



THE
GREAT
COURSES®

Topic
Philosophy

Subtopic
Modern Philosophy

An Introduction to Formal Logic

Course Guidebook

Dr. Steven Gimbel
Gettysburg College



PUBLISHED BY:

THE GREAT COURSES
Corporate Headquarters
4840 Westfields Boulevard, Suite 500
Chantilly, Virginia 20151-2299
Phone: 1-800-832-2412
Fax: 703-378-3819
www.thegreatcourses.com

Copyright © The Teaching Company, 2016

Printed in the United States of America

This book is in copyright. All rights reserved.

Without limiting the rights under copyright reserved above,
no part of this publication may be reproduced, stored in
or introduced into a retrieval system, or transmitted,
in any form, or by any means
(electronic, mechanical, photocopying, recording, or otherwise),
without the prior written permission of
The Teaching Company.



Steven Gimbel, Ph.D.

Professor of Philosophy
Gettysburg College

Dr. Steven Gimbel received his Ph.D. from Johns Hopkins University before joining the faculty at Gettysburg College, where he is a Professor of Philosophy. At Gettysburg, he has received the Luther W. and Bernice L. Thompson Distinguished Teaching Award and was named to the Edwin T. and Cynthia Shearer Johnson Distinguished Teaching Chair in the Humanities. He also serves as Chair of the Philosophy Department.

Dr. Gimbel's research focuses on the philosophy of science, exploring the nature of scientific reasoning and the ways in which science and culture interact. He has published numerous articles and four books: *Defending Einstein: Hans Reichenbach's Writings on Space, Time and Motion*; *Exploring the Scientific Method: Cases and Questions*; *Einstein's Jewish Science: Physics at the Intersection of Politics and Religion*; and *Einstein: His Space and Times*.

Dr. Gimbel's previous Great Course is *Redefining Reality: The Intellectual Implications of Modern Science*. ■

TABLE OF CONTENTS

INTRODUCTION

Professor Biography	i
Course Scope.	1

LECTURE GUIDES

LECTURE 1

Why Study Logic?	4
----------------------------	---

LECTURE 2

Introduction to Logical Concepts	12
--	----

LECTURE 3

Informal Logic and Fallacies.	22
---------------------------------------	----

LECTURE 4

Fallacies of Faulty Authority	31
---	----

LECTURE 5

Fallacies of Cause and Effect.	39
--	----

LECTURE 6

Fallacies of Irrelevance.	47
-----------------------------------	----

LECTURE 7

Inductive Reasoning.	55
------------------------------	----

LECTURE 8	
Induction in Polls and Science	64
LECTURE 9	
Introduction to Formal Logic.	73
LECTURE 10	
Truth-Functional Logic	86
LECTURE 11	
Truth Tables	98
LECTURE 12	
Truth Tables and Validity.	112
LECTURE 13	
Natural Deduction.	121
LECTURE 14	
Logical Proofs with Equivalences.	133
LECTURE 15	
Conditional and Indirect Proofs	144
LECTURE 16	
First-Order Predicate Logic	152
LECTURE 17	
Validity in First-Order Predicate Logic	162
LECTURE 18	
Demonstrating Invalidity	172
LECTURE 19	
Relational Logic	180

LECTURE 20	
Introducing Logical Identity	190
LECTURE 21	
Logic and Mathematics	199
LECTURE 22	
Proof and Paradox	208
LECTURE 23	
Modal Logic	217
LECTURE 24	
Three-Valued and Fuzzy Logic.	224

SUPPLEMENTAL MATERIAL

Bibliography	232
Image Credits.	234

An Introduction to Formal Logic

Logic is the study of rational argumentation. A belief is rational if we have good reason to believe that it is true. An argument is a set of sentences such that one sentence, the conclusion, is claimed to follow from the other sentences, the premises. That means that logic is the study of the support we need to have good reason to believe the things we believe.

We are not wired to be logical. All kinds of bad reasons for believing things seem to us like perfectly good reasons. But the situation is not hopeless. It turns out that logic is a skill that can be learned. We can be trained to spot the reasoning errors, look for the proper kind of support, and analyze arguments we hear from others to assess their strength and weaknesses.

We'll begin by examining the human mind and seeing the ways in which logic is and is not a natural part of the way we think. We'll look at some of our cognitive biases, ways in which social psychologists have demonstrated that the brain naturally works against good inferences. Humans can be rational beings, but it takes work to realize the pitfalls we need to avoid.

Then, we'll introduce a wide range of logical concepts. We will rigorously introduce the notion of an argument and examine both the types of arguments—deductive and inductive—and the criteria by which we assess an argument—validity and well-groundedness. We will learn that arguments have two parts: conclusions (that which is being argued for) and premises (the support given for the conclusion).

Next, we'll focus on informal logic—that is, considerations of well-groundedness, the criterion of assessment that considers the truth of an argument's premises. We'll learn to spot common fallacies, reasoning errors that sound good to the ear but that actually undermine the support for the conclusion.

Inductive arguments are those that begin with observation and lead to broader, generalized conclusions. We'll study the basis for strong inductive arguments and the ways in which inductive arguments can go wrong, and then we'll examine the use of inductive arguments in science and polls.

Formal deductive logic is the part of logic concerned with the forms of deductive arguments. An argument is deductive if the content of the conclusion is contained in the content of the premises. If a deductive argument has good form, regardless of the truth or falsity of its premises, we say that the argument is valid.

To handle different types of sentences, we will develop two logical languages. The first, truth-functional logic, looks at the logical behavior of truth-functional connectives—words such as “not,” “and,” “or,” and “if”—which produce sentences that we can determine the truth or falsity of simply by knowing the truth or falsity of the sentences combined. We will learn how to determine whether or not such arguments are valid using both truth tables and natural deduction proofs.

The second logical language is first-order predicate logic, which builds on top of the elements of truth-functional logic in a way that allows us to account for the logical content within sentences as well as between them. In this way, a more powerful logical system can be constructed that can handle everything in truth-functional logic, everything in Aristotle's logic, and more.

The pinnacle of first-order predicate logic for us will be the ability to account for sentences that contain quantities. This means that we can

begin to account for mathematical truths in our logical system. The possibility of giving logical justifications for mathematical propositions is a major element in the history of the development of logic. We'll discuss the rise and fall of logicism, the view that logic provides the ultimate foundation for all of mathematics.

We'll examine modal logic, which deals with the concepts of possibility and necessity, distinguishing between those sentences that just happen to be true and those sentences that must be true. We will see how modal logic was reinterpreted to develop a logic for ethical claims.

Finally, we'll explore three-valued logic and fuzzy logic. The basis for all of the logical considerations is that all sentences must have one of two truth values: true and false. But what if we include a third truth value—or even a continuum, ranging from definitely true to definitely false, possibly taking any value in between?

In this course, we'll see that logic is a tool for helping us think in a clearer, more rigorous fashion. Logic allows us to distinguish between proper and improper forms of reasoning but never leads us to conclusions on its own. The content of our arguments come from our experiences, beliefs, and hypotheses. Logic does not control life, but it is an important part of a well-lived life.

Why Study Logic?

Since the times of classical Greece, the definition of “human being” has hinged on the notion of people being rational. Our brain is what sets us apart from the rest of the animal kingdom, and the nature of that brain is reason. The claim has been that we are built for logic. But contemporary psychology and sociology have shown otherwise. We have built in cognitive biases, and we have internalized social facts from our cultural background, and these both often lead us to unwarranted conclusions. We cannot rely on our natural instincts, but must train ourselves in the way of rational thought.

INNATE HUMAN IRRATIONALITY

- ▶ We have arranged our entire civil society around the idea that we are logical beings naturally capable of rational thought. But the world is complex. In the same way that there are optical illusions, images that will trick our senses into thinking they perceive something they do not, there are cognitive illusions—that is, ways of thinking that convince us that we are right about something we believe when we are, in fact, wrong.
- ▶ Contrary to our beliefs, we are not wired to be completely rational. Indeed, there are multiple factors that lead us to be irrational. Some of these are psychological, some are sociological, and some are cultural.
- ▶ But understanding that they are out there is the first step to realizing that we need to think carefully about thinking carefully and to teach ourselves how to think in ways that are the most likely to avoid errors.

- ▶ One of the more well-known examples of being led to irrational belief comes from psychologist Solomon Asch's experiments in the early 1950s on conformity. The subject of the experiment would sit at a table with several other people, who he or she was told were also participating in the experiment but unbeknownst to the subject were really confederates working with the experimenter.
- ▶ It was explained that this was a study in perception and that they would be shown a series of charts with line segments of different length (labeled A, B, and C) and that they just needed to say which one was the same length as a fourth line.



- ▶ The table was arranged so that the subject would answer last, after the others said aloud which line they thought was of the same size as the comparison line. The first chart was shown, and all of the confederates answered correctly, and so did the test subject. The same thing happened with the next chart.
- ▶ Then, on the third try, the confederates all gave what was obviously the wrong answer. C was the same size as the comparison line, but they all responded that A was. When it got to the test subject's turn, 12 out of 18 times in the original study (and it has been reconfirmed many times since), the test subject gave the wrong answer so as not to stick out.
- ▶ Not conforming with the group is something we fear, and we do things that we know are wrong in order to get along with the majority. What Asch shows is interesting: that we will act in such a way as to do something contrary to our reasonable beliefs.
- ▶ Interviewed after the experiment, most subjects said that they knew the answer was wrong, but they didn't want to embarrass themselves or mess up the researcher's data by being different.

More interestingly, others said that because the task was so easy and because the others answered so quickly and confidently, they started to doubt themselves.

- ▶ Asch's study was about what people would do—whether they would be willing to act in a way that they knew was wrong. But what came out was that, in some cases, it was not just about their actions, but also about their thoughts, about the way in which humans will change what we believe in order to fit in with those in our environment.
- ▶ This is what psychologist Irving Janis called groupthink, which occurs when we sacrifice critical capacities of thought in order to adopt a consensus belief. All other things being equal, we would not have come to hold this belief, but because we are in a group that believes it, we do, too. It is not just that we act as if we believe it—we do believe it.
- ▶ Janis, Asch, and a few generations of researchers following them have documented all kinds of factors that amplify or diminish groupthink. If those around you are unified in expressing the belief and do so passionately, you are more likely to get swept up in it.
- ▶ But if even one person expresses skepticism, you are more likely to be skeptical yourself.
- ▶ If there is a strong, charismatic leader espousing the view, people are more likely to come to believe it.
- ▶ If there are rewards for believing and punishments for disbelieving, people will respond.
- ▶ Our beliefs are affected by the beliefs, actions, and personalities of those around us, even if they violate what we otherwise know to be good reasoning.

- ▶ Documenting this has become one task of social psychology, and these researchers have documented many ways that we naturally reason badly, called cognitive biases.
- ▶ Many cognitive biases have been discovered and replicated in multiple experiments. Some involve behaviors—effects such as hyperbolic discounting, where we will irrationally choose less of a reward now instead of delaying gratification for more of a reward later, even if the greater reward is certain to be granted.
- ▶ There is also what has come to be known as irrational escalation effect, where we will judge an action less risky than we know it is because we have already invested in it.
- ▶ Humans are not purely rational beings. Our ultimate interest in cognitive biases involves the ways in which we are innately wired to acquire false beliefs.

RATIONAL ACTIONS AND THOUGHTS

- ▶ We know that we *act* contrary to what we know to be rational, but just as interesting—or even more so—are the ways in which we are regularly wrong about what *is* rational.
- ▶ Researchers Rolf Reber and Norbert Schwarz have shown that the font a sentence is written in will have a measurable effect on how likely we are to think it is true. The clearer the words are visually, the more stock we put in their truth.
- ▶ Psychologists define the halo effect as follows: If someone is successful or has a virtue in one way—for example, the person is successful in business—then we will take that as evidence that the person is an authority in whatever he or she does or says. Furthermore, if someone we think is attractive tells us something and someone unattractive tells us the same thing, we are much

more likely to believe it when told to us by the good-looking person, even with the same or less evidence.

- ▶ We have a positive outcome bias—that is, we think that we are more likely to succeed at tasks, win at games, or have good things happen to us than is actually the case. We imagine ourselves after the positive outcome, and because we can envision it so clearly, we come to think that it is more probable than it is.
- ▶ We overestimate our abilities. The overconfidence effect is the irrational degree of belief we put in our own choices, even when we know they are random. Interestingly, this effect is amplified when the person is the most ignorant. We are most overconfident about our abilities to do that which we know the least about.
- ▶ The Dunning-Kruger effect is the name given to the fact that the less people know about an area or how to do something, the more likely they are to overestimate their ability to do it or understand it. The more ignorant we are, the more brazen we are in our belief about our abilities and knowledge.
- ▶ In the early 1980s, psychologist Benjamin Libet ran a series of experiments that showed that the parts of our brain that determine our actions are sometimes triggered before the parts of our brain that make decisions. We act first, and then we use our brain to concoct a story that justifies our belief that we acted because we consciously and rationally decided to, even if that justification has nothing to do with the real reason for the action.
- ▶ We find this kind of “backfilling a justification” behavior on different cognitive levels. We suffer greatly from what social psychologists call confirmation bias—that is, when we hold a particular belief, we search out that which we believe supports the belief and explain away or outright ignore that which undermines rational belief in the proposition.

- ▶ When faced with a small amount of supporting evidence and a huge amount of disconfirming or falsifying evidence, we focus on that which backs us up and use it to swamp the overwhelming evidence against us—especially when the belief is core to our worldview. We see this often in the world of politics.
- ▶ We will do whatever we need to in order to save our preexisting beliefs. Where do these come from? That is the purview of sociology. We are acculturated into a belief system, and the worldview it gives us—the basic categories and presuppositions that come with it—is notoriously difficult to change.
- ▶ Emile Durkheim, one of the founding fathers of sociology, discussed what he called “social facts,” which we acquire from being part of a society. Social facts are ways of thinking or acting that originate outside the individual, are enforced by the society, and become a part of the individual.
- ▶ These are beliefs that become invisible to us because they are the lens through which we see the world, and questioning their truth strikes us as absurd, if not dangerous.
- ▶ Where are our culturally generated blind spots? What do we believe is necessarily true that is, in fact, completely false? We don’t know, but figuring it out is the job of philosophers and cultural critics, and among the most important tools they employ is logic.

LOGIC AND REASONING

- ▶ We are not rational animals by nature, but that does not mean that we cannot be or should not be. Cognitive biases from our psychology and social facts from our cultural upbringing and political context might give us false beliefs, but we are capable of excavating and rigorously analyzing them.

- ▶ We can question what we believe, even our most deeply held beliefs, and objectively determine whether there are legitimate grounds for rational belief. For this, we need logic.
- ▶ It might turn out that there is good reason to believe what we believe. In this case, our belief will be strengthened. If it turns out that there is good reason to disbelieve it, then we will have saved ourselves from having a false belief, which could not only affect other beliefs, but also cause us to act in ways that have negative consequences for us and those around us.
- ▶ When we are aware of our psychological and cultural biases, we are able to step away from them and begin to think rationally. But it does not happen automatically. We might be aware of some cognitive biases and still fall into them.
- ▶ Logical thinking is not an innate talent, but it is a skill that one can develop. It takes training and practice. You must see what makes some forms of reasoning effective at determining what is likely to be true and what makes other forms of reasoning likely to lead you into error, even though it seems attractive to us.

READINGS

Durkheim, *The Rules of Sociological Method*.

Fine, *A Mind of Its Own*.

Sutherland, *Irrationality*.

QUESTIONS

1.

If the human brain is the result of evolution, which selects for properties that have an advantage in survival, then why do we have these cognitive biases built into us?

2.

How do we best go about finding the cultural blind spots that we have? If there are assumptions we have been given our whole lives, how do we find them?

Introduction to Logical Concepts

The word “logic” means rational argumentation. A belief is rational if we have good reason to believe it, and an argument is a set of sentences such that one sentence, the conclusion, is claimed to follow from the other sentences, the premises. The first skill we have to master is finding the parts of an argument—that is, recognizing what is the conclusion and what are the premises. Then, we need to figure out what type of argument we have: deductive or inductive. Finally, we can begin to analyze the argument. For this, there are two criteria that have to be satisfied: validity and well-groundedness.

THE DEFINITION OF “LOGIC”

- ▶ For the purposes of this course, logic is the study of rational argumentation. By “rational,” we mean that which we have grounds to show is likely true. A belief is rational if we have good reason to believe it is at least probably the case in reality.
- ▶ The other word, “argumentation,” is the important one. The central term in all of logic is “argument.” For logicians, an argument is a set of sentences such that one sentence, the conclusion, is claimed to follow from the other sentences, the premises.
- ▶ Arguments have two parts: a conclusion and premises. The conclusion is the point of the argument. It is the thing being argued for—that which we are trying to convince ourselves or others of. We give arguments in order to provide legitimate reasons to believe the conclusion. The premises are those

reasons. The premises are the grounds that are being proposed to support rational belief in the conclusion.

- ▶ Every argument has one and only one conclusion. Premises, on the other hand, can come in any number. There are mathematical arguments with an infinite number of premises and weird logical arguments that have no premises.
- ▶ Consider an example: All men are mortal. Socrates is a man; therefore, Socrates is mortal. In this case, the conclusion is “Socrates is mortal.” That is what the argument is trying to convince us of. Why should we believe that Socrates is mortal? Because he is a man, and all men are mortal.
- ▶ The first thing you do when you approach an argument is to find the conclusion and then set out the premises. It is crucially important that you do this correctly.
- ▶ Consider what would happen if we misidentified the conclusion in this case. If it is true that all men are mortal and that Socrates is one of these men, then it turns out that it is absolutely the case that Socrates is mortal. But suppose we wrongly thought that the conclusion is “all men are mortal” and that the premises are “Socrates is a man” and “Socrates is mortal.” Just because one man is mortal, it doesn’t necessary follow that they all are.
- ▶ By misidentifying the conclusion and premises, we have taken a good argument, one that gives us good reason to believe something, and turned it into a flawed argument that does not.
- ▶ So, the first tasks that are necessary for us to develop are figuring out when we have an argument and determining what the conclusion is and what the premises are. We often have help with these tasks—called indicator words. There are certain words we use to point out conclusions, and there are certain words we use to point out premises.

- ▶ A conclusion is supposed to be the thing that is established by the argument, so we use words that indicate this. The most obvious one is “therefore.” My heart is beating; therefore, I am alive. We use other words for this function as well, including “thus,” “hence,” “so,” and the Latin “ergo.” But be careful: Not every use of these words indicates a conclusion.
- ▶ We also use words and phrases that indicate premises, such as “because,” “since,” and “given that.” Again, not every use of these words is an indication of a premise in an argument.
- ▶ Indicator words are the easiest way to determine whether we have an argument and, if so, what the conclusion is and what the premises are. But we don’t always have indicator words. How, then, do we determine if we are looking at an argument?
- ▶ The easiest way is to try to insert your own indicator words—for example, “therefore” or “because.” If you look at a passage and the word “therefore” can be naturally inserted in a way that maintains the meaning of the passage, you are probably looking at an argument, and what immediately follows “therefore” is your conclusion.
- ▶ Similarly, if you can insert “because” into a passage without changing the meaning, you are likely looking at an argument, and what comes right after your inserted “because” is probably a premise.

TYPES OF ARGUMENTS

- ▶ The first two steps are to figure out if a passage contains an argument and then to identify the structure—that is, to pick out its conclusion and fully lay out its premises (both stated and implied). The third step is to determine what kind of argument it is.

- ▶ Arguments come in two kinds: deductive and inductive. An argument is deductive if its conclusion is no broader than its premises—that is, if the conclusion only refers to that which is mentioned in the premises. The technical term we use is to say that deductive inferences are “non-ampliative”—that is, the conclusion is contained within the content of the premises.
- ▶ Reconsider the following argument: “All men are mortal. Socrates is a man. Therefore, Socrates is mortal.” In the premises, we have all men being mortal, but then move in the conclusion to just one of these men being mortal. Deduction moves from broad to narrow.
- ▶ Inductive arguments are ampliative—that is, they do have a conclusion that is broader than the premises. For example, the Scottish philosopher David Hume famously considers the following argument: “The sun has risen every other morning, so the sun will rise tomorrow as well.”
- ▶ Notice that the premises contain a lot of data: that the sun has risen every day in the past. But the conclusion is not about one of those days; it is about a day outside of the data set. We are using facts about millions of other days to expand our reasonable belief to one new one.



DAVID HUME
(1711–1776)

- ▶ In the case of the Socrates being mortal argument, we had a claim about all men and Socrates was one of them, but here we have a collection of data and the day in the conclusion is not one of them. It is a further instance not contained in the content of the premises.
- ▶ With deduction, we are arguing from broad to narrow, but with induction, we are arguing from narrow to broad.
- ▶ The important result of this difference is in the degree of belief we can reasonably put in the conclusions of these different types of arguments. Because the content of a deductive argument is already contained in the premises, then in a successful deductive argument, the conclusion will be absolutely certain. If all men are indeed mortal and Socrates is one of those men, then there is no other alternative but that Socrates is mortal.
- ▶ But with an inductive argument, because the conclusion outruns the premises, we have no guarantees. Just because the sun has always risen, there is no guarantee that it will do so again. It probably will, but not definitely. Successful inductive arguments, because they are ampliative, only give us high probability, not absolute certainty.
- ▶ With an inductive argument, because the conclusion lies outside the scope of the premises, there is a risk that even a very good inductive argument might have a false conclusion. This is not to say that we should not believe it. We should believe that which is probably true. It is wonderful to achieve the absolute certainty we get from deduction, but in most real-life cases, we are restricted to the high probability of induction.
- ▶ It will turn out that when we evaluate arguments—when we determine whether an argument is successful or not, whether it has flaws or not—we will need different logical tools for deductive and inductive arguments.

CRITERIA FOR ACCEPTABLE ARGUMENTS

- ▶ Identifying the parts of arguments and the types of arguments are the initial steps, but the whole point is to determine which ones are good arguments. Which ones give you good reason to believe in the truth of the conclusion?
- ▶ To evaluate arguments, we use two criteria: validity and well-groundedness. An argument is valid if and only if, assuming the truth of the premises for the sake of argument, the conclusion follows from them.
- ▶ The important thing to notice here is that we are assuming the premises are true for the sake of argument. Maybe they are true; maybe they are false. We don't care. Validity does not concern the content of the premises.
- ▶ All we are looking at is whether the premises, *if* true, would lead you to the conclusion. Validity is not about the content of the argument, but about the form of the argument. Validity looks at the skeleton of the argument and determines if it is strong enough to support the weight of the conclusion.
- ▶ That we are assuming the truth of the premises in the first criterion—validity—should bother you a little bit. After all, that is a huge assumption to make. What justifies our ability to make such an assumption is the second criterion: well-groundedness.
- ▶ An argument is well-grounded if and only if all of its premises are true. Well-grounded arguments have true premises. Maybe the conclusion is true, or maybe it is false, but what is important for us in looking at the well-groundedness of an argument is just the truth or falsity of the premises.

- ▶ An argument that satisfies both of the criteria—an argument that is both valid and well-grounded—is considered to be sound. A sound argument gives us good reason to believe its conclusion.
- ▶ In order to determine which arguments are sound, we need to develop tests for validity and well-groundedness. Validity looks at the structural elements of the argument, and its study is called formal logic, because it is an examination of the form of arguments. Validity for deductive and inductive arguments are completely different matters, and we need different tools.
- ▶ Well-groundedness concerns look at the acceptability of the argument other than the form and are called informal logic. We will begin there as we work to build a complete account of the ways we determine what arguments give us good reason to believe their conclusions.

READINGS

Barker, *The Elements of Logic*, chap. 1.

Copi, *Introduction to Logic*, chap. 1.

Kahane, *Logic and Philosophy*, chap. 1.

QUESTIONS

1.

Which of the following passages contains an argument?

- a Groundhogs burrow in the winter to protect themselves from the cold. They acquire a layer of fat through the summer and fall months, which allows them to stay warm and nourish themselves as they hibernate. They flourish in areas where food is plentiful, making it easier for them to survive the winter.

- b The groundhog population here is likely to die off in the near future because residential construction is eliminating easy access to food. Without the food, they will not have sufficient fat for them to make it through the winter, and the population will decrease until it can no longer sustain itself.

2.

Find the conclusion and premises in the following arguments.

- a You should buy the most expensive model of this computer. I know that it is a lot of money, but the speed and increased functionality are worth it. If you skimp on the model, then it will not be that much better than your current computer, and that would make it an even bigger waste of money.
- b I don't like Indian food. You can't eat the gluten that is in most Italian food. The Persian restaurant is very expensive. The only other restaurant on this block is the Irish pub, which we both like and is affordable. Let's eat there.

3.

Are the following arguments deductive or inductive?

- a If you don't water your plants, they will die. If we don't drink water, we will die. If you take a fish out of water, it will die. It seems that it is true of everything living that it needs water.
- b That recipe calls for a tablespoon of sugar. It says that it serves 12. There's only four of us, so if I scale back the recipe, I'll only use a third of a tablespoon of sugar. That's just a teaspoon.

ANSWERS

1.

- a This is not an argument; it's simply a collection of facts about groundhogs.
- b This is an argument, because it is providing reasons why one should believe that the groundhog population in a specific region will die off.

2.

- a Premise: The most expensive model of computer has increased speed and functionality, which make it worth the extra money.

Premise: The less expensive model is not much better than your current computer, so buying it would be a waste of money.

Conclusion: You should buy the most expensive model of computer.

- b Premise: I don't like Indian food.

Premise: You cannot eat at the Italian restaurant.

Premise: The Persian restaurant is very expensive.

Premise: The Irish pub is the only other restaurant on the block and is affordable.

Conclusion: We should eat at the Irish pub.

3.

- a This is inductive, because it generalizes from plants, humans, and fish to all living things.

- b This is deductive, because four is one-third of 12, and one-third of a tablespoon is a teaspoon, so it necessarily follows that the scaled-back version of the recipe will call for a single teaspoon of sugar.

Informal Logic and Fallacies

An argument is well-grounded if all of its premises are true. In assessing the well-groundedness of arguments, we need to be on the lookout for a number of standard reasoning errors, called fallacies. One category of fallacy concerns what logicians call begging the question. These are unwarranted assumptions that depend on accepting the conclusion, which is a problem because the point of the argument is to provide independent warrant for belief in the conclusion. We can commit this error by creating various fallacies, such as circular argument, begging the question, equivocation, and distinction without a difference.

FALLACIES AS REASONING ERRORS

- ▶ There are common reasoning errors that create unsound arguments but that are nonetheless attractive to our minds. They sound like good arguments, but in fact they are unsound arguments. These reasoning errors are called fallacies. A fallacy is an identifiable category of argument that does not support its conclusion.
- ▶ One of the most well-known fallacies is circular argument. A circular argument is one in which the conclusion is identical to the premise. Consider the following argument: “My name is Steve; therefore, my name is Steve.” Is it valid? If we assume the truth of the premise, are we led to also accept the truth of the conclusion?
- ▶ Yes. It is a perfectly valid argument. But is it a good argument? Does the premise give us independent warrant for rational belief in the conclusion? No.

- ▶ The point of an argument is to provide independent support for the conclusion because the conclusion is in doubt. But if the conclusion is in doubt, then so is the premise, because they are one and the same proposition.
- ▶ If we do not have independent reason to think that the premise is true, then the well-groundedness of the argument is in doubt, and the argument cannot be said to be sound, and we do not have good reason to believe the conclusion. It is a bad argument logically.
- ▶ Unfortunately, it is an effective argument psychologically. It works well, if done right. If we say something clearly, slowly, loudly, confidently, or forcefully enough, people will believe it. Repeat it and it seems more likely to be the case. Circular argument is a fallacy, but it is effective.
- ▶ Does the fact that it is logically flawed mean we have reason to reject the conclusion of a circular argument? If we have a valid, well-grounded argument, then we have good reason to accept the likely truth of the conclusion. But if we find a flaw, a fallacy committed in the argument—if the argument is invalid or not well-grounded—do we have reason to think it is likely false? No.
- ▶ Just because someone makes a bad argument for a conclusion does not give us rational justification for thinking that the conclusion is not the case. Remember that an argument provides us good reason to believe in the likely truth of the conclusion. If an argument fails, then it means that those specific premises do not give us reason to believe that conclusion.
- ▶ But it doesn't mean that there can be no other set of premises that does. There could be a good argument for that conclusion, and this one just isn't it. Every true proposition can be made the conclusion of a terrible argument.

- ▶ If we find a fallacy—if an argument we are analyzing turns out to be flawed—what do we know about the truth or falsity of the conclusion? Nothing. Might it be true? Yes, for other reasons. Might it be false? Maybe. So, what do we know about it? Nothing. Are we justified in rejecting it? No. We have to suspend belief. Keep in mind that a lack of presentation of a good argument is not a refutation.

CIRCULAR ARGUMENT

- ▶ An argument is circular when the premise and the conclusion are the same proposition. Notice the use of the term “proposition” here and not the term “sentence.” In philosophy, a sentence is a grammatically correct string of words. A proposition is the content of the sentence.
- ▶ Different sentences can be used to express the same proposition. The easier-to-spot version of circular argument is where the premise and the conclusion are the same sentence, as in the example, “My name is Steve because my name is Steve.”
- ▶ But more often, the premise and conclusion will be different sentences that express the same proposition, thereby camouflaging the fact that the argument is circular. For example, “My name is Steve because Steve is what I am called.”
- ▶ While this example might seem contrived, we find examples often in real life. One place you will frequently come across the easier-to-identify version is when you question a deeply held belief. We will often feel the need to justify our more peripheral beliefs, but certain things we take as foundational and either beyond the need for evidence or so central to the way we see the world that we never thought to question them.

- ▶ Why are people of your class deserving of benefits not given to those beneath? Because we are. Why is gold more valuable than silver? Because it is. These circular arguments will often be given with a sense of disbelief that anyone could even think to request justification for the belief.
- ▶ Frequently, the premise and the conclusion will be different sentences expressing the same proposition. Because they sound so different to the ear, the fact that they say the same thing is overlooked, and the argument is wrongly thought to be sound.
- ▶ Consider a naïve version of a standard ethical argument for vegetarianism that we call the argument from sentience: “It is wrong to kill animals because it is wrong to kill anything that can feel pain.”
- ▶ What makes this circular is that the only things that can feel pain are animals. To feel pain, you need a central nervous system. But anything that has a central nervous system would be an animal. So, the argument really just says that it is wrong to kill animals because it is wrong to kill animals. It restates the conclusion as the premise in a fashion that is not obvious at first glance.

BEGGING THE QUESTION

- ▶ Circular arguments are one instance of a larger class of fallacies called begging the question. This is a phrase that is often misused. One will frequently hear someone say, “That begs the question that...” when they mean, “That leads one to ask the question...”
- ▶ But what begging the question really means is arguing unfairly in a way that tries to use the conclusion in support of itself. Anytime we try to get the conclusion to pull itself up by its own logical bootstraps instead of giving independent reasons for belief, getting its support from propositions distinct from itself, we are looking at begging the question.

- ▶ What is a question? Is it anything of a certain grammatical form where our voice goes up at the end? No. There are non-questions that look like questions and sound like questions but are not questions. A question is a request for information. A pseudo-question, on the other hand, seems like a request for information but really isn't. For example, "You're not going wear that, are you?"
- ▶ A so-called leading question is a question that is not fairly asked to elicit an honest response from the listener, but rather a sentence that looks like a question but is designed to lead the listener to a particular desired response. "Does this make my butt look big?" is a famous leading question that one might receive from a significant other. It is not a question; there is only one correct answer.
- ▶ So, one way to beg the question is to use questions that are not questions. Another way is to use the connotative power of language to frame questions unfairly. Words have both denotation—that is, they pick out certain things—and connotation—that is, they lend an emotional weight to those things.
- ▶ Companies vying for our business do this all the time: Looking to relocate to another part of the country, but you don't have a lot of money? Try Budget Movers. Are they cheaper than other companies? We have no evidence, but the name implies it. The words are chosen because of their ability to sway the customer.

EQUIVOCATION

- ▶ The basis of the fallacy known as equivocation is ambiguity. Words can mean more than one thing. This is not a logical problem, just a feature of language. The fallacy of equivocation occurs when we change the meaning of a word in the middle of an argument. Consider the following argument.

- ◆ Tables are furniture.
 - ◆ My statistics book has tables in it.
 - ◆ Therefore, there is furniture in my statistics book.
- ▶ It certainly looks valid, doesn't it? And all of the premises are true. So, it must be sound, right? We have good reason to believe the conclusion—except, of course, that the conclusion is absurd. What's the flaw?
- ▶ The meaning of the word “table” has changed from a raised flat surface on which to place things to a rectangular array of numbers. There is nothing inherently wrong with the ambiguity as long as the meaning of operative ambiguous terms is maintained throughout the argument, but in this case, the meaning has changed, and the result is the absurdity of thinking that there is furniture inside of a book.
- ▶ Unfortunately, not all examples of equivocation are this clear-cut. The usual cases are trickier because many ambiguous terms will have distinct meanings that are related. As such, the shift from one to the other is subtler and, as a result, easier to overlook. Consider the following example.
- ◆ We have a right to vote.
 - ◆ One should always do what is right.
 - ◆ Therefore, one should always vote.
- ▶ In this case, the word “right” is being equivocated upon. In one case, we have one sense of “right,” meaning a legally protected action, and in the second premise, the same word means an act that is morally necessary. These are two different meanings, but ones that can be confused if you are not careful.

DISTINCTION WITHOUT A DIFFERENCE

- ▶ A distinction is a linguistic separation of two concepts that are different. Distinction without a difference is a fallacy that occurs when we try to draw a distinction between two things that are not, in fact, distinct. “I didn’t steal it; I just didn’t ask before I borrowed it.”
- ▶ One place we will often find people trying to draw distinctions without a difference is when they are trying to justify holding views they know they should not.
- ▶ We will often try to justify our misdeeds or our problematic views by trying to distance them from the categories to which we know they really belong. “I’m not racist; I just think that members of that particular minority are less intelligent.”
- ▶ Notice the difference between distinction without a difference and circular argument: In the case of distinction without a difference, the form is “It’s not *A*; it’s *A*,” whereas in a circular argument, we are saying, “*A* because *A*.”
- ▶ In distinction without a difference, we are trying to say that *A* is both *A* and not *A*, whereas in a circular argument, *A* certainly is *A*, and we should believe *A* because of itself. Both repeat *A*, but in making very different logical mistakes.

READINGS

Copi, *Introduction to Logic*, chap. 3.

Damer, *Attacking Faulty Reasoning*, chap. 5.

Kahane, *Logic and Contemporary Rhetoric*, chap. 3.

QUESTIONS

Identify the fallacies in the following passages from the following list: circular argument, question-begging language, equivocation, distinction without a difference.

1.

That dish can't be too spicy; after all, she said she made it with chili peppers. Chilly is the opposite of hot, and she didn't say she used hot peppers.

2.

Do you want to go to the same old boring beach this year for vacation, or should we be adventurous and try something new, such as the mountains?

3.

It's not that I don't care about you; I'm just not concerned with your life choices.

4.

Fried foods are bad for you because they are not part of a healthy diet.

ANSWERS

1.

equivocation

2.

question-begging language

3.

distinction without a difference

4.

circular argument

Fallacies of Faulty Authority

No one knows everything. In many situations, we have to rely on others to learn what we need, or we have to infer something from other things we know. These are both fine ways of acquiring rational beliefs, if done properly. In this lecture, you will continue learning about ways that arguments can go wrong by examining five new reasoning errors that are associated with flawed appeals to authority and flawed inferences.

APPEAL TO AUTHORITY

- ▶ When we want to know something we don't know, it is perfectly rational to ask someone who does. We call this an appeal to authority, and arguing from authority is a legitimate means of reasoning. We do it all the time, and we should. When you are sick, hopefully it is your physician you seek out, because he or she is the expert you need.
- ▶ What makes someone a legitimate authority, someone whose answer to a question we would have good reason to believe? There are three requirements. First, the purported authority must have actual physical presence in the material world. The authority must exist.
- ▶ Surprisingly, we hear arguments from authority that violate this simple criterion all the time. Someone who says, "I heard somewhere that..." or "I read somewhere that..." is making an argument from authority, but doing so in a way that does not tell you the name of the authority. Whenever someone cites an authority, you have every right to demand to know the name of the authority.

- ▶ The second step in a successful argument from authority is to make sure that the person cited as an authority is, in fact, an expert—that is, someone we have good reason to believe would know the correct answer to the relevant question.
- ▶ What makes someone an expert? It depends on the case. On one hand, if the question were about basic general science, for example, then someone who has a basic education in the area would be enough to qualify. On the other hand, if it were an intricate question about a cutting-edge aspect of science, then an expert would be someone with more education who is involved in that particular subfield.
- ▶ The third criterion in considering someone a legitimate authority is that the person must not only be in a position to know the answer to the question, but the person must also not have a stake in your believing one way or the other. The expert must be objective or disinterested.
- ▶ If you are buying a used car from Joe, someone you know to be a mechanic, should you take his word on the condition of the engine? Joe is the sort of person who would know, but he also stands to profit if he convinces you that the car is in better shape than it is. You should not consider Joe to be an authority in this case, even though he is indeed an expert. You should take the car to another mechanic to give you an objective appraisal.
- ▶ If someone meets all three of these requirements, that person can be considered an authority, and his or her word can be reasonably taken. If someone fails to meet one or more of these, then the use of that person's word in an argument from authority would be a reasoning error.

APPEAL TO COMMON OPINION

- ▶ A related fallacy is called appeal to common opinion, also known by the Latin, *ad populum*. Simply because many people believe something does not make it true. It was widely believed that the world was flat and that slavery was morally acceptable.
- ▶ At some level, we all know that simply because something is a common belief, that is not sufficient for rational belief, but what we are dealing with here is not mere intellectual laziness but a deeper cognitive bias: groupthink.
- ▶ We like to think of ourselves as independent minded, but we are deeply influenced by the views of those around us. We don't like to stick out, and when we believe something the majority of others do not, it will often lead to doubt and insecurity. When we surround ourselves with others of like mind, it is comforting to us, leading us to be surer of our views than is rationally warranted.
- ▶ In a deep way, we are programmed to commit the fallacy of appeal to common opinion. To be rational, we need to learn to keep this proclivity in check. This does not mean that it is always rational to reject common opinion; sometimes appealing to common opinion is a good inference.
- ▶ Suppose that you are attending a performance at a venue where you have never been before, and after the show is over, you need to find your way to your car, which is in the garage where many of those in the theater have also parked. If you just follow the crowd, you will probably end up exactly where you need to be. In this case, there is nothing wrong with believing that what everyone else believes is likely true. But this is not the usual case.

APPEAL TO TRADITION

- ▶ Widely held beliefs might be true, or they might be false. We need independent reason to believe them rationally, and simply appealing to the fact that everyone else thinks so is not enough. But there might be reasons given why everyone else thinks so.
- ▶ Some of these might be good, but others are not, despite the fact that they are often cited. One of these problematic justifications is called an appeal to tradition: “But we’ve always done it that way....” Some traditions are good; some traditions are not.
- ▶ If a belief, or belief system, or way of doing things has been honed and crafted over generations, there might be good reason to keep doing something in a certain way.
- ▶ We learn from mistakes, both our own and those of others. And it is certainly true that being an apprentice to someone with much more experience can be a wonderful way of learning—that is the argument from authority. But it requires evidence that the way we have been doing it is the right way, or the thing we have always believed is, in fact, true.
- ▶ Intellectual inertia is not rational justification. Just because it has always been believed or always been done a certain way does not make that belief or that method justified.
- ▶ An interesting version we find uses an appeal to other people’s traditions. “It is an ancient Chinese cure” is something one hears as an attempt to justify approaches to alternative medicine. Are ancient Chinese medical practices effective? This may or may not be true, and we need independent evidence. Merely appealing to tradition is not sufficient.

THE FALLACY OF NOVELTY

- ▶ In the same way that looking backward in time does not give us legitimate warrant for rational belief, neither does looking forward. The converse of the fallacy of appeal to tradition is the fallacy of novelty. Just because something is the latest does not necessarily make it the greatest. There is a reason we say “new and improved”; just because it is new does not entail that it is improved.
- ▶ We regularly see this fallacy in advertising. Manufacturers of products have a problem: They want you to be satisfied with the product you have bought from them so that you will be loyal to their brand, but they also want you to be dissatisfied so that you will look to buy a replacement for it.
- ▶ One way of making you want something you have already bought is to convince you that while what you have is good, there is a new one that is even better. If it is better—if it has capabilities the old one does not that would be helpful, or if its new design makes it more efficient or easier to use—then these might be legitimate reasons to decide to purchase a new one.
- ▶ But manufacturers try fallacious versions of this approach as well. Just think of the terms that are often used for new versions: “upgrades” or “updates.” They’re explicitly saying that new is better just because it is new. This may or may not be true and thus requires evidence, not just assertions of its status as the most recent addition.

ARGUING BY ANALOGY

- ▶ Arguing by analogy is like faulty authority in that it’s a flawed version of a good form of reasoning. It is a perfectly fine way to arrive at reasonable beliefs.

- ▶ We use arguments from analogy in science all the time. We use computer models to predict the weather. Over the past few decades, these computer analogues of the actual weather systems have become increasingly accurate.
- ▶ What makes arguments from analogy successful is that the systems selected to be analogues do share certain structural similarities to the system being modeled. The fallacy of faulty analogy occurs when we argue by analogy using a flawed analogy, where the system and the analogue are not alike in the ways used to draw the inference.
- ▶ Successful analogies involve structural similarities between two systems. An analogy can go wrong if that structural similarity is missing.
- ▶ Some political candidates cite business experience as a central reason to suppose them competent for public office. In both cases, they would hold positions of authority in an organization. But the goal of running a business is to defeat competition in a marketplace and create monetary profits for only your shareholders. By contrast, the goal in politics is to create laws in a way that brings more than monetary benefit to all of society.
- ▶ These are radically different tasks, and one cannot infer competence in one from success in the other. The analogy between the two does not hold because the structures of the two do not share the requisite commonalities.
- ▶ But even if there is a commonality, that is not sufficient. For example, someone might argue that having a drink of alcohol is like shooting up heroin: Both alter the brain, both are addictive, and lives have been shattered by both. While that is certainly true, the effects of a single glass of wine are different from heroin use. The analogy fails despite the similarity because of a radical difference in degree.

READINGS

Copi, *Introduction to Logic*, chap. 3.

Damer, *Attacking Faulty Reasoning*, chap. 6.

Kahane, *Logic and Contemporary Rhetoric*, chap. 5.

QUESTIONS

Identify the fallacies in the following passages from the following list: faulty authority, appeal to common opinion, faulty analogy, fallacy of novelty.

1.

My GPS says that we should take route 70 to route 97. You think we should take route 85 to Buckminster Road after looking at a map. I'm thinking that the GPS directions are better because the GPS is equipped with the latest route-finding algorithm.

2.

Look at the line for that new movie! It is out of the theater door, down the street, and around the corner. That film must be great.

3.

I was reading the expert reviews for this product on its website. All of them were outstanding. I think that this could be the product we've been waiting for.

4.

She must wear dentures. He said that her teeth were like the stars, and we know that the stars come out at night.

ANSWERS

1.

fallacy of novelty

2.

appeal to common opinion

3.

faulty authority

4.

faulty analogy

Fallacies of Cause and Effect

Much of the reasoning we do involves the relation of cause and effect. What will happen if you take a new medication while also taking your pill for high blood pressure? But cause-and-effect reasoning, called causal reasoning, can be quite difficult. As the old dictum goes, correlation does not imply causation. Just because two events are frequently seen together does not mean that one necessarily causes the other. As you will learn in this lecture, there are several regularly recurring mistakes that people make when asserting cause and effect.

CAUSE-AND-EFFECT REASONING

- ▶ Some of the most interesting and important claims we make are about cause and effect. So much of what we want to know about the world, about each other, and about ourselves involves the kinds of “why” questions that assert cause-and-effect relations. Unfortunately, they are also among the most difficult to establish logically.
- ▶ Most people have heard the old logical dictum that correlation does not entail causation—that is, that simply because you can find two things occurring together, it does not mean that we can assert with any certainty that one caused the other.
- ▶ It may be the case, but simply identifying a correlation does not itself establish in any way that there is a cause-and-effect relation. We need to understand the mechanism by which event *A* brings about event *B* if we want to be justified in asserting that *A* causes *B*.

- ▶ We will look at different ways that this kind of erroneous inference can be made—that is, different ways in which arguments that attempt to establish that event *A* caused event *B* miss their mark.
- ▶ It is certainly true that for event *A* to cause event *B*, *A* must precede *B* in time. Causes come before their effects. First, we observe the cause and then the effect. But just because we saw event *A* before event *B* does not itself give us good reason to infer that *A* caused *B*.

THE POST HOC FALLACY

- ▶ Causes come before their effects, but time order is not sufficient warrant to assert cause-and-effect relations. To do so is to commit the error known by the Latin phrase “post hoc, ergo propter hoc,” which means “after this, therefore because of it.” Just because it comes after, it does not mean that it happened because of it. Most logicians abbreviate this as the “post hoc” fallacy.
- ▶ We often make post hoc arguments around individual events when the effect is something unusual or significant. Your car, which has run beautifully for 10 years with no need of repairs, is suddenly making strange noises, and you lent it to Ben just last week. What did he do to your car?
- ▶ When something odd occurs, we try to determine why. And if we see something different in the antecedent context—if there was something else unusual beforehand—we might jump to the conclusion that whatever was different before must be the cause of what was different after. But that prior difference might be completely unrelated to the posterior difference.

- ▶ Unbeknownst to you, your timing belt was on its last legs, a part that usually only lasts 10 years. Ben's driving had nothing to do with it.
- ▶ The post hoc inferences that we cling to the strongest are those that do not correlate single events, but when we notice that most of the time when we observe events of type *A*, we also tend to observe thereafter events of type *B*. Seeing the repetition of the correlation strengthens our belief in the causation.

NEGLECT OF A COMMON CAUSE

- ▶ Correlation does not necessarily mean causation. Again, if we want to assert a cause-and-effect relation between *A* and *B*, we need more than time order. We need the mechanism. Making the inference based only on time order, only on correlation, can lead to a flawed inference in a few different ways.
- ▶ By taking correlation to imply causation, we commit the fallacy called neglect of a common cause. Just because whenever we see *A*, we also see *B* does not mean that *A* causes *B* or that *B* causes *A*; there might be a third thing, *C*, that causes both *A* and *B*.
- ▶ It might be the case that whenever you see your normally mild-mannered coworker taking an aspirin, he is uncharacteristically cranky and short-tempered. This does not mean that aspirin has a negative side effect that has psychological ramifications, nor does it mean that he takes the pills because of his change in disposition. Indeed, there might a third factor—for example, a headache or a toothache—that causes both the foul mood and the taking of the medicine.

CAUSAL OVERSIMPLIFICATION

- ▶ Real-world phenomena are complex and often brought about by a convergence of a multiplicity of factors. By misconstruing or ignoring this complexity, we can commit one of two cause-and-effect reasoning errors. The first is called causal oversimplification, which involves picking out one part of a complex causal web and ignoring the web.
- ▶ We see this fallacy committed often when we are dealing with complex social issues. The claim has been made that the reason the divorce rate in America is so high is because of the women's rights movement.
- ▶ It wasn't until feminism removed women from their traditional gender roles that the divorce rate started climbing significantly. The feminist movement caused the breakdown of American marriage.
- ▶ There is certainly no doubt that the advance of women's rights is a part of the story about the increased number and rate of divorce.
- ▶ Given the social context of most women before the advance of feminism, many women would not have had the financial means to live on their own, because of a lack of educational opportunities or well-paying jobs, and as a result, some women were forced to stay in unsatisfying or abusive marriages.
- ▶ But once a greater range of possibilities became normalized for women, such circumstances were less common, and bad marriages dissolved more frequently.
- ▶ But the complete story is more than just desertion of traditional gender roles. There are many other aspects of the changing culture that are intricately wrapped up in the story. To simply pick out one element, one causal factor, and say it is *the* cause is to commit the fallacy of causal oversimplification.

CONFUSION OF A NECESSARY WITH A SUFFICIENT CONDITION

- ▶ An error that is related to causal oversimplification is the confusion of a necessary with a sufficient condition. A condition *A* is necessary for *B* if you cannot have *B* without first having had *A*. In other words, *A* is necessary for *B* if *A* is required for even the possibility of *B*. *A* doesn't bring about *B* by itself, but if there is no *A*, there is no *B*.
- ▶ Oxygen, for example, is necessary for fire. If there is no oxygen, there can be no fire. This doesn't mean that everywhere there is oxygen there will also be fire, but take away the oxygen and you remove the chance for fire. Oxygen is necessary for fire.
- ▶ A condition *A* is sufficient for *B* if *A*, by itself, is enough to bring about *B*. Winning a high-stakes lottery, for example, is sufficient to become a millionaire. But while it is sufficient—that is, it is by itself enough to make you a millionaire—it is not necessary. There are other ways to become a millionaire. You could start a tech company that gets bought by Google, or you could be born or married into it.
- ▶ While thinking that a necessary condition is sufficient is the most common version of this fallacy, the converse can be found occasionally as well. Sometimes we take a sufficient condition and wrongly assert it to be necessary.
- ▶ This will often be the result of focusing on a particular way of doing something that has become habitual for us, and we allow ourselves to be blinded to other ways of accomplishing the same task.
- ▶ For example, you might ask your child why a dirty dish that was just used is sitting in the sink. Your child might answer, "The dishwasher is running." While putting a dirty plate in the dishwasher is sufficient for cleaning it, it is not necessary. There are other ways to clean it, such as washing it by hand.

THE SLIPPERY SLOPE FALLACY

- ▶ The final causal fallacy is perhaps the most famous: the slippery slope fallacy. There is no doubt that there are causal chains—that is, an event A causes B , but then the effect of B becomes the cause of C , which in turn causes D .
- ▶ We can have causally related chains of events. This is what leads to an alternate name one will sometimes see for the slippery slope fallacy, the domino fallacy, which refers to the tipping of dominoes. They are arranged until the first one goes, and then it hits the second, which hits the third, and so on.
- ▶ While such chains of cause-and-effect relations exist, the fallacy is where one asserts the existence of such a chain without giving full causal arguments for each step in the chain. Arguing for cause-and-effect relations is tricky.
- ▶ We need to show the underlying mechanism at work, and in the complex world of intervening causes, often A would bring about B , all other things being equal, but when rubber meets road, not everything is always equal.
- ▶ When warning people of an act we think is imprudent, we will often neglect to do all of the logical heavy lifting and simply assert a causal chain or the likelihood of it and leave it at that.
- ▶ “I wouldn’t take that first sip of beer. It always starts with beer, but then it goes to wine and then hard liquor, which paves the way for marijuana, and then addictive drugs like cocaine and heroin. That little sip might seem harmless, but it is the first step on a slippery slope to addiction, losing your house and your family—everything will be gone.”

- ▶ Of course, there are stories of people who have followed this unfortunate path to ruin. But the question for us is whether there is good reason to believe that each of these steps down the slippery slope awaits the particular person warned. If so, make the case for each step. If you cannot, then the argument fails.

READINGS

Copi, *Introduction to Logic*, chap. 3.

Damer, *Attacking Faulty Reasoning*, chap. 8.

Kahane, *Logic and Contemporary Rhetoric*, chaps. 2 and 4.

QUESTIONS

Identify the fallacies in the following passages from the following list: post hoc fallacy, neglect of a common cause, causal oversimplification, confusion of a necessary and sufficient condition.

1.

People who drive nice cars also tend to have large homes. I guess that if you have a car that nice, you don't want to park it in front of a small house.

2.

Zydeco music always has an accordion in it. This polka band has an accordion player, so I guess they play Zydeco.

3.

I had eggs for breakfast, and then I played the best round of golf in my entire life. It's eggs every Saturday for me. I want to bring down my handicap.

4.

The reason drug use is down is because they have been showing those “just say no” commercials on television. The message must really be getting through to people.

ANSWERS

1.

neglect of a common cause

2.

confusion of a necessary and sufficient condition

3.

post hoc fallacy

4.

causal oversimplification

Fallacies of Irrelevance

When we disagree with others, the process is not merely logical, but emotional. When someone thinks that we are wrong about something, we feel attacked and often feel justified in attacking in return. As you will learn in this lecture, the fallacies of irrelevance frequently occur when our emotions lead us to focus on some aspect of the disagreement other than the actual disagreement. We have to make sure that we are on guard at all times to avoid committing—and getting sidetracked by—these diversionary fallacies.

FALLACIES OF IRRELEVANCE

- ▶ One of the most difficult aspects of engaging in passionate discourse is keeping the discussion focused on the question at hand. When someone is disagreeing with us, especially if it involves a proposition we take to be important, we can feel attacked.
- ▶ The result is that the fight-or-flight portion of the brain becomes engaged, and this can overtake the functioning of the part of the brain associated with our rational faculties. We feel that we have to defend ourselves through any means possible, not necessarily the ones that will lead to open-minded consideration.
- ▶ The outcome is often logically unfortunate. Frequently, the conversation gets hijacked by irrelevant appeals that cause the discussion to veer off in directions that do not serve the central point, but serve only to obscure it.

AD HOMINEM

- ▶ One of the most common diversionary fallacies is where instead of attacking the argument, we instead focus on attacking the arguer. This fallacy is known by its Latin name, “ad hominem,” which translates as “to the man.” The idea is that we are focused on the person instead of the case the person is making.
- ▶ Arguments are acceptable if they are sound—that is, if their form is valid and their premises are well-grounded. If an argument is valid and has true premises, then that is true regardless of whose mouth it comes out of.
- ▶ Arguments stand or fall on their own merits. Whose mouth it comes out of is irrelevant. The argument is valid because of its logical structure and is well-grounded because of the truth of its premises. The identity, background, or motivation of the speaker has nothing to do with the satisfaction of either criterion.
- ▶ To argue that we should not accept—or, indeed, even consider—the argument because of the source of the argument is to commit an ad hominem fallacy.
- ▶ Ad hominem attacks tend to come in three general categories. The first is the “you’re a jerk” version. There are horrible, immoral people in this world who do nothing to make the world a better place and who often serve their own petty desires at the cost of the well-being of others. But if that person makes an argument, we need to analyze it objectively. We need to evaluate the validity of the argument and assess the likelihood of the premises’ truth.
- ▶ The second version is a form of guilt by association where we discount an argument not for objective reasons, but because the person offering it belongs to some identifiable group. Don’t listen to her; she’s a feminist. You can’t take his argument seriously; he’s a conservative.

- ▶ A common variation of this kind is to point out that the speaker is not among those who follow the advice the speaker is giving. Known by its Latin name, “tu quoque,” the “but you do it, too” objection is just an illegitimate ad hominem attack. It might be true that the person telling you not to drink is an alcoholic, but that doesn’t mean it is not good advice.
- ▶ The third class of ad hominem attacks is where we focus on the motivations of the speaker. “Of course, you’d say that. You stand to profit if it’s true.” Again, maybe that is correct, or maybe it isn’t, but the argument stands or falls on its own merits, regardless of who, where, when, or why the argument is made.

ATTACKING A STRAW MAN

- ▶ Another diversionary tactic that we must be on guard against is called attacking a straw man. The strange name comes from the fact that it is easier to beat the stuffing out of a scarecrow than it is to take on an actual human being. It is a metaphor for arguments that do not address the actual argument made but rather a weaker, easier-to-refute version.
- ▶ Logicians have something called the principle of charity, according to which, when one analyzes an argument, one must assess the strongest-possible version of that argument. To defeat a weak version does nothing in terms of demonstrating the given argument to be unsound. It can only be rejected as not providing legitimate grounds for rational belief in its conclusion if the strongest version, the best understanding, is seen to be flawed.
- ▶ Think of prize fighting. If a particular boxer is the reigning champion, then he or she has to take on all challengers. One cannot keep the title of heavyweight champion of the world and fight only, for example, 12 year olds. While there are surely some tough preteens out there, the point of being the champion of the

world is that you are the top of the top—that you can defeat the toughest competition anywhere.

- ▶ It is the same thing with argumentation. If we are to refute an argument—find a flaw in it that leads us to reject it as providing good reason to believe—then, like the heavyweight champion of the world, we as critical thinkers need to take on the strongest version of the argument. To take on a weaker version and then assert that we have done anything is to attack a straw man.
- ▶ There are two main varieties of attacking a straw man. One version is to alter the scope of the premises offered, making them broader or narrower than the ones offered to weaken the argument while keeping the rest of the premises intact. The hint that this is what you are hearing is the phrase “Oh, so what you are saying is...”
- ▶ The other kind of straw man argument is more radical. It is where the interlocutor replaces all of the premises wholesale. When you hear the phrase “the real reason...,” you are likely looking at a straw man argument.
- ▶ Why would someone say “the real reason”? Because what they are doing is replacing the original reasons—that is, the premises—with new premises, and odds are that these new premises are going to be a whole lot easier to undermine. But in undermining them, the interlocutor has done nothing with regard to the soundness of the original argument, because the original argument is gone.

RED HERRING

- ▶ Where attacking a straw man is the error wherein we replace the premises of an offered argument, the fallacy known as a red herring is where we change the conclusion.

- ▶ When we replace premises with those that are easier to attack but maintain the conclusion, we are still talking about the same thing, only talking about it differently. But when we change the conclusion, we are completely changing the topic of conversation. That is a red herring—the ultimate in argumentative diversion.
- ▶ Anyone who has ever been in a serious interpersonal relationship knows all about red herrings.
 - ◊ “You really need to clean those dishes in the sink. You make yourself a snack and just clutter the kitchen and leave it for me. That is not respectful or fair to me.”
 - ◊ “Well, if you want to talk about messes and respect, what about the fact that you never pick up your dirty clothes in the bathroom? You just throw them on the floor before you get in the shower and leave them there.”
- ▶ This began as a discussion about dishes. The first partner made the following argument.
 - ◊ One should clean up one’s own messes, because not to do so is disrespectful and unfair.
 - ◊ The dishes in the sink are your mess that is not cleaned up.
 - ◊ Therefore, you should do those dishes out of respect and fairness.
- ▶ That seems like a sound argument. The conclusion follows from premises that certainly seem to be true. How does the other partner respond to this argument? Not by showing that the argument is flawed, but by giving a new argument.
 - ◊ One should clean up one’s own messes, because not to do so is disrespectful and unfair.
 - ◊ The clothes on the bathroom floor are your mess that is not cleaned up.

- ◊ Therefore, you should pick up the clothes on the bathroom floor out of respect and fairness.
- ▶ Notice that while the form of these two arguments is the same and there is some overlap in content, the conclusions are different propositions—that is, they are completely different arguments. Both are worth assessing, but they need to be considered one at a time.
- ▶ What we have here is a combination of a red herring and tu quoque. Well, you do it, too, or something so much like it that you can't criticize me for doing what you do.
- ▶ If I do it, then I should be criticized and I should change my ways, but that is a different question from what we are talking about, which is the stack of dirty dishes in the sink.
- ▶ The thing about people is that while we certainly have trainable rational capabilities, we are also bundles of insecurities and dedicated to agendas of our own which we take to be crucially important. Arguments between two people in a relationship can display insecurities.
- ▶ Agenda-based red herrings are often seen in political discussions. Consider a conversation like the following.
 - ◊ “Your gun control proposal is an affront to gun owner rights. We are talking about liberty being seized by an overinvasive government here.”
 - ◊ “Oh, that’s funny coming from the person who proposed such draconian abortion regulations. If you want to talk about rights and liberty being stripped by an overinvasive government, there is example A.”
- ▶ Notice what happened. We started with a conversation about the political benefits and flaws of a proposed piece of legislation

about firearms, and instead of evaluating the argument, we shifted to a completely different topic: the permissibility of abortion.

- ▶ Both are important issues. We should give both careful, thoughtful attention. But we need to do so one at a time. “I understand that abortion rights is an important issue to you, and we will give it our due attention, but right now we are talking about gun control.”
- ▶ When there is an issue that is important to us, we will often see traces of it everywhere. Just because something reminds you of a topic you want to discuss doesn’t mean that we cannot first finish the discussion we have started.

READINGS

Copi, *Introduction to Logic*, chap. 3.

Damer, *Attacking Faulty Reasoning*, chap. 9.

Kahane, *Logic and Contemporary Rhetoric*, chap. 4.

QUESTIONS

Identify the fallacies in the following passages from the following list: ad hominem, tu quoque, attacking a straw man, red herring.

1.

Don't listen to him. He can't even speak proper English, so you know his argument is also nonsense.

2.

You say that the changes to the tax code would promote fairness, but the real reason you are in favor of it is that you want to punish the rich.

3.

You say that we need to help the homeless, but what about the working poor who have a place to live? Do you think we should just ignore them?

4.

You know that famous celebrity who is always going on and on about the need to care about the environment and leave a small carbon footprint? It turns out that she has a mansion, and you know that thing uses a ton of electricity in the summer when she runs the air conditioner. So, if she can use a lot of power, so can I.

ANSWERS

1.

ad hominem

2.

attacking a straw man

3.

red herring

4.

tu quoque

Inductive Reasoning

This lecture begins the move from thinking about well-groundedness to validity. Each of the two different types of arguments, inductive and deductive, require different means of determining validity. Inductive arguments start with a set of observed instances and use that information to infer beyond it. Inductive inferences do not give us the certainty of deduction, but in the messiness of the real world, they are the inferences we most often make. Because their conclusions outrun the content of their premises, inductions are incapable of giving us complete confidence in their conclusions, but they give us high probability, and for reasonable belief, that is good enough.

INDUCTIVE ARGUMENTS

- ▶ Arguments come in two different kinds: deductive and inductive. A deductive argument is one in which the scope is non-ampliative—that is, the scope of the conclusion is no broader than the scope of the premises.
- ▶ In other words, deductive arguments go from broad to narrow; the conclusion doesn't talk about anything that wasn't already covered in the premises. Because there is no new information in the conclusion—because it is just milking a specific result out of the premises—one nice thing about deductive arguments is that if they are sound, or well-grounded, then their conclusions must be true.

- ▶ Inductive arguments, on the other hand, are ampliative—that is, their conclusions do move beyond the scope of the premises to give us rational belief about something we have not yet observed. Induction is ampliative in that it amplifies our rational beliefs; it takes us from narrow to broad.
- ▶ Inductive arguments are wonderful because they give us new knowledge about the world. They take what we already know and give us logical permission to believe new things that we did not know before.
- ▶ Deduction only rearranges our previous knowledge into new forms we may or may not have considered, but induction actually generates completely novel beliefs about the world. This does not come without a fee, and the cost of this growing of our stockpile of rational beliefs is certainty.
- ▶ Because deductive inferences are only rearranging what we knew before, if what we knew before was known to be true, then the conclusions that come from deduction will also be true. But with induction, because we are making a logical leap beyond the content of the premise set, there is no absolute guarantee that the result must be true.
- ▶ The best we get from induction is likely truth. If you have a good inductive argument, the conclusion is probably true. While “probably true” is less desirable than “definitely true,” “probably true” is sufficient for rational belief. We should believe that which is probably true.
- ▶ Deductive certainty in all of our beliefs would be wonderful, but it is not available to us. We need inductive inferences because in most real-life situations, it is all we have. And it is good enough. We make inductive inferences all the time, and we should. High probability is all we need for rational belief.

TYPES OF INDUCTIVE ARGUMENT

- ▶ While there are many different kinds of inductive arguments, there are three forms that are the most important because they are the most frequently used. In one form, called inductive analogy, we move from a set of data showing universal adherence of a particular property in a particular population and assert it of a future individual.
- ▶ An inductive analogy is an inductive argument of the following form.
 - ◊ P_1 has the property A .
 - ◊ P_2 has the property A .
 - ◊ P_3 has the property A .
 - ◊ ...
 - ◊ P_n has the property A .
 - ◊ I have only seen n instances of P .
 - ◊ Therefore, P_{n+1} has the property A .
- ▶ “I have seen some number n different P s, and every one of them has had the property A ; therefore, I believe that the next P I see will also have the property A .”
- ▶ We make these kinds of inferences all the time. “No, I don’t want to go see that scary movie, because it’ll give me nightmares like all the other ones.”
- ▶ When we apply what we’ve learned from all other instances in the past to one in the present or future, inductive analogy is the form of the argument we are using.
- ▶ A related but stronger inductive inference is the universal generalization. It has the following form.
 - ◊ P_1 has the property A .
 - ◊ P_2 has the property A .

- ◊ P_3 has the property A .
- ◊ ...
- ◊ P_n has the property A .
- ◊ I have only seen n instances of P .
- ◊ Therefore, all P s have the property A .

- ▶ While both inductive analogy and universal generalization have the same premises, notice the difference in the conclusions. One is an analogy in that it says that a future instance will be like all past instances, whereas the other makes a much broader claim, saying something about *all* members of the observed population.
- ▶ Universal generalizations are also common in everyday life. “Every time I’ve tried something with cilantro in it, I have thought it’s disgusting. I don’t like cilantro.” Notice that in this example, we are making a generalization about an entire category of things, not just a prediction about a single upcoming instance.
- ▶ The third form of inductive inference is called statistical generalization. Its form is as follows.
 - ◊ X percent of all observed P s have the property A .
 - ◊ Therefore, X percent of all P s have the property A .
- ▶ Like universal generalization, we are generalizing over the entire set of P s from some limited sample of P s, but here we are not attributing the property to all of them, but to some percentage—either an explicit percentage or a vaguer amount—of the population.
- ▶ “Four out of five dentists surveyed recommend sugarless gum for their patients who chew gum, so likely the overwhelming majority of dentists overall agree.”
- ▶ In some cases, we are generalizing a more specific statistic, such as a percentage of the population as a whole; sometimes we are taking a specific statistic, such as 80 percent of dentists

surveyed, and generalizing it broadly; and sometimes we are taking a vague sense, such as “a lot of the time” or “most people surveyed,” and making equally vague generalizations.

- ▶ We need to be careful that we don't fall prey to one of the standard fallacies we see with inductive inferences: exaggerated accuracy. If a certain basketball player has made 75 percent of his free throws for the year, that could mean that we think he will hit most of his free throws in the next game—that's a perfectly fine statistical generalization.
- ▶ But it is not good reason to think he will make three out of four. There is a lot of mathematical machinery in the study of statistics that tells you exactly how big a sample you need and what percentage needs to have the property in order to have a 90 percent, 95 percent, or 99 percent degree of confidence that your statistic can be accurately generalized to the entire population.
- ▶ When you have a statistic you want to generalize, you are almost always going to have to soften your generalized conclusion a bit.

INDUCTIVE FALLACIES

- ▶ One of the fallacies associated with inductive reasoning is exaggerated accuracy. There are several more we need to be aware of. One involves premises that appear in the first two inductive argument forms and implicitly in the third.
- ▶ Note the last premise in inductive analogy and universal generalization: “I have only seen n instances of P ,” and note the form of the premise in statistical generalization: X percent of all observed P s have the property A .” In both cases, we see what is called requirement of complete information.

- ▶ Induction is unlike deduction in that deduction has, but induction lacks, the property logicians call monotonicity. An inference is monotonic if adding more premises will not turn a valid argument into an invalid one. This is true for deductive arguments.
- ▶ As long as the argument is valid to begin with, new premises—no matter what they are—will always keep the argument valid. While this is true for all deductive arguments, it is not true for inductive arguments. Inductive arguments are non-monotonic because adding a new premise can turn a perfectly good inductive argument into a bad one. One new piece of information can completely destroy an argument in a way that cannot happen for deductive arguments.
- ▶ So, for an inductive inference to be valid, we need a guarantee that the evidence given in the premises is all of the relevant observations we have. We need complete information. To fail to do so is called cherry-picking.
- ▶ This is a common error where when someone wants to support a point, he or she will present lots of evidence—so much evidence that we can't help but believe the conclusion it leads to. But if the person was careful to select only the evidence that supports his or her position and excludes the counterevidence, we will be led to believe something that is not well supported.
- ▶ The key to avoiding this error is in the selection procedure for the sample we use to collect our observation. The sample is the set of n individuals mentioned in the premises. How do we go about getting our data from which to make the inductive inference?
- ▶ There are two keys to having an acceptable sample for a good inductive inference. The first is sample size. We need the sample to be large enough to support our inference. How large is large enough?

- ▶ This is something that statisticians have studied in great detail and have tables dedicated to showing for varying degrees of confidence in our inferences. For the purposes of this course, one should be on the lookout for absurdly small samples, called anecdotal evidence.
- ▶ If someone is generalizing on the basis of just a few examples or often just a single experience, then we have the fallacy of insufficient sample. If this error is committed while making a universal generalization, we sometimes call it a hasty generalization, but the term “insufficient sample” covers all of the forms of inductive reasoning from too small of a sample.
- ▶ The size of the sample isn’t the only thing we need to be concerned about. Bad samples could be quite large. To make a good induction, samples have to be representative of the population over which we are making the inference.
- ▶ The sample needs to look like the population in miniature because there might be aspects of different subgroups within the population that affect the distribution of property we are examining. If you have a poorly distributed sample, you commit the fallacy of unrepresentative data.
- ▶ Good samples model the heterogeneity of the population. If your population is homogeneous, then you don’t need to worry so much, but the key is that you cover the relevant subgroups in proportion to their general representation. This assumes, of course, that you know what subgroups are relevant and their proportion of the general population beforehand.
- ▶ But sometimes you don’t. In those cases, the key is a random sample. The idea is that if we pull enough individuals out of the population without a bias toward or away from any subgroup, then relevant subgroups will show up in the sample roughly the proportion they occupy in the whole.

- ▶ The key to the random sample, then, is to make sure that your sample is large enough that small subgroups will appear and that your selection procedure does not accidentally bias your selection toward or away from such subgroups.
- ▶ The gambler's fallacy is one in which we try to make an inductive inference from data that bear no cause-and-effect relationship to each other, where there is not a probabilistic relation between the members of the sample and the property being observed. Arguing that past experience in any way affects the next instance is to commit the gambler's fallacy.

READINGS

Barker, *The Elements of Logic*, chap. 7.

Copi, *Introduction to Logic*, chap. 3.

Damer, *Attacking Faulty Reasoning*, chap. 5.

Kahane, *Logic and Contemporary Rhetoric*, chap. 3.

QUESTIONS

1.

If inductive arguments do not give us conclusions with absolute certainty, why should we believe the conclusion of a sound inductive argument? What inductive arguments do you accept in day-to-day life?

2.

Identify the fallacies in the following passages from the following list: cherry-picking, insufficient sample, unrepresentative data, gambler's fallacy.

- a Al Capone: Italian and mafia. Lucky Luciano: Italian and mafia. John Gotti: Italian and mafia. Sammy the Bull: Italian and mafia. I could go on and on, naming hundreds of Italian mobsters. So, it must be the case that all Italians are connected.
- b I always hold my breath when crossing a bridge to make sure that nothing bad happens.
- c My Ford has 300,000 miles on it and has never needed a single major repair. Ford makes quality cars.
- d A survey conducted by Christian televangelist Pat Robertson polled more than a thousand of his followers and overwhelmingly showed that people want creationism instead of evolution taught in public schools.

ANSWERS

1.

Answers will vary.

2.

- a cherry-picking
- b gambler's fallacy
- c insufficient sample
- d unrepresentative data

Induction in Polls and Science

Most of what we believe traces back to inductive reasoning. We learn from experience, and the reason that works is because of the successfulness of induction. Inductive arguments are everywhere in life. There are two places, however, where we conspicuously find formalized inductive inferences that deserve extra discussion: Inductive inferences are made whenever a poll is taken and reported, and induction is a crucial element of scientific reasoning. In this lecture, you will learn about induction in these two contexts. You'll learn how to read and understand what information a poll is giving and how to understand the roles of induction in scientific results.

INDUCTIVE INFERENCE IN POLLS

- ▶ We have looked at the basic forms of inductive reasoning and the ways in which some inductive arguments can go wrong. With respect to polls, those are precisely the errors we need to be on the lookout for, with some interesting twists and a few additions.
- ▶ Virtually all polls taken are of the form of the inductive argument called statistical generalization.
 - ◆ X percent of all observed P s have the property A .
 - ◆ Therefore, X percent of all P s have the property A .
- ▶ We are generalizing a statistic we find in our observed sample to the entire population. What we want to examine with respect to polling is what makes for a successful statistical generalization and how much reason to believe is achieved as a result.

- ▶ Inductive arguments, because they are ampliative—that is, because their conclusions are broader in scope than their premises—come with risk. Even the best inductive arguments do not provide us with the certainty of deductive arguments.
- ▶ Successful inductions give us high probability, a level of reasonable belief. We want to see, then, how a single poll or collection of polls should affect what we believe and how deeply.
- ▶ Several of the problems that can plague inductive arguments come from flawed samples. This is the case with polls, too. Indeed, one of the biggest concerns of legitimate pollsters is sampling.
- ▶ For a good sample, we need both size and distribution—that is, we need for the poll to have asked enough people, and we need for those people to resemble the population being generalized in miniature. All of the relevant subgroups must be present in the sample in roughly the proportion in which you find them in the general population.
- ▶ With respect to size, we looked at the fallacy of insufficient sample in its simplest form: anecdotal evidence, where we generalized from just one or two experiences. But in a poll, where we are trying to generalize over an entire electorate or an entire national population, how big is big enough?
- ▶ It is actually surprising how small a sample can be and still be of an acceptable size to make meaningful claims about so large a population. About 1,000 respondents is sufficient for a nationwide poll in the United States, which has more than 150 million registered voters.
- ▶ The other concern, once we have enough people, is that the sample is well distributed—that it looks like the population as a whole in miniature. With polls, we have two concerns: that we are sampling the right population and that the sample is properly representative.

- ▶ But even when we have determined the proper population and can develop filters to screen out those who do not belong, we have the question of creating a well-distributed sample.
- ▶ There are a few different ways of trying to develop a well-distributed sample. First, we could identify all of the important demographic groups and their proportions in the population beforehand and then make sure that our sample is shaped to resemble it.
- ▶ Second, we could take a random sample. If we select enough people at random from a large group, the sample will come to resemble the population as long as there were not selection biases toward or against given subgroups.
- ▶ Clearly, the first option is preferable. If we know what properties are germane, we can find people with the desired profiles in the desired proportions. But how do we know this?
- ▶ Often, it is done inductively. We look at past polls and see how accurate their predictions came out and see what their sample looked like in terms of different subgroups. Pollsters call this making a model of the population. They collect their data but then have to decide how to weight the contributions of various parts of the population.
- ▶ Once the sample worries are behind us, there is a new concern: Can we trust what people tell the pollster? There are two worries here. The first is that people sometimes intentionally mislead pollsters.
- ▶ Especially on questions that have a moral element, people don't want to be seen as being on the unpopular side of the question and will tell pollsters what they think will lead the pollster to think better of their character, even if it does not reflect the way they will actually vote or what they actually believe.

- ▶ The name for this is the social desirability bias and with respect to polls is often referred to as the Bradley effect, named for Los Angeles mayor Tom Bradley, who when running for governor of California had a sizeable lead in the polls but lost the election. Some have contended that this was the result of people falsely telling pollsters that they would vote for Bradley when they would not, so that they would not appear to be racist.
- ▶ Second, it turns out that very different results can be achieved by asking the same question in different ways. How the pollster chooses to ask the question will affect how people answer it. This is a cognitive bias that social psychologists call the framing effect, and it relates to the fallacy of begging the question.
- ▶ Words not only denote—that is, pick out objects to refer to—but also have connotative power; they convey emotional or value-laden judgments as well. Questions framed differently will allow the connotative power of the language employed to steer the listener toward or away from particular viewpoints in ways that are subtler than leading questions.
- ▶ Polls not only report on the public mood and beliefs, but also influence them. This means that we need to be cautious. Questions might have been asked intentionally or unintentionally in a way that biases them.
- ▶ One's own ideological filter can lead one to frame a question in a leading way that seems perfectly fair to you. The key, then, to making good inferences from polling data is to take a step back. Polls themselves are inductive arguments—that is, they take data and extrapolate a general result.
- ▶ But we are now able to use the polls themselves as data in an inductive argument at a higher level. We can now take polls of polls. There are websites that work as poll aggregators, collecting

all of the polls on a particular topic and displaying their results, the size of their samples, etc. By bringing all the polls together, we can form a rational basis for reasonable belief based on them.

INDUCTION IN SCIENCE

- ▶ Induction is used in two different ways in science: in supporting hypotheses and in testing theories. We need to keep these two contexts separate.
- ▶ A hypothesis is a proposed statement of purported fact. Scientific theories, by contrast, are sets of general axioms, which together form a system of thought that provides a full picture of the workings of some part of nature.
- ▶ Hypotheses are proposed individual statements of possible truth; they are more specific than the axioms, and we get evidence for them individually. The axioms work together as a group, and what we test in that very different context is the theory as a whole. We might be able to derive hypotheses when working within the theory, but the parts of the theory are not themselves hypotheses.
- ▶ There are different inductive processes for hypotheses and theories. The philosopher Karl Popper pointed out that a hypothesis is scientific only if it is falsifiable—that is, only if there are observable circumstances that would render the statement false.
- ▶ How do we go about our scientific inductive inferences? The first step is to identify the independent and the dependent variables. The independent variable is the thing we adjust or administer—the thing under our control. The dependent variable is what we measure. It is the thing not under our control, and it may or may not change as a result of our adjusting the independent variable.

- ▶ For example, to determine whether high doses of vitamin C reduce the symptoms of the common cold, we can give someone high doses of vitamin C—the independent variable—and we can then check to see if they receive relief from their cold symptoms—the dependent variable.
- ▶ One problem is that the world is a complex, messy place. While we are giving someone vitamin C, they might also be eating something else that is decreasing their symptoms, and we will wrongly attribute the effect to our independent variable. So, we need to conduct the experiment in a way that controls for other independent variables—that is, we need to do our best to create a fixed environment in which we can screen off intervening causes.
- ▶ When we get the data, we analyze it to see if the results are statistically significant. If so, we have reason to think that the results are consistent with the hypothesis. What that means is that the hypothesis has not been falsified, but we also do not yet have enough evidence to think it is probably true.
- ▶ Where do we see induction in science with respect to theories? The philosopher of science Hans Reichenbach drew a distinction between the context of discovery and the context of justification. What this distinction has come to mean is the context in which scientists come up with their theories and the context in which they provide good reason to believe they are true.



HANS REICHENBACH
(1891–1953)

- ▶ The context of discovery is generally thought to be free—that is, there is no specific logic of discovery, no turn-the-crank method of coming up with scientific theories. But while there is no set method, there is induction, because scientists are working from their experiences and the data.
- ▶ They have a question about how a certain system works, and they consider what they know and make inductive leaps. They look for models—analogy where the system could be thought to work like a different system that is better understood.
- ▶ The most important place in scientific reasoning that we find induction is in the context of justification. Once a theory has been proposed, why should we believe it? Theories are testable; they have effects, results, and predictions that come from them. These observable results of a theory are determined deductively—that is, if the theory is true, then in this given situation, a particular observable consequence should result.
- ▶ We go into the lab, set up the situation, and see if we observe or measure the result as expected. If not, then the theory has failed and, as it stands, is not acceptable. It will either have to be rejected or fixed.
- ▶ But if the theory says to expect a particular result and we observe it, now we have evidence in favor of the theory. That evidence is inductive. It might be that a particular theory predicts the result, but there will also be other theories that are different from the first one but are also supported by the result.
- ▶ As such, none of the theories are certain, in the way that deductive inferences are, but rather they receive inductive support. They are more likely to be true than they had been. The probability of truth has been increased.

- ▶ To go from supporting evidence, which makes a theory more likely, to conclusive evidence, which makes a theory probably true, we need lots of evidence as well as evidence of different types.
- ▶ It is good for a theory if it accounts for everything we already knew. We call this retrodiction. This is especially true if what we already knew was previously unexplained. But better than explaining what we already knew, prediction is taken as strong evidence.
- ▶ The best evidence brings about what scientist William Whewell termed “consilience,” which is when a theory designed to account for phenomena of type *A* turns out also to account for phenomena of type *B*. If you set out to explain one thing and also are able to explain something completely different, that is strong evidence that your theory is probably true.

READINGS

Bradburn, Sudman, and Wansink, *Asking Questions*.

Gimbel, *Exploring the Scientific Method*.

Kelley, *The Art of Reasoning*, chaps. 15–17.

QUESTIONS

1.

Push polls are fake polls in which people believe that they are being asked for their opinions by a reputable pollster but are actually being lobbied by an interested party that is trying to change their opinion. Political campaigns, for example, will have people pose as pollsters and ask biased questions designed to influence the respondents' answer. Because the person believes the pollster to be legitimate, the respondent thinks that the questions are not skewed and therefore

have their own thoughts unfairly affected. Are all polls, in some sense, push polls, or can a fair poll be conducted that takes the temperature of a population without affecting it?

2.


Because science is inductive in the inference it requires for evidence, science proves nothing. If the results of science are never proven—if scientific progress might require us to surrender even the most deeply held current beliefs—then is it ever rational to believe the results of science?

Introduction to Formal Logic

You have learned that there are two criteria for acceptable arguments: well-groundedness and validity. You have also seen that there are two types of arguments: inductive and deductive. So far, you have examined some aspects of informal logic—that is, analysis of the well-groundedness of arguments—and you have looked at validity concerns for inductive arguments. What remains is deductive validity, the aspect of logic that has historically commanded a great amount of attention from logicians. In this lecture, you will be introduced to formal logic.

FORMAL LOGIC

- ▶ In part because of its use in mathematics and in part because it is the study of what we can know with absolute certainty, there is a long trail of work in formal logic. The first figure in the history of thought to give us a functional formal account of reasoning was Aristotle.
- ▶ Possibly the smartest person ever to have lived, Aristotle provided us with the starting point of virtually every academic discipline, from economics and physics to literary theory and ethics. But for this course, it is his work on categorical logic that is important.



ARISTOTLE
(384 B.C.E.–322 B.C.E.)

CATEGORICAL PROPOSITIONS

- ▶ Aristotle's logic begins with a little bit of grammar. Formal logic examines what propositions necessarily follow from what other propositions because of their forms.
- ▶ Propositions need to be expressed as declarative sentences, and declarative sentences, Aristotle contends, have a specific form. All declarative sentences that we will consider have a subject and a predicate.
- ▶ The subject is what the sentence is about, and the predicate is what we are asserting about the subject. A sentence is true if and only if the subject does have the property asserted by the predicate, and it is false if the subject does not.
- ▶ These categorical sentences come in four types. First is the universal affirmative, sentences of the following form: All *As* are *B*.
 - ◆ All people have noses.
- ▶ Second is the universal negative: No *As* are *B*.
 - ◆ No circles have angles.
- ▶ Third is the particular affirmative: Some *As* are *B*.
 - ◆ Some people are blonde.
- ▶ Finally, there is the particular negative: Some *As* are not *B*.
 - ◆ Some wines are not expensive.
- ▶ We arrange these sentences into what is called the square of opposition.

UNIVERSAL

AFFIRMATIVE	All As are <i>B</i>	No As are <i>B</i>	NEGATIVE
	Some As are <i>B</i>	Some As are not <i>B</i>	

PARTICULAR

- ▶ The top line is for universal sentences. The bottom line is for particular sentences. The left side is for affirmative sentences. The right side is for negative sentences.
- ▶ There are abbreviations for each. Universal affirmative sentences are called *A* sentences. Universal negations are called *E* sentences. Particular affirmative sentences are *I* sentences, and particular negative sentences are called *O* sentences.

UNIVERSAL

AFFIRMATIVE	A All As are <i>B</i>	E No As are <i>B</i>	NEGATIVE
	I Some As are <i>B</i>	O Some As are not <i>B</i>	

PARTICULAR

- ▶ The reason we call this the square of opposition is because opposite corners—sentences that are connected diagonally on the square—are negations of each other. That is, an A sentence is true if and only if the O sentence is false.
- ▶ All Boy Scouts are boys if and only if it is not true that some Boy Scouts are not boys. Similarly, if it is false that all Boy Scouts are boys, then there must be some Boy Scout that isn't a boy. A and O sentences will always have different truth-values.
- ▶ And it is the same for E and I sentences. If it is true that no clowns are happy, then it is false that some clowns are happy. And if it is true that some clowns are happy, then it must be false that no clowns are happy. E and I sentences must have different truth-values.
- ▶ The sentences on opposite corners are called contradictories because they contradict each other: One and only one can be true at a time. But this is not the case with A and E or I and O sentences.
- ▶ A and E sentences are called contraries. They cannot both be true, but they can both be false. It is false that all paintings use the color blue and false that no paintings use the color blue.
- ▶ I and O are called subcontraries. They can both be true, but they cannot both be false. It is true that some people are having a birthday today and some people aren't. But if it is false that some A s are B , then it must be true that some A s are not B .
- ▶ In these I and O sentences, the word “some” means at least one, maybe all. By using the word “some” we are not saying that only some, but not all. It might be simultaneously true that all people have mothers and that some people have mothers.

- ▶ The word “all” is slightly more complex. It means every member of the category. But what about the peculiar categories that are empty? Consider the following *A* sentence: “All unicorns have a horn.” Is this sentence true or false? It seems to be true by definition—to have a horn is part of what it is to be a unicorn. But at the same time, there are no unicorns for the sentence to be true of.
- ▶ How do we make sense of this? The answer is that we have two different meanings: the hypothetical viewpoint and the existential viewpoint.
- ▶ In the hypothetical viewpoint, the sentence “All unicorns have a horn” is true, because it means that all unicorns, if there are any (and there might not be), have a horn.
- ▶ From the existential viewpoint, the sentence “All unicorns have a horn” is false, because it now means that there are unicorns and all of them have a horn. We need to know whether we are working with the hypothetical or existential viewpoint before we start doing our logical manipulations.
- ▶ Notice one important difference. From the existential viewpoint, *A* sentences imply *I* sentences and *E* sentences imply *O* sentences. If no square has five sides, then some squares do not have five sides.
- ▶ As long as we know that the subject exists, then if the predicate holds for all, it must hold for some. But from the hypothetical viewpoint, just because an *A* sentence is true does not mean that the corresponding *I* sentence will be true.
- ▶ In the hypothetical viewpoint, it is true that all griffins have the body of a lion, but that does not mean that some griffins do, because the word “some” means that there is at least one, and the griffin is a mythical beast. There aren’t any.

- ▶ So, with vacuous subjects, we can have true *A* sentences and false *I* sentences or true *E* sentences and false *O* sentences. Both viewpoints are legitimate, but for clarity's sake, let's assume the hypothetical viewpoint unless otherwise noted.

CATEGORICAL SYLLOGISMS

- ▶ Recall that our interest here is deductive validity. We want to know when we have to believe a conclusion if we also believe the premises. For Aristotle, the key to reasoning is a type of argument called a syllogism, which is an argument with two premises.
- ▶ Categorical syllogisms are arguments with a categorical sentence as a conclusion and two categorical sentences as premises.
 - ◊ All humans are mortal.
 - ◊ All Greeks are human.
 - ◊ Therefore, all Greeks are mortal.
- ▶ The conclusion has two terms: a subject and a predicate. The subject of the conclusion is called the minor term (**S**). The predicate of the conclusion is called the major term (**P**).
- ▶ In the example, "Greek" is the minor term, and "mortal" is the major term. There is another term that appears in both premises, but not in the conclusion. This is called the middle term (**M**). "Human" is the middle term in the example.
- ▶ The premise with the minor term and the middle term is called the minor premise, and the one with the middle term and the major term is called the major premise. We always write out the major premise first.

- ▶ The example has a major premise that is an *A* sentence, a minor premise that is an *A* sentence, and a conclusion that is yet another *A* sentence. Reading top to bottom, we say that this argument has the mood *AAA*. *AAA* syllogisms are nice because if they turn out to be invalid, you can call for a free tow into the nearest garage.

A	$\begin{array}{cc} M & P \\ \hline \end{array}$	All humans are mortal.	Major Premise
A	$\begin{array}{cc} S & M \\ \hline \end{array}$	All Greeks are human.	Minor Premise
A	$\begin{array}{cc} S & P \\ \hline \end{array}$	Therefore, all Greeks are mortal.	Conclusion

- ▶ But notice that we can have different variations of *AAA*. Compare our example with a new one.
 - ◊ All Greeks are human.
 - ◊ All humans are mortal.
 - ◊ Therefore, all mortals are Greeks.
- ▶ These two arguments both have *A* sentences for premises, but because the terms are switched around, they are different arguments. Indeed, the first example is valid, but the second is not. The point, though, is that the order of the terms matters.
- ▶ There are four possible arrangements of the terms. Each arrangement is called a figure.
 - ◊ The first figure has a major premise that starts with the middle term and a minor premise that starts with the minor term.
 - ◊ The second figure has a major premise starting with the major term and a minor premise starting with the minor term.
 - ◊ The third figure starts both premises with the middle term.
 - ◊ The fourth figure starts the major premise with the major term and the minor premise with the middle term.

- ▶ This is a complete catalogue of possible figures.

1		2		3		4		
M	P	P	M	M	P	P	M	Major Premise
S	M	S	M	M	S	M	S	Minor Premise
S	P	S	P	S	P	S	P	Conclusion

- ▶ Combining mood with figure, we get a complete list of possible forms of categorical syllogisms. For example, “Some dogs have four legs, and some four-legged beings are not cats; therefore, no cats are dogs.”
- ▶ The major premise is an *I* sentence: “Some dogs have four legs.” The minor premise is an *O* sentence: “Some four-legged beings are not cats.” And the conclusion is an *E* sentence: “No cats are dogs.”
- ▶ So, the mood is *IOE*. The middle term appears second in the major premise and first in the minor premise, so it is of the fourth figure. We have an argument that is *IOE-4*.
- ▶ The first example argument has a mood of *AAA*, and the middle term is second in the major premise and first in the minor premise, so it is an *AAA-1*.
- ▶ There are four possibilities for the major premise—*A*, *E*, *I*, *O*—and the same four possibilities for the minor premise and the conclusion. That gives us $4 \times 4 \times 4$, or 64 possible moods. We then have four figures for each mood, giving us 64×4 , or 256 categorical syllogisms whose validity status we need to determine. How do we do that?

ARISTOTLE'S FIVE RULES

- ▶ There are five simple rules that detect the 15 valid forms from the hypothetical viewpoint. But before we work with the rules, we need to understand one more concept: a distributed term.
- ▶ A term in a categorical sentence is distributed if that sentence says something about the entire category the term refers to. So, in an *A* sentence like “All Greeks are human,” the subject is distributed because it tells us something about all Greeks.
- ▶ Likewise, in an *E* sentence like “No cowboys are werewolves,” we are told something about the entire class of cowboys, so again the subject is distributed.
- ▶ In an *I* sentence like “Some people are allergic to peanuts,” we are not told anything about any entire class, so no term is distributed.
- ▶ In an *O* sentence like “Some computers are not made by Apple,” it seems intuitively that, like with *I* sentences, neither term is distributed. But, in fact, the predicate of an *O* sentence is distributed. We know of the entire category of products made by Apple that it fails to include some computers.
- ▶ So, *O* sentences distribute the predicate, *A* and *E* sentences distribute their subjects, and *I* sentences distribute nothing.
- ▶ Having this concept, we can now set out the five rules.
 - 1 In all valid syllogisms, the middle term is distributed in at least one of the premises.
 - 2 In all valid syllogisms, any term distributed in the conclusion is also distributed in the premises.
 - 3 In all valid syllogisms, at least one of the premises must be affirmative.


- 4 In all valid syllogisms, if the conclusion is negative, one premise must be negative.
 - 5 In all valid syllogisms, if the conclusion is particular, at least one of the premises must be particular. (This one is for the hypothetical viewpoint only; for the existential viewpoint, just use the first four.)
- ▶ A syllogism is valid if and only if it satisfies all five of these rules. If it violates even one, it's an invalid argument.

VENN DIAGRAMS

- ▶ These rules work, but they are not terribly intuitive. We can determine which categorical syllogisms are valid and which are not, but we come away without a sense of why. It would be nice to have a method that allows us to see why a categorical syllogism is or is not valid.
- ▶ Fortunately, we have one. It was provided by the 19th-century English logician John Venn. He came up with a way of visually representing the content of each type of categorical propositions in diagrams that forever bear his name: Venn diagrams.
- ▶ We start with a pair of overlapping circles, with the one on the left representing the subject class and the one on the right representing the predicate class. The area in the overlap is the part of the subject class that also is in the predicate class.
- ▶ The parts of the circles outside of the overlap are for members of the classes that belong only to one or the other class. We indicate that a region is empty by shading it in, and we represent that an area is not empty by putting an *X* in it.

- ▶ For categorical syllogisms, we just enter the information from the premises into a larger diagram that has three circles, one for the minor term on the bottom left, one for the major term on the bottom right, and one for the middle term, which is placed above and appropriately in the middle.
- ▶ We enter in the premises and see if what comes out contains the information within the conclusion. If it is, then the argument is valid. If not, it is invalid.

VENN DIAGRAM



READINGS

Barker, *The Elements of Logic*, chap. 2.
Copi, *Introduction to Logic*, chaps. 5 and 6.

QUESTIONS

1.

Identify the form of the following categorical syllogisms.

- No dogs are blue. Some blue things are not fruit. Therefore, some fruit are dogs.
- All desserts are sweet. Some children are sweet. Therefore, some children are desserts.

2.

Use Aristotle's rules to determine if the following syllogism is valid.

Some movies are not comedies.

Some movies are not romances.

Therefore, some romances are not comedies.

3.

Use Aristotle's rules to determine if the following syllogism is valid.

No food is poisonous.

Some mushrooms are poisonous.

Therefore, some mushrooms are not food.

4.

Use Venn diagrams to show whether the two categorical syllogisms in questions 3 and 4 are valid.

ANSWERS

1.

a *EOI-4*

b *AII-2*

2.

In *OOO-1*, the predicate of each sentence is distributed. The first rule requires that the middle term be distributed in one of the premises. *O* sentences distribute the predicate, but the middle term is the subject in both premises and therefore is undistributed. Because the middle term is not distributed, by the first rule, the argument is invalid.

3.

In *EIO-1*, the minor premise is an *E* sentence, so both terms, “food” and “poisonous,” are distributed. The major premise is an *I* sentence, and nothing is distributed in an *I* sentence. The conclusion is an *O* sentence, which distributes the predicate—in this case, “food.” The first rule requires that the middle term be distributed in one of the sentences, and it is in the minor premise. The second rule requires that any term distributed in the conclusion be distributed in the premises. The term “food” is distributed in the conclusion and also in the minor premise. The third rule requires that at least one of the premises be affirmative, and the major premise satisfies that one. The fourth rule requires that if the conclusion is negative, then one of the premises is, and we do have a negative conclusion and a negative minor premise. Finally, according to the fifth rule, if the conclusion is particular—which it is—then one of the premises must be particular, and the major premise is. So, the argument is valid.

Truth-Functional Logic

Aristotle's categorical logic was the first attempt at creating a rigorous calculus of human thought, a surefire means of determining when a deductive argument is valid. The problem with Aristotle's system is that it only works for sentences of categorical form; for non-categorical sentences, the tool kit of Aristotelian logic will not let us do what we need to do. We need a different system—one that can handle a wider range of propositions that we use in deductive arguments. This system is called truth-functional logic.

THE NEED FOR TRUTH-FUNCTIONAL LOGIC

- ▶ At the end of the 19th century and beginning of the 20th century, truth-functional logic was developed by philosophers who found a need for a new language. Philosophy has to be done with words. But the words we use in everyday, ordinary language are ambiguous and vague. They are not defined precisely enough to do the work that philosophers need.
- ▶ Instead of doing philosophy in ordinary language, they thought, they would develop a new artificial language, one that the exactitude they required was built into its very grammatical structure.
- ▶ It would resemble spoken language enough that we could translate our philosophical questions and intuitions into it, but it would be strict enough that we could finally answer some of these questions.
- ▶ We could see whether they were real questions or merely pseudo-questions, and if they were real questions, we could determine

their truth conditions—that is, what we would have to look for to know if they were true or false. The German mathematician/logician/philosopher Gottlob Frege called his attempt at framing such a language *Begriffsschrift*, or concept writing. It ultimately became truth-functional logic.

- ▶ In the phrase “truth-functional,” we use the word “function” in the same way that mathematicians do. Mathematicians look at an algebraic equation—for example, $y = x + 2$ —and say that y is a function of x . What this means is that for every value of x , there is a single, unique value for y . For example, if x is 1, then y is 3.
- ▶ A function takes in something and then spits out a well-defined, completely determined something else. Numerical functions take in numbers or pairs of numbers and spit out a single number. Instead of numbers, truth-functions take in truth-values and spit out a single truth-value.
- ▶ In classical logic, there are two truth-values: true and false. Every sentence has one and only one truth-value. If a sentence is true, it is not false. If it is not true, then it has to be false. Truth-functional logic is a two-valued system, and every sentence has one or the other of these values.
- ▶ For now, it doesn't matter how we know which it is or whether there are some sentences whose truth-values we don't know how to ascertain. All we are concerned with is *if* these sentences are true, does that mean some other sentence also has to be true?
- ▶ In fact, the content of the sentence will be completely irrelevant. All we are concerned with here is the form. When we translate spoken language into this language, we will remove everything about the content of the propositions and strip it down to the bare skeleton—the logical structure. Whether one sentence follows from another depends only on its form, so content will go away for us.

THE ELEMENTS OF TRUTH-FUNCTIONAL LANGUAGE

- ▶ Truth-functional language has two elements: atomic sentences and truth-functional connectives. Atomic sentences are simple declarative sentences that are either true or false.
 - ◆ “The sky is blue” is atomic.
- ▶ This is a simple sentence that has a truth-value. We may or may not know the truth-value, but again, that is of no consequence. In our language, we use lowercase letters to represent atomic sentences.
- ▶ To our atomic sentences, we add truth-functional connectives. A connective is a word that joins atomic sentences together to form a more complex molecular sentence. The word “and,” for example, is a connective.
- ▶ We can take the sentence “The sky is blue” and the sentence “I am 12 feet tall” and use the connective “and” to create a whole new sentence: “The sky is blue and I am twelve feet tall.” Connectives just join atomic sentences to form new sentences.
- ▶ A connective is truth-functional if and only if the truth-value of the molecular sentence is completely and uniquely determined, knowing only the truth-values of the constituent atomic sentences and the definition of the connective. “And” is truth-functional.
- ▶ We define truth-functional connectives in terms of what is called a truth table. In a truth table, we start by setting out all of the atomic sentences involved, listing all of the possible combinations of truth-values beneath.
- ▶ “And” is a dyadic connective—that is, it joins two sentences together—so we need two sentences (it doesn’t matter what they say, just that they’re atomic): Let’s call them p and q . We start our

truth table by writing down a column for p and a column next to it for q .

p	q
T	T
T	F
F	T
F	F

- ▶ We then list all of the combinations of truth-values they could have. Because we have two sentences, each of which could have one of two truth-values, there are four possibilities (shown at right).
- ▶ After our atomic sentences are all entered in our truth table, we add a new column in which appears nothing but what has come before in the table and one new connective. We will use the symbol “&” as our symbol for “and,” although some logicians prefer the wedge (\wedge) or a dot (\cdot).
- ▶ Next is the important step: filling in the value for the molecular sentence. In truth-functional logic, what we mean by “and” is “and.”

- ▶ Think about how we use the word “and.” Suppose that you have a friend Bob over at the coffee machine and you say, “Bob has sugar in his coffee and Bob has cream in his coffee.” When is that sentence true, and when is it false? All the cases are already in our truth table, so let’s take them one by one.

p	q	$p \& q$
T	T	T
T	F	F
F	T	F
F	F	F

- ▶ The first case has both sentences true. If it is true that Bob has sugar in his coffee “and” it is true that Bob has cream in his coffee, then what do we know about “Bob has sugar in his coffee and Bob has cream in his coffee”? We know it is true.
- ▶ “And” sentences are true when both sentences joined are true. So, in the first row of the third column in our truth table, we put a “T.”

- ▶ Suppose that he does have sugar but doesn't have cream? What do we say about the sentence "Bob has sugar in his coffee and Bob has cream in his coffee"? It is false.
- ▶ Similarly, if he has cream but no sugar, for an "and" sentence to be true, both constituent atomic sentences must be true. And if neither are true, then it's false.
- ▶ "And" is truth-functional because it is possible to construct this truth table; that is, knowing just the truth-values of p and q , we can uniquely determine the value for $p \& q$ in every case.

p	q	$p \& q$
T	T	
T	F	
F	T	
F	F	

- ▶ Not all connectives are truth-functional. Consider the word "because." It is a connective. But it is not truth-functional. We cannot construct a truth table for it. Let's try. We'll use a greater than ($>$) symbol to mean "because."
- ▶ What goes in the first row of the third column? If p is true and q is true, do we know if " p because q " is true? Let p be "The sky is blue" and let q be "Albert Einstein was a physicist."
- ▶ Both p and q are true. Is it also true that the sky is blue because Albert Einstein was a physicist? No. To determine the truth-value of "because" sentences, we need to know more than just the truth-values of the constituent atomic sentences. So, "because" will have no place in our truth-functional language as a connective.
- ▶ We will, in fact, have four truth-functional connectives. The first is negation, which is just our way of saying "not." It is our only monadic, or one-place, connective.

- ▶ We will use the minus sign (\neg) to represent it, although others use the tilde (\sim) or a minus sign with a little nib on the end ($\bar{\neg}$). Its truth table is shown at right.

p	$\neg p$
T	F
F	T

- ▶ Negation flips the truth-value of the sentence it is applied to.
- ▶ In spoken language, in addition to the simple word “not,” we also use phrases like “it’s not the case that” or “it is false that” to denote negation.

- ▶ The second is conjunction, which is logic terminology for “and.” We have already seen its truth table.

p	q	$p \& q$
T	T	T
T	F	F
F	T	F
F	F	F

- ▶ In spoken language, we use the word “and” as well as “as well as,” “in addition to,” and “also.” One that fools a lot of people is “but,” which in many contexts is equivalent to “and.” “I intended to call you back, but I got distracted.” In this sentence, “but” just means “and.” The sentence is true if and only if both atomic sentences are true.

- ▶ The third connective is the disjunction, “or.” For the disjunction, we will use the universal symbol \vee , which comes from the word *vel*, which means “or” in Latin.

- ▶ Let’s see if we can construct the truth table. It is another dyadic connective, so the table starts the same.

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

- ▶ Let’s see if we can fill in the last column. We’ll start with the middle two rows because they are the easiest. Let’s think about our friend Bob at the coffee machine.

- ▶ Suppose that the sentence we want to consider is “Bob has sugar or cream in his coffee.” Sentence p is “Bob has sugar in his coffee,” sentence q is “Bob has cream in his coffee,” and the disjunction is “Bob has sugar in his coffee or Bob has cream in his coffee.”
- ▶ Let’s take the second case, in which he does have sugar but has no cream. What do we know about “Bob has sugar or cream in his coffee”? It is true. Similarly, if he has cream, but no sugar, the disjunction “Bob has sugar or cream in his coffee” is true. If one or the other is true, then the “or” sentence is true.
- ▶ What about the last row? Suppose that Bob has neither sugar nor cream in his coffee. What do we know about the sentence “Bob has sugar or cream in his coffee”? It is false.
- ▶ For the top row, suppose that Bob has both sugar and cream in his coffee and you say, “Bob has sugar or cream in his coffee.” Are you right or wrong? Is the sentence true or false?
- ▶ We use the word “or” to mean both “one, the other, but not both” and “one, the other, or both.” If you are having coffee, you might ask, “Sugar or cream?” and your friend could reasonably reply, “Yes, both please.” But if your friend is having tea and you offer, “Lemon or cream?” then clearly you mean which of the two, not both.
- ▶ One is what we call the inclusive sense of “or,” and the other is the exclusive “or.” Which one do we use? It is arbitrary, because once we pick one, we can use the other connectives to build the other.
- ▶ It doesn’t go away; it just becomes a little more cumbersome to write. So, purely as a convention, we will choose the inclusive “or” because we prefer to be inclusive. Logic is for everyone.

- ▶ So, the truth table for disjunction, then, is as shown at right.

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

- ▶ In spoken language, we often include the term “either” when we are signaling to our audience that we mean the exclusive “or” and use just “or” to mean the inclusive.

- ▶ “We should either go all the way or not start at all” implies that there are two mutually exclusive possibilities and we have to select only one. “If you like that shade of purple, paint the upstairs or the downstairs powder room in that color” is inclusive, leaving open the chance to do both.

- ▶ Another way of saying “or” is “unless.” Just like “or,” we have inclusive and exclusive senses of “unless.” “When the call comes in, page me unless I’m in a meeting” is exclusive. “I’ll have just salad, unless they have that soup I love” is inclusive. But in both cases, unless is “or.”

- ▶ The last connective is the conditional, which is our “if, then” sentences. The sentence that is the “if” clause is called the antecedent, and the sentence that is the “then” clause is the consequent. We will write the conditional using the right-pointing arrow (\rightarrow); others use the sideways horseshoe (\supset).

- ▶ In this truth table, we have a dyadic connective—two atomic sentences are being joined—so it starts off like the last two.

p	q	$p \rightarrow q$
T	T	
T	F	
F	T	
F	F	

- ▶ Let’s address the last column. It will help to think of this in terms of an example. Suppose that you make a claim to your friend Jane in the form of

a conditional—for example, “If you beat me in backgammon, I will give you my car.” When do you lie to Jane? When can she take you to court for failing to uphold a contract?

- ▶ In the first case, the antecedent p and the consequent q are both true. Jane did beat you at backgammon, and you did give her your car. Did you lie? No. So, the first row of the third column gets a “T.”
- ▶ For the other easy case, skip down to the bottom line. The antecedent and the consequent are both false. Jane did not beat you at backgammon, and you did not give her your car. Did you lie? No. In two-valued logic, any sentence that is not false must be true. Because this is not false, it is thus true.
- ▶ Consider the third line. Jane did not beat you at backgammon, but you gave her your car anyway. Did you lie? No. You never said that beating you at backgammon was the only way to get your car. Maybe Jane bought it from you, or maybe you were feeling generous and just gave it to her. Regardless of the how and why, Jane did not beat you at backgammon but did get the car.
- ▶ You did not violate your agreement, so your “if, then” sentence is not false and therefore must be true.
- ▶ Consider the second line. Here, Jane did beat you at backgammon, and you refused to give her your car. Does she have warrant for complaint? Absolutely.
- ▶ This is the only case in which the conditional is false: when the antecedent is true and the consequent is false.
- ▶ In spoken language, we have a number of ways to express the conditional. “If you put your hand on a hot stove, then you will feel pain”: $s \rightarrow p$. We drop the

p	q	$p \rightarrow q$
T	T	T
T	F	F
F	T	T
F	F	T

word “then”: “If you put your hand on a hot stove, you will feel pain”: $s \rightarrow p$. We change the order: “You will feel pain if you put your hand on a hot stove”: still $s \rightarrow p$.

- ▶ The word “if” picks out the antecedent wherever it is: in front, in the middle, etc.; wherever you see the word “if,” you have found the antecedent.
- ▶ It is crucial when translating a conditional to correctly identify the antecedent and consequent. Notice a difference in the truth tables of the conjunction and disjunction, on the one hand, and the conditional on the other.
- ▶ Notice that for “and” and “or,” they are symmetric. Paper “or” plastic; plastic “or” paper. Logically, there is no difference. The truth-value won’t change, regardless of the order. But this is not so with the conditional, where the order does matter.
- ▶ Consider the following two sentences: “If you get shot, then you bleed” and “If you bleed, then you get shot.” These are two very different sentences. The first is true, and second is false. So, for a conditional, unlike for a conjunction or disjunction, order matters.
- ▶ So, we look for the word “if” to pick out our antecedent. There are other stylistic variants for “if.” “When” sometimes means “if.” “Yes, I will go to dinner and a movie with you...when pigs fly.” In other words, if pigs fly, feel free to make a reservation for two. “Given that” and “on the condition that” are other ways of saying “if.”
- ▶ The interesting one is “only if.” We said that wherever you find the word “if,” it picks out your antecedent. The only exception is when the word “if” is paired with “only.” “Only if” always picks out the consequent.
- ▶ Think about the difference between the sentences. “There is fire only if there is oxygen.” This means that if there is fire, then there

is oxygen: $f \rightarrow o$. It does not mean that if there is oxygen, then there is fire.

- ▶ The sentence is true; there is oxygen around you but no fire. So, it could not be translated as $o \rightarrow f$. So, “if” picks out the antecedent wherever it is, unless it is “only if,” in which case it picks out the consequent.

READINGS

Barker, *The Elements of Logic*, chap. 3.

Copi, *Introduction to Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 2.

QUESTIONS

1.

Consider the connective “it is the case that,” for which we can use the symbol $=$. $=$ is a one-place connective—that is, it works on a single sentence (for example, $=p$). Is $=$ truth-functional?

2.

Translate the following into truth-functional logic.

- a If I am here and you are here, then all is well in the world.
- b Unless it is raining or too hot, we will go to the park and have a picnic.
- c Because I am hungry and you are tired, things will go badly. If we go home, I will not be hungry and you will not be tired. Things will not go badly this time. We are going home.

ANSWERS

1.

We can construct a truth-table for $=$, as shown at right.

p	$=p$
-----	------

That means that we can determine whether $=p$ is true knowing only whether p is true or false. So, $=$ is a truth-functional connective (albeit not a terribly interesting one).

T	T
---	---

F	F
---	---

2.

- Let i represent "I am here," u represent "You are here," and a represent "All is well in the world": $(i \& u) \rightarrow a$.
- Let r represent "It is raining," h represent "It is too hot," p represent "We will go to the park," and e represent "We will have a picnic": $(r \vee h) \vee (p \& e)$.
- Let h represent "I am hungry," t represent "You are tired," b represent "Things go badly," and m represent "We go home":

$$(h \& t) \rightarrow b$$

$$m \rightarrow (-h \& -t)$$

$$-b$$

Therefore, m .

Truth Tables

Truth-functional logic is an artificial logical language designed to show us what sentences follow from what other sentences. And as you will discover in this lecture, truth tables are amazing logical instruments. We can use them to answer a wide range of questions that we might have about truth-functional sentences. And the answers they give are absolute and rigorous. With truth tables, we have an algorithmic method for determining logical properties.

TRUTH TABLES FOR CONNECTIVES

- ▶ Atomic sentences are basic sentences that are either true or false, such as “The sky is blue.” Connectives are words that connect atomic sentences together into larger, more complex molecular sentences.
- ▶ Connectives are said to be truth-functional if and only if the truth-value of the molecular sentence created using the connective can be completely and uniquely determined knowing nothing more than the truth-values of the constituent atomic sentences and the definition of the connective.
- ▶ This is equivalent to saying that we can construct a truth table for the connectives. A truth table contains a row for each of the possible arrangements of truth-values for all of the constituent atomic sentences and a listing of the resulting truth-values for molecular sentences.

- ▶ It is very nice that we can define our connectives using truth tables, but this is actually their most trivial use. We can construct truth tables for any truth-functional sentence, and when we do, they tell us the truth conditions for that sentence—meaning what must be the case in the world for the sentence to be true and what can be the case in the world that renders the sentence false.
- ▶ For example, consider the following sentence in our truth-functional language: $(p \vee q) \& \neg p$. The first step in constructing a truth table is to list all of the atomic sentences in the object sentence. For this example, it's just p and q .
- ▶ Each new column must contain only what already appears in the truth table plus one new truth-functional connective. We do this by finding the main connective and building up the parts it connects.
- ▶ In our sentence, $(p \vee q) \& \neg p$, it is the conjunction that applies to the rest of the sentence. The first conjunct is $(p \vee q)$. It is not in the table, but p is and q is, so by adding one new connective, \vee , we can make it.
- ▶ Once one side of the conjunction is in the table, we can work on the other side, $\neg p$, which is not in the table, but p is, so by adding one connective, \neg , we can build it.
- ▶ We now have what is on the left side of the “and” and what is on the right side of the “and” in our table, so we can add a column that joins them.
- ▶ Then, we can add the possible combinations of truth-values for p and q : T-T, T-F, F-T, F-F.
- ▶ We know that $p \vee q$ is true when p is true, q is true, or both are true. And $\neg p$ always has the opposite truth-value of p .

- ▶ Finally, we can fill in the values for the last column using the values from the second and third columns.

p	q	$p \vee q$	$\neg p$	$(p \vee q) \& \neg p$
T	T	T	F	F
T	F	T	F	F
F	T	T	T	T
F	F	F	T	F

- ▶ It is $p \vee q$ “and” $\neg p$, and we know that “and” sentences are only true when both conjoined sentences are true. So, we see that the only time the sentence $(p \vee q) \& \neg p$ is true is when p is false and q is true.

TYPES OF TRUTH-FUNCTIONAL SENTENCES

- ▶ By determining the truth conditions of a truth-functional sentence, we can tell what type of sentence it is. Truth-functional sentences come in three categories: tautologies, contradictions, and contingencies.
- ▶ A tautology is a sentence that is always true. For example, “It is raining or it is not raining.” That is a sentence that is true no matter what the weather is doing. We can tell when a sentence is a tautology using truth tables by looking at the column beneath it and seeing if it is a complete line of nothing but Ts.
- ▶ By contrast, a contradiction is a sentence that is always false. For example, “It is raining and it is not raining.” Because an “and” sentence is only true when both conjoined sentences are true and because the negation of a sentence always has the opposite truth-value, “It is raining and it is not raining” must be

false, because when one side is true, the other will be false. In a truth table, this is seen when the column beneath the sentence is nothing but Fs.

- ▶ A contingency is a sentence whose truth-value is contingent on how the world is. Sometimes it is true; sometimes it is false. A sentence is a contingency if it is neither a contradiction nor a tautology—that is, if there is at least one set of truth-values for the constituent atomic sentences that renders the sentence true and at least one set of truth-values for the constituent atomic sentences that renders the sentence false.
- ▶ The sentence we just used as an example, $(p \vee q) \& \neg p$, is a contingency because its column has at least one T and at least one F.

CONSTRUCTING TRUTH TABLES FOR MULTIPLE SENTENCES

- ▶ We can construct truth tables for any sentence that allows us to determine the sentence's truth conditions, and this tells us what type of sentence it is: tautology, contradiction, or contingency. But we can construct truth tables that include multiple sentences as well. This allows us to use truth tables to determine whether certain relations between sentences hold.
- ▶ One such relation is truth-functional equivalence. Two truth-functional sentences are logically equivalent if and only if they have the same truth conditions; that is, they always have the same truth-value—both are true or both are false—no matter the state of the world.
- ▶ We test for this by constructing a truth table that includes both sentences and see if the arrangement of Ts and Fs in the two columns are exactly the same.

- ▶ Is the sentence $\neg a \& \neg b$ equivalent to $\neg(a \& b)$? Using algebra, can you distribute the negation through the sentence? We need a truth table that includes both sentences.
- ▶ First, list the atomic sentences: a and b .
- ▶ Next, let's build our first sentence. The main connective in $\neg a \& \neg b$ is the conjunction. It conjoins $\neg a$ with $\neg b$, neither of which is in the table, but both of which can be made by adding one connective to something that is already in the table.
- ▶ Once we have $\neg a$ and $\neg b$, we can make $\neg a \& \neg b$. That is one of our sentences. Now, let's build the other. The main connective in $\neg(a \& b)$ is the negation. It applies to the rest of the sentence. So, we need $a \& b$. It isn't in our table, but a is and b is, so we can make it.
- ▶ Once we have $a \& b$, we just negate it in the last column and we have both sentences we are looking to compare.
- ▶ Next, we enter the truth-values, just like last time: T-T, T-F, F-T, F-F.
- ▶ Because $\neg a$ is the negation of a , and $\neg b$ is the negation of b , the next two columns are as follows.
 - ◆ $\neg a$: F, F, T, T.
 - ◆ $\neg b$: F, T, F, T.
- ▶ The next column is the conjunction of the previous two. Note that it is only true when both $\neg a$ and $\neg b$ are true: F, F, F, T.
- ▶ The next column is just the conjunction of a and b : T, F, F, F.
- ▶ The last column is just the negation of the values we just determined: F, T, T, T.

a	b	$\neg a$	$\neg b$	$\neg a \& \neg b$	$a \& b$	$\neg(a \& b)$
T	T	F	F	F	T	F
T	F	F	T	F	F	T
F	T	T	F	F	F	T
F	F	T	T	T	F	T

- ▶ So, to see if the two sentences are equivalent, let's see if they have exactly the same arrangement of Ts and Fs beneath them. Look at the second and third rows of $\neg a \& \neg b$ and $\neg(a \& b)$. They are not equivalent.
- ▶ But what would happen if instead of $\neg(a \& b)$, we compare $\neg a \& \neg b$ to $\neg(a \vee b)$? Let's change the conjunction from $\&$ to \vee in the appropriate columns. We know what $a \vee b$ looks like. It's only false when both a and b are false. And the last column is just the negation of the one we just determined.

a	b	$\neg a$	$\neg b$	$\neg a \& \neg b$	$a \vee b$	$\neg(a \vee b)$
T	T	F	F	F	T	F
T	F	F	T	F	T	F
F	T	T	F	F	T	F
F	F	T	T	T	F	T

- ▶ This time, we compare the two columns, $\neg a \& \neg b$ and $\neg(a \vee b)$, and they are exactly alike. Therefore, $\neg a \& \neg b$ is truth-functionally equivalent to $\neg(a \vee b)$.
- ▶ Intuitively, this makes sense. If we say, "Bob doesn't have sugar *and* Bob doesn't have cream in his coffee," that is synonymous to saying, "Bob doesn't have sugar *or* cream in his coffee."

- ▶ This is exactly what we want. The results of truth-functional logic match those of our natural way of reasoning, but we have a precise method of demonstrating it.
- ▶ In addition to equivalence, another relation between sentences is consistency. Two sentences are consistent if it is possible for them both to be true at the same time—that is, if there is some arrangement of the constituent atomic sentences such that both sentences are rendered true. Notice that this is an incredibly weak relation: “The sky is blue” is consistent with “I am a warthog.”
- ▶ We can see this using a truth table. “The sky is blue” is an atomic sentence. Let’s use s to abbreviate it. “I am a warthog” is also atomic, and we’ll use w for it.
- ▶ The truth table is trivial to set up. There are two atomic sentences, and we put them in the table.
- ▶ Next, we add the truth-values.
- ▶ Because the first line is a case in which both sentences are true, they are consistent. “The sky is blue” is also consistent with “I am not a warthog.”
- ▶ Next, we need an additional column for $\neg w$. We know that $\neg w$ has the opposite truth-value as w , which gives us what is shown at right.
- ▶ In the second row, s is true and $\neg w$ is true. So, the two are also consistent.
- ▶ Indeed, “The sky is blue” is consistent with pretty much everything except “The sky is not blue.”

s	w	$\neg w$
T	T	F
T	F	T
F	T	F
F	F	T

- ▶ To tell if two truth-functional sentences are consistent, we construct a truth table with both of them in it and see if there is any row—in which both have the truth-value T.
- ▶ There is a third relation between sentences called implication. A sentence S_1 truth-functionally implies sentence S_2 if and only if whenever S_1 is true, S_2 is also true—that is, there is no case in which S_1 is true and S_2 is false. This does not mean that they both have all of the same truth-values—that is, equivalence. Rather, S_1 implies S_2 when S_1 's truth is enough to guarantee S_2 's truth.
- ▶ Implication is unlike equivalence and consistency in that it is not necessarily symmetric—that is, just because S_1 implies S_2 does not mean that S_2 implies S_1 . Indeed, the only time we will have mutual implication is when we have two sentences that are equivalent.
- ▶ For example, consider the following two sentences: “If today is Sunday, then I’ll visit my parents, unless today is not Sunday and I will visit my parents.” “Today is Sunday, or if I visit my parents, then it is not Sunday.”
- ▶ The first sentence is composed of two atomic sentences: “Today is Sunday” and “I will visit my parents.” Let’s use x for the first and y for the second: “If x then y , unless it is not the case that x and y . The comma shows us that the main connective is “unless”—which is another way of saying “or.” So, we have the following.
 - ◊ $(\text{if } x, \text{ then } y) \vee (\text{it is not the case that } x \text{ and } y)$
- ▶ The first disjunct is just a conditional.
 - ◊ $(x \rightarrow y) \vee (\text{it is not the case that } x \text{ and } y)$
- ▶ The second disjunct is a conjunction.
 - ◊ $(x \rightarrow y) \vee (\text{it is not the case that } x \& y)$

- ▶ It is not the case that x is just $\neg x$. So, the first sentence is as follows.
 - ◊ $(x \rightarrow y) \vee (\neg x \& y)$
- ▶ The second sentence is “Today is Sunday, or if I visit my parents, then it is not Sunday.” Substitute in the x and y while leaving the logical terms in place: “ x , or if y , then not x .”
- ▶ The main connective is the disjunction.
 - ◊ $x \vee (\text{if } y \text{ then not } x)$
- ▶ The second disjunct is a conditional.
 - ◊ $x \vee (y \rightarrow \text{not-}x)$
- ▶ The consequent is a negation, giving us the following.
 - ◊ $x \vee (y \rightarrow \neg x)$
- ▶ So there are our sentences: $(x \rightarrow y) \vee (\neg x \& y)$ and $x \vee (y \rightarrow \neg x)$.
- ▶ What type of sentence is each? Are they equivalent? Are they consistent? And does either imply the other?
- ▶ First, we need a truth table for two atomic sentences, x and y .
- ▶ The first sentence has the disjunction, the “or,” as its main connective. On the left side is $x \rightarrow y$. The right side is $\neg x \& y$. We have y , but not $\neg x$, but we do have x , so we can make $\neg x$.
- ▶ Next, we can make $\neg x \& y$.
- ▶ By adding the \vee , we can complete our first sentence: $(x \rightarrow y) \vee (\neg x \& y)$.

- ▶ In the second sentence, $x \vee (y \rightarrow \neg x)$, the main connective is the “or.” The left side is just x , which is in the table already, and the right side is the conditional $y \rightarrow \neg x$, and we already have y and $\neg x$ in the table. So, we need one more column for that, as well as a final column for the second sentence.
- ▶ The next step is to add the truth-values.
- ▶ We know what the conditional ($x \rightarrow y$) looks like.
- ▶ The fourth column is just the negation of the first.
- ▶ The fifth column is just the fourth conjoined with the second.
- ▶ The sixth column, the first sentence, is the third column “or” the fifth column.
- ▶ The seventh column has the second as the antecedent and the fourth as the consequent.
- ▶ The last column, the second sentence, is the first column “or” the seventh column.

x	y	$x \rightarrow y$	$\neg x$	$\neg x \& y$	$(x \rightarrow y) \vee (\neg x \& y)$	$y \rightarrow \neg x$	$x \vee (y \rightarrow \neg x)$
T	T	T	F	F	T	F	T
T	F	F	F	F	F	T	T
F	T	T	T	T	T	T	T
F	F	T	T	F	T	T	T

- ▶ When we look at the first sentence, we see that underneath it is T, F, T, T.

- ▶ The second sentence has nothing but Ts beneath it. It is a tautology. Those two columns are not the same, so the two sentences are not equivalent.
- ▶ In the first, third, and fourth rows, both sentences have the truth-value T; remember that we only need one, so they are consistent.
- ▶ Let's think carefully about implication: Is there a case in which the first sentence is true and the second sentence is false? No, because the second sentence is never false.
- ▶ The first sentence does imply the second. Indeed, if the second sentence is a tautology, then any sentence would imply it, because it would be impossible for the first sentence to be true in a case where the second sentence is false, given that the second sentence is never false.
- ▶ What about the other direction? Is there a case in which the second sentence is true and the first one is false? Yes, this occurs in the second row. So, the second sentence does not imply the first.

READINGS

Barker, *The Elements of Logic*, chap. 3.

Copi, *Introduction to Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 2.

QUESTIONS

1.

Use a truth table to determine whether the sentence $p \rightarrow [q \vee (-q \rightarrow p)]$ is a tautology, contradiction, or contingency.

2.

Translate the following sentence, construct a truth table for it, and determine if it is a tautology, contradiction, or contingency.

“If you get a pie, then pick up ice cream, but if you don’t get ice cream, don’t get a pie.”

3.

Consider the following two sentences:

“I will pick up the kids from school, you will pick up the kids from school, or we will both pick up the kids from school”

and

“If you don’t pick up the kids from school, I will.”

Are they consistent? Are they equivalent? Does one imply the other?

4.

Consider the following two sentences:

“If you are not in love, don’t get married”

and

“If you are in love, get married.”

Are they consistent? Are they equivalent? Does one imply the other?

ANSWERS

1.

tautology

p	q	$\neg q$	$\neg q \rightarrow p$	$q \vee (\neg q \rightarrow p)$	$p \rightarrow [q \vee (\neg q \rightarrow p)]$
T	T	F	T	T	T
T	F	T	T	T	T
F	T	F	T	T	T
F	F	T	F	F	T

2.

contingency

$(p \rightarrow i) \& (\neg i \rightarrow \neg p)$

p	i	$p \rightarrow i$	$\neg i$	$\neg p$	$\neg i \rightarrow \neg p$	$(p \rightarrow i) \& (\neg i \rightarrow \neg p)$
T	T	T	F	F	T	T
T	F	F	T	F	F	F
F	T	T	F	T	T	T
F	F	T	T	T	T	T

3.

The sentences are consistent, equivalent, and imply each other. $(i \vee u) \& (i \& u)$ (this could also be translated simply as $i \vee u$), $\neg u \rightarrow i$.

i	u	$i \vee u$	$i \& u$	$(i \vee u) \vee (i \& u)$	$\neg u$	$\neg u \rightarrow i$
T	T	T	T	T	F	T
T	F	T	F	T	T	T
F	T	T	F	T	F	T
F	F	F	F	F	T	F

4.

The sentences are consistent. They are not equivalent, and there is no implication.

$\neg l \rightarrow \neg m$, $l \rightarrow m$

l	m	$\neg l$	$\neg m$	$\neg l \rightarrow \neg m$	$l \rightarrow m$
T	T	F	F	T	T
T	F	F	T	T	F
F	T	T	F	F	T
F	F	T	T	T	T

Truth Tables and Validity

Truth-functional logic provides us with the tools to assess arguments surrounding many of the things we believe. Does the evidence we cite really lead us with absolute certainty to the conclusions we hold? To find out whether they do, we can use truth tables. As you know, they can be used to test for properties of individual sentences and relations between sentences. But our ultimate goal in logic is to test for validity—that is, to see when we are logically obligated to believe an argument's conclusion based on the information contained in the argument's premises. As you will learn in this lecture, truth tables are able to do this, too.

VALIDITY AND INVALIDITY FOR DEDUCTIVE ARGUMENTS

- ▶ In order to understand the means of testing for validity and invalidity, recall a relation between truth-functional sentences that you have already learned how to test for: implication.
 - ◆ A sentence S_1 implies a sentence S_2 if and only if in every case in which S_1 is true, S_2 is also true—that is, there is no case in which S_1 is true and S_2 is false. We test for implication by constructing a truth table containing both S_1 and S_2 . We then look at every row in which S_1 is true and see if S_2 is also true in every case.
- ▶ Validity is just a generalization of implication. An argument is valid if and only if its premise set implies its conclusion. As such, we can just generalize our truth table test to turn it into a validity test—and, interestingly, an invalidity test.

- ▶ It might seem obvious that we should be able to use a single test to demonstrate that an argument is valid if it passes and invalid if it fails, but there is another validity test that is not an invalidity test. But truth tables are able to do both. How?
 - 1 Construct a big truth table that includes all of the premises and the conclusion.
 - 2 Fill in all of the truth-values.
 - 3 Look for every case—that is, every row—in which all of the premises are true.
 - 4 See if the conclusion is true in all of those cases. If it is, then the argument is valid. If there exists even one case in which all of the premises are true and the conclusion is false, then the argument is invalid.

MODUS PONENS

- ▶ Consider the famous argument called modus ponens, which is Latin for “the way that affirms.” It is any argument of the form $m \rightarrow n$; m , therefore, n .
- ▶ For example, consider the following argument: “If you are drunk, then your cognitive abilities are impaired. You are drunk. Therefore, your cognitive abilities are impaired.”
- ▶ Let m be the sentence “you are drunk” and n be the sentence “your cognitive abilities are impaired.”
- ▶ We need a truth table that includes the premises and the conclusion. The first step in building a truth table is to list all of the atomic sentences we will need. In this case, it is m and n .
- ▶ Next, let’s look at the first premise: $m \rightarrow n$. It’s not in our table, but we have both m and n , so by adding just one connective, we can make it.

- ▶ We look at the second premise, m , and see that it is already in the table, as is the conclusion, n .
- ▶ Next, add the truth-values. This is just the truth table for the conditional.
- ▶ Once we have the complete truth table, we can test the argument for validity. In which rows are both of the premises true? It is only the first row. Is the conclusion also true in this case? Yes. So, the argument is valid.
- ▶ If it is true that you are drunk and if it is true that being drunk leads to impaired cognitive abilities, then it must also be true that your cognitive abilities are diminished.

m	n	$m \rightarrow n$
T	T	T
T	F	F
F	T	T
F	F	T

THE FALLACY OF AFFIRMING THE CONSEQUENT

- ▶ Modus ponens's evil twin is known to logicians as the fallacy of affirming the consequent. It is any argument of the form $m \rightarrow n$; n , therefore m .
- ▶ Using the same sentences as the previous example, the argument would be as follows: "If you are drunk, then your cognitive abilities are impaired. Your cognitive abilities are impaired. Therefore, you are drunk."
- ▶ The name probably gives away the validity status, but let's see whether this so-called fallacy really is an invalid argument form.
- ▶ We need a truth table that includes both of the premises and the conclusion. This one is easy for us because it is just the truth table we constructed for the previous example. The

only difference is that we have switched the conclusion and one of the premises.

<i>m</i>	<i>n</i>	<i>m</i> → <i>n</i>
T	T	T
T	F	F
F	T	T
F	F	T

- ▶ To test for validity, let's find every row in which both of the premises are true. In both of the first and the third rows, the conditional "if you are drunk, then your cognitive abilities are decreased" is true, and in both of these cases you do have decreased cognitive capacities.
- ▶ In the first case, you are drunk, but in the third case, you are not drunk but have decreased cognitive capacities for some other reason—perhaps lack of sleep, for example.
- ▶ Is the conclusion also true in both of those cases? It is for the first row, but not for the third row.
- ▶ Note what this means: If we are presented with this argument and we know for a fact that both of the premises are true—it is the case that being drunk diminishes your cognitive abilities, and you do, in fact, have diminished cognitive abilities—do we know with absolute certainty whether you are drunk? That is, do we know if the conclusion is also true?
- ▶ No, because we don't know if our world is the first case or the third case. In both of these rows, the premises are true, but this is not enough for us to know which one is our case and, therefore, not enough to know whether the conclusion is true or false.
- ▶ This argument is invalid. It does not give us reason to either believe or deny the conclusion. This argument gives us reason to believe absolutely nothing. The fallacy of affirming the consequent lives up to its name.

THREE ATOMIC SENTENCES

- ▶ The following example is a bit more complicated than the previous two. “If there is precipitation and the temperature is below freezing, then there is snow. There is not snow or the temperature is below freezing. Therefore, there is precipitation or there is snow.”
- ▶ Let a be the sentence “there is precipitation.” Let b be the sentence “the temperature is below freezing.” And let c be the sentence “there is snow.”
- ▶ This gives us the argument form $(a \& b) \rightarrow c; \neg c \vee b$, therefore $a \vee c$. There are three atomic sentences. How do we build a truth table for this? The same way as we do with two. The first step is to list all of the atomic sentences.
- ▶ Next, let’s work on the first premise. The main connective is the conditional. The antecedent is $a \& b$. We already have both a and b in the table, so we can make $a \& b$ by adding one connective.
- ▶ We have c , the consequent, already in the table, so by adding the arrow to what is already in the table, we can make the first premise: $(a \& b) \rightarrow c$.
- ▶ In the second premise, the main connective is the “or.” On the left side, we see $\neg c$. We have c , so one connective added makes it.
- ▶ Once we have the left side and the right side of the disjunction in the table, we can make the second premise: $\neg c \vee b$.
- ▶ Finally, we have a and c already in the table, so we can add one connective to make the conclusion: $a \vee c$.
- ▶ Once we set up the columns, we can add the truth-values. We have three atomic sentences that can take two possible truth-

values, so the number of combinations is $2 \times 2 \times 2 = 8$. There are 8 rows in each column. The first three columns are as follows.

<i>a</i>	<i>b</i>	<i>c</i>	<i>a</i> & <i>b</i>	(<i>a</i> & <i>b</i>) → <i>c</i>	¬ <i>c</i>	¬ <i>c</i> ∨ <i>b</i>	<i>a</i> ∨ <i>c</i>
T	T	T					
T	T	F					
T	F	T					
T	F	F					
F	T	T					
F	T	F					
F	F	T					
F	F	F					

- ▶ The fourth column is just the conjunction of the first two.
- ▶ The fifth column is a little tricky. Remember that the only time the conditional is false is when the antecedent is true and the consequent is false. But notice that in this case, the antecedent is listed after the consequent in the table, so we need to be mindful.
- ▶ For the sixth column, ¬*c* is just the negation of *c*.
- ▶ The seventh column is a disjunction of the sixth and second columns.
- ▶ Finally, the conclusion in the eighth column is the disjunction of the first and third columns.

a	b	c	$a \& b$	$(a \& b) \rightarrow c$	$\neg c$	$\neg c \vee b$	$a \vee c$
T	T	T	T	T	F	T	T
T	T	F	T	F	T	T	T
T	F	T	F	T	F	F	T
T	F	F	F	T	T	T	T
F	T	T	F	T	F	T	T
F	T	F	F	T	T	T	F
F	F	T	F	T	F	F	T
F	F	F	F	T	T	T	F

- Now that our truth table is filled in, let's find every case in which both of the premises are true. It is rows 1, 4, 5, 6, and 8. Is the conclusion true in all of these cases? Lines 6 and 8 show us cases in which the premises are both true and the conclusion is false. So, this is an invalid argument.

READINGS

Barker, *The Elements of Logic*, chap. 3.

Copi, *Introduction to Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 2.

QUESTIONS

1.

Is the following argument valid?

$p \rightarrow q, \neg p \vee \neg q$. Therefore, $p \vee q$.

2.

Determine whether the following argument is valid using a truth table.

Your mother will stay with us only if we kennel the dog and get all of the carpets in the entire house steam cleaned. I am not going through the whole hassle of getting the house prepared for steam cleaning again. So, tell your mother she is not staying with us.

ANSWERS

1.

invalid (row 4)

p	q	$p \rightarrow q$	$\neg p$	$\neg q$	$\neg p \vee \neg q$	$p \vee q$
T	T	T	F	F	F	T
T	F	F	F	T	T	T
F	T	T	T	F	T	T
F	F	T	T	T	T	F

2.

valid

$m \rightarrow (k \& c)$. $\neg c$. Therefore, $\neg m$.

<i>m</i>	<i>k</i>	<i>c</i>	<i>k & c</i>	$m \rightarrow (k \& c)$	$\neg c$	$\neg m$
T	T	T	T	T	F	F
T	T	F	F	F	T	F
T	F	T	F	F	F	F
T	F	F	F	F	T	F
F	T	T	T	T	F	T
F	T	F	F	T	T	T
F	F	T	F	T	F	T
F	F	F	F	T	T	T

Natural Deduction

Logic determines what we have good reason to believe. We have legitimate warrant to accept the conclusion of an argument if that argument is both well-grounded and valid. To protect ourselves, we need a validity test. For arguments in our truth-functional language, we have truth tables. In this lecture, you will be introduced to a different validity test: natural deduction proofs. Why introduce a second validity test when we already have a perfectly good one? Truth tables are not always user-friendly, and our truth-functional language is insufficient for all of the logical work we want to do.

NATURAL DEDUCTION PROOFS

- ▶ Think of a natural deduction proof as a game. First, look at the playing board. Then, learn the goal we need to accomplish in order to win and the rules we have to follow in seeking that goal. Finally, learn strategies for successful playing.
- ▶ When it comes to the playing board, a natural deduction proof has three columns. The first column is comprised of positive integers in ascending order. Every line has a number, which is the name of the line.
- ▶ The second column includes nothing but true sentences written in our truth-functional language. If a single false sentence gets into the second column, it could create logical havoc. The second column is, therefore, a very exclusive logical club.

- ▶ For a sentence to move past the velvet rope and gain admittance to the second column, it needs an identification card that vouches for its undeniable truth. That identification card is a justification, and it is placed in the third column.
- ▶ So, the game of natural deduction proof is played on a board that has numbered and justified true sentences written in our truth-functional language. First comes the number, then the sentence, and then the justification.
- ▶ There are several categories of justification, but the first one is premise.
- ▶ Natural deduction proof is a validity test, and an argument is valid if and only if, assuming the truth of the premises for the sake of argument, the conclusion must follow from them.
- ▶ We are trying to show that if the premises are true, then the conclusion has to be, too. So, we begin by admitting the premises to the second column. Being a premise is like being born to a wealthy family: You get privileges you did not have to earn, unlike others.
- ▶ All other sentences will need to earn the justification that will let them into the second column, but the premises are simply granted justification by virtue of the argument whose validity is in question.
- ▶ The game begins when all of the premises are entered onto the board as numbered lines, justified by premise. The game ends when the conclusion appears as a justified line of the proof.
- ▶ Because only true sentences are allowed in the second column, if we start with the premises and show that the conclusion follows, then we will have, indeed, proven the validity of the argument.

THE RULES OF INFERENCE

- ▶ To get us from the premises to the conclusion, we need some moves. The first set is rules of inference, which are simple valid argument forms. If we have lines in the proof that have the form of the premises, then we can write a new line in the proof of the form of the conclusion of this valid argument form.
- ▶ Because the form is valid, we know that if the premises are true, the conclusion must be and therefore can be entered into the second column.
- ▶ We have already seen one of these rules of inference in the previous lecture. Recall the argument form modus ponens—that is, $a \rightarrow b$; a , therefore b .
- ▶ If we have a conditional that we know is true and we also know that the antecedent is true, then we can conclude that the consequent must be true as well.
- ▶ We can use modus ponens if we have a conditional on some line of the proof and on any other line of the proof we have the antecedent of that conditional.
- ▶ If these two lines appear, then on a subsequent line, we can write the consequent of the conditional in the second column and next to it, in the third column, write “MP” for modus ponens and write the numbers of the two lines on which the conditional and antecedent are found.
- ▶ Let’s work out a proof for the following argument.

- 1 $a \rightarrow b$
- 2 $b \rightarrow c$
- 3 $c \rightarrow d$
- 4 $d \rightarrow e$

- 5 $e \rightarrow f$
- 6 $f \rightarrow g$
- 7 a therefore, g

- ▶ We start by listing all of the premises as numbered lines of the proof, justified by premise. The proof therefore begins as follows.

- 1 $a \rightarrow b$ premise
- 2 $b \rightarrow c$ premise
- 3 $c \rightarrow d$ premise
- 4 $d \rightarrow e$ premise
- 5 $e \rightarrow f$ premise
- 6 $f \rightarrow g$ premise
- 7 a premise

- ▶ We look at line 1, and we have a condition $a \rightarrow b$, and then we see on line 7 that we have a . If a , then b , and we know a , so on line 8 we can write the following.

- 8 b MP 1,7

- ▶ The justification must include the name of the rule of inference used and the lines on which it is being used. This says that we used modus ponens on lines 1 and 7 in order to get this line, line 8.
- ▶ But look at line 2. It says that if we have b , we get c , and our new line 8 gives us b . So, we can add the following.

- 9 c MP 2,8

- ▶ And this gives us lines 10 through 13.

- 10 d MP 3,9
- 11 e MP 4, 10
- 12 f MP 5, 11
- 13 g MP 6, 12

- ▶ At this point, let's look at what we have. The conclusion g showed up as a justified line of the proof. This means that the proof is complete. We win the game. We declare the argument valid by writing beneath the proof "QED," which stands for *quod erat demonstrandum*, or "which was to be demonstrated." With this proof, we have proven the validity of the argument.
- ▶ The less optimistic sibling of *modus ponens*—the way of affirming—is *modus tollens*, or the way of denying. If we have, on any line of a proof, a conditional, and if on any other line of the proof, we have the negation of the consequent, then we may write on a further line of the proof the negation of the antecedent, justified by *modus tollens* (MT). In other words, we are using the valid argument form.

$$\begin{array}{l}
 a \rightarrow b \\
 \neg b \\
 \text{Therefore, } \neg a
 \end{array}$$

- ▶ If a is true, then b must be true. But we know that b is not true, so a cannot have been, or else b would be. This is valid.
- ▶ The other rule for sentences that are conditionals is the hypothetical syllogism (HS). Recall that a syllogism is an argument with two premises. A hypothetical syllogism is an argument with two conditionals as premises.

$$\begin{array}{l}
 a \rightarrow b \\
 b \rightarrow c \\
 \text{Therefore, } a \rightarrow c
 \end{array}$$

- ▶ If we have two conditionals in a proof and the antecedent of one is the consequent of the other, we can combine them, eliminating the middle term. We justify the new line with "HS" and the line numbers of the operative conditionals.

- ▶ This gives us another way of producing a proof from the original example.

1 $a \rightarrow b$ premise
2 $b \rightarrow c$ premise
3 $c \rightarrow d$ premise
4 $d \rightarrow e$ premise
5 $e \rightarrow f$ premise
6 $f \rightarrow g$ premise
7 a premise
8 $a \rightarrow c$ HS 1,2

- ▶ So, we took the conditional on line 1 and the conditional on line 2 and squished them together. If we have a , then we get b . But b gives us c . So, a gives us c .
- ▶ But this can be done again and again.

1 $a \rightarrow b$ premise
2 $b \rightarrow c$ premise
3 $c \rightarrow d$ premise
4 $d \rightarrow e$ premise
5 $e \rightarrow f$ premise
6 $f \rightarrow g$ premise
7 a premise
8 $a \rightarrow c$ HS 1,2
9 $a \rightarrow d$ HS 3,8
10 $a \rightarrow e$ HS 4,9
11 $a \rightarrow f$ HS 5,10
12 $a \rightarrow g$ HS 6,11

- ▶ But we now have line 12 that tells us that if a , then g , and line 7 that gives us a . So, by modus ponens, we get the following

13 g MP 7,12

- ▶ This is not uncommon. Many arguments will have multiple ways of constructing a proof.
- ▶ So far, all of the examples have used atomic sentences as the antecedent and consequent, but this is not necessary. This was only done to make the examples easier.
- ▶ You can use modus ponens, modus tollens, and the hypothetical syllogism on lines that are as complex as you want, as long as the main connective is the conditional and the form holds.
- ▶ Modus ponens just requires a conditional on one line and the antecedent on another. They don't have to be atomic sentences but could be molecular sentences of any complexity.
- ▶ With the hypothetical syllogism, we match up the middle term, and the antecedent stays the antecedent and the consequent stays the consequent, even if the terms are complex molecular sentences.
- ▶ If we have some sentence on line m and another on line n , then we know that both are true because all of the sentences in the second column are true. But if they are both true, then the conjunction of the two must also be true.
- ▶ So, we can take the sentences on any two lines and put an “and” between them, justifying the move with the rule of inference called conjunction (Conj). Suppose that we have the following.

8 $(d \rightarrow h) \vee (e \& t)$

9 $i \& (f \vee g)$

- ▶ Then, we can make the following line.

10 $[(d \rightarrow h) \vee (e \& t)] \& [i \& (f \vee g)]$ Conj 8,9

- ▶ The conjunction rule allows you to join any two lines with an “and.”
- ▶ Conjunction introduces an “and.” To get rid of one, you use the rule called simplification (Simp).
- ▶ We know that the only time a conjunction $a \& b$ is true is when a is true and b is true. So, if you have a conjunction in your proof, you can break it down and write both conjuncts on their own lines, justified by simplification. Suppose that we are given the following.

8 $(s \vee r) \& (x \rightarrow t)$

- ▶ Then, we can write the following.

9 $s \vee r$ Simp 8

10 $x \rightarrow t$ Simp 8

- ▶ Use simplification anytime you can. You will often be able to do things with the parts that you cannot do with the whole. Break it down and free up as much logical fodder as you can.
- ▶ Be careful, however: We can only use simplification if the conjunction is the main connective. You cannot simplify parts of lines, only entire lines. In general, rules of inference must only be used on entire lines—that is, operate on the main connective of the sentences.
- ▶ We also have rules for introducing and removing disjunctions. To introduce one, we use the rule called addition (Add). Given any line, we can turn it into a disjunction with any sentence. Suppose that we are given the following.

23 t

- ▶ We can form the following.

24 $t \vee m$ Add 23

- ▶ Where did m come from? Why m ? If you don't like m , you can make it w or even $(c \& f) \vee (q \& \neg j)$. You can add any sentence to any other sentence just because you want to.
- ▶ But we said that the sentences in the second column have to be true or else really bad things happen. How can we just add any sentence?
- ▶ Think about disjunctions: "Or" sentences are true when one or the other (or both) disjuncts are true.
- ▶ While we can introduce any sentence we want—true or false—it will be stuck in the disjunction. It will not be able to be freed and appear on its own.
- ▶ To get a sentence out of a disjunction, we use the disjunctive syllogism (DS). If we know that Bob has sugar or cream in his coffee and we know he doesn't have cream, then we know he has sugar. If we know that Bob has sugar or cream in his coffee and we know he doesn't have sugar, then we know he has cream. That is the disjunctive syllogism.
- ▶ If we have a line whose main connective is an "or," and we know that one of the disjuncts is false, then the other one has to be true. It looks like the following.

5 $(c \& d) \vee f$
 6 $\neg f$
 7 $c \& d$ DS 5,6

- ▶ You can use the disjunctive syllogism if either disjunct is negated. The order does not matter.

- ▶ The last rule has three premises and is called the constructive dilemma (CD). It is an argument of the following form.

$a \rightarrow b$
 $c \rightarrow d$
 $a \vee c$
 Therefore, $b \vee d$

- ▶ The idea is that a takes you to b , and c takes you to d . You know that either a or c is true, which means that you will be led, then, to either b or d . So, if you have two unrelated conditionals, see if there is a disjunction of the antecedents. If so, you can use the constructive dilemma.

THE RULES OF INFERENCE

MP	$a \rightarrow b$; a , therefore b
MT	$a \rightarrow b$; $\neg b$, therefore $\neg a$
HS	$a \rightarrow b$; $b \rightarrow c$, therefore $a \rightarrow c$
Conj	a , b , therefore $a \& b$
Simp	$a \& b$, therefore a , therefore b
Add	a , therefore $a \vee b$
DS	$a \vee b$, $\neg a$, therefore b or $a \vee b$, $\neg b$, therefore a
CD	$a \rightarrow b$, $c \rightarrow d$, $a \vee c$, therefore $b \vee d$

READINGS

- Barker, *The Elements of Logic*, chap. 3.
 Copi, *Introduction to Logic*, chap. 8.
 Hurley, *Logic*, chap. 7.
 Kahane, *Logic and Philosophy*, chap. 4.

QUESTIONS

1.

Provide a proof for the following argument.

$p \vee (r \& q)$. $\neg p \& (q \rightarrow s)$. Therefore, $r \& s$.

2.

Translate the following argument into truth-functional logic and construct a proof to show that it is valid.

If you have an identification card showing that you are older than 21, then either you are of legal age or this ID is a fake. I see your card, but I know that you are not 21. Hence, this must be a fake ID.

ANSWERS

1.

1	$p \vee (r \& q)$	premise
2	$\neg p \& (q \rightarrow s)$	premise
3	$\neg p$	Simp 2
4	$q \rightarrow s$	Simp 2
5	$r \& q$	DS 1,3
6	r	Simp 5
7	q	Simp 5
8	s	MP 4,7
9	$r \& s$	Conj 6,8

2.

c: You have an identification card showing that you are older than 21.

o: You are older than 21.

f: The identification card is fake.

If you have an identification card showing that you are older than 21, then either you are of legal age or this ID is a fake.

$$c \rightarrow (o \vee f)$$

I see your card, but I know that you are not 21.

$$c \& \neg o$$

Hence, this must be a fake ID.

Therefore, *f*.

1	$c \rightarrow (o \vee f)$	premise
2	$c \& \neg o$	premise
3	c	Simp 1
4	$\neg o$	Simp 1
5	$o \vee f$	MP 1,3
6	f	DS 4,5

Logical Proofs with Equivalences

The system of proof we have is sound, but it is not complete. A system of proof is sound if and only if you can only construct a proof for valid arguments. Our system as it is so far must be sound because the only moves we can use are rules of inference, which are themselves valid arguments. But the system is not complete. A system of proof is complete if and only if every valid argument formable in the language can be proven in the system. As the system stands now, there are valid arguments that we are unable to prove. We need more moves. We call these equivalences.

THE EQUIVALENCES

- ▶ Equivalences are sentence forms in our truth-functional language that are truth-functionally equivalent—that is, sentences that must always have the same truth-value. For any of these, if you doubt that they are really equivalent, work out a truth table, and you will see that they are.
- ▶ The first equivalence is called double negation (DN). It asserts that a sentence p and the double-negated sentence $--p$ are equivalent. “It is not the case that you didn’t call your mother” means that you did call your mother.
- ▶ Two negations yield an affirmation. Negation changes the truth-value of a sentence from true to false or false to true. So, doing it twice just brings you back to the original truth-value. Therefore,

you could always add two negations or remove a pair of adjacent negations without changing the truth-value of the original sentence. For us, this makes the two equivalent.

- ▶ Suppose that we have a proof with the following lines.

6 $-c \rightarrow h$
7 $-h$

- ▶ We have a conditional and the negation of the consequent, so we can use the rule of inference modus tollens to get the negation of the antecedent in the proof. The antecedent is $-c$, so we get the following.

8 $--c$ MT 6,7

- ▶ But if the conclusion we are looking for is c , then we would need one more step.

9 c DN 8

- ▶ Because c and $--c$ are equivalent, if c appears in the second column, then it must be true, and therefore $-c$ must be true and thus can also be entered in the second column. If there are two negations directly next to each other, not separated by parentheses, we can wipe them away.
- ▶ However, in some circumstances, we will need them. In these cases, we can use DN to add them. Suppose that we have a disjunction in a proof.

12 $b \vee -d$

- ▶ Suppose that we also have the sentence d in the proof.

13 d

- ▶ This means that we know that b is true or d is false. But we know that d cannot be false because line 13 says that it is true. We would like to use the disjunctive syllogism here, but to use the rule DS, we need a disjunction and the negation of one of the disjuncts. We do not have that.
- ▶ But by using DN, we can manufacture it.

14 $\neg\neg d$ DN 13

- ▶ Next, we do have the negation of one of the disjuncts. One is b and the other is $\neg d$. The negation of $\neg d$ is $\neg\neg d$. And that is what we have on line 14. So, we can now do the following.

15 b DS 12,14

- ▶ So, double negation can be used to add or take away two negations that are directly next to each other anywhere in a proof.
- ▶ There are two important differences between the rules of inference and the equivalences. First, equivalences work in both directions. We can substitute $\neg\neg c$ for c , or c for $\neg\neg c$. Modus ponens, for example, is an inference, so it only goes one way.
- ▶ The second difference is that we can use equivalences on parts of lines. Rules of inference can only be used on entire lines. The reason for this is that while rules of inference are truth-preserving—that is, they always take you from a true sentence to another true sentence—equivalences always maintain the truth-value of the sentence, true or false.
- ▶ Truth-functional sentences derive their truth-values solely from the values of the constituent parts. Some of these parts may be true; some may be false. Equivalences always leave the truth-value unchanged and therefore can be used on parts of lines, where rules of inference cannot.

- ▶ One similarity is that although we will be using atomic letters like a and b to show the equivalences, they may be used on molecular sentences or any level of complexity. We could take any line in a proof, atomic or not, and apply any of our equivalences.
- ▶ The second equivalence is a pair of relations called De Morgan's theorem (DeM), named for the 19th-century logical pioneer Augustus De Morgan. It says that the negation of a disjunction is a conjunction of the disjuncts and the negation of a conjunction is a disjunction of the conjuncts.

$$\neg(a \vee b) :: \neg a \wedge \neg b$$

$$\neg(a \wedge b) :: \neg a \vee \neg b$$

- ▶ When you negate an “or,” you get an “and,” and when you negate an “and,” you get an “or.”
- ▶ Think of this in terms of Bob at the coffee machine. If we say that Bob has sugar or cream in his coffee, how could we be wrong? The only way we are wrong about Bob having sugar or cream is when he has neither—that is, when he doesn't have sugar *and* he doesn't have cream.
- ▶ If we say that Bob has sugar and cream in his coffee and again are wrong, what could be the case? We would be wrong about his having both sugar and cream in his coffee if he didn't have sugar *or* he didn't have cream, *or* he had neither. This is De Morgan's theorem.
- ▶ We cannot distribute a negation, but we do have an equivalence called distribution where we can distribute an “and” or an “or.”

$$a \vee (b \wedge c) :: (a \vee b) \wedge (a \vee c)$$

$$a \wedge (b \vee c) :: (a \wedge b) \vee (a \wedge c)$$

- ▶ We can distribute the $a \vee$ or the $a \wedge$ through the other sentence.

- ▶ Notice an important aspect of this equivalence. One side is a disjunction and the other is a conjunction. A helpful hint when doing proofs is that anytime you have the negation of a disjunction, use De Morgan's theorem as quickly as possible, because it results in a conjunction upon which we can use the rule of inference simplification and get new pieces we might be able to use.
- ▶ Distribution has two sibling equivalences, association and commutation. Where distribution is used when you have mixed conjunctions and disjunctions, association works where you have two “ands” or two “ors.”

$$a \& (b \& c) :: (a \& b) \& c$$

$$a \vee (b \vee c) :: (a \vee b) \vee c$$

- ▶ This is association (Assoc). For an “and” sentence, the only time it is true is when both of the conjuncts are true. So, $a \& b$ is true when a is true and b is true, and $(a \& b) \& c$ is true only when $(a \& b)$ is true and c is true—that is, when all three are true.
- ▶ Similarly, when we regroup them, the only time the sentence is true is when all three are true. They are equivalent. In the case of the disjunction, the only time the sentence is false is when a , b , and c are all false, no matter how they are grouped. So, if we have two “ands” or two “ors,” we are free to shift the parentheses how we see fit.
- ▶ The last of the trio is commutation (Comm), where we are taking the terms of a conjunction or disjunction and moving them somewhere else.

$$a \vee b :: b \vee a$$

$$a \& b :: b \& a$$

- ▶ Remember that in basic truth tables, conjunction and disjunction are symmetric. This is just a result of that fact.
- ▶ But remember that the conditional was not. The only time a conditional is false is when the antecedent is true and the consequent is false.
- ▶ So, let's suppose that we have two sentences, a , which is true, and b , which is false: $a \rightarrow b$ is true then false, and thus is false. But $b \rightarrow a$ gives us false then true, and that case is true. So, if we try to take the antecedent and the consequent and swap their places, we change the truth-value: a then b and b then a are not equivalent.
- ▶ We can't switch around the antecedent and the consequent and maintain the truth-value of the resulting sentence, so while commutation holds for conjunction and disjunction, it does not for the conditional.
- ▶ The conditional gets its own version of this, called contraposition (Contra), as in "contrary" (a change in truth-value) and "position" (a change in location).

$$a \rightarrow b :: \neg b \rightarrow \neg a$$

- ▶ If we know that $a \rightarrow b$ is true, then if you get a , b must follow. So, if there is no b to be found, there could not have been an a in the first place. If there is fire, then there is oxygen. In other words, if there is no oxygen, then there is no fire. We can swap the antecedent for the consequent in a conditional as long as we negate both in the move.
- ▶ The conditional also has its own version of distribution, which is called exportation (Exp).

$$(a \& b) \rightarrow c :: a \rightarrow (b \rightarrow c)$$

- ▶ Suppose that we need two things to happen to bring about a third—for example, we need for the temperature to be below freezing and for there to be precipitation in order to get snow. That means that if the temperature is below freezing, then if there is precipitation, it will snow.
- ▶ So, we can shift those parentheses the way we do with conjunctions and disjunctions using the rule association, but only if we change the first arrow to a conjunction or vice versa.
- ▶ That is the equivalence that the conditional and the conjunction share.
- ▶ The disjunction and the conditional share one as well. It is called implication (Impl).

$$a \rightarrow b :: -a \vee b$$

- ▶ If you say to one of your teenage children, “If you are going to go to that party, you will complete your chores,” then what you are saying is that the teenager’s chores will be complete or the teenager will not be attending that party.
- ▶ The key here is to think of the truth table for the conditional. The only time a conditional is false is when the antecedent is true and the consequent is false. That means that a conditional is true if that case is not the case—in other words, if the antecedent is false or the consequent is true.
- ▶ And that is just what implication says. This means that we can switch back and forth between conditionals and disjunctions as long as we negate the first term.

- ▶ The last equivalence is given the misleading name tautology (Taut).

$a \vee a :: a$

$a \& a :: a$

- ▶ It is fairly obvious that a “or” a or a “and” a is the same thing as saying a : Tonight, let’s have either Indian food or Indian food. Actually, let’s have both Indian food and Indian food.
- ▶ So, anytime you have a conjunction or a disjunction of the same term, you can collapse it.
- ▶ Of course, $a \vee a$ and $a \& a$ are not tautologies. They are contingencies. If a is true, then $a \& a$ and $a \vee a$ are true. If a is false, then $a \& a$ (false and false) and $a \vee a$ (false or false) are both false. On the other hand, $a \rightarrow a$ is not equivalent to a . Why?
- ▶ Because it *is* a tautology. If we have a sentence a , then $a \rightarrow a$ will always be true, regardless of the truth-value of a . If a is true, then $a \rightarrow a$ is true, then true, and that is true. If a is false, then $a \rightarrow a$ becomes false, then false, and this is also true. The only time a conditional is false is when the antecedent is true and the consequent is false.
- ▶ If the antecedent and the consequent are the same sentence, then this could never be the case. As such, $a \vee a$ and $a \& a$ are equivalent to a , and using the equivalence tautology, we can substitute one for the other at will.
- ▶ But a and $a \rightarrow a$ are not equivalent because a could be false and $a \rightarrow a$ would be true. So, the equivalence only holds for “and” and “or.”

THE EQUIVALENCES

DN	$a::\neg\neg a$
DeM	$\neg(a\vee b)::\neg a\&\neg b; \neg(a\&b)::\neg a\vee\neg b$
Dist	$a\vee(b\&c)::(a\vee b)\&(a\vee c); a\&(b\vee c)::(a\&b)\vee(a\&c)$
Assoc	$a\vee(b\vee c)::(a\vee b)\vee c; a\&(b\&c)::(a\&b)\&c$
Comm	$a\vee b::b\vee a; a\&b::b\&a$
Contra	$a\rightarrow b::\neg b\rightarrow\neg a$
Exp	$(a\&b)\rightarrow c::a\rightarrow(b\rightarrow c)$
Impl	$a\rightarrow b::\neg a\vee b$
Taut	$a\vee a::a; a\&a::a$

READINGS

Copi, *Introduction to Logic*, chap. 8.

Hurley, *Logic*, chap. 7.

Kahane, *Logic and Philosophy*, chap. 5.

QUESTIONS

1.

Construct a proof for the following argument.

$(p\vee\neg q)\vee r, \neg p\&\neg r$. Therefore, $q\rightarrow s$.

2.

Translate the following argument and construct a proof to show that it is valid.

I need to be in Atlanta for work, but we could also visit Boston or Chicago. If we fly through Denver, then it makes sense to go to Chicago. I suggest that we give Chicago a miss this time. So, we'll go to Atlanta and Boston and not fly through Denver.

ANSWERS

1.

1	$(p \vee \neg q) \vee r$	premise
2	$\neg p \& \neg r$	premise
3	$\neg p$	Simp 2
4	$\neg r$	Simp 2
5	$p \vee (\neg q \vee r)$	Assoc 1
6	$\neg q \vee r$	DS 3,5
7	$\neg q$	DS 4,6
8	$\neg q \vee s$	Add 7
9	$q \rightarrow s$	Impl 8

2.

I need to be in Atlanta for work, but we could also visit Boston or Chicago.

$$a \& (b \vee c)$$

If we fly through Denver, then it makes sense to go to Chicago.

$$d \rightarrow c$$

I suggest we give Chicago a miss this time.

$\neg c$

So, we'll go to Atlanta and Boston and not fly through Denver.

Therefore, $a \& (b \& \neg d)$.

1	$\&(b \vee c)$	premise
2	$d \rightarrow c$	premise
3	$\neg c$	premise
4	$\neg d$	MT 2,3
5	$(a \& b) \vee (a \& c)$	DeM1
6	$\neg c \vee \neg a$	Add 3
7	$\neg a \vee \neg c$	Comm 6
8	$\neg(a \& c)$	DeM 7
9	$a \& b$	DS 5,8
10	$(a \& b) \& \neg d$	Conj 4,9
11	$a \& (b \& \neg d)$	Assoc 10

Conditional and Indirect Proofs

We have been developing a system of natural deduction. We want a system that is both sound and complete. A system of proof is sound if every argument for which there is a proof is valid. Our system so far is sound. A system of proof is complete if every valid argument has a proof in the system. Our system is not yet complete. There are valid arguments that our rules of inference and our equivalences are not enough to complete a proof. We need two new proof strategies and a new justification to reach that ultimate goal. The justification is assumption.

ASSUMPTION AS A JUSTIFICATION

- ▶ At this point in the course, you should be flabbergasted at the thought of assumption being a justification. How can we enter a sentence in the second column that is a mere assumption?
- ▶ After all the time we spent discussing why it was essential to only put sentences in the second column that we know to be absolutely true, do you mean to tell us that now we can just enter any sentence in the second column and justify it as an assumption? Surely, there are limits to what can be assumed, and those limits are based on rational inferences about what we can know.
- ▶ No. You are free to assume any sentence you want and enter it into the second column of the proof. We did say clearly and explicitly that if even one false sentence shows up in the second column, logical chaos could result.

- ▶ But it is okay. We have protection. We have boxes. When we introduce an assumption (Assumpt) into a proof, we put it in a box, and that box is logical quarantine. Nothing inside of the box is allowed to come out into the general population.
- ▶ Anything can be brought into the box, but once a sentence has been in the box with the assumption, anything that inferred from it is to be deemed to be contagious in that it could be infected with the possible falsity of the assumption.
- ▶ When we open a box and insert an assumption, we create a new logical world—one that resembles our logical world in certain ways, so we can infer true propositions in the real world from what we observe in the hypothetical box world. One way to draw such inferences is called conditional proof.

CONDITIONAL PROOF

- ▶ We use conditional proof when we want to prove a conditional. We use this form of reasoning all the time, especially if we have children. There are two ways one can learn lessons in life: the easy way and the hard way. The sentence “You should not put your hand on a hot stove” can be learned the hard way by putting your hand on the stove.
- ▶ The easy way is conditional proof. We say to the person who has never yet put a hand on the hot stove, “Before you do, let’s think about this.” Let’s start by assuming that you put your hand on the hot stove. Don’t really do it in the real world, but create a hypothetical world in which you did. What would be the result?
- ▶ The temperature of the stove is significantly higher than that of your hand. By the first law of thermodynamics, heat would flow rapidly from the stove into your hand. Your hand heats up very quickly—so quickly that it would cause damage to the skin.

You would suffer burns. The nerves in your hand would send message of those burns to your brain, which would register it as incredible amounts of pain.

- ▶ What do we know from this thought experiment? We know that if you put your hand on a hot stove, then you will experience incredible amounts of pain. By assuming the antecedent and then deriving the truth of the consequent, we can assert the conditional “If you put your hand on a hot stove, you will feel incredible amounts of pain.” We learned the truth of the conditional the easy way.
- ▶ This is how conditional proof works. To demonstrate a conditional by conditional proof, the first step is to enter all of the premises into the proof. The second step is to open a box and insert the antecedent into the box as an assumption justified by “Assumpt.”
- ▶ Next, pull your premises into the box as needed. Then, use rules of inference and equivalences. Proceed as if everything were normal inside of the box, until such time as the consequent of the conditional appears as a justified line inside of the box.
- ▶ At that point, we know that if the antecedent is true, then the consequent has to be true. So, close the box, and in the general proof, write down the conditional justified by CP m,n ., where m is the first line inside the box and n is the last line inside the box.

INDIRECT PROOF

- ▶ One use of assumptions is in conditional proof. The other is called indirect proof.

- ▶ When we think of proofs in mathematics, we think of Euclidean geometry, which are proofs of the type we started with, where you assume the premises and demonstrate the conclusion. This is called direct proof.
- ▶ But most mathematical proofs do not take this form. Most are indirect proofs, or to use the Latin name, *reductio ad absurdum*, or reduce to absurdity.
- ▶ The “absurdity” is a contradiction. Recall how much we worry about contradictions. If even one contradiction is true, then everything is true, and truth goes away. We fear contradictions. But, like early humans who learned to harness the power of otherwise dangerous fire for cooking, we, too, will learn to control the power of contradictions.
- ▶ The fundamental principle behind traditional logic is the law of the excluded middle—that all sentences are either true or false. A sentence that isn’t true is false, and a sentence that isn’t false is true.
- ▶ So, if we want to prove that a sentence has to be true, that is the same thing as proving that it can’t be false. This is the trick. We want to prove that p cannot be false. But how?
- ▶ We know that contradictions are always false, so if we can show that the negation of p , in conjunction with the premises we are asserting to be true, necessarily leads to a contradiction, then we have logical grounds on which to reject the negation of p .
- ▶ But if $\neg p$ is false, then p must be true. If we can show that denying p leads to absurdity, then we have no choice but to accept p . That is indirect proof.

- ▶ It begins when we enter the premises into the proof, then open a box and insert the negation of what we are trying to prove into the box. We use our premises, the assumption, and our rules of inference and equivalences until a contradiction, any sentence of the form $a \& \neg a$, appears in the box.
- ▶ It does not have to be a contradiction involving the assumption. Any contradiction will do, because, as you have learned, if you grant even a single contradiction, all sentences—including all contradictions—are true.
- ▶ At this point, we have shown that the negation of the conclusion, when added to the premises, necessarily results in contradiction, so we close the box, and in the general proof, we write down the conclusion, justified by IP with the line numbers from the opening to the closing of the box.
- ▶ In a conditional proof, the box was a hypothetical logical world in which we posit something that might not be true in order to see what else would have to be true.
- ▶ In an indirect proof, inside the box is a bizarre logical world in which we put propositions together that we believe cannot be put together in order to observe the nonsense that results.
- ▶ If you are particularly clever, you might have realized that it might not be the case that the negation of the conclusion is responsible for the contradiction. All we have shown is that the combined set of the premises and the negated conclusion lead to a contradiction.
- ▶ How can we assert that the negation of the conclusion is necessarily false when it might not be the conclusion that is to blame for the derived contradiction? If the fault is with the premises and not the negation of the conclusion, what gives us the right to assert the conclusion?

- ▶ The answer is that if it is not the assumption that is to blame for the appearance of the contradiction, then it has to arise from the premises alone. That means that the premises are inconsistent—that they cannot all be true at the same time.
- ▶ An argument with inconsistent premises must be valid, because it would be impossible for all of its premises to be true and its conclusion to be false. So, on a technicality, we know that the argument is valid. So, whether the contradiction comes from the assumption or not, the derivation of the contradiction guarantees that the argument is valid.

READINGS

Barker, *The Elements of Logic*, chap. 3.

Hurley, *Logic*, chap. 7.

Kahane, *Logic and Philosophy*, chap. 6.

QUESTIONS

1.

Use conditional proof to show that the following argument is valid.

$s \rightarrow (y \& z)$. $(w \vee q) \rightarrow (y \& x)$. Therefore, $(s \vee q) \rightarrow (y \vee w)$.

2.

Use indirect proof to show that the following argument is valid.

$(a \vee l) \& (a \vee m)$. $a \rightarrow (t \& -l)$. Therefore, $(t \vee m) \vee -l$.

3.

Translate the following argument and construct both a conditional and an indirect proof to demonstrate its validity.

I can't eat turkey or pasta without overeating. So, if I eat turkey, I will eat turkey and overeat.

ANSWERS

1.

1	$s \rightarrow (y \& z)$	premise
2	$(w \vee q) \rightarrow (y \& x)$	premise
3	$s \vee q$	Assumpt
4	$(s \vee q) \vee w$	Add 3
5	$s \vee (q \vee w)$	Assoc 4
6	$s \vee (w \vee q)$	Comm 5
7	$(y \& z) \vee (y \& x)$	CD 1,2,6
8	$y \& (z \vee x)$	Dist 7
9	y	Simp 8
10	$y \vee w$	Add 9
11	$(s \vee q) \rightarrow (y \vee w)$	CP 3-10

2.

1	$\neg l \rightarrow n$	premise
2	$(n \& \neg t) \rightarrow m$	premise
3	$\neg [(t \vee m) \vee \neg l]$	Assumpt
4	$\neg (t \vee m) \& \neg l$	DeM 3
5	$\neg (t \vee m)$	Simp 4
6	$\neg l$	Simp 4
7	$\neg t \& \neg m$	DeM 5
8	$\neg t$	Simp 7
9	$\neg m$	Simp 7
10	n	MP 1,6

11	$n \& \neg t$	Conj 8,10
12	m	MP 2,11
13	$m \& \neg m$	Conj 9,12
14	$(t \vee m) \vee \neg l$	IP 3–13

3.

t : I eat turkey; p : I eat pasta; o : I overeat.

$(t \vee p) \rightarrow o$; therefore, $t \rightarrow (t \& o)$.

Conditional proof

1	$(t \vee p) \rightarrow o$	premise
2	t	Assumpt
3	$t \vee p$	Add 2
4	o	MP 1,3
5	$t \& o$	Conj 2,4
6	$t \rightarrow (t \& o)$	CP 2–5

Indirect proof

1	$(t \vee p) \rightarrow o$	premise
2	$\neg[t \rightarrow (t \& o)]$	Assumpt
3	$\neg[\neg t \vee (t \& o)]$	Impl 2
4	$\neg\neg t \& \neg(t \& o)$	DeM 3
5	$t \& \neg(t \& o)$	DN 4
6	t	Simp 5
7	$\neg(t \& o)$	Simp 5
8	$\neg t \vee \neg o$	DeM7
9	$t \vee p$	Add 6
10	o	MP 1,9
11	$t \& o$	Conj 6,10
12	$(t \& o) \& \neg(t \& o)$	Conj 7,11
13	$t \rightarrow (t \& o)$	IP 2–12

First-Order Predicate Logic

Truth-functional logic allowed us to account for arguments where the relevant logical structure occurred between sentences. The truth-functional connectives showed us the truth conditions of complex combinations of sentences. But some of the logically relevant elements of our arguments occur inside of the sentences. Truth-functional logic is too broad to get at this. We need a new, sharper logical language that builds on truth-functional logic but allows us to pull the logical content from within sentences. This new language is first-order predicate logic.

THE NEED FOR FIRST-ORDER PREDICATE LOGIC

- ▶ Think about the early lectures about logic. The standard argument we used as an example was the following classic argument.
 - ◆ All men are mortal.
 - ◆ Socrates is a man.
 - ◆ Therefore, Socrates is mortal.
- ▶ We said that it is valid. Can we use our truth-functional apparatus to demonstrate this? Let a represent “All men are mortal,” b represent the sentence “Socrates is a man,” and c stand for “Socrates is mortal.” Can we give a proof from the premises a and b to the conclusion c ? No.
- ▶ Truth-functional logic is also known as sentential logic because it is the logic of sentences and the connectives that connect them. It is, in a certain sense, a very blunt weapon, because it can only deal with the logical relations between sentences.

- ▶ What we see in the Socrates example is the need to apply logical machinery inside of sentences. We do need to be able to evaluate the relations between sentences, but we also need the much finer work done of illuminating the logical relationships within sentences.
- ▶ Think about Aristotle's categorical logic, in which we have inferences that concern the content inside of the grammatical structure of the individual sentences. Truth-functional logic cannot get at the aspects that Aristotle did.
- ▶ So, for a period, there was a running argument as to which is the real logic: Aristotle's categorical approach or the truth-functional approach.
- ▶ That debate ended when we formulated a new logical language that combined the two together. This new hybrid language is first-order predicate logic. It contains all of the elements of truth-functional logic—all of our connectives, exactly as we have been using them—but adds some new elements that allow us to investigate the logical content of the grammar of individual sentences instead of simply giving them a truth-value.

ELEMENTS OF THE FIRST-ORDER PREDICATE LANGUAGE

- ▶ We begin our new language with the four truth-functional connectives: \neg , $\&$, \vee , and \rightarrow . But now, instead of sentences being represented by a single lowercase letter, individual sentences require a subset of three new categories of elements.
- ▶ As we said when we were starting Aristotelian logic, sentences have a subject and a predicate. Sentences say something about something else. The thing that we are talking about we will call the individual. What we are saying about it we will call the predicate.

- ▶ We will start with the easiest version in which our predicates are properties, such as “is tall,” “is blue,” or “is yummy.” We will represent these properties in our language with uppercase letters, such as T , B , or Y . So, uppercase letters are our properties.
- ▶ One restriction on properties is that they must be world referring, such as being tall, blue, or yummy. They cannot be logical properties, such as “is true,” “is valid,” or “is logically equivalent to.” This is why we call this language “first-order” logic. The first-order language talks about things, and the second-order language talks about the first-order language.
- ▶ The subject, that which is being said to have the property, will be represented by lowercase letters. But we need to be careful because there are two different types of individuals we must distinguish between. Borrowing terms from mathematics, there are constants and variables.
- ▶ A constant is a specific thing, such as Bob, the sky, or a particular sandwich. Constants are not classes of things, but particular things. For these, we use the lowercase letters in the first part of the alphabet—the letters a through t .
- ▶ So, if we want to say in our new language that Bob is tall, we select an uppercase letter for the property “is tall”—for example, T . And we need a lowercase letter to represent Bob—for example, b .
- ▶ In this language, we put the predicate first and the individual second. So, “Bob is tall” is written Tb . For “Larry is tall,” we use l to represent Larry, and we get Tl . Bob and Larry are both tall. That is a truth-functional combination of the sentences “Bob is tall” and “Larry is tall,” giving us $Tb \& Tl$. We cannot write Tbl , because T is a property and only applies to one thing.
- ▶ What about “If Bob is tall, you’re a monkey’s uncle”? Here, again, is a truth-functional combination of predicate sentences. We have

seen that “Bob is tall” is Tb , and now we need a capital letter for the property of being a monkey’s uncle. Let’s use M . And for you, we’ll use the lowercase letter y . So, “If Bob is tall, you’re a monkey’s uncle” gets written $Tb \rightarrow My$.

- ▶ We need to discuss sentences that are not about specific entities that we can point to. Think of the end of a mystery novel. All of the suspects have been gathered in the study and the detective says, “Someone in this room is a murderer.” That someone is a particular individual, but we don’t know whom. So, we can’t use a constant.
- ▶ For unknown individuals or for all of the individuals in some category, we need something else in our language. We need variables. Variables will be lowercase letters at the end of the alphabet—the letters u through z . These will represent arbitrary members of some class or unknown members of a class.
- ▶ Suppose that we want to write the title of the song “Everything Is Awesome” in our new language. How would we do that? You might think that it would be—using A for “awesome”— Ax . That is “some x has the property A .” But then how would we write “Something is awesome”? Wouldn’t we also write it Ax ?
- ▶ But these are clearly different propositions, and the whole point of our artificial logical language is to avoid the ambiguity we find in ordinary spoken language. We need to do something about this.
- ▶ That something is a new pair of elements, called quantifiers. We have two of them. The first is the existential quantifier (\exists). We read it as “for some.”
- ▶ Quantifiers have to quantify over a variable name. So, we always pair a quantifier with a variable. $\exists x$ reads “for some x .” We cannot quantify a constant: $\exists b$ makes no sense in our language. “For

some Bob”? There is no need for “some Bob”; there is simply “Bob.” Whenever you use a quantifier, you need to use a variable and only a variable.

- ▶ To say, “Something is awesome,” we would write $\exists xBx$.
- ▶ Suppose that Bob is quite distinguished looking, but not beautiful. We might want to say, “Something is beautiful, but it isn’t Bob.” Here is a truth-functional combination of sentences. Remember that “but” is a way of saying “and,” so we have the following.
 - ◆ Something is beautiful, and it is not the case that Bob is beautiful.
- ▶ That would translate into our new language as follows.

$$\exists xBx \& \neg Bb$$

- ▶ That is the existential quantifier. It is called this because it asserts the existence of something.
- ▶ The other quantifier is the universal quantifier (\forall), which is read “for all.” If we want to say, “Everything is beautiful,” we pick a variable to quantify over—for example, y —and we assert B of all y —that is, $\forall yBy$.
- ▶ Suppose that we want to say, “If Bob is not beautiful, then not everything is beautiful.” We have a conditional, as follows.
 - ◆ If Bob is not beautiful, then it is not the case that all things are beautiful.
- ▶ Substituting in our truth-functional symbols of the connectives:
 - ◆ \neg Bob is beautiful \rightarrow \neg all things are beautiful

- ▶ We now know how to translate these first-order pieces.

$$-Bb \rightarrow -\forall yBy$$

- ▶ Notice again that all quantifiers have to quantify over something, and that something has to be a variable. All quantifiers need a variable.
- ▶ The relation goes in the other direction as well. All variables need a quantifier. If T is the property “is tall,” what does Tz mean? It doesn’t mean “something is tall” or “everything is tall,” because those both need quantifiers. In our language, it does not mean anything. It is a violation of the grammar of first-order predicate logic.

FUNDAMENTAL FORMS

- ▶ Instead of the sentence letters of truth-functional logic, we now have a more complex representation that allows us to get at the content contained inside of sentences the way that Aristotelian logic did. Indeed, we can subsume Aristotle’s categorical logic within this broader approach. To do that, we need to translate the four categorical forms into our new first-order language.
- ▶ Let’s start with A sentences, sentences of the form “All A s are B ,” such as “All people have noses.” We need one capital letter to represent the property of being human—for example, H —and another for having a nose—for example, N . It is saying that everything that is in the H group is also in the N group, so we are looking at a universal quantification: It is “all somethings are something.”
- ▶ So, we need a universal quantifier and a variable name. It is going to be $x(Hx \supset Nx)$, where the question mark will be one of our truth-functional connectives. Which one?

- ▶ Some people intuitively think that it should be “and.” But it isn’t. Note what the sentence says if we include a conjunction: $\forall x(Hx \& Nx)$ says that everything is both human and has a nose. That is not what we are trying to say. We are trying to say that if something were to be a human, then it would be a thing with a nose. So, we need the connective to be a conditional: “All humans have noses” is written $\forall x(Hx \rightarrow Nx)$ in our language.
- ▶ What about *E* sentences, sentences of the form “No *As* are *B*”?
- ▶ The flawed intuition many people have in this case is to slap a negation in front of the *A* sentence form. But remember the square of opposition: The negation of an *A* sentence is an *O* sentence, not an *E* sentence. Just because it is false that all people love cilantro does not mean that no people love cilantro, just that some people don’t.
- ▶ Remember that an *E* sentence is a universal sentence. As such, we will need a universal quantifier. If we say, “No people live on Mars,” we are saying something about all people—that they are not Martians. What we are saying when we say that no people are living on Mars is that, for all things, if that thing is a person, we know that that thing does not live on Mars. In symbols, $\forall x(Hx \rightarrow \neg Mx)$.
- ▶ So, the negation does not go in front of the sentence, but rather in front of the second property. An *E* sentence, “No *As* are *B*,” is still a claim about the entire set of things in the *A* class, but instead of saying that they are all in the *B* class, we are saying that they are not in the *B* class.
- ▶ What about the particular sentences? Let’s start with *I* sentences: “Some *As* are *B*.” Clearly, this is an existential claim. We are saying that there is something and that thing is both *A* and *B*.

- ▶ We translate *I* sentences as $\exists z(Az \& Bz)$. Unlike with the universal sentences, we do use the conjunction in the particular. We are not saying that there is something that if it were an *A*, it would be a *B*. Instead, we are saying that there is something, and it is both *A* and *B*.
- ▶ What about *O* sentences? Let's think about sentences of the form "Some *As* are not *B*." Again, it is making an existential claim. We are saying that there is at least one thing, and it has certain properties. In this case, the properties are "being *A*" and "not being *B*." So, translating it fully, we get $\exists x(Ax \& \neg Bx)$; there is a thing, and it is *A* and not *B*.
- ▶ We now have all four of our categorical sentence forms translated into our new first-order predicate language.

<i>A</i> : $\forall x(Ax \rightarrow Bx)$	<i>E</i> : $\forall x(Ax \rightarrow \neg Bx)$
<i>I</i> : $\exists x(Ax \& Bx)$	<i>O</i> : $\exists x(Ax \& \neg Bx)$

- ▶ Add to these our categorical forms, the five basic forms, and we have the building blocks to translate most sentences into our first-order predicate language.

The sky is blue	<i>Bs</i>
Everything is blue	$\forall x Bx$
Nothing is blue	$\forall x \neg Bx$ or $\neg \exists x Bx$
Something is blue	$\exists x Bx$
Something is not blue	$\exists x \neg Bx$

- ▶ We will call these nine sentence forms our fundamental forms. Many of the sentences we come across in normal conversation are either one of these fundamental forms or a truth-functional combination of them.

- ▶ When translating, then, the key is to go step by step.
 - ◊ First, look at the sentence and figure out if it is one of the nine fundamental forms or if it is a truth-functional combination of the fundamental forms.
 - ◊ If it is a truth-functional combination, find the main connective and insert the symbol for it, while leaving the rest of the sentence in spoken language.
 - ◊ Look at the parts connected by the connectives and again ask whether these are fundamental forms or truth-functional combinations. Keep going until you are left with nothing but truth-functional symbols and spoken language versions of fundamental forms.
 - ◊ At this point, or if the sentence was not a truth-functional combination, ask whether the fundamental form is categorical or basic. Once you know that, figure out which form, and the translation from there is straightforward.

READINGS

Barker, *The Elements of Logic*, chap. 4.

Copi, *Introduction to Logic*, chap. 10.

Hurley, *Logic*, chap. 6.

Kahane, *Logic and Philosophy*, chap. 7.

QUESTIONS

Translate the following sentences into first-order predicate logic. Use the following abbreviations: e = Emmett Kelley, C = is a clown, A = is an artist, P = is a person, L = loves clowns, F = is afraid of clowns.

1.

Emmett Kelley was a clown and an artist.

2.

All the world loves a clown.

3.

Some people love clowns, but some people are afraid of them.

4.

Clowns are artists, but only some artists are clowns.

ANSWERS

1.

$Ce \& Ae$

2.

$\forall x(Px \rightarrow Lx)$

3.

$\exists x(Px \& Lx) \& \exists y(Py \& Fy)$

4.

$\forall x(Cx \rightarrow Ax) \& \exists y(Ay \& Cy)$

Validity in First-Order Predicate Logic

Because the universal quantifier potentially applies to an infinite number of things, truth tables are not an option for us as a validity test for first-order predicate logic. We can, however, extend our system of natural deduction proofs to work for this stronger language. To do so, we need to add four new rules and one new equivalence—all of which you will learn about in this lecture.

NEW RULES OF INFERENCE

- ▶ In first-order predicate logic, we need our other method of demonstrating validity, natural deduction proof. We will, however, need to augment it in order to deal with quantifiers. This will mean four new rules of inference and one new equivalence.
- ▶ The four new rules of inference will include two rules for taking away quantifiers and two rules for adding them on. The idea is to strip the quantifiers away and reduce things back to truth-functional sentences. Then, once we have a sentence of the form we want, we add the quantifiers back in to get the first-order predicate sentence we were looking for. In essence, these rules allow us to reduce things to what we already know how to do.
- ▶ These rules have personalities. Two are laid-back, relaxed rules that just do their job, no matter what is going on around them. The other two are finicky, picky rules that refuse to do their job unless very specific demands are met.

- ▶ The first rule, universal instantiation (UI), allows us to remove a universal quantifier that applies to an entire line and replace the quantified variable with either a free variable or any constant.
- ▶ Consider the sentence $\forall x(Bx \& Cx)$; everything is both B and C . If everything is B and C , what do we know about your favorite pair of socks? It has to be B and C . If everything is B and C , we know that the sentence formed by replacing the quantified variable x with any constant must also be true. If we know it is true, then it can appear in the second column of a proof.
- ▶ But we can also replace the quantified variable with an unquantified, or free, variable. That is, suppose that we had the following.

12 $\forall x(Bx \& Cx)$

- ▶ Then, we could write the following.

13 $By \& Cy$ UI 12

- ▶ Previously in this course, we made the seemingly unequivocal claim that free variables are not allowed in our language—that $By \& Cy$ has no meaning. It doesn't mean that some y are $B \& C$ and it didn't mean that all y are $B \& C$. It means that it doesn't have a meaning in our language—that it is grammatically incorrect.
- ▶ This is true, outside the context of a proof. In a proof, the use of UI gives the sentence $By \& Cy$ a meaning. It means that for any arbitrary thing—for example, y — y is $B \& C$. We can reason about a universal property of the whole group if we have reason to think that it holds for each member of the group.
- ▶ While universal instantiation allows us to take away a universal quantifier that applies to an entire line in a proof and replace the quantified variable with either a constant or any free variable name, existential instantiation (EI) allows us to remove an

existential quantifier and replace it with a free variable, but only with a free variable that appears free nowhere earlier in the proof.

- ▶ There are two restrictions on EI that we need to be clear about. First, we cannot use EI to instantiate a constant. Consider the sentence $\exists zBz$. It says that something is B . What is that thing? We don't know what exact thing it is, only that something—that is, at least one thing, maybe all things—is B .
- ▶ If we used a constant to instantiate an existential, that is saying that we *do* know what the thing is, and it is c . Maybe it is c ; maybe it isn't. If all we know is that it is something, then asserting that some constant *is* the thing is to say more than we know, and this means that we could be wrong.
- ▶ If a sentence could be wrong, then it cannot be entered in the second column of the proof. So, EI cannot be used to take away a quantifier and substitute a constant for the quantified variable. It has to be a variable that gets put in.
- ▶ Again, when we use EI we have a sentence with a free variable, something that we generally don't allow. But this free variable means something slightly different from the free variable that results from UI, in which the free variable means “any arbitrary thing.”
- ▶ By contrast, EI means “the thing whose identity we do not yet, and may never, know.” Think of the victim of a murder who is discovered with no identifying information. If the victim is male, we use the term “John Doe.” This name is a placeholder. It refers to the person who was killed, even though we don't know who the person is. The free variable instantiated with EI is the logical equivalent of “John Doe.”
- ▶ Suppose that we have the following.

8 $\exists xBx$

- ▶ Then, we can use EI to get the following—as long as z does not appear free earlier in the proof.

9 Bz EI 8

- ▶ In this sentence, the free variable z means whatever the thing is that is z . From line 8, we know that it exists; we just have no clue about its identity, so we will use z to refer to whatever that thing is that is B .
- ▶ The reason that we have to instantiate with a new variable name is that all we know about the thing that is B is that it is the thing that is B . If we reused a variable that appeared free earlier in the proof, then we would be saying that the same thing has several properties—maybe it does, and maybe it doesn't, but we cannot be sure. For example, suppose the following.

10 $\exists xBx$

11 $\exists xCx$

- ▶ If we used EI on both of these lines with the same variable name, we would get the following.

12 By EI 10

13 Cy EI 11

- ▶ And then we could form the following.

14 $By \& Cy$ Conj 12,13

- ▶ Maybe there is a thing that is both B and C , but we don't know that. Maybe one thing is B and a completely different thing is C . It is true that some numbers are even and that some numbers are odd, but if we make the mistake of re-instantiating with the same free variable, we will be proving that some numbers are both even and odd.

- ▶ It is therefore crucial that every time we use EI, we introduce a new free variable name.
- ▶ While the first two rules are instantiation rules that let us remove quantifiers, the other two are generalization rules that let us reinsert them. And just as one came with no restrictions and the other had constraints, it is the same with generalization rules. The only difference is that it is the existential that is the easy one.
- ▶ Existential generalization (EG) is the rule by which we add an existential quantifier to a line. We can always take a line with a free variable or a constant in it and add an existential quantifier to bind that variable. If the line has a constant, then we know that the property asserted of the constant is true of something—namely, that thing.
- ▶ We might know that Jane is happy. If this is true, then it must necessarily be true that something is happy. Because we are being less specific, we are guaranteed that the move will be truth-preserving, and therefore, we can write the new existentially quantified sentence in the second column of the proof.
- ▶ The same is true if we use EG to generalize a free variable, regardless of whether it was instantiated universally or existentially. If it is true that all books have pages, then it is true that some books have pages. You can existentially generalize any free variable anytime and anywhere.
- ▶ The same is not true of universal generalization (UG), which cannot be used on a constant. Just because John's mom let him go to the party doesn't mean that all moms will. Similarly, we cannot use UG to generalize a variable that appears free in a line justified by EG. Just because some men are jerks does not mean that we can conclude that all are.

- ▶ We can only use UG to generalize a free variable that was introduced into the proof using UI and that has never appeared free in a line on which EI has been used. EI stamps the sentences it operates on with the quantificational version of guilt by association.
- ▶ Suppose that we have the sentence “Everything has mass, but if that thing is heavier than the earth, then there is something else that orbits it.” It is a truth-functional combination. The main connective is “and.”

(Everything has mass) & (If that thing is heavier than the earth,
then there is something that orbits it)

- ▶ It is important to see that both conjuncts in this sentence refer to the same thing. In the first conjunct, we are saying that any given thing you choose has a property—having mass. In the second conjunct, we are saying that the thing we picked for the first conjunct also has the property of being orbited by something else if it has a mass greater than that of the earth. Because both conjuncts refer to the same thing, the quantifier we use must have both sentences in its scope. So, we say the following.

$\forall x[(x \text{ has mass}) \& (\text{if } x \text{ is heavier than the earth, then there is something that orbits } x)]$

- ▶ Having mass is just a property, so let's use M to represent that.

$\forall x[Mx \& (\text{if } x \text{ is heavier than the earth, then there is something that orbits } x)]$

- ▶ The second conjunct is clearly a conditional.

$\forall x[Mx \& (x \text{ is heavier than the earth} \rightarrow \text{there is something that orbits } x)]$

- ▶ Being heavier than the earth is a property, and so is “orbits x .” Let’s use H and O for those. So, what we are saying is that there is some thing y such that if x is heavier than the earth, then y orbits x . This gives us the following.

$$\forall x[Mx \& \exists y(Hx \rightarrow Oy)]$$

- ▶ If we had that as a premise in an argument, we could use UI on it because the scope of the universal quantifier is the entire rest of the sentence.

1	$\forall x[Mx \& \exists y(Hx \rightarrow Oy)]$	Premise
2	$Mx \& \exists y(Hx \rightarrow Oy)$	UI 1

- ▶ Next, we can simplify.

3	Mx	Simp 2
4	$\exists y(Hx \rightarrow Oy)$	Simp 2

- ▶ On line 4, we see a sentence with an existential quantifier whose scope is the entire rest of the line. So, we can use EI. We cannot instantiate it with the free variable x because it already appears free previously in the proof.
- ▶ If we did use x again, we would have something orbiting itself. Clearly, that’s not what we want. But we could use any other free variable name for this thing—whatever it is—that is orbiting x . Let’s keep it as y .

5	$Hx \rightarrow Oy$	EI 4
---	---------------------	------

- ▶ Now we are in the situation we wanted to illustrate. We introduced the free variable named x into the proof using UI in line 2. So, you would think that we would be free to generalize it using UG. But you would be wrong.

- ▶ Because of line 5, x appears free in a line justified by EI, and this means that x can no longer be generalized using EG, even though it was introduced via a universal quantifier. By simply being free on a line justified by EI, even though it was not being instantiated, that is enough to disqualify it from being generalizable using UG.

A NEW EQUIVALENCE

- ▶ Finally, we have one new equivalence: quantifier negation (QN). If we have a negation next to a quantifier, we can reverse the order if we switch the quantifier.

$$-\exists::-\text{ and } -\forall::\exists-$$

- ▶ If it is false that some people are 12 feet tall, then it is true that all people are not 12 feet tall. If it is false that all grass is green, then it is true that some grass is not green. We can always move a negation from one side of a quantifier to another if we switch the quantifier.
- ▶ These are the only adjustments necessary to have a validity test for first-order predicate logic. We take an argument in our first-order predicate language, enter the premises into the proof, use our original eight rules of inference, our new four rules of inference, our original nine equivalences, our new equivalence, and our proof strategies until such time as the conclusion appears as a justified line in the proof, and we have a complete demonstration of the validity.

READINGS

- Barker, *The Elements of Logic*, chap. 3.
Copi, *Introduction to Logic*, chap. 10.
Hurley, *Logic*, chap. 7.
Kahane, *Logic and Philosophy*, chap. 5.

QUESTIONS

1.

Show that the following argument is valid by constructing a proof.

$\exists x(Fx \& Gx)$. $\forall x(Gx \rightarrow Hx)$. Therefore, $\exists x(Fx \& Hx)$.

2.

Translate the following argument and show that it is valid by constructing a proof.

All magical unicorns are immortal. There are no immortal beings.
Therefore, no unicorns are magical.

ANSWERS

1.

1	$\exists x(Fx \& Gx)$	premise
2	$\forall x(Gx \rightarrow Hx)$	premise
3	$Fx \& Gx$	EI 1
4	$Gx \rightarrow Hx$	UI 2
5	Fx	Simp 3
6	Gx	Simp 3
7	Hx	MP 4,6
8	$Fx \& Hx$	Conj 5,7
9	$\exists x(Fx \& Hx)$	EG 8

2.

All magical unicorns are immortal.

$$\forall x[(Mx \& Ux) \rightarrow Ix]$$

There are no immortal beings.

$$\forall x \neg Ix$$

Therefore, no unicorns are magical.

$$\forall x(Ux \rightarrow \neg Mx)$$

1	$\forall x[(Mx \& Ux) \rightarrow Ix]$	premise
2	$\neg Ix$	premise
3	Ux	Assumpt
4	$(Mx \& Ux) \rightarrow Ix$	UI 1
5	$\neg Ix$	UI 2
6	$\neg(Mx \& Ux)$	MT 4,5
7	$\neg Mx \vee \neg Ux$	DeM 6
8	$\neg \neg Ux$	DN 3
9	$\neg Mx$	DS 7,8
10	$Ux \rightarrow \neg Mx$	CP 3-9
11	$\forall x(Ux \rightarrow \neg Mx)$	UG 10

Demonstrating Invalidity

One of the wonderful aspects of truth tables is that they are both validity and invalidity tests. But in first-order predicate logic, we cannot construct truth tables. Natural deduction proof is our only validity test. Successful construction of a proof guarantees that the argument is valid, but an inability to find a proof does not establish an argument's invalidity. So, we'll need a different way of demonstrating invalidity. We will formulate two different ways of demonstrating that an argument in first-order predicate logic is invalid: the method of counterexample and the method of expansion. Both will show that an argument is invalid by introducing semantic elements into our purely syntactic approach.

SEMANTICS, SYNTAX, AND PRAGMATICS

- ▶ There are three levels to language: syntax, semantics, and pragmatics. Syntax deals with form and structure. Questions of grammar are syntactic, as are word order and internal structure.
- ▶ Semantics looks at meaning. Questions of what a word refers to is a semantic concern. Semantic issues are important issues that are complex and philosophically interesting.
- ▶ Syntax looks at what we say, semantics looks at what it means, and pragmatics deals with how we say it and how that can change the meaning. Issues of tone or inference from that which is unsaid are pragmatic aspects.

- ▶ Our logical languages are thin because they are purely syntactic languages. In a deep sense, they are languages with no meaning. When we translate spoken language arguments into truth-functional or first-order predicate logic, we remove the semantic content.
- ▶ Why are we ignoring semantics with our logical languages if it is important? Because our interest is determining deductive validity, which is a purely syntactic concept.
- ▶ Whether a deductive argument is valid has nothing to do with what the argument is about. It only matters whether the conclusion follows from the premises. As a result, our artificial languages remove the flesh from the spoken language arguments and rigorously display the inner skeleton, because that is sufficient to answer our questions about validity.
- ▶ But in demonstrating invalidity, we are going to restore semantic elements. The idea is that, in one sense, these semantic aspects are irrelevant. Consider the invalid argument affirming the consequent: If a , then b ; b , therefore a .
- ▶ This is a flawed argument form, regardless of what we fill in for a and b . We could say “If you are a dog, you have a nose. You have a nose. Therefore, you are a dog.” Or we could say, “If you won the lottery, you would be happy. You are happy. Therefore, you won the lottery.”
- ▶ Whether you are a dog or have won the lottery is irrelevant here. The content does not matter because it is the underlying form of the arguments that is to blame. But the flaw in the structure becomes clearer to us when we add some semantic content to that structure.

- ▶ We are not declaring individual arguments to be invalid, but argument forms. When we add semantic content to the argument forms, we are not merely saying that just the filled-in version of the argument is invalid; rather, we are saying that the skeleton itself is flawed. We are saying that any argument with that same form is invalid because we are showing that it is possible to construct a bad argument with that form.

THE METHOD OF COUNTEREXAMPLE

- ▶ An argument is invalid if it is possible for its premises to be true while its conclusion is false. In truth-functional logic, we could construct a truth table in which every possible combination of truth-values for the constituent sentences is considered.
- ▶ But another way is to simply come up with an example of an argument that has the same form but has true premises and a false conclusion. This is called the method of counterexample, and because we don't have the ability to construct truth tables for first-order predicate logic, we will have to use it.
- ▶ In order to give semantic content to an argument form in first-order predicate logic, there are three things we need to do.
 - 1 We need to give meaning to the quantifiers. (When we say "all x " or "some x ," all what's or some what's are we talking about?)
 - 2 We need to say which specific elements of this universe are picked out by our constants.
 - 3 We need to give meaning to the properties.
- ▶ These combine to give what logicians call an interpretation. An interpreted argument has the semantic flesh put back on the syntactic skeleton, and thus we can talk about the truth or falsity of the interpreted sentences.

- ▶ The goal is to generate an interpretation of our argument that results in the premises all being true and the conclusion false. This will conclusively show that the argument form is invalid.
- ▶ If we want to construct an interpretation, we need first to give meaning to the quantifiers. We do this by specifying what is called a domain of discourse, or a universe.
- ▶ When we use a universal quantifier to say “all y ,” what is the “all” we are talking about? It can be any set. It could be the set of all politicians or the set of all brown things. It could be the set of all things whatsoever—what is called an unrestricted domain.
- ▶ Often, we will pick sets of numbers to be our domain of discourse because their properties are well defined and well behaved.
- ▶ Next, we assign a specific member of the universe to each constant. If the domain of discourse is the set of positive integers, each constant becomes a particular number. If the domain is brown things, then each constant becomes a specific brown thing—not, for example, mud, but rather a particular mud stain on a particular shirt.
- ▶ Each constant gets mapped onto a particular member of the universe, but two different constants can get mapped onto the same thing. People have multiple names. Some might call you by your name, while others might call you “mother” or “father,” for example. The same can be true for constants.
- ▶ Finally, we take each property and assign it a meaning by selecting a subset of our domain of discourse. This can be done in two ways. We can simply choose members and put them together to form a subset. The property is then being a member of that subset. Or—and this is the more common way—we specify an actual property.

- ▶ Note that the empty set, the set with no members, is a subset of every set. That means that it is acceptable to have empty categories. Indeed, sometimes it will be quite advantageous. For example, we can use “being a unicorn” as a property.
- ▶ Once we have specified a domain of discourse, chosen a member for each constant, and assigned subsets to our properties, we can then translate our formal propositions into spoken language and determine if the premises of our argument are true and if our conclusion is false on this interpretation. If so, we have found a counterexample, and our argument is invalid.

THE METHOD OF EXPANSION

- ▶ In addition to the method of counterexample, the other invalidity test is the method of expansion. Coming up with counterexamples requires that we fully endow our argument with semantic content. Expansions only require some semantic content, but not full meaning.
- ▶ We do this by starting with a small universe that contains only two things: a and b , for example. We then see if we can make the premises true and the conclusion false for this small universe.
- ▶ We do this by eliminating the quantifiers, reducing the problem to truth-functional logic. If there are only two things, then our universal quantifier becomes a conjunction.
- ▶ The sentence $\forall xFx$ says that everything is F , but everything in our small universe is just a and b . So, $\forall xFx$ is equivalent to $Fa \& Fb$. Similarly, $\exists xFx$ says that something is F . That thing has to be either a , b , or both, because that is all there is. So, $\exists xFx$ is equivalent to $Fa \vee Fb$. This holds for any sentence in our first-order predicate language.

- ▶ Consider the sentence $\forall x[(Dx \vee Nx) \rightarrow Jx]$. We see that we have a universal quantifier, so we start by writing down a conjunction. On the left side, we replace every instance of the quantified variable—in this case, x —with a . On the right side, we replace every instance of x with b . This gives us $[(Da \vee Na) \rightarrow Ja] \& [(Db \vee Nb) \rightarrow Jb]$.
- ▶ Suppose that we have a complex existential quantification—for example, $\exists y[(Gy \& \neg Fy) \vee (\neg Gy \& Fy)]$. It is an existential quantification, so we start by writing down a disjunction, and on the left side, we replace every instance of the quantified variable—in this case, y —with a . On the right side, we replace every instance of y with b . This gives us the following.

$$[(Ga \& \neg Fa) \vee (\neg Ga \& Fa)] \vee [(Gb \& \neg Fb) \vee (\neg Gb \& Fb)]$$

- ▶ But what if it is not a quantification, but a truth-functional combination of quantified propositions? We can handle that, too. Recall that the scope of a quantifier ends at a truth-functional connective that is not shielded within parentheses. So, what we do is keep all of the truth-functional connectives in place and individually expand the quantifications.
- ▶ How is this an invalidity test? When we expand the sentences, the quantifiers are gone. All we have are sentences that can be true or false and truth-functional connectives.
- ▶ We have successfully reduced a first-order predicate question to a truth-functional question, and we know how to answer the truth-functional question of determining whether an argument is invalid.
- ▶ We could construct a truth table and find a row in which all the premises are true and the conclusion is false. It turns out that the truth tables will quickly become unwieldy and that the preferred

approach is to plug truth-values into the atomic sentence to see if we could arrange them to give us true premises and a false conclusion.

READINGS

Barker, *The Elements of Logic*, chap. 3.

Copi, *Introduction to Logic*, chap. 9.

Hurley, *Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 5.

QUESTIONS

1.

Give an interpretation to show that the following argument is invalid.

$\forall x(Fx \rightarrow Gx)$. $\forall x(Fx \rightarrow Hx)$. Therefore, $\forall x[(Gx \& Hx) \rightarrow Fx]$.

2.

Create an expansion to show that the following argument is invalid.

$\exists x(Fx \& Gx)$. $\forall x(Fx \rightarrow Hx)$. Therefore, $\forall x(Fx \& Hx)$.

ANSWERS

1.

Domain = animals; F = human, G = warm-blooded, H = has a heart.

$\forall x(Fx \rightarrow Gx)$ = All humans are warm-blooded. (true)

$\forall x(Fx \rightarrow Hx)$ = All humans have a heart. (true)

$\forall x([(Gx \& Hx) \rightarrow Fx])$ = All warm-blooded animals with a heart are humans. (false)

2.

Domain = {*a, b*; *Fa* = T, *Fb* = F, *Ga* = T, *Gb* = F, *Ha* = T, *Hb* = F.

$$\exists x(Fx \& Gx) = (Fa \& Ga) \vee (Fb \& Gb) = (T \& T) \vee (F \& F) = T \vee F = T.$$

$$\forall x(Fx \rightarrow Hx) = (Fa \rightarrow Ha) \& (Fb \rightarrow Gb) = (T \rightarrow T) \& (F \rightarrow F) = T \& T = T.$$

$$\forall x(Fx \& Hx) = (Fa \& Ha) \& (Fb \& Hb) = (T \& T) \& (F \& F) = T \& F = F.$$

Relational Logic

First-order predicate logic exposes the logical structure within sentences. But the language as we have developed it so far can only handle certain kinds of internal logical connections: the attributions of properties to individuals. In order to be able to more fully express the kind of content we put in our arguments, we need to add another piece of machinery: relations between individuals. In doing this, translation becomes significantly trickier because when we need to assert a relation between variables, we will need a different quantifier to bind each variable. Multiple quantifiers as well as mixtures of properties and relations combine to make translating into our extended first-order language more difficult.

RELATIONS WITH CONSTANTS

- ▶ Just as we did for properties, we will use capital letters to represent relations. The only difference is that the number of lowercase letters that follow the capital letter will increase. We translated “Phil is tall” into our first-order predicate language as Tp , where T is the property of being tall and p is the constant for Phil.
- ▶ If we take T to be the relation “taller than” and j to be the constant for Jose, then we will translate “Phil is taller than Jose” as Tpj . Relations will look like properties, except that they will have more constants or variables.
- ▶ We can combine these to create truth-functional combinations, as in the sentence “Phil is tall, but he is not taller than Jose.” In this case, the main connective is “and.” Remember that “but” is just another way of saying “and.”

- ▶ So, we have the following.

$(\text{Phil is tall}) \& (\text{It is not the case that Phil is taller than Jose})$

- ▶ “Phil is tall” is just Tp .

$Tp \& (\text{It is not the case that Phil is taller than Jose})$

- ▶ “It is not the case that Phil is taller than Jose” is a truth-functional combination with the main connective “not.”

$Tp \& \neg(\text{Phil is taller than Jose})$

- ▶ “Phil is taller than Jose” is Tpj , so the final formulation of the sentence in our first-order language is as follows.

$Tp \& \neg Tpj$

- ▶ You might think that we should have translated the sentence as follows, because Jose is taller than Phil if Phil is not taller than Jose.

$Tp \& Tjp$

- ▶ But this is problematic. Suppose that Jose and Phil are the same height; then, neither is taller than the other. The negation of Txy is not necessarily Tyx .
- ▶ Some relations are symmetric—that is, if x bears the relation to y , then y bears the relation to x . Consider the relation “is the sibling of.” If you are Bob’s sibling, then he is yours. It is a symmetric relation.
- ▶ Other relations are asymmetric. If x bears the relation to y , then y does not bear it to x . Think of the relation “is the parent of.” If x is the parent of y , then y is not the parent of x .

- ▶ Yet other relations are nonsymmetric. If x bears the relation to y , then y might or might not bear the relation to x . Think of the relation “is in love with.” When x loves y and y loves x , life can be happy, but just because x loves y does not mean that y loves x .
- ▶ To return to our example, when we translate, we need to translate the sentence itself. To take “Phil is not taller than Jose” and turn it into “Jose is taller than Phil” is to import external knowledge about the “is taller than” relation, knowledge that is not in the structure of the sentence—that is, not part of the syntax.
- ▶ But, rather, this new knowledge is in the semantic content. We know what “is taller than” means and that (with the rare exception of people of the same height) it is antisymmetric. But when we look at the validity of a relational argument, we need to ask what the syntax alone tells us. If the relation is antisymmetric, then we will be able to formulate that as a sentence that will be a premise in the argument.
- ▶ We will be able to say that if one person is taller than a second, then the second is not taller than the first. That is content we will need to make explicit as a premise. It cannot be simply assumed in the translating because it is not part of the syntactic structure of the sentence being translated.
- ▶ So far, we have looked only at relations between two individuals, but we are in no way limited to them. Think of the relation “between.” It requires three things, one to be in the middle and two to be on either side. We could write “Kim is between Sharleen and Gregory” as *Bksg*. This is the usual way of writing it—that we put the subject first.
- ▶ The objection is that it is a between-ness relation, yet the thing between is not visually between in the representation in our language. One could define the symbol for between-ness such that “Kim is between Sharleen and Gregory” is represented *Bskg*.

That is fine, but it's not how we often do it. We will have to be clear how we define the relations in this extension of our language.

RELATIONS WITH VARIABLES

- ▶ Relations with constants are straightforward. Define the relation, and part of that definition will be which place after the capital letter stands for which place in the relation. But aside from that, it is simple. This is not the case once we start in with variables and quantifiers.
- ▶ Suppose that we want to say, "Phil is taller than something or other." We need a constant for Phil but a variable for the unknown thing that Phil is taller than. It is something, not everything, so the quantifier for that variable will be existential.

$$\exists x T p x$$

- ▶ This, of course, is different than "something is taller than Phil," which would be the following.

$$\exists x T x p$$

- ▶ The order matters.
- ▶ The move to the universal quantifier should be just as simple. Phil is taller than everything: $\forall x T p x$. Everything is taller than Phil: $\forall x T x p$.
- ▶ The next move is the Pandora's box of first-order logic. What about "something is taller than something"? We need multiple quantifiers. We are saying that there is a thing and there is another thing, and the first thing bears the "taller than" relation to the second thing.

$$\exists x \exists y T x y$$

- ▶ Notice that this is exactly the same sentence as the following.

$$\exists y \exists x T y x$$

- ▶ Remember that bound variables—that is, variables that are in the scope of a quantifier—are dummy variables. They are just placeholders that show us which quantifier they are connected to. With properties, this never really came into play, but with relations, it will be a relevant fact with great regularity.
- ▶ For intuition's sake, let's switch relations to the loving relation and restrict ourselves to talking about people. Suppose that we want to say, "Someone loves someone or other." We know that this is the following.

$$\exists x \exists y L x y$$

- ▶ There exists a thing such that there is another thing, and the first thing loves the second one.
- ▶ Suppose that we want to say, "Someone loves everyone." Now, we are saying that there is a thing, and that thing bears the loving relation to all other things.

$$\exists x \forall y L x y$$

- ▶ This is a different sentence from "Everyone loves someone or other." That is saying that for all things, there is a thing that the first thing loves.

$$\forall x \exists y L x y$$

- ▶ We need to be careful here, because there are two different propositions we need to keep separate in our mind. One is "Everyone loves someone or other." This means that if you picked any given person, that person loves some person.

- ▶ This is different from the sentence “There is a person everyone loves.” Now, we are saying that there is a thing, and that thing is loved by all other things.

$$\exists y \forall x Lxy$$

- ▶ Notice how changing the order of the quantifiers has a radical effect on the meaning of the sentence. $\forall x \exists y Lxy$ means that for each and every person, there is some person out there whom that person loves. But switch around the quantifiers and we get $\exists y \forall x Lxy$, which says that there is some person who is universally loved. That is not something that follows from the first sentence. The order of the quantifiers matters just as much as the order of the variables.
- ▶ Instead of just the two quantified basic forms that we had with properties, we now have eight.

1	$\exists x \exists y Lxy$	Someone loves someone or other.
2	$\exists y \exists x Lxy$	Someone is loved by someone or other.
3	$\forall x \forall y Lxy$	Everyone loves everyone.
4	$\forall y \forall x Lxy$	Everyone is loved by everyone.
5	$\exists x \forall y Lxy$	Someone loves everyone.
6	$\forall x \exists y Lxy$	Everyone loves someone or other.
7	$\exists y \forall x Lxy$	There is someone who is loved by everyone.
8	$\forall y \exists x Lxy$	Everyone is loved by someone or other.

- ▶ There is one more. All of these presume that the one loving and the one loved could be different. But what about the sentences “Someone loves himself/herself” and “Everyone loves himself/herself”? For these, the one doing the loving, the first variable, and the beloved, the second variable, are the same thing. As such, they get the same variable and are bound by the same quantifier.

$\exists x Lxx$	Something loves itself.
$\forall x Lxx$	Everything loves itself.

- ▶ So, when we go through the process of asking whether something is a truth-functional combination of quantified propositions or is a fundamental form, and we find that it is fundamental, and then ask whether it is categorical or basic, and we answer basic, we have 10 new basic forms we need to know.
- ▶ With the categorical forms, we move up another notch of complexity. We will keep the basic form of the *A*, *E*, *I*, and *O* sentences, but the placement of quantifiers and the placement of parentheses will become very important.
- ▶ Consider the *A* sentence “All who love someone are loved by someone or other.” The sentence makes clear that if person *A* loves person *B*, then there is a person *C*, who might not be *A* or *B*, who loves *A*.
- ▶ We know that this is an *A* sentence, so it will have the general form $\forall x(Ax \rightarrow Bx)$, but it will need to be made more intricate to handle the relations. The key is to go step by step, partially translating as we go. So, we start with the following.

All people who love someone or other are loved by someone.

- ▶ We identify it as an *A* sentence and add in the structure, while keeping the content in place not related to the general form.

$\forall x(x \text{ loves someone or other} \rightarrow \text{someone loves } x)$

- ▶ Because the person who *x* loves in the antecedent might be different from the person who loves *x* in the consequent, they will need different quantifiers. And because neither appears in both terms, the quantifiers will be a part of each term. Because *x* appears in both, the quantifier for *x* must appear outside of the parentheses in the order that both terms lie within its scope.

$\forall x[y(x \text{ loves } y) \rightarrow \exists z(z \text{ loves } x)]$

$$\forall x(\exists yLxy \rightarrow \exists zLzx)$$

READINGS

Barker, *The Elements of Logic*, chap. 5.

Hurley, *Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 8.

QUESTIONS

1.

Translate the following sentences into first-order relational logic. Use $Hx = x$ is human, $Mxy = x$ and y are married, $Pxy = x$ is a parent of y , $Gxy = x$ is a grandparent of y , $j = \text{John}$, $m = \text{Mary}$, $r = \text{Roberta}$.

- John and Mary are married and are the parents of Roberta.
- Everyone has a person who is his/her parent.
- Anyone who is the parent of one's parent is one's grandparent.

2.

Translate the following argument into first-order relational logic and construct a proof for it to show that it is valid.

Everyone has a person who is his/her parent. Anyone who is the parent of one's parent is one's grandparent. Therefore, everyone has a grandparent.

ANSWERS

1.

- a $Mjm \& (Pjr \& Pmr)$
- b $\forall x[Hx \rightarrow \exists y(Hy \& Pyx)]$
- c $\forall x \forall y \forall z [(Pyx \& Pzy) \rightarrow Gzx]$

2.

Everyone has a person who is his/her parent.

$$\forall x[Hx \rightarrow \exists y(Hy \& Pyx)]$$

Anyone who is the parent of one's parent is one's grandparent.

$$\forall x \forall y \forall z [(Pyx \& Pzy) \rightarrow Gzx]$$

Therefore, everyone has a grandparent.

$$\forall x(Hx \rightarrow \exists yGyx)$$

1	$\forall x[Hx \rightarrow \exists y(Hy \& Pyx)]$	premise
2	$\forall x \forall y \forall z [(Pyx \& Pzy) \rightarrow Gzx]$	premise
3	Hx	Assumpt
4	$Hx \rightarrow \exists yPyx$	UI 1
5	$\exists yPyx$	MP 3,4
6	$Hw \& Pwx$	EI 5
7	Hw	Simp 6
8	Pwx	Simp 6
9	$Hw \rightarrow \exists y(Hy \& Pyw)$	UI 1
10	$\exists y(Hy \& Pyw)$	MP 7,9
11	$Hz \& Pzw$	EI 11
12	Hz	Simp 11
13	Pzw	Simp 11

14	$Pwx \& Pzw$	Conj 12,13
15	$\forall y \forall z [(Pyx \& Pzy) \rightarrow Gzx]$	UI 2
16	$\forall z [(Pwx \& Pzw) \rightarrow Gzx]$	UI 15
17	$(Pwx \& Pzw) \rightarrow Gzx$	UI 16
18	Gzx	MP 14,17
19	$\exists y Gyx$	EG 18
20	$Hx \rightarrow \exists y Gyx$	CP 3–19
21	$\forall x (Hx \rightarrow \exists y Gyx)$	UG 20

Introducing Logical Identity

There is one final addition that needs to be made to our first-order relational logical language, a new relation called identity. Two things are identical only if they are the same thing—that is, two individuals are identical if and only if they are different names for the same thing. Identity is an equivalence relation—that is, it is reflexive, symmetric, and transitive. Once we have it, we can start to translate sentences that include quantities such as “I have two arms” and “I would like at least three slices of pizza.”

IDENTITY AS AN EQUIVALENCE RELATION

- ▶ Identity is an equivalence relation. We talked about truth-functional equivalence when we were examining everything that could be done with truth tables. We saw that two sentences were truth-functionally equivalent if and only if they always shared the same truth-value—that is, the arrangement of Ts and Fs in the columns in a truth table were identical.
- ▶ We made use of this notion to give us our nine equivalences in our system of natural deduction proof. This provided us with new tools that we could use to substitute one sentence form for another in a fashion that still guaranteed truth. These tools greatly enlarged the number of arguments we could construct proofs for.
- ▶ But just as we expanded everything else in first-order predicate logic so that we could move beyond just talking about truth-functional relations between sentences, when we talk about identity in first-order logic, we will move from the notion of equivalent sentences and place our equivalence relation inside of sentences.

- ▶ With identity, we now want to talk about the equivalence of individuals. We want to be able to say that two things are the same thing. This is a different notion of equivalence. Two sentences can be truth-functionally equivalent and say different things—have different meanings. So, we need a new, different, more robust concept of equivalence.
- ▶ Equivalence relations are relations that have three important properties: *r*, *s*, and *t*. Equivalence is reflexive. That means that a thing is always equivalent to itself. Think of reflexive as reflective: When you look in a mirror, you never have to wonder whom you are seeing. Your reflection is always you. That is *r*.
- ▶ Next is *s*, which stands for symmetric: If *a* is equivalent to *b*, then *b* is equivalent to *a*. Some relations, such as “is the sibling of,” are symmetric. Some, such as “is the parent of,” are antisymmetric. And others, such as “is the brother of,” are nonsymmetric. For equivalence, the relation must be symmetric.
- ▶ Finally, we have *t*, which stands for transitive. If *a* is equivalent to *b* and *b* is equivalent to *c*, then we know that *a* is equivalent to *c*. This should remind us of the kind of reasoning behind the hypothetical syllogism, the rule of inference in natural deduction proofs where if we have $a \rightarrow b$, and $b \rightarrow c$, we can infer $a \rightarrow c$.
- ▶ If you have all three of these properties, then you get the most important of all logical relations: If you have a relation that satisfies *r*, *s*, and *t*—that is, reflexive, symmetric, and transitive—then you have equivalence.

TRANSLATING WITH IDENTITY

- ▶ When we apply this relation to individuals, we have identity. For identity, we will use in our language the universal symbol for equivalence, the equal sign ($=$). It will not be a truth-functional

connective that applies to sentences, but rather a relation within sentences that shows the equivalence between individuals—that is, constants and variables.

- ▶ What makes identity the most important of all logical relations is that it allows us to translate into our logical language all kinds of sentences that we could not translate before.
- ▶ Consider the following two sentences: “Dave is taller than everyone” and “Dave is taller than everyone else.” The first sentence is false. Dave is someone. So, if it is true that Dave is taller than everyone, then it would have to be true that Dave is taller than Dave.
- ▶ But we know that the “is taller than” relation is anti-reflexive—that is, because everything is always the same size as itself, nothing can be taller than itself. It is not a contradiction because this knowledge about the “is taller than” relation is something semantic that we are importing. It is false because of the meaning of “is taller than,” not because of the form of the sentence. But with the interpretation we have given that form, we know that the sentence is always false.
- ▶ The other sentence, “Dave is taller than everyone else,” however, could be true. It says that Dave is taller than everyone who is not Dave. To know whether that sentence is true or false, we need to know not just the form of the sentence and the nature of the “is taller than” relation, but now we need to know how tall Dave is and how tall everyone else is. These are completely different sentences.
- ▶ The first sentence could be translated into our language already: “Dave is taller than everyone.” It is of the fundamental form, and it is categorical. It is an *A* sentence. “All people are things that Dave is taller than.” We know that an *A* sentence, then, takes the following logical form.

$$\forall x(x \text{ is a person} \rightarrow \text{Dave is taller than } x)$$

- ▶ We need a letter for the property “is a person.” We use H for that. We need a letter for the “is taller than” relation. Let’s use T . And we need a constant for Dave; d would be the obvious candidate. This, then, gives us the following.

$$\forall x(Hx \rightarrow Tdx)$$

- ▶ Without equivalence, we cannot translate the second sentence without some philosophical weirdness. The sentence is “Dave is taller than everyone else.” It is of the fundamental form, and it is categorical. It is still an A sentence, but the slight difference is in the antecedent.

$$\forall x(x \text{ is a person other than Dave} \rightarrow \text{Dave is taller than } x)$$

- ▶ To say that “ x is a person other than Dave” is to assert a conjunction; we are saying that x is a person, and x is not Dave.

$$\forall x[(x \text{ is a person} \ \& \ x \text{ is not Dave}) \rightarrow \text{Dave is taller than } x]$$

- ▶ We know how to do most of this already.

$$\forall x[(Hx \ \& \ x \text{ is not Dave}) \rightarrow Tdx]$$

- ▶ But what do we do with “ x is not Dave”? The intuitive move many people make is to make a property of Dave-ness and assert that x does not have this property.

$$\forall x[(Hx \ \& \ -Dx) \rightarrow Tdx]$$

- ▶ Making a property out of an entity, though, opens some strange philosophical problems. How do we define Dave-ness? Dave is always changing. Are there essential properties that belong always and only to Dave?

- ▶ Give us any set of properties and we can usually point to either other things that have them and are not Dave, or we can point to times of Dave's life or after his death when Dave would have to be categorized as not-Dave. But Dave is always Dave.
- ▶ But even if you could come up with a satisfactory set of properties that cluster to form Dave-ness, we still want to be able to meaningfully talk about hypothetical cases in which Dave is still Dave but lacks certain aspects that Dave has.
- ▶ Would Dave have been so smart and successful if he had different parents or if he had gone to different schools? We could have interesting discussions about these types of questions and come up with rational arguments that seem to be sensible. But what we are doing is adding and subtracting properties from Dave in order to do it. As such, taking Dave not only to be an individual but also a property seems to be something we want to avoid.
- ▶ Recall the purpose of developing this artificial language: to have a tool that allows us to reason clearly and rigorously without the ambiguity and sloppiness of ordinary spoken language. By allowing individuals to become properties in this way, we are inviting exactly the kinds of problems we are trying to avoid. To escape this, we use identity.
- ▶ That is not to say that there are no problems with self-identity. What does it mean to be the same thing over time? But, while interesting, these are not new problems we are introducing through our language.
- ▶ Using equivalence, then, we can finish our translation.

$$\forall x([Hx \& \neg(x=d)] \rightarrow Tdx)$$

- ▶ Dave is taller than everyone else.

- ▶ Perhaps the most important sentences we can now translate are ones involving quantities. Consider the sentence “Steve’s class has at least two students registered for it.” Previously, we might have wrongly thought that this is an *I* sentence: “Some students are registered for Steve’s class.” We know how to translate an *I* sentence. In the following, *S* is the property of being a student, and *R* is the property of being registered for Steve’s class.

$$\exists x(Sx \& Rx)$$

- ▶ This says that there is at least one student registered for Steve’s class. So, we might have wrongly thought that this is a hidden truth-functional combination and that it really is a conjunction of two *I* sentences.

$$\exists x(Sx \& Rx) \& \exists y(Sy \& Ry)$$

- ▶ That says, “There is one student *x* who is registered for Steve’s class and a student *y* who is also registered for Steve’s class.”
- ▶ The problem here is that we do not know who *x* or *y* is. All we know from these sentences is that *x* is some member of our domain of discourse and *y* is some member of our domain of discourse.
- ▶ Could both of these be true and *x* and *y* be the same thing? Absolutely. What we need is an additional claim, an assurance that *x* and *y* refer to different members of the universe. And for that, we need equivalence.

$$\exists x \exists y [(Sx \& Rx) \& (Sy \& Ry)] \& \neg (x=y)$$

- ▶ Now that says, “There are at least two students registered for Steve’s class.”

- ▶ Notice what we had to do with the quantifiers, grouping them at the front of the sentence. The reason is that the equivalence (or, in this case, the negation of the equivalence) relates x to y , so the equivalence needs to be within the scope of both existential quantifiers.
- ▶ The jurisdiction of a quantifier ends at a truth-functional connective. So, in the second formulation of the sentence, $\exists x(Sx \& Rx) \& \exists y(Sy \& Ry)$, the $\exists x$ logically disappears once we hit the central conjunction.
- ▶ In order to have the equivalence within the scope of both quantifiers, we need to put the entire rest of the sentences inside the braces and put the quantifiers outside so that they reach the equivalence portion.

EQUIVALENCE SUBSTITUTION

- ▶ We can now translate sentences that require equivalence into our language, and that means that we can translate arguments. Can we determine if they are valid?
- ▶ We will need one new rule for our proof structure: equivalence substitution (ES). It is, perhaps, the most obvious rule. If, on some line of a proof, we have, as the entire content of the line, a statement of equivalence between two individuals, then on any subsequent line in which either individual appears in a sentence, we can substitute the other equivalent individual.
- ▶ Note that this only applies to free variable names—that is, variables that are not in the scope of a quantifier. We cannot change a bound variable name—that is, a quantified variable—without changing the meaning of the sentence.

- ▶ So, we can use ES on free variables and constants. When we say that Mark Twain is Samuel Clemens, we are asserting the equivalence of two names. We can equate two variables, as in cases in which we claim that whoever is stealing your lunch is also the person who is spray painting graffiti on the walls. And when we say that the butler did it, we are equating a named individual, Jeeves, with a free variable, the previously unnamed murderer. We can use ES on an equivalence that connects any free variable(s) or constant(s).

READINGS

Hurley, *Logic*, chap. 8.

Kahane, *Logic and Philosophy*, chap. 8.

QUESTIONS

1.

Translate the following sentences, using $Ax = x$ is an apple; $Hxy = x$ has y .

James has an apple, and Florence has a different apple.

2.

Translate the following argument and construct a proof to show that it is valid.

James has at least two apples. Therefore, James has an apple.

ANSWERS

1.

$\exists x\exists y\{[(Ax \& Hjx) \& (Ay \& Hfy)] \& -(x=y)\}$.

2.

$\exists x\exists y\{[(Ax \& Hjx) \& (Ay \& Hjy)] \& -(x=y)\}$; therefore, $\exists x(Ax \& Hjx)$.

- | | | |
|---|--|---------|
| 1 | $\exists x\exists y\{[(Ax \& Hjx) \& (Ay \& Hjy)] \& -(x=y)\}$ | premise |
| 2 | $\exists y\{[(Ax \& Hjx) \& (Ay \& Hjy)] \& -(x=y)\}$ | EI 1 |
| 3 | $\{[(Ax \& Hjx) \& (Ay \& Hjy)] \& -(x=y)\}$ | EI 2 |
| 4 | $(Ax \& Hjx) \& (Ay \& Hjy)$ | Simp 3 |
| 5 | $Ax \& Hjx$ | Simp 4 |
| 6 | $\exists x(Ax \& Hjx)$ | EG 5 |

Logic and Mathematics

As you learned in the previous lecture, equivalence relation allows us to translate into our language numerical sentences. This means that we can now ask questions about the form of mathematical reasoning. This is the question that launched the study of deductive logic in the first place. If our study of deductive logic looked like work you have seen in mathematics classes, that's because the two studies are inextricable. We cannot discuss mathematics without discussing logic, and we cannot discuss logic without discussing mathematics.

LOGIC IN EUCLIDEAN GEOMETRY

- ▶ Aristotle was the first to develop a systematic approach to formal reasoning with his categorical logic, but the hallmark of formal reasoning, the book that brought deduction to the forefront of the Western intellectual tradition, was written about 50 years later. It was Euclid's *The Elements*.
- ▶ Euclid's use of deductive proofs started with five axioms and five postulates that were all seemingly self-justifying—that is, they were so obvious that no one could possibly deny their truth. He starts with general mathematical ideas, such as equals added to equals yields equals.
- ▶ But what was amazing is where he went from there. He used these basic truths, and then added some definitions. Definitions have to be true because they are true by definition. We are free to define any term in any way we want. So, Euclid gives us definitions of “point,” “line,” “circle,” and other geometric terms.

- ▶ But then he combines all of this and, using strict deductive reasoning, demonstrates the necessary truth of hundreds of theorems—that is, propositions he derives from the axioms and postulates. These are not obvious. Some of them are incredibly intricate, and some of them are counterintuitive. But after reading Euclid’s work, one is convinced that each must be true.
- ▶ What Euclid did that made his book one of the most important in human history is his organization of the theorems of plane geometry—his use of a strict deductive system to derive them all from first truths.
- ▶ Euclid’s use of deduction was so elegant and powerful that it made thinkers who came after him long to do for every aspect of human thought what he did for geometry. So, Euclid’s *Elements* reigned as the unchallenged epitome of rigorous thought—that is, until some tried to improve it.
- ▶ Mathematicians love elegance: doing the most with the least, not making any more assumptions than are absolutely necessary. Euclid had made 10 such assumptions—that is, he had five axioms and five postulates.
- ▶ But one of the postulates bothered mathematicians—the fifth one. It seemed different. The following are the five postulates.
 - 1 A line segment can be drawn connecting any two points.
 - 2 A line segment can be extended indefinitely far in either direction.
 - 3 A circle of any radius can be drawn around any point.
 - 4 All right angles are equal to one another.
 - 5 Take any two straight lines and have a third line fall across them. If, on one side of the crossing line, the other two lines are approaching each other, they will eventually meet on that side of the crossing line.

- ▶ The fifth postulate seems much more like one of the many theorems that are proven in the book than the other postulates, which are more about what can be constructed. So, some thought, maybe it is a theorem. If it is, then it could be proven from the definitions, the five axioms, and the other four postulates.
- ▶ If you could create a deductive argument with the fifth postulate as a conclusion, then it would not need to be assumed—it would be omitted from the list of postulates. Diminishing this list would create a more elegant system, and that would be astounding. Over the centuries, many mathematicians tried. All failed.
- ▶ Despite the universal failure, the nagging sense persisted. So, some changed tactics. Instead of constructing what we have called a direct proof, maybe they could instead proceed indirectly, in which one sets out the premises and the negation of the conclusion and reasons until a contradiction rises.
- ▶ Many attempts were made, but no one could show that this process led to explicit contradictions, and this is what is needed to conclude an indirect proof. After a while, several mathematicians thought that perhaps the contradictions would never arise. If that were true, then what they were playing with was not a self-contradictory set of premises, but rather propositions that were consistent.

THE DISCOVERY OF NON-EUCLIDEAN GEOMETRY

- ▶ But if that were true, what they had in their hands was a new geometry—a geometry other than that of Euclid, a non-Euclidean geometry. While a handful of mathematicians nearly simultaneously came to this conclusion, the name most associated with this move to asserting the existence of a geometry different from that of Euclid is the Russian mathematician Nicolay Lobachevsky, and the system is now widely referred to as Lobachevskian geometry.

- ▶ The introduction of Lobachevskian geometry set off alarm bells in the world of mathematics. Non-Euclidean geometry seemed like a contradiction in terms. Euclid had used deduction. Deductive reasoning guarantees the truth of its conclusions. How could there be alternative mathematical truths?
- ▶ Mathematical truths are not like physical truths. Things in the world could be otherwise, but mathematical truths seem to be necessary. They have to be true.
- ▶ So, we have one geometry with one set of axioms and a second geometry with a different set. Which one is true? Which set of axioms should we believe?
- ▶ Surely, nearly everyone thought, it must be Euclid. How could we possibly reject Euclid? So, the push continued to find the contradiction that lay hidden among the results of Lobachevskian geometry. If a single contradiction could be found, the whole non-Euclidean threat could be rejected.
- ▶ But it was a pointless search. This fact was proven by three different mathematicians: the German Felix Klein, the Frenchman Henri Poincaré, and the Italian Eugenio Beltrami. Each constructed what we call a relative consistency proof for non-Euclidean geometry.
- ▶ A consistency proof is an argument that establishes that the axioms of a system are non-contradictory. A full-out consistency proof for non-Euclidean geometry would have been wonderful and established beyond any doubt that non-Euclidean geometries were non-contradictory. But we got something only half as good: a relative consistency proof.
- ▶ Klein, Beltrami, and Poincaré took the language of the non-Euclidean universe and erased all of the meanings. They then gave the non-Euclidean terms a new meaning in the Euclidean world and showed that the relations among the reinterpreted

axioms are true in the Euclidean universe. They made Euclidean models of non-Euclidean geometry.

- ▶ The radical result is that the only way the axioms could form an inconsistent set—that is, the only way non-Euclidean geometry would ever produce a contradiction—is if Euclidean geometry did.
- ▶ Because we could show a formal relation between the sentence forms and because consistency and inconsistency are formal properties, we could prove that the two systems are relatively consistent. We did not show that non-Euclidean geometry is consistent, but rather that it is only inconsistent if Euclidean geometry is inconsistent.
- ▶ But is Euclidean geometry consistent? Mathematicians thought, if it is, we should be able to prove it. So, how can we prove that Euclidean geometry is consistent?
- ▶ We had been working with it for centuries, and no one had found a single contradiction yet. But just because a contradiction had not been found yet does not mean that one will not surface eventually. We needed a proof.
- ▶ But before we could prove Euclid's geometry consistent, a small detail needed to be taken care of. It was long known that Euclid had played fast and loose with some of his proofs and had made use of some assumptions that he had not accounted for in his axioms and postulates. If we were going to prove that Euclidean geometry was consistent, then we first needed to clean up Euclid's work.

LOGIC IN MATHEMATICS

- ▶ We needed a complete re-axiomatization of Euclidean geometry and a rigorous formalization of the logic to be used. It was a big project that would require a big brain and intellectual courage.

The man for the job was one of the greatest minds in the history of mathematics, David Hilbert.



DAVID HILBERT
(1862–1943)

- ▶ Unlike some mathematicians of his time, Hilbert saw logic as the most powerful tool of the mathematician. He embraced it and the problems that came from it. Contrary to constructivism—which held that logical proofs were not to be admitted in mathematics—Hilbert became the founding father of formalism, the view that mathematics was about the forms of equations and mathematical entities.
- ▶ This formalist project needed a formal methodology. With his student Wilhelm Ackermann, Hilbert wrote a book called *Principles of Mathematical Logic*. Logic was to be the method, and it needed to be made rigorous, so Hilbert developed what we have come to know as natural deduction proof.
- ▶ It was Hilbert who first developed the system we have been studying, although his version is a bit different. Hilbert’s system has many fewer rules, but philosophically, the foundation is the same. We have a formal process that demonstrates validity, a formal property using purely formal means.
- ▶ He took this view and attacked many of the major mathematical problems of his day—among them, the question of geometry.

- ▶ Poincaré, Beltrami, and Klein had produced relative consistency proofs. They had shown that the only way that the non-Euclidean geometry of Lobachevsky could be inconsistent, self-contradictory, is if Euclidean geometry was as well.
- ▶ No one doubted the consistency of Euclid, but no one had a proof of it either. And if we are mathematicians, we have warrant for rational belief only if we have a proof. So, is Euclidean geometry consistent? We need a proof.
- ▶ But before we could prove Euclid's system consistent, we needed a reformulation of Euclid. We needed for all and only the assumptions, the real axioms, to be set out. We needed a strict means of logical process—natural deduction.
- ▶ And we needed the theorems of Euclid to be re-derived. Hilbert did exactly this in his book *The Foundations of Geometry*, published in 1899. In it, he replaces Euclid's 30 definitions, five postulates, and five axioms with 20 axioms. He then proceeds to do the unthinkable—he improves Euclid. Hilbert had created a completely logically precise version of Euclidean geometry.

SAVING GEOMETRY BY REDUCING IT TO ARITHMETIC

- ▶ The question could now be approached. Is Euclid consistent? For that, Hilbert could not produce a proof. But he could take the next step down the road begun by Klein. Hilbert could produce a relative consistency proof for Euclid, showing that the only way that Euclidean geometry held inconsistencies is if arithmetic did.
- ▶ Math is about the internal structure and relations of the concepts within the axioms, and Hilbert showed how we could produce a coherent interpretation of the Euclidean axioms inside the language of arithmetic.

- ▶ We could take the sentence forms of the axioms he used to reformulate Euclid and find their equivalences inside of the results of arithmetic. That means that the consistency of all of geometry now relied on the consistency of arithmetic.
- ▶ But is arithmetic consistent? Mathematicians doubted this less than they doubted the consistency of Euclid. But this certainty is unwarranted, or at least ungrounded, without a proof. Could we prove arithmetic consistent?
- ▶ In the year 1900, Hilbert gave the keynote address at the Second International Congress of Mathematics, the most prestigious gathering of mathematicians from all around the world. Hilbert's address laid out the 23 problems that required the combined efforts of the mathematical world in the 20th century. It is one of the most important addresses ever given in the history of mathematics.
- ▶ Of those 23 problems, about 12 were solved in that century—there is still debate about a few of them. For this course, the important one is the second problem: “to investigate the consistency of the arithmetic axioms.” Geometry now rested on arithmetic, and it was incumbent on mathematicians to prove a logical property of arithmetic in order to provide mathematics with the firm foundation it required.
- ▶ Logic was not only a part of mathematics, but the truth of all mathematics itself now rested on the ability to produce a proof of a logical property. Logic and mathematics had become inextricably entwined.

READINGS

Euclid, *The Elements*.

Gray, *The Hilbert Challenge*.

Kramer, *The Nature and Growth of Modern Mathematics*, chaps. 29–32.

Yandell, *The Honors Class*.

QUESTIONS

1.

David Hilbert's logical approach to mathematics upset mathematicians like Paul Gordan because the purpose of mathematics, they argued, is to produce what you prove. By proving that something cannot not exist, they contended, is not really showing us the thing, and the thing is what we want. Was Gordan correct that Hilbert was not really doing mathematics, or did they just have an old-fashioned idea of what mathematics is? Does mathematics have revolutions that change how we think of the mathematical realm like scientific fields do?

2.

The results of Euclidean geometry are easy to visualize, and they seem to resemble our actual observations. Could observations be used as a foundation for mathematical truth? Doing so would mean that they would be approximate and not necessary truths. Does this diminish mathematics? Does the fact that we have a difficult time envisioning non-Euclidean geometry mean that it is or has to be false? What role does and should our ability to envision something play in determining when a mathematical proposition is true or false?

Proof and Paradox

At the turn of the 20th century, the foundations of mathematical truth hinged on logic. The question of the consistency of Euclidean geometry had been reduced to the question of the consistency of arithmetic. So, at the beginning of the 20th century, we were left with the logical question of the consistency of arithmetic. Logic and mathematics had become entangled, but to a group of thinkers, including Gottlob Frege and Bertrand Russell, the connection would run even deeper than others had previously suspected.

ARITHMETIC IN THE 20TH CENTURY


- ▶ Traditionally, mathematics was thought to have two separate realms: geometry (shapes in space) and arithmetic (numbers and equations). Non-Euclidean geometry had the mathematical world in a tizzy, but it would be wrong to say that life was peaceful in the arithmetic realm.
- ▶ The consistency of arithmetic, the ultimate grounding of numerical truths, was something that was considered indubitable. But at the end of the 19th century, nothing was beyond doubt.
- ▶ The great German mathematician Richard Dedekind had considered that while non-Euclidean geometry had been reduced to Euclidean geometry and Euclidean geometry had been reduced to arithmetic, perhaps the can needed to be kicked one more step down the road and arithmetic needed to be reduced to set theory.

- ▶ Think about numbers as conceptual representations of sets and addition, subtraction, multiplication, and division as operations on sets—for example, merging sets or looking for members that sets have in common. This way, we would just need a complete and consistent theory of sets, and that could produce for us the long-sought foundation for mathematical truth.
- ▶ Sharing Dedekind's interest was a younger German mathematician, Georg Cantor, whose results are among the strangest in mathematics. Cantor thought about numbers as the size of sets. Two sets are of the same size if we can draw a line from each member of one set and have it end at a member of the other set, and there is no member of either set that is not paired.
- ▶ We can, therefore, talk about the size of sets without counting. If everyone who buys a concert ticket sits in one seat and we want to see if the show is sold out, undersold, or oversold, we don't need to count the seats and count the tickets sold, we just need to see if there are empty seats, if there are people without seats, or if every person has a seat and every seat has a person.
- ▶ We call this mapping. A mapping is “one-to-one and onto,” what mathematicians call a homeomorphism, if all of the members of each set are mapped onto one member of the other set. Homeomorphic sets are the same size.
- ▶ The ability of being able to say that two sets have the same number of members without having to count means that we can talk about infinite sets. Some infinite sets are the same size even though one might be a proper subset of the other—for example, even numbers and the positive integers. Some infinite sets are larger than other infinite sets—for example, real numbers and rational numbers.

- ▶ Mathematicians became suspicious of thinking about infinity. We cannot construct infinite sets, so we should not be able to speak of them. They are not part of mathematics.
- ▶ The use of logic and the use of sets in mathematics were joined not only in the minds of opponents, but also in the minds of those who saw it as the path to ultimate basis for certainty in mathematics.

QUESTIONING THE BASIS OF MATHEMATICAL TRUTH

- ▶ The question of the consistency of non-Euclidean geometry was reduced to the consistency of Euclidean geometry and that reduced to the consistency of arithmetic, which was in turn reduced to the consistency of set theory. Was this a never-ending—and therefore ultimately unjustified—path? What is the intellectual bedrock on which it all ultimately rests?
- ▶ Some philosophers, such as Plato, thought that numbers are mystical, metaphysical things, and that is where the ultimate justification resides. Other philosophers, such as Immanuel Kant, argued that it was in the human mind, in our own intuitions, that we find the justification.



IMMANUEL KANT
(1724–1804)

- ▶ But many mathematicians were loath to take these philosophical paths. They wanted absolute rigor on purely mathematical terms. For Hilbert, the answer was his view, formalism, in which we just started with an axiomatized theory that we hypothetically take to be true.
- ▶ This meant that all of mathematics was ultimately ungrounded and that math was just a game like basketball, with arbitrary rules whose results do not mean anything away from the playing field. Many mathematicians were scandalized by such trivialization of their work.
- ▶ Another view was called logicism. According to this view, there is an ultimate end to the reduction. The last stop was logic itself. Logic is necessarily consistent. If we could reduce set theory to pure logic, then there could be no doubt that it is not only consistent, but necessarily true because logic is based on definitions and tautologies.
- ▶ According to the logicists, all of mathematics is just complex logic when you get all the way down. Logic is not merely a method to be used in doing mathematics; logic is the reason for and the justification of mathematics.

THE LOGICIST PROJECT

- ▶ The project of coming up with an axiomatized theory of sets that could be reduced to logic was most famously begun by the Italian mathematician Giuseppe Peano. He used nine axioms framed in an early version of formal logic that he developed for the work.
- ▶ It involved natural numbers and a function called successor. The idea is that we can posit basic truths about a number and its successor, the number that comes next. Using this, he showed


how the truths of addition and multiplication could be worked out within the system, which necessarily obeyed the laws of logic.

- ▶ Impressed with this work, but having the sense that it needed to be taken further, was the German mathematician Gottlob Frege. He wanted to reconceptualize all of mathematics in a fashion that would be able to show how every mathematical truth could be traced back to logical concepts via absolutely rigorous deductions. There would be no gaps in the reasoning, no appeals to intuition; it would all be strict and formal.
- ▶ He knew that such rigor in deduction would require a logical language. He appreciated Peano's work but thought that an even more powerful logical language would be needed, so he constructed what he called his *Begriffsschrift*, or concept writing.
- ▶ It would be his logical language within which his completely rigorous work could be done. It is a progenitor of our first-order predicate logic, although one that is much more difficult to work in.
- ▶ Having his logical language, he began to work on setting out a version of set theory, which would be a step more abstract than Peano's. Indeed, he sought to create a logical, conceptual framework from which Peano's axioms would emerge as theorems—as results.
- ▶ For Frege, mathematics must be completely reducible to the concepts of logic. But what, then, are numbers? Think about the two seemingly quite different ways we use numbers. On the one hand, we talk about numbers abstractly. On the other hand, numbers are amounts of things.
- ▶ How can we make sense of numbers in such a way that they are not like all other normal, observable properties? How can numbers be both in the world and be abstract concepts?

- ▶ Sets are conceptual things, but we can map abstractly generated concepts onto concepts we apply to the world. But that mapping is also conceptual, so mathematics should be reducible to the underlying logic if the system can be created properly. The ultimate concept of “set” would be the basis for everything in mathematics, and it would be abstract yet applicable.
- ▶ Frege worked for years to develop this system. He set out a version of his view in his earlier book, *The Foundations of Arithmetic*, and then his full system in his masterwork, *Basic Laws of Arithmetic*. He thought he had done it.

RUSSELL’S PARADOX

- ▶ The first volume of *Basic Laws of Arithmetic* came out in 1893, and he had continued to work on the project for the better part of the following decade. In 1902, he was ready with the second volume.
- ▶ But before he could publish it, he received a letter from an unknown philosopher and logician in Britain, the young Bertrand Russell. He was one of the few who had read Frege’s work and admired it to the degree Frege thought it warranted.



BERTRAND RUSSELL
(1872–1970)

Russell agreed that logic was the basis of all mathematical truth, and Frege’s work took us along the path to successfully showing it.

- ▶ However, Russell came to realize, Frege's work could not be the last word on the matter. Frege thought he had succeeded in developing a system in which set theory and thus arithmetic derived from pure logic, thereby guaranteeing that set theory would be free of contradictions.
- ▶ But Russell informed him that, in fact, the opposite was true: Frege's system generated a contradiction, what has come to be known as Russell's paradox.
- ▶ Russell thought that Frege's problem was that his axiom for set construction was far too permissive. It allowed sentences to assert properties of sentences of its own order. This was a problem. We need a logical caste system, a hierarchy with strict regimentation and an absolute rule against semantic fraternization.
- ▶ This approach was what came to be known as Russell's theory of types, and he thought that it was sufficient to save Frege's logicist project. He thought that this sort of logical segregation would make the paradoxes disappear while still saving the idea that all math could be derived from nothing but logic and abstract concepts added in as definitions.
- ▶ Russell worked for years developing a system from which an axiomization of arithmetic would emerge. He joined forces with one of his former professors, Alfred North Whitehead, and together they built the first-order logical language that we learned to work in, and they used it to produce a mammoth work of logic: *Principia Mathematica*. It is three huge volumes from which the basic propositions of arithmetic are painstakingly derived.
- ▶ Unlike Frege's effort, it was a work that was received with great fanfare. The logicist project seemed to have been carried out. Logic saved mathematical truth and gave mathematics a firm foundation from which to emerge. For an equation to be true is for it to be provable in the system of axioms set out in *Principia Mathematica*.

GÖDEL'S INCOMPLETENESS THEOREM

- ▶ In 1931, Austrian mathematician Kurt Gödel did what Russell thought could not be done. Gödel produced a paradox from the axiom set of *Principia Mathematica*, a paradox not unlike the one Russell had pointed out could be formed in Frege's system. This result is famously known as Gödel's incompleteness theorem.
- ▶ Gödel created a kind of relative consistency proof of his own. But instead of finding a model of non-Euclidean geometry in the Euclidean language, he found a model of second-order mathematical sentences in the first-order language.
- ▶ In this way, he considers the second-order sentence "this sentence is unprovable." Either it has a Gödel number or it doesn't—that is, either it has a proof or it doesn't.
- ▶ Russell and the other logicians, to be true is to be provable in the system of logic and to be false is to be unprovable. If "this sentence is unprovable" can be proven, then it is false but has a proof. This is not allowed. If "this sentence is unprovable" does not have a proof, then it is true, but does not have a proof. Either way, the logicist definition of truth as provability within the system is flawed.
- ▶ But the problem is that this flaw is mirrored in the arithmetic itself through the Gödel numbering. The mirroring connects second-order truths to arithmetic truths.
- ▶ So, Russell and Whitehead's axioms of arithmetic necessarily either let in true sentences that cannot be proven (this makes the system incomplete) or false sentences that can be proven (this makes the system unsound). Either way, logicism as Russell worked it out is dead. Russell's attempt to save Frege's project fails.

READINGS

Coffa, *The Semantic Tradition*.

Kramer, *The Nature and Growth of Modern Mathematics*, chaps. 30–33.

Nagel and Newman, *Gödel's Proof*.

QUESTIONS

1.

Cantor's work on trans-infinite numbers requires an axiom of infinity, an assumption that posits the existence of a first infinite number. From there, all of the strangeness of his results follows. Mathematicians who were bothered by his results contended that we should not accept his initial assumption—that it was based on a misunderstanding of the infinite. This seems to be a philosophical objection. On what grounds should we accept or reject basic assumptions in mathematics? Is math ultimately philosophical? Is there a logical basis? Is it a matter of applicability to science? Can we accept any basis we want?

2.

If logicism has failed, which of its competitors is a better foundation for mathematical truth? Is Plato correct that mathematics is not about this world, but rather about an ideal world of concepts? Is Kant correct that mathematics is really a form of psychology that investigates how the human mind thinks about numbers and shapes? Is Hilbert correct that mathematics does not give us absolute truths, but only results of different axiom games we can choose to play if we want? Are empiricists correct that mathematical truths are really just a kind of physical truth—that one apple and one pear together are two pieces of fruit and that mathematical truth is no longer exact and necessary but just another variety of plain old truths?

Modal Logic

We hold there to be two different senses of truth: those sentences that happen to be true and those sentences that must be true. Our logic so far does not distinguish between garden-variety truth and necessary truth. Can we develop a logic that does account for the difference between these kinds of true sentences? This is what we call modal logic. “Modality” is the term that philosophers use for the concepts “possible” and “necessary.” Modal logic is an augmented version of our first-order relational logical language that includes new operators for possibility and necessity.

MODALITY

- ▶ We have been using a particular word repeatedly in these lectures as if it is clear what it means. The word is “true.” What does it mean for a sentence to be true? Philosophers have debated this for centuries.
- ▶ One standard definition, what we call the correspondence theory of truth, holds that a sentence is true if what it says about the universe is actually the case in the universe.
- ▶ Philosophers have long realized that there are two different types of true sentence. One type of true sentence happens to be true, but the world could be otherwise. These contingent truths are possibilities.
- ▶ But then there are sentences we think must be true no matter how the world is. We call these necessary truths. We have thought that these necessary truths are sprinkled throughout

human endeavors—in such areas as mathematics, science, ethics, and metaphysics.

- ▶ The concept of necessary truth is stronger than mere truth or that which happens to be true. The term “modality” is used by philosophers to refer to the concepts of necessity and possibility.
- ▶ The logic we have been exploring is only geared to a single concept of truth. If we want to expand it to be able to look at arguments involving possible and necessary truths, then we need to create an augmented system. We need modal logic.
- ▶ We have seen some necessary truths already. Think back to truth-functional logic. We saw that there are tautologies, sentences like “I have a brother or I don’t have a brother” that are always true. We can extend this to our first-order predicate language and come up with necessary truths like “All dogs are dogs.”
- ▶ There are sentences like these that are necessarily true. These sentences are necessarily true because of their form. Any sentence of the form $a \vee \neg a$ or of the form $\forall x(Dx \rightarrow Dx)$ will be necessary truths.
- ▶ There are some examples from mathematics, science, ethics, and religion that are supposedly necessary truths because of their content, not because of their form. This means that we will need some additional machinery to deal with the logic of modal concepts.
- ▶ We need to add two new operators to our first-order predicate language. The box (\Box) will be our symbol for “necessary,” and the diamond (\Diamond) will be our symbol for “possible.” They will function syntactically sort of like quantifiers in that they will be put in front of the sentences that they are intended to modify.
- ▶ We might say that the sentence “Everything is F ” is true—that is, we can assert $\forall xFx$. But if we want to make a stronger claim

and say that it is no accident of the universe that everything is F , but rather that it could not be otherwise—that everything must be F —then we can write $\Box \forall xFx$. This says that the sentence $\forall xFx$ is a special sentence; it is a necessary truth.

- ▶ In the same way, we could add a special notation to indicate that we don't know if $\forall xFx$ is true or not, but it is certainly possible that all things have the property F . In this case, what we are asserting is $\Diamond \forall xFx$. This says that the sentence $\forall xFx$ is like a contingency in truth-functional logic. It is a sentence that could be true.
- ▶ There is a logical relation between these two. In the same way that we have the equivalence quantifier negation that allows us to switch the order of a negation and a quantifier—for example, “not everything” is the same as “something is not”—the same holds true with our modal operators.
- ▶ We can employ the equivalence modal negation (MN) and switch the order of a modal operator and a negation. If we say that it is not the case that “Everything is G ” is necessary, then we have the sentence $\neg \Box \forall xGx$. But to say that something is not necessarily true is to say that it is possible that it is false. But this is to say that $\Diamond \neg \forall xGx$.
- ▶ Similarly, if we say that it is not possible that everything is F , then we are asserting the sentence $\neg \Diamond \forall xGx$. But to say that it is not possible for it to be true, we are saying that it is necessarily false—that is, we are asserting the sentence $\Box \neg \forall xGx$. If we move the negation to the other side of the modal operator and switch the operator, we end up with an entirely equivalent proposition.
- ▶ One way that philosophers make sense of these new operators is to give them an interpretation in terms of possible worlds. By the term “possible world,” we mean a way the world could be. Our world, the real world, is just one among the infinite possible worlds.

- ▶ When we write a diamond in front of a sentence, we are saying that the sentence is possibly true. By that, we mean that there is a possible world (or some set of possible worlds) in which this sentence is true. When we then use that sentence in valid arguments, we are seeing what else must also be true in this particular possible world.
- ▶ When we put a box in front of a sentence, we are saying that the sentence is necessarily true. A sentence is necessarily true if it is true in every possible world. No matter what world you consider, the necessary truths will be true in them. If we put a box in front of a negated sentence—that is, we say it is necessarily false—that means that it is true in no possible world.

THE MECHANICS OF MODAL LOGIC

- ▶ The translation of spoken-language modal sentences into our expanded first-order modal logical language is not different from our regular translation in our first-order language, except for the new machinery.
- ▶ With all of the other additions we have made, we first learned to translate and then jumped right into how to use them in proofs. What are the new rules and equivalences associated with the introduced elements? With modality, however, this is more complicated because there is not single logical meaning for them. Rather, there are several possible meanings.
- ▶ There are multiple viewpoints for modal logic. They are called modal systems of different strengths because each stronger level commits us to a new rule, where the weaker ones do not. We account for different meanings of modality by having different modal languages that use different additional axioms.

- ◇ K (named for the philosopher-logician Saul Kripke) is the weakest modal language, according to which a sentence is necessary only if we can prove it using no premises. In this language, the axiom $\Box(P \rightarrow Q) \rightarrow (\Box P \rightarrow \Box Q)$ is true and can be inserted into a proof for any sentences P and Q .
 - ◇ M or T (named independently by Kurt Gödel and Georg Henrik von Wright), the next stronger language, adds the claim that if a sentence is necessarily true, then it must also be true—that is, $\Box P \rightarrow P$.
 - ◇ S3 (where “S” stands for “strength”) adds the additional claim to M or T that if a sentence is necessarily true, then it must also be possibly true—that is, $\Box P \rightarrow \Diamond P$.
 - ◇ S4 adds to S3 the additional claim that if a sentence is necessarily true, then it is necessarily true that it is necessarily true—that is, $\Box P \rightarrow \Box \Box P$.
 - ◇ S5, the strongest modal language, adds to S4 the claim that all modal claims are necessary—that is, not only are necessary truths necessarily necessary, but that sentences that are possible are necessarily possible. In other words, $\Diamond P \rightarrow \Box \Diamond P$.
- ▶ With modal logic, we are left with the question that mathematicians faced when suddenly they had multiple geometric systems: Which one is right? Which of these five modal systems is the real one? The stance that logicians take is similar to the one David Hilbert took with respect to mathematical systems—take whichever one you want and then ask the philosophers questions.

DEONTIC LOGIC

- ▶ Philosophers have made another interesting use of modal logic, turning it into what is called deontic logic. The term “deontology” refers to an approach to logic that is based on duties, absolute rules that must be followed in order to act morally.

- ▶ Moral theory utilizes three different statuses for possible acts.
 - ◊ An action could be morally necessary, meaning that you have to do it.
 - ◊ An action could be morally prohibited; these are things that you should never do.
 - ◊ An action could be morally permissible, meaning that you could do it if you want or don't do it if you don't want—either way is fine.

- ▶ Deontological ethicists thought that perhaps modal logic could be modified to become deontic logic. Morally necessary actions would be represented with the box (\Box). Instead of $\Box P$ representing the necessary truth of some sentence, it would represent the necessity of a person's undertaking the act P . Morally impermissible acts would be represented as $\Box \neg P$. This means that it is necessary to avoid doing P . The diamond (\Diamond) would be used for morally permissible acts.

- ▶ The hope was that in adapting modal logic to a moral logic, we could take the mushiness and hurt feelings that often accompany arguments surrounding questions of ethics and replace these feelings with the rigor and objectivity that we see in logical argumentation.

- ▶ Unfortunately, the project has not taken off as much as early advocates had hoped. A few snags have surfaced. Unlike with the logic of truth in which every sentence is either true or false and we never have to choose between truths, there can be situations in which we have to choose between morally required acts. Additionally, are ethical sentences intrinsically necessary or only necessary in a context?

READINGS

- Epstein, *Propositional Logics*, chap. 6.
Layman, *The Power of Logic*, chap. 12.

QUESTIONS

1.

Is S5 too strong? Is everything necessarily the way it has to be? Could things be otherwise? Are the alternative possible worlds really possible? Are things the way they are because that is the way they have to be?

2.

Some philosophers have argued that to say that a sentence is true is to say that it describes reality. For example, to say that the sentence “Madagascar is an island” is true, it must be the case in reality that Madagascar is, indeed, an island. So, if we hold that there are true sentences about possibility—for example, if we think it is true that “It is possible that Madagascar wouldn’t have been an island”—then there is a possible world in which Madagascar is not an island. But if to be true is to describe an aspect of reality, then if we hold sentences about possibility to be true, then the possible worlds must be part of reality—in other words, the possible worlds must all actually exist. Do we want to give up on having true sentences about possibility? Do we want to allow for reality to include not only our world but every possible world, or is there another way to understand the truth of sentences about possibility?

3.

Is a deontic logic possible? Could there be a logic that helps us determine what we should do if we want to act ethically?

Three-Valued and Fuzzy Logic

All of our lectures thus far have elucidated certain elements of classical logic. What defines classical logic is the adherence to a central proposition: the law of the excluded middle. This is the claim that all sentences have one of two truth-values: true or false. But what happens if we deny this axiom? Because this is the foundation of all the logical inferences we have examined so far, the fear is that by eliminating the law of the excluded middle, we make logic impossible. It turns out that we don't; we just make it more complicated.

THREE-VALUED LOGIC

- ▶ By denying that all sentences must be either true or false, we create what we call multivalued logic. The simplest of these is three-valued logic. We are used to having just two truth-values: T and F. But we are going to augment this set with one new member. The name of this new truth-value is a completely arbitrary choice for which logicians have no unified position.
- ▶ Some use M for “middle”—that is, it is some middle value between true and false. Others use I for “indeterminate” or U for “unknown” or N for “neither true nor false.” We'll use M, partially because it does not presuppose an interpretation for this new truth-value.
- ▶ One could develop a full-blown first-order predicate logic, but for simplicity and to see some of the results more clearly, we will look at what introducing a new truth-value does for truth-functional logic.

- ▶ We can still use truth tables, but they will have to be a bit larger because of the new truth-value. For negation, for example, instead of having two rows, we will now have three. For a sentence p , we have the truth table shown at right.

p	$\neg p$
T	F
M	M
F	T

- ▶ The only new information is the middle line. The negation of a sentence with truth-value M remains M. At first glance, this should strike us as odd. Negation changes truth-values. But if we think of M as indeterminate or unknown, then if we don't know whether p is true, then we also don't know if it is false. So, if p has the value M, then so should $\neg p$.
- ▶ What about the other connectives? Let's make one big truth table for all three remaining connectives.
- ▶ We know what the columns are. One each for p and q , and columns for "or," "and," and "if-then." We have three possible values for each of the constituent sentences, so we will need nine rows to capture all of the combinations.
- ▶ Let's start with the disjunction, p "or" q . All of our truth-values from two-valued logic remain the same, so we can enter all of those into the table.
- ▶ All that is left are the ones that involve the new value M. We know that a disjunction is true whenever one of the disjuncts is true. So, let's put in T for $T \vee M$ or $M \vee T$.
- ▶ What about the remaining cases: $M \vee F$, $F \vee M$, and $M \vee M$? If, again, we think of M as indeterminate or unknown, then all three of these will be unknown as well, because if they turn out to be true, so will the sentence, but if they turn out to be false, so will the sentence. So, we can complete the entire column for disjunction.

- ▶ What about the conjunction, p “and” q ? Again, the two-valued values remain. So, we can add those to the table.
- ▶ We know that a conjunction is false when one conjunct is false, so $M \& F$ and $F \& M$ will both give us F .
- ▶ “True or unknown” remains unknown, and “unknown or unknown” is quite unknown, so we can now complete the column.
- ▶ Last is the conditional: if p , then q . Once again, the two-valued values can be plugged in.
- ▶ We know that the only time a conditional is false is when the antecedent is true and the consequent is false. So, if the antecedent is false or the consequent is true, then we know that the conditional is true regardless of the truth-value of the remaining sentence. This was true for two-valued logic and remains true for three-valued logic.
- ▶ If a sentence is $T \rightarrow M$, then there is a chance that it could end up $T \rightarrow F$, which would make it false, and it is possible that it could end up $T \rightarrow T$, which would have it end up true. So, $T \rightarrow M$ would remain M . The same kind of reasoning holds for $M \rightarrow F$. So, we can plug these values into the table.
- ▶ The only remaining case is $M \rightarrow M$. Different three-valued systems treat this case differently. Some take this to give us a value of T , while others take it to give us a value of M . For the sake of being more intuitive, let's go with $M \rightarrow M$ as M . So, we now have a complete truth table, as follows.

p	q	$p \vee q$	$p \& q$	$p \rightarrow q$
T	T	T	T	T
T	M	T	M	M
T	F	T	F	F
M	T	T	M	T
M	M	M	M	M
M	F	M	F	M
F	T	T	F	T
F	M	M	F	T
F	F	F	F	T

- ▶ One consequence of our choice of $M \rightarrow M$ being M is that the entire row for $p=M$ and $q=M$ has a value of M. If we also think of negation, $\neg p$ and $\neg q$ with values of M also remain M.
- ▶ As a result, no matter what combination of truth-functional connectives we use in creating molecular sentences and no matter how complex we make those sentences, the row with all truth-values M will always end up with the truth-value M.
- ▶ So, there can be no tautologies or contradictions as we traditionally think of them. We can have modified or pseudo-tautologies that have no truth-value of F, or we can have modified or pseudo-contradictions that have no truth-values of T. But neither of these will function as tautologies or contradictions did in two-valued truth-functional logic.
- ▶ We can still use truth tables to determine validity and invalidity in exactly the same way. We are still making sure that there are

no cases in which the premises are all true and the conclusion is false. But now, we have a stronger and weaker sense of validity.

- ▶ Recall that an argument is invalid if there is even a single case in which all of the premises are true and the conclusion is false. But now there are two ways to avoid that. The first is validity, in which for every case in which all of the premises are true, the conclusion is also true. We'll call this strong validity.
- ▶ But then there will be the arguments in which there are no cases in which the premises are true and the conclusion is false, but there are some cases in which the premises are true and the conclusion is true and some cases in which the premises are true and the conclusion is middle. We can think of this as weak validity.
- ▶ But what about the case in which the premises are all true and the conclusion is sometimes T and sometimes M? We need a new third value, an M version of validity for these arguments.
- ▶ What does this M value really mean? We have been saying, for the sake of intuition, that it means unknown, but unknown is not a truth-value.
- ▶ There is a difference between what is and what we know about what is. Three-valued logic does not seem apt for describing what is, but it does seem appropriate for talking about our knowledge—and our lack thereof.
- ▶ We can be in one of three states with respect to the truth-value of a sentence: we can know it is true, we can know it is false, or we can be in a state of doubt, not knowing whether it is true or false. As such, three-valued logic does seem to be a good model of our reasoning when we see ourselves as being in a state of knowledge about some propositions and a state of ignorance about others.

- ▶ With three-valued logic, we deny the law of the excluded middle by positing a third truth-value. But we could add more. We could create a four-valued logic or an 18-valued logic. It is not clear what uses we would have for them, but such utilitarian, pragmatic questions hold little weight with logicians and mathematicians.
- ▶ But in all of these cases, we are making a particular assumption that we could deny and create an even more general logic. Why must these truth-values be discrete—that is, why is it an absolute partitioning of sentences into one of these truth-values? Why not make a smooth continuum of truth? In this case, we have an infinite number of truth-values.

FUZZY LOGIC

- ▶ It is tempting to generalize our use of three-valued logic as a measure of our knowledge or lack thereof and see fuzzy logic as a measure of our certainty about a sentence's truth. But this is not what fuzzy logic is about. That is induction. Fuzzy logic is not concerned with likelihood, but rather deals with fuzzy sets.
- ▶ A non-fuzzy set is one where all members of the universe are either in or out of the set. Suppose that the domain is the set of humans, and the set we want to examine is the set of all pregnant people. This is non-fuzzy; either you are or are not pregnant. There is a clear and absolute line between in and out, and no one straddles that line.
- ▶ But some sets are not like that. Some sets are fuzzy; objects in the universe can be in the set to varying degrees. Again, take the domain to be humans, and now consider the property of hairiness. For example, when men grow a beard, they move further into the fuzzy set. But over the years, as their hairline recedes, they move further out of the fuzzy set. Hairiness ranges from zero (hairless) on one extreme to one on the other extreme (werewolf-like).

- ▶ All humans have some degree of hairiness. So, while all people are hairy, some sets are fuzzy as well. Logicians have devised operations on these sets, called fuzzy connectives, that allow us to have a fuzzy version of truth-functional logic and even a fuzzy version of predicate logic. The values plugged in are numbers between zero and one, and the result of the connectives is a new value between zero and one.
- ▶ There are several versions of fuzzy logic, each defining these connectives in slightly different fashion. Which of the definitions we use depends on the needs that the fuzzy logic serves.
- ▶ The standard example of a system that utilizes fuzzy logic is the heating and cooling system in your house. There is a thermostat connected to a switch that controls the heating and cooling system. You ask your system to keep your home at some particular temperature.
- ▶ If you have an old system, then this is governed by classical logic. Let p be the sentence “The temperature is the one I desire.” If p is true, then the system does nothing. If p is false, then the system turns on—for example, the heater is on full blast until the temperature is reached, at which point it is turned off.
- ▶ But more modern systems use fuzzy logic to be more energy efficient. If the temperature is the one desired—if p has a value of one—then the system does nothing. But if the temperature deviates from the desired temperature, then it springs to action in a more intricate fashion.
- ▶ If it is too cold, it turns on the heat to an appropriate degree, depending on the degree of deviation from the desired temperature. If it is way too cold, it cranks the heat way up. If it is just a little too cold, it gives the heater just a small shot of power to let it do a little bit of work. This is a different logic governing the operation of the system—it is a fuzzy logic.

- ▶ Fuzzy logic is a generalization of multivalued logic that we created by seeing what would happen if we denied the central claim of classical logic, the law of the excluded middle. We found that if we were clever about it, we could find a way of reasoning that is helpful and useful. And that is what logic is really all about.

READINGS

Epstein, *Propositional Logics*, chap. 8.

Frielberger and McNeill, *Fuzzy Logic*.

QUESTIONS

1.

In three-valued logic, the negation of the new value middle is still middle. Does it make sense to make the negation of a value the same value? Is that not the opposite of what we mean by negation?

2.

Fuzzy logic might be useful—that is, it might be helpful to think of truth as if it varies smoothly—but does it really? Isn't truth more like being pregnant, where you are or you aren't? Is there really a sense of “in between” when it comes to truth?

BIBLIOGRAPHY

Barker, Stephen. *The Elements of Logic*. New York: McGraw-Hill, 1989. A classic textbook that introduces logical concepts for the non-technician.

Bradburn, Norman, Seymour Sudman, and Brian Wansink. *Asking Questions: The Definitive Guide to Questionnaire Design*. Hoboken, NJ: Jossey-Bass, 2004. A clear and comprehensive discussion on how to conduct a poll and the pitfalls to be avoided.

Coffa, J. Alberto. *The Semantic Tradition: From Kant to Carnap*. New York: Cambridge University Press, 1991. A detailed discussion of the emergence of analytic philosophy from the advances of science, mathematics, and logic in the 19th and 20th centuries.

Copi, Irving. *Introduction to Logic*. New York: Macmillan, 1972. One of the best-selling and longest-lived college-level logic texts out there.

Damer, T. Edward. *Attacking Faulty Reasoning*. Boston: Wadsworth, 2009. A guidebook to the fallacies that can be committed in making arguments.

Durkheim, Emile. *The Rules of Sociological Method*. New York: Free Press, 1966. One of the founding fathers of the field of sociology carefully considers what is the subject matter of sociology and how sociological research should be done.

Epstein, Richard. *Propositional Logics*. Belmont, CA: Wadsworth, 2001. A college-level text that examines a full range of logical languages, their structures, and usages.

Euclid. *The Elements*. New York: Dover, 1952. The classic axiomatization of plane geometry, where all of the results are derived through deductive proofs.

Fine, Cordelia. *A Mind of Its Own: How Your Brain Distorts and Deceives*. New York: Norton, 2006. A popular account of the ways in which our neural wiring leads us into errors while thinking it is correct.

Frielberger, Paul, and Daniel McNeill. *Fuzzy Logic*. New York: Simon & Schuster, 1993. A popularly accessible account of the development of fuzzy logic and its applications to real-world systems.

Gimbel, Steven. *Exploring the Scientific Method*. Chicago: University of Chicago Press, 2011. An investigation into the various accounts that have been proposed by scientists and philosophers of science to account for the logic behind the scientific method.

Gray, Jeremy. *The Hilbert Challenge*. New York: Cambridge University Press, 2001. An examination of the 23 problems David Hilbert charged mathematics with solving in the 20th century. It discusses the problems, the attempts to solve them, and whether the attempts were successful.

Hurley, Patrick. *Logic*. Belmont, CA: Wadsworth, 1996. A college-level textbook that examines both formal and informal logic.

Kahane, Howard. *Logic and Philosophy*. Belmont, CA: Wadsworth, 1978. A college-level textbook that examines deductive logic and the philosophical questions that gave rise to the interest in the field.

Kelley, David. *The Art of Reasoning*. New York: W. W. Norton, 1998. A college-level textbook that covers both inductive and deductive reasoning.

Kramer, Edna. *The Nature and Growth of Modern Mathematics*. Princeton: Princeton University Press, 1981. An encyclopedic account of the history of mathematical thought from ancient times through the first half of the 20th century.

Layman, Stephen. *The Power of Logic*. Toronto: Mayfield, 1999. A college-level logic textbook that covers deduction, induction, probability, and modal logic.

Nagel, Ernest, and James Newman. *Gödel's Proof*. New York: New York University Press, 1958. A popularly accessible account of the development of Gödel's incompleteness theorem and the structure of the argument itself.

Sutherland, Stuart. *Irrationality: The Enemy Within*. London: Constable and Company, 1992. A popular book that examines the cognitive biases that humans possess and the ways in which they lead us into error.

Yandell, Ben. *The Honors Class: Hilbert's Problems and Their Solvers*. Natick, MA: A K Peters, 2002. An examination of the 23 problems David Hilbert charged mathematics with solving in the 20th century. It discusses the problems, the attempts to solve them, and whether the attempts were successful.

IMAGE CREDITS

15. © GeorgiosArt/iStock/Thinkstock.
- 69 Vienna museum/Wikimedia Commons/Public Domain.
73. © denisk0/iStock/Thinkstock.
- 204 Göttingen State and University Library/Manuscripts
and Scholarly Collections/flickr/Public Domain.
210. © wynnter/iStock/Thinkstock.
213. Conquistador/Wikimedia Commons/Public Domain.