so poorly the last few nights,” “how come there are so many accidents at this intersection,” and so on. Scientists and researchers pursue the same kinds of questions in a much more formal way, looking for the causes of a fish die-off, foods that contribute to heart disease, or whether a tax cut stimulated the economy, for example.

Either way, when looking for causes a common place to start is to look for statistical patterns, and finding the right kind of pattern can indeed be a good jumping off point towards this end. But what counts as the right kind of pattern, and why can it be a good jumping off point? The answer to both questions is grounded in the fact, noted above, that *causes make their effects more likely*. This fact means that the effect occurs more commonly when the cause is present than when it is absent. That is, when one thing is the cause of another, then this statistical pattern will hold between them (though this pattern holds in other cases as well).

To illustrate, smoking causes lung cancer, and consequently we find that the percentage of smokers who develop lung cancer is higher than the percentage of non-smokers who do. Similarly, studying for exams causally contributes to getting good grades on them, and we see that the percentage of students who get good grades is higher among people who study than people who do not. Finally, consuming a lot of sugar causes people to get cavities, and we accordingly observe that the percentage of people with many cavities is higher among those who consume a lot of sugar than it is among those who do not. The statistical pattern common to all of these causal relationships is called a positive correlation. This relation is also sometimes called an association or link (particularly in news reporting). Regardless of how we refer to it, in general we can say:

■ If X causes Y, then X is positively correlated with Y.

We will take two points from this. First, when you are trying to identify the cause of some phenomena it makes sense to look for a correlation involving it, because the cause—whatever it turns out to be—will be something that is correlated with the effect you are wondering about. Importantly, however, a correlation does not imply or entail a causal relationship (for reasons we will discuss in a bit). Nevertheless, a correlation is a place to start—it is a clue to causation.

Second, when you are looking for causes, and you find that a suspected cause is not correlated with the effect, then you can rule it out. Say that you get a raised taste bud after you eat sometimes, and you are trying to figure out what is causing it. You read online that for some people spicy foods can cause raised taste buds, and you’ve been watching what you eat lately. However, you’ve found you don’t get a raised taste bud when you eat spicy foods any more commonly than when you don’t eat spicy foods. Given this, spicy food is not correlated with your raised taste buds, and you can rule it out and move on (by *Modus Tollens!*).

Before we move on, it is important to clarify the relation a correlation expresses. In general terms, a positive correlation is a comparative statistic. We will use X and Y to stand for different categories or groups. Although correlations can hold between other kinds of variable as well (e.g. quantitative or continuous variables), we will focus on categorical correlations as they are both common and somewhat easier to understand. Given this, we will say:

X is **positively correlated** with Y when (and only when) the percentage of Xs that are Ys is *greater than* the percentage of non-Xs that are Ys.

To illustrate, let us look at a couple of new cases so you can clearly see the pattern. Being a professional basketball player is positively correlated with being 6'4 or taller, in that the percentage of professional basketball players who are 6'4 or taller is *greater than* the percentage of people who are not professional basketball players who are 6'4 or taller.

Importantly, a positive correlation is not saying that the number of Xs that are Ys is higher than the number of non-Xs that are Ys. There are, after all, only around 600 professional basketball players in the United States, although most of them are 6'4 or taller, there are hundreds of thousands of people who are not professional basketball players who are 6'4 or over within United States (pop. 330 million). Again, a positive correlation is saying that the percentage is higher—not the number. Here is one more example: in the United States being a Democrat is positively correlated with being a vegetarian. That is, the percentage of Democrats who are vegetarians is *greater than* the percentage of non-Democrats (Independents and Republicans) who are vegetarians.¹

As we noted at the beginning of the chapter, prevention is a causal relation as well. This means that we might be interested in identifying cases where the presence of some object or event makes the occurrence of another less likely, and this would show up in a statistical pattern called a negative correlation. We will say that:

X is **negatively correlated** with Y when (and only when) the percentage of Xs that are Ys is *less than* the percentage of non-Xs that are Ys.

For example, getting the vaccine for COVID-19 is negatively correlated with getting a severe infection. That is, the percentage of people who got the vaccine and suffered severe infections by the virus is *less than* the percentage of people who did not get the vaccine and suffered severe infections. Importantly, this is not inconsistent with the point made earlier that causes make their effects more likely. In the case of prevention, a cause produces, and therefore makes more likely, an effect that subsequently interferes or disrupts a distinct set of causes. By interfering in this second chain of events, the cause thereby makes the effect in question less likely.

It is important to add that correlations can be stronger or weaker. In everyday informal thinking, it is difficult to be precise about the degree of correlation. We tend to have only a vague idea of how much more commonly Y occurs given X than non-X. However, scientists and researchers are normally in a position to be much more precise, and measure the degree of correlation on a scale between -1 and 1. On this scale, 0 says that there is no meaningful relationship between the two variables, whereas 1 means there is perfect positive correlation between two variables, and -1 refers to a perfect negative correlation. This matters because the weaker the correlation the more likely it is mere coincidence, whereas the stronger the correlation the more likely it captures a genuine relation between two things.

Finally, although we have been thinking about correlation in the context of identifying causes, they are useful for other purposes as well. This is because a correlation allows you to make predictions. If you know that X is correlated with Y, and you know that X is present, then you know that Y is more likely *than not* to be present as well. For example, say you work in marketing. Knowing that consumers who are interested in your product tend to watch the game show *Jeopardy!* at a greater rate than other consumers, gives you important information about where to spend your advertising dollars. In this case, you aren't looking to identify cause/effect relationships, but simply to predict what potential consumers are watching. In any case, now that we have a sense for what a correlation says, let us turn to its use in causal inferences.

SECTION 4: INFERRING CAUSE FROM CORRELATION

Suppose that you have discovered a correlation. You learn from a reputable source that teenagers that spend more than 3 hours a day on social media tend to have higher rates of anxiety and depression than teenagers who spend less time on social media. This makes you start thinking about your own social media usage, and to wonder whether you should try to limit it as a way to prevent depression or anxiety. This train of thought embodies a move from correlation to cause, and implies that extended social media use is a partial cause of depression/anxiety. But can we really draw such a conclusion from the information provided?

Well, we need to be careful. Let's take a look at this inference:

1. Teens who spend more than 3 hours per day on social media have higher rates of anxiety and depression than those that spend less time.
2. So probably extended use of social media causally contributes to anxiety and depression among teens.

This argument begins with an observation about the world, and draws a conclusion about the likely cause or explanation for this state of affairs. That is, this is an *Inference to the Best Explanation*, and the link between the stated premise and stated conclusion is the assumption that the best explanation for this correlation is a causal relation.

[picture—teens and social media]

Given this, let us recall and apply the steps for evaluating an *Inference to the Best Explanation*.

Three Questions to Ask of *Inferences to the Best Explanation*:

- How likely is the proposed explanation?
- Are there other plausible explanations? (Failure: Hasty Explanation)
- Would the truth of the proposed explanation be less surprising than the truth of any competitor? (Failure: Poor Explanation)

Step one: does the proposal that social media use causally contributes to depression and anxiety seem like a likely explanation for the correlation? Well, if one really causes the other, then we would certainly find a correlation between the two. Whether a person regards this as plausible will depend on their relevant background knowledge, but this will likely seem plausible to many people, at least on its face.

Step two: are there other plausible explanations for the correlation? This is a little more difficult to answer, but yes there are (after all, there are many possible explanations for any phenomena—including correla-

tions). For example, it might be the other way around—it might be that anxiety and depression among teens leads them to spend a lot of time on social media. Alternatively, it may be that there is no causal relation at all—perhaps there is some distinct third factor that causes anxiety and depression among teens and pushes them toward extended social media use at the same time. For example, perhaps interpersonal conflicts with friends or family can have these effects. Finally, this correlation may be sheer coincidence. In this case, there will be no deeper explanation for the correlation—it is just a statistical anomaly. This is important to emphasize. Although coincidences are not common (that is part of what makes them coincidental!), they do happen—especially in cases where correlation is based on relatively few observations.

Given this, we are now in a position to understand the warning you might have heard that “correlation does not imply causation.” The idea behind this catchphrase is that we cannot jump from the observation of a correlation to a conclusion about causation without first considering and comparing alternative explanations. That is, because there are always many explanations for any given correlation, we cannot simply pick out one cause and claim it is THE explanation unless we have reason to believe that it is the best one available. To do otherwise is to draw a *Hasty Explanation*.



"Jumping to conclusions" by Cross-stitch ninja BY-NC-ND 2.0

This brings us to step 3: what is the most likely explanation among the options out there? In order to think through this, we will need to evaluate and ultimately rule out, or at least cast doubt upon, some of these explanations. How do we begin? Luckily, the possible explanations for a correlation between X and Y can be separated into only 4 types.

- Type 1: X causally contributes to Y.
- Type 2: Y causally contributes to X.
- Type 3: This is some underlying causal factor or factors which relate X to Y.
- Type 4: The correlation is accidental or coincidental.

Again, we can only conclude that one of these explanations is correct to the extent that we have ruled out, or at least cast doubt on its alternatives.

This raises the question: how do we rule out alternative causes? The answer to this question is complicated, but we can offer some general strategies. First, we know that causes come before their effects, thus if we

know that X occurred before Y, we can rule out the Type 2 explanation that Y caused X. Returning to the example, we'd want to know which came first—extensive social media use or depression/anxiety? Second, the more evidence we have supporting the correlation, the less likely the correlation is merely coincidental (Type 4). That is, a correlation that is based on *more* observations or data is less likely to be coincidental than one based on fewer observations or less data. So, we might ask: how many teens did researchers look at in conducting the study? This correlation is more likely to be mere coincidence if they looked at 10 teens than if they looked at 1000. A third, and related, point is that a correlation is less likely to be accidental or coincidental to the extent that we can envision a possible mechanism connecting X and Y (though our ability to do depends a lot on our background knowledge). Thus we might ask: are there some known connections between extended social media use and people's emotional state that might account for the connection? Fourth, what about the possibility of underlying causal factors (Type 3)? This is the most difficult possibility to rule out or cast doubt upon. After all, there are always factors we might not be aware of. In this case, we can only proceed by considering possible underlying causes and trying to rule them out one by one (although we will discuss this point further in the next chapter).

In any case, once we've identified one explanation as the most likely by at least casting doubt on the others, we can claim to have identified a correlation that is likely to capture a causal relation. Such correlations are called **significant correlations**. There are other techniques that we can use to rule out possible causes, and thereby identify significant correlations. In the next chapter we will take a look at one of the best methods: a controlled experiment.

EXERCISES

Exercise Set 20A:

#1:

Give an example of a reliable cause (other than the examples given in the reading).

#2:

Give an example of a probabilistic cause (other than the examples given in the reading).

#3:

Give an example of a correlation that you think is significant or important (other than the examples given in the reading).

#4:

Suppose that being a Bollog is positively correlated with being a Trollog. Does it thereby follow that most Bollogs are Trollogs? Why or why not? (*Note: you do not need to know what a Bollog or Trollog are in order to answer this question.*)

Exercise Set 20B:

Directions: Using the following chart to answer the questions below.

A college is thinking about changing its mascot, and wants to gauge community attitudes. So the school conducts a poll. As the poll is set up, students must reply either 'In Favor' or 'Not in Favor'. Here is the part of the information they collected that compares responses from students involved in athletics at the college and students who are not.

	IN FAVOR	NOT IN FAVOR
ATHLETES	15	85
NON-ATHLETES	80	120

Exercise 20B Poll Data

#1:

Is there a correlation between being an *athlete* and being *not in favor* of the change? If so, explain, and say whether it is positive or negative.

#2:

Is there a correlation between being an *athlete* and being *in favor* of the change? If so, explain, and say whether it is positive or negative.

#3:

Is there a correlation between being *not in favor* and being an *athlete*? If so, explain, and say whether it is positive or negative.

Exercise Set 20C:

Directions: for each of the following (i) identify the correlation on the basis of which the causal inference is drawn, (ii) the causal conclusion, and (iii) using appropriate questions think through the argument, and comment on the strength of the inference. In each case, assume the correlation is true.

#1:

Do you want healthy teeth? Buy a speedboat! According to a recent study, people who own speedboats have fewer serious dental problems than those who do not.

#2:

People with allergies to animal dander are less likely to have pets than those who do not have these allergies. Therefore, people should avoid getting pets if they want to avoid developing allergies to pet dander.

#3:

We should seriously consider putting limits on the amount of cheese a person can buy, since per capita cheese consumption almost perfectly matches the number of people who die by becoming tangled in their bedsheets each year!²

#4:

College entrance exams aren't fair since you can effectively *buy* a high score. Look at the stats: students from high income families disproportionately score in the top 10% of all SAT scores.

#5:

You should encourage your grandpa to walk more. I read the other day that elderly adults who walk for at least 180 minutes per week have lower rates of dementia than those who walk less.

#6:

Teenagers who are classified as heavy users of marijuana are more likely to have strained relations with at least one parent than students who use marijuana less often or not at all. We should recommend limiting marijuana use among teenagers as a way to improve relationships within families.

#7:

During the 20th century, Washington D.C.'s professional football team had an amazing degree of control over presidential elections. Over the interval from 1940-2000 the outcome of the final Washington D.C football team's home game perfectly matches presidential election results. If they won, the incumbent party held on to the presidency; if they lost, the opposition party took it.³

#8:

You have slept very poorly the last few nights. You are wondering what is causing it, and notice that there have been poor air quality advisories the last few days too. Moreover, prior to this stretch of poor air quality you were sleeping fine. You reason that the poor air quality is probably disrupting your sleep and think about getting an air purifier for your room.

Notes

1. "American Dietary Preferences are Split Across Party Lines." (2018, Nov. 22). *Economist*. <https://www.economist.com/graphic-detail/2018/11/22/american-dietary-preferences-are-split-across-party-lines>
2. Source: Spurious Correlations. <https://tylervigen.com/old-version.html>.
3. Allen, Scott. (2016, Oct. 20). "Redskins Rule used to predict election, but the guy who discovered it now says it's a crock." *Washington Post*. <https://www.washingtonpost.com/news/dc-sports-bog/wp/2016/10/20/redskins-rule-used-to-predict-elections-but-the-guy-who-discovered-it-now-says-its-a-crock/>.

Significant Correlations and Controlled Studies

SECTION 1: INTRODUCTION

Science is in the business of offering explanations—more specifically, it is in the business of trying to discover the causal relations in nature. Identifying a correlation can be one step in this process, and in the last chapter we looked at correlations and how they relate to causal thinking more broadly. As we saw, although a correlation can be a good place to start, by itself a correlation never justifies drawing a causal conclusion. After all, there are many possible explanations for the existence of a correlation. Moreover, we briefly talked about some ways of ruling out alternative explanations, and thereby identifying **significant correlations**—or correlations that are suggestive of a causal relation. In this chapter we will look at some more sophisticated techniques for identifying significant correlations—techniques often used in science and social science. We will look at two of the most common: the observational study and the randomized controlled study. The following is not intended to prepare you to conduct your own studies—conducting a study requires expertise and resources (and time!) most of us do not have. Rather, the aim is to show the reasoning behind these studies, and to give us the tools and vocabulary to be more informed and critical readers of scientific studies.

SECTION 2: OBSERVATIONAL AND EXPERIMENTAL STUDIES

Suppose that you have spotted a pattern. Suppose that among your friends you have noticed that those who regularly eat breakfast have higher grade point averages (GPAs) than those who do not. This observation leads you to think that maybe you ought to eat breakfast, and maybe other students should too. Put in our terminology: among your friends, regularly eating breakfast is correlated with GPA, and this correlation may be indicative of a causal relation between the two. So, is that right—does this correlation establish that eating breakfast is causally relevant to GPA? Well no. To conclude that the eating breakfast is causally relevant to GPA is to risk a hasty explanation. Not only is this correlation based on a relatively small sample (your friends), but there are many possible explanations for any correlation, and we need to investigate the plausibility and adequacy of these alternatives before identifying a particular explanation as the best.

Nevertheless, observing this pattern gives you a place to start. In order to determine whether this pattern is suggestive of a causal relation, you would need to scale up your inquiry in a way that would allow you to (i) find a genuine correlation if there is one, and (ii) rule out, or at least cast doubt upon, some of the possible explanations for the correlation. We will discuss two techniques a researcher might use. The first technique is called an observational study. To conduct an **observational study** in this case, a researcher would

systematically identify and examine a group of students to see whether there is a significant correlation between the two factors in question. Alternatively, a researcher might choose to conduct an experimental study. Unlike an observational study, in an **experimental study** researchers go beyond mere observation to intervene by systematically exposing subjects to the suspected cause. In this case, the researcher would (among other things) identify a specific group of students, give them breakfast (or not), and see what happens. How do these techniques take up or cast doubt upon alternative explanations for the correlations? Let us take a closer look one-at-a-time.

SECTION 3: OBSERVATIONAL STUDIES

There are different kinds of observational study. A retrospective observational study begins with the effect you are interested in and looks backwards in time to try to isolate a cause. To conduct a retrospective study on the Breakfast/GPA hypothesis, you'd start with students' GPAs and then look to see whether they eat breakfast or not. A prospective observational study, on the other hand, begins with a suspected cause and follows it forward in time to see if the effect follows. So, in the case at hand you'd start by looking at breakfast eating patterns and look for patterns in GPAs. There are other differences between prospective and retrospective studies (e.g. retrospective studies tend to be much easier to conduct), but in discussing the reasoning behind observational studies we will focus on a prospective study, though our conclusions will largely apply to retrospective studies as well.

Let us return to the case at hand. In order to dig deeper into your tentative view that breakfast contributes to GPA, you've decided to conduct a prospective observational study. How will this work? You'll have to start by identifying a group of subjects. These people will constitute your sample. The larger your sample is, the less likely your results will be mere coincidence. So, you'll want to work with the largest sample that is feasible given the limitations of time, resources, and so on. How should you choose the people in your sample? Here is a strategy that might seem intuitive: find as many students as you can who regularly eat breakfast. Once you've got this group, you could then simply look at their GPAs. This is too quick, however. This strategy for identifying a sample won't work since, as we've seen, correlations are comparative claims. If we are looking to see whether eating breakfast is significantly correlated with GPA, we'll need to compare the percentage of people who eat breakfast and have high GPAs to the percentage of people who do not eat breakfast who have high GPAs. That is, we need to look at two groups with respect to GPAs: the **experimental group** that exhibits the suspected cause and the **control group** that does not. Let's call the experimental and control groups in this case the *Breakfast Group* and the *No-Breakfast Group*.

Say you have identified a reasonably large sample of students, half of which tend to eat breakfast, half of which do not. This makes it more likely you'll find a genuine correlation if there is one, and mitigates against the possibility of coincidence. Are we ready to look at GPAs? Not yet. Remember, we want to know whether breakfast is a causal contributor to GPA. Finding a correlation between the two will not, all by itself, suggest this, since there are many possible explanations for any correlation. Recall the four types of explanation from Chapter 20. It could be that breakfast contributes to GPA as you suspect (Type 1), but it could also be the other way around (Type 2). Moreover, there could be some underlying factor that explains the correlation (Type 3), or it could be coincidence (Type 4).

The set-up of the study mitigates the possibility of coincidence, and background knowledge casts doubt on Type 2, since it seems unlikely that a GPA could influence breakfast habits. This leaves Type 3: an underlying cause. Could there be some underlying factor that accounts for a person's breakfast habits and their GPA? Sure. Of particular importance is the possibility of one or more confounding factors (sometimes called "lurk-

ing factors”). A **confounding factor** for a particular study is a factor which is (i) correlated with the suspected cause under investigation and (ii) a partial cause of the effect.



“Saffron” by *albastrica mititica* CC BY 2.0

Here is an example. Suppose you were interested in the possible health benefits of the spice, saffron. To investigate this question, you might examine people who regularly eat saffron to see if they are healthier than those who never (or rarely) do. There are confounding factors here, however. Saffron is an extremely expensive spice; as a consequence, people who regularly consume saffron also tend to be relatively wealthy. Thus, there is a correlation between wealth and the suspected cause (saffron consumption). In addition, there is also a correlation between wealth and the suspected effect, since wealthy people have access to regular preventative health care at a higher rate than non-wealthy people do. Thus, if there is a correlation between saffron consumption and health, this correlation may be explained by people’s wealth—not because saffron has any health benefits. That is, wealth is a confounding factor in the effort to determine whether saffron has health benefits.

Let’s return to the question of whether eating breakfast is causally relevant to GPA. How could we conduct our study to avoid confounding factors? In brief, we will need to (i) think about the kind of factors that may be causally relevant to GPA, and (ii) make sure we take this into account in deciding how to construct the *Breakfast* and *No-Breakfast* groups. Given this, let us consider factors that might influence a person’s GPA. One prominent determinant of a student’s GPA, for example, is the difficulty of their course schedule. In some classes it is more difficult to get an ‘A’ than in others. To simplify, let us pretend that there is only one class like this—say Organic Chemistry. In addition, imagine that 10% of the students in the *Breakfast Group*, but none of the students in the *No-Breakfast Group*, are enrolled in Organic Chemistry. In this case, we would have a confounding factor, since being enrolled in Organic Chemistry is causally relevant to a student’s GPA *and* is correlated with the suspected cause (eating breakfast).

Why does this matter? Well, if we don’t somehow account for students’ course schedules, our results might not tell us anything definitive about whether eating breakfast is relevant to GPA—even if it is! Here is why: say that we find no difference in GPA between the groups. The natural inference here would be that eating breakfast—at least within this group of students—had no effect on GPA. But if the groups differ in this way, it may be that that eating breakfast *really did* have an effect, but that this effect has been hidden or masked by the unusually low grades coming from the students enrolled in Organic Chemistry. Put in different terms, if eating breakfast is relevant, then you wouldn’t know it, since its effects on GPA are mixed together or con-

founded with the effects of course difficulty. In order to determine if eating breakfast contributes to GPA we need to isolate it from other causes, so that we can spot its effects (if there are any).

In order to prevent this kind of confounding, we need to take account of, or control for, the difficulty of students' courses. In general terms, to **control for** a particular factor or variable is to ensure that there is no difference with respect to that variable between the two groups you are comparing. This allows you to isolate the potential effects of the factor you are interested in. To control for the difficulty of students' courses in this case would mean making sure that the *Breakfast* and *No-Breakfast* groups are as similar as possible with respect to the number of students enrolled in Organic Chemistry. Similarly, to control for wealth in the saffron study discussed above would mean looking at two groups of equally wealthy people—some of whom regularly eat saffron, some of whom do not.

SECTION 4: LIMITATIONS OF OBSERVATIONAL STUDIES

Observational studies, even when they control for known confounding factors and study large samples, do not always give accurate results. Perhaps the chief problem is that observational studies like this cannot account for the effects of **unknown confounding factors**. In other words, an observational study *cannot* rule out all the competing explanations for the phenomena in question.

Take the case discussed above; suppose we have conducted our study in a way that controls for a variety of relevant factors (including course difficulty), and that our original suspicion has been confirmed—in this much larger and more diverse group there is a correlation between eating breakfast and GPA. While this gives us stronger reasons to think that this correlation is not mere coincidence, and further that eating breakfast is causally relevant—we cannot be sure this is the case. The problem is that there may be *unknown* confounders out there. For example, it may have not occurred to you that getting up early might be causally relevant to a person's GPA. Let us assume for a moment that this is right, and that, further, people who get up early tend to eat breakfast. If we haven't controlled for getting up early, then the results of our observational study will suggest that eating breakfast is causally relevant when, in reality, it is not.

A well-known case illustrating the limitations of observational studies involves hormone replacement therapy. In the 1990s a number of professional observational studies correlated hormone replacement therapy with a decreased risk of heart disease in post-menopausal women. Scientists had a good sense of how this replacement therapy might work, and as a result many doctors recommended this to their postmenopausal patients. However, in the early 2000s it became clear that there must be a confounding factor at work in these studies. A large randomized controlled experiment showed that though hormone replacement therapy could have positive health benefits for some women, it did not prevent heart disease and was actually associated with a number of negative health outcomes. As it turned out, overall health was the confounding factor. On balance, healthy women were more likely to pursue and be prescribed hormone replacement therapy than their unhealthy counterparts. Since a person's overall health is causally relevant to the chance of developing heart disease, this was a confounding factor.

This shows that observational studies have their limits, and that there are alternative methods for ruling out competing explanations. Let us take a closer look at the kind of study which ultimately revealed the flaw in this observational study—the randomized controlled experiment.

SECTION 5: RANDOMIZED CONTROLLED EXPERIMENTS

In well-conducted observational studies researchers make careful choices about what populations to observe. While this is also true in an experimental study, a researcher conducting an experimental study goes one step further by manipulating the suspected causal variable. Consider the breakfast/GPA case discussed above. In an observational version of this study, you divide students into two carefully chosen groups—the *Breakfast* and *No Breakfast* groups—and then check to see whether students who eat breakfast end up having better GPAs than those who do not. In contrast, in an experimental study the researcher chooses which students will eat breakfast (the experimental group) and which will not (control). There are different kinds of experimental study, but here we will focus on a particularly important one: the randomized controlled experiment or trial. As one author recently put it, “The randomized controlled trial...is one of the simplest, most powerful, and revolutionary tools of research.”¹ In this section we will explain what a randomized trial is, and why it is so powerful.

By their very design experimental studies can cast doubt on many of the possible explanations for a correlation. First, like an observational study, an experimental study takes a systematic look at a wide body of information, and in doing so limits the possibility of sheer coincidence (Type 4 explanations). Second, given a correlation between Xs and Ys, an experimental study can rule out the possibility that Ys are causing the Xs instead of vice versa (Type 2), since causes precede effects. In an experimental study the researcher introduces the suspected causal factor (X) to subjects in which the suspected effect (Y) is absent. So if the effect is subsequently observed in the population you can be sure that this is not a case of Ys causing Xs.

As we have seen, observational studies are always subject to the possibility of unknown confounders. However, the possibility of confounders, known and unknown, is greatly limited by a process of randomization—hence the value of the randomized controlled experiment. But wait. What, exactly, is randomized, and how does this mitigate worries about confounders?

What makes an experiment like this randomized is that the individuals or members of the experimental group and the control group are chosen randomly from a targeted population. To do so, a researcher uses a procedure that gives each member of the population an equal chance of being chosen. So, in this case you might assign students a number and use a random number generator that is available online to assign individuals to each group. Ok, but how does randomizing who is chosen for each group mitigate against confounders?

Let us return to the unknown confounder considered in the breakfast case: for the sake of argument let us assume that it is not eating breakfast, but getting up early that is causally relevant to GPA (perhaps because they are more alert during morning classes). Because there are more people who get up early in the *Breakfast Group* than in the *No-Breakfast Group* your observational study will mistakenly suggest that breakfast is causally relevant to GPA. However, in a randomized controlled experiment this is much less likely. How so?

By randomly assigning students to either the experimental group or the control group, you would likely distribute early risers evenly into both groups (at least roughly). Randomizing in this way breaks the problematic relation between breakfast eaters and early risers, and allows us to see more clearly whether eating breakfast, in itself, has any effect. The same will go for other unknown confounders; after all, arbitrarily splitting subjects into the control and experimental groups will likely distribute other possible confounders evenly (at least roughly) between the groups. In other words, by randomly assigning subjects to the experimental and control groups the researcher will thereby automatically control for unknown confounders.



"Start The Day With A Smile" by outdoorPDK CC BY-NC-SA 2.0

This is the chief benefit of a randomized controlled experiment—it reduces the chances that an unknown confounder will bias your results. Unfortunately, however, a study of this kind does not completely eliminate this possibility. It is always possible—even with a random assignment of subjects—that some causally relevant factor will end up accidentally associated with the experimental group. Even so, a randomized controlled experiment is much more likely to give you an accurate picture of the relationship between two (or more variables), and these studies are considered the single best way to determine causal relationships.

SECTION 6: LIMITATIONS ON EXPERIMENTAL STUDIES

The comparison above between different kinds of experimental techniques raises a question: if randomized controlled experiments are the best way to identify significant correlations and causal relations, why do researchers ever do any other kind of study? First off, randomized controlled experiments are time consuming and can be quite expensive to conduct. Second, in some cases it is not practical or ethical to do a randomized controlled experiment. To illustrate, consider the disease known as Ebola. Ebola is a viral hemorrhagic fever that carries a high risk of death for which there is no known cure. However, there are a number of experimental treatments. The problem is that in order to do a randomized controlled study for one of these treatments, researchers would need to set up a control group that received only a placebo. In this case, the control group would not get the possible benefits of the treatment, which in this case might realistically include survival.

This point generalizes, in that when a person's life is at stake, an experimental study can be simply unethical. This point was made in a (somewhat) humorous way in an article from the *British Medical Journal*:

"As with many interventions intended to prevent ill health, the effectiveness of parachutes has not been subjected to rigorous evaluation by using randomized controlled trials. Advocates of evidence based medicine have criticised the adoption of interventions evaluated by using only observational data. We think that everyone might benefit if the most radical protagonists of evidence based medicine organized and participated in a double-blind randomized placebo controlled, crossover trial of the parachute."²

Conducting a randomized study of the effectiveness of parachutes using real people would be obviously unethical, and so too when it comes to other potentially life-saving or life interventions.

A third consideration is that simply because randomized controlled studies are the single best way to uncover significant correlations, does not mean that observational studies have no value. Observational studies tend to be more economical in terms of time, effort, and resources than experimental studies, and so are useful (among other things) for studying a hypothesis in a preliminary way—as a way of deciding whether it is worth using the resources to conduct a randomized controlled experiment. Moreover, when multiple observational build upon one another and largely concur in their conclusions, then they can give us good reason to endorse causal conclusions.

EXERCISES

Exercise Set 21A:

Directions: Consider the following possible studies. What confounding factors might you need to control for?

#1:

A study to assess whether being an early reader (a person who learns to read before age 5) causally contributes to academic success.

#2:

A study to assess whether fluency in a second language improves scores on standardized tests.

#3:

A study to assess whether taking multi-vitamins prevents chronic illness.

#4:

A study to assess whether listening to classical music while studying makes it easier to remember information.

Exercise Set 21B:

#1:

Comment on the following experiment:

A vitamin company has developed a new pill intended to prevent strep throat (a kind of bacterial infection). The company claims that they pill was given to 2000 people daily for a month, and only 3% of subject came down with strep throat during this time. On the basis of this experiment, the company sells the vitamin as a preventative.

#2: What kind of study is the following, and what do you make of the study itself?

A researcher is interested in investigating whether pet ownership contributes to lower blood pressure. The researcher identifies 30 people who have pets and 30 people who do not. She takes everybody's blood pressure several times, and then averages the results. It turns out that the pet owners in the researcher's sample have lower blood pressure than the non-pet owners. The researcher takes this to be good preliminary evidence of a causal connection.

#3:

In his book *Exercised: Why Something We Never Evolved to Do is Healthy and Rewarding*, Daniel Lieberman comments on a study of physical activity and health. He summarizes and evaluates the study as follows.

“Researchers put accelerometers “on a diverse sample of eight thousand Americans above the age of forty-five and then tallied up who died over the next four years—about 5 percent of the sample. Predictably, those who were more sedentary died at faster rates, but these rates were lower in people who rarely sat for long uninterrupted bouts....One flaw with this study is that people who are already sick are inherently less able to get up and be active.”³

What kind of study is Lieberman commenting on, and what correlation did it find? Also, explain Lieberman’s criticism of the study.

#4:

What kind of study is the following, and what do you make of the study itself?

A researcher wants to study whether taking notes by hand is more effective than typing notes when it comes to remembering them. He identifies 30 students who are willing to participate, and allows them choose whether they’ll be in the experimental group or the control group, and the groups end up even (15 people in each). The researcher then requires students to attend the same hour-long lecture. One group takes notes by hand, the other takes notes by typing on their laptops. Students are allowed to study their notes and are given a test over the lecture 3 days later. When the researcher crunches the numbers, it turns out that students in this population who took notes by hand got higher grades than those who took notes on a laptop.

#5:

Come up with one experiment you’d love to know the results of—if you had the time, money, and expertise to do so. What would it be, how you set it up, and why?

Notes

1. Jadad, Alejandro R. and Enkin, Murray W. (2007). *Randomized Controlled Trials: Questions, Answers, and Musings 2nd ed.* Malden, MA: Blackwell, 1.
2. Smith G. C. , Pell, J. P. (2003). "Parachute use to prevent death and major trauma related to gravitational challenge: systematic review of randomised controlled trials." *BMJ*. 327(7429):1459-61.
3. Lieberman, Daniel E. (2020). *Exercised: Why Something We Never Evolved to Do is Healthy and Rewarding*. New York: Pantheon Books, 66-67.

Unit #7 Summary

In this unit we looked at some forms of reasoning commonly used in the sciences. We began by looking at some important components of statistics and their visual representations. In Chapter 20 we turned to an examination of causal language, and a particularly important kind of comparative statistic: the correlation. Although correlations can be especially useful starting point for identifying likely cause/effect relationships, not every correlation captures a causal relation. Before drawing an inference like this, we need to check to see whether other conditions have been met. Chief among them is whether there are other possible explanations for the correlation besides the hypothesized causal relation. As we saw, there are all kinds of techniques for identifying significant correlations, though some techniques are better than others.

A Key Question for a Statistic:

- What, exactly, was counted?

Questions to Ask of Inferences from correlation to cause:

- How likely is the proposed explanation?
- Are there other plausible explanations for the correlation?
- Would the truth of the proposed explanation be less surprising than the truth of any competitor?

KEY TERMS:

- Statistical Claim
- Comparative Statistics
- Outliers
- Scale
- Relevantly Similar Statistics
- Causal Inference
- Complete Cause
- Partial Cause
- Reliable Cause
- Probabilistic Cause
- Positive Correlation
- Negative Correlation
- Significant Correlation
- Observational Study
- Experimental Study
- Experimental Group
- Control/Control Group
- Confounding Factor
- Controlling for a Variable
- Unknown Confounders
- Experimental Group
- Control Group

FURTHER READING

There are a number of short readable books about the pitfalls of dealing with statistics. A famous, but dated, source is *How to Lie with Statistics* by Darrell Huff and Irving Geis. For more contemporary examples, you could check out Joel Best's *Damned Lies and Statistics* or *Is that a Fact? A Field Guide to Statistical and Sci-*

entific Information by Mark Battersby. For a deeper discussion of experimental structure and set up, see *Randomized Controlled Trials: Questions, Answers, and Musings* by Alejandro Jadad and Murray Enkin.

Concluding Thoughts

In conclusion, let us briefly look back over the text as a whole. Taking this wider angle will allow us to see how the text has woven several themes together, and will put us in a position to think about reasoning and argument more broadly. As noted at the outset, the goal of the text is to help its readers reflectively enhance their general reasoning skills. We have pursued this goal by walking through the basics of argument identification, analysis, and evaluation within the social and psychological contexts in which they normally occur. Along the way, we have gradually built up a vocabulary for describing the elements, standards, and circumstances for reasoning.

The most important running theme in the text has been *asking the right question*. As we have seen, uncovering the intended structure of an argument means asking about the author's main point and reasons, working to cooperatively resolve ambiguity, and so on. Turning to argument evaluation, we isolated three main questions to ask of any (inductive) argument.

Evaluation: The Three Main Questions

- *Are the premises likely to be true? (A check for Factual Correctness)*
- *Would the truth of the premises make the truth of the conclusion probable? (A check for "internal" Logical Strength)*
- *Is there any other relevant information available? (A check for "external" Logical Strength)*

We discussed the process of asking each one of these questions. While checking for factual correctness can be difficult in some cases (e.g. conditionals), we spent most of our time thinking about how to determine logical strength. Indeed, when we know what kind of argument we are talking about, there are more specific questions we can ask to determine an argument's logical strength. We focused on four of the most common types of (inductive) argument, and identified a series of more specific questions to help us determine logical strength for each type.

Two Questions to Ask of *Arguments from Analogy*:

- Is the noted similarity relevant to the inferred similarity?
- Are there differences that are relevant?

Three Questions to Ask of *Inferences to the Best Explanation*:

- How likely is the proposed explanation?
- Are there other plausible explanations?
- Would the truth of the proposed explanation be less surprising than the truth of any competitor?

Two Questions to Ask of *Inductive Generalizations*:

- Is the sample large enough?
- Is the sample diverse enough?

Two Questions to Ask of *Inductive Applications*:

- Is the individual in question a member of the subject class *or* not a member of the predicate class?
- Is the individual in question a member of other relevant classes?

Another main theme at work throughout the text has to do with the influence of social factors on our thinking. Reasoning is not the solitary task it is sometimes imagined to be, and almost always takes place within a broader social context. We began by focusing on cooperative dialogue, and cooperative disagreement specifically. In addition, we have seen how social factors can activate biases *and* help mitigate against them. We have discussed the extent to which we depend on other people for information, and use them to check our own thinking. Moreover, we saw how important trust is for learning from other people, and took a close look at when trust is warranted, and when it is not. Lastly, we have seen how the internet and social media give us access to an incredible variety of new voices and information, while simultaneously amplifying misinformation and making it difficult to know who to trust.

The final theme is psychological. At the outset, we distinguished between largely automatic and implicit reasoning processes, on the one hand, and consciously directed ones, on the other. Although we have focused on the latter, we have considered how our conscious reasoning can be influenced by these implicit processes. Myside bias, for example, can arise from largely intuitive responses and preferences. In addition, we saw that in some contexts we are subject to a variety of biases or cognitive illusions. We often read conditionals as if they were symmetric even though they are not, and we tend to overestimate how common something is when it is especially interesting or provocative. Further, we underestimate the influence of situational factors when we explain other people's behavior, and we tend to think other people are more like us than they really are.

So where does this leave us? Hopefully, taking the time to reflectively work through this text has made you a more careful and active thinker, and has given you greater confidence to navigate the important disagreements, challenges, and decisions you will face.