

***Introduction to Philosophy - Classical and
Contemporary Readings 9th edition Perry Test Bank***

SKU: 9780197543818-TEST-BANK

Instructor's Manual and Test Bank

to accompany

Introduction to Philosophy:

Classical and Contemporary Readings

Ninth Edition

by Perry, Bratman, and Fischer

John Martin Fischer

Neal A. Tognazzini

D. Justin Coates

David Beglin

Andrew Law

New York • Oxford

Oxford University Press

2021

TABLE OF CONTENTS

Preface	6
Sample Syllabi	8
Logical Toolkit	14
Writing Philosophy Papers	40
Reading Summaries	
PART I: Philosophy	44
Bertrand Russell, “The Value of Philosophy”	44
Plato, “Apology: Defence of Socrates”	45
PART II: God and Evil	46
A. Why Believe?	
Saint Anselm, “The Ontological Argument”	46
Saint Thomas Aquinas, “The Existence of God”	47
William Paley, “Natural Theology”	48
William James, “The Will to Believe”	49
B. The Problem of Evil	
David Hume, “Dialogues Concerning Natural Religion”	50
Gottfried Leibniz, “God, Evil and the Best of All Possible Worlds”	51
William Rowe, “The Problem of Evil and Some Varieties of Atheism”	52
Marlin Adams and Stewart Sutherland, “Horrendous Evils and the Problem of Evil”	53
Eleonore Stump, “The Mirror of Evil”	54
Louise Antony, “For the Love of Reason”	55
PART III: Knowledge and Reality	56
A. Plato, Descartes, and the Problems of Skepticism	
Plato’s “Allegory of the Cave”	56
Robert Nozick, “The Experience Machine”	57
René Descartes, “Meditations on First Philosophy”	58
Keith DeRose and Ted A. Warfield, “Responding to Skepticism”	59
Peter Graham, “Beginning to Respond to Skepticism”	60

B. Hume's Problems and Some Solutions	
David Hume, "An Enquiry Concerning Human Understanding"	61
W. C. Salmon, "The Problem of Induction"	62
PART IV: Minds, Bodies, and Persons	63
A. The Traditional Problem of Minds and Bodies	
David M. Armstrong, "The Nature of Mind"	63
Paul M. Churchland, "Eliminative Materialism"	64
Frank Jackson, "What Mary Didn't Know"	65
Patricia Churchland, "Neurophilosophy"	66
B. Minds, Brains, and Machines	
A. M. Turing, "Computing Machinery and Intelligence"	67
John R. Searle, "Minds, Brains, and Programs"	68
C. Personal Identity	
John Perry, "A Dialogue on Personal Identity and Immortality"	69
Derek Parfit, "Personal Identity"	70
J. David Velleman, "So It Goes"	71
Daniel Dennett, "Where Am I?"	72
Marya Schechtman, "Personhood and Personal Identity"	74
Agnieszka Jaworska, "Respecting the Margins of Agency: Alzheimer's Patients and the Capacity to Value"	75
D. Freedom, Determinism, and Responsibility	
Roderick M. Chisholm, "Human Freedom and the Self"	76
David Hume, "Of Liberty and Necessity"	78
Harry G. Frankfurt, "Alternate Possibilities and Moral Responsibility"	79
John Martin Fischer, "Responsiveness and Moral Responsibility"	80
Harry G. Frankfurt, "Freedom of the Will and the Concept of a Person"	81
Gary Watson, "Free Agency"	82
Christian Helmut Wenzel, "Free Will and Zhuangzi: An Introduction"	83
PART V: Ethics and Society	84
A. Utilitarianism	
Jeremy Bentham, "The Principle of Utility"	84
John Stuart Mill, "Utilitarianism"	85

E. F. Carritt, "Criticisms of Utilitarianism"	86
Mozi, "Excerpts from Mozi"	87
Peter Singer, "Famine, Affluence, and Morality"	88
B. Kantian Ethics	
Immanuel Kant, "Groundwork of the Metaphysic of Morals"	89
Onora O'Neill, "Kantian Approaches to Some Famine Problems"	90
C. Aristotelian Ethics	
Aristotle, "Nichomachean Ethics"	91
Rosalind Hursthouse, "Right Action"	92
D. Justice and Equality	
John Rawls, "A Theory of Justice"	93
Robert Nozick, "Justice and Entitlement"	94
John Stuart Mill, "The Subjection of Women"	95
Annette Baier, "The Need for More Than Justice"	96
Elizabeth Anderson, "What's Wrong with Inequality?"	97
E. Contemporary Moral Problems	
Judith Jarvis Thomson, "A Defense of Abortion"	98
Debra Satz, "Markets in Women's Reproductive Labor"	99
Kwame Anthony Appiah, "Racisms"	100
Linda Martin Alcoff, "Racism and Visible Race"	101
PART VI: Existential Issues	102
A. Meaning in Life	
Albert Camus, "The Myth of Sisyphus"	102
Thomas Nagel, "The Absurd"	103
Richard Taylor, "The Meaning of Human Existence"	104
Susan Wolf, "The Meanings of Lives"	105
B. Death	
Thomas Nagel, "Death"	106
James Baillie, "Existential Shock"	107
John Martin Fischer, "Excerpt from <i>Death, Immortality and Meaning in Life</i> "	108
Amy Olberding, "Sorrow and the Sage: Grief in the <i>Zhuangzi</i> "	109
Jenann Ismael, "The Ethical Importance of Death"	110

Sample Exam Questions	112
Part I: Philosophy	112
Part II: God and Evil	117
Part III: Knowledge and Reality	143
Part IV: Minds, Bodies, and Persons	158
Part V: Ethics and Society	202
Part VI: Existential Issues	244

Preface

Welcome to the Instructor's Manual for our textbook, *Introduction to Philosophy*. This Instructor's Manual was introduced with the 4th edition, but we have updated it for the 9th edition to reflect the slightly revised contents. We hope that this will continue to be a useful resource for instructors who already use the text, and we hope that new instructors will also discover its benefits.

The Book

The textbook is arranged topically and divided into seven parts: Philosophy; God and Evil; Knowledge and Reality; Minds, Bodies, and Persons; Ethics and Society; Meaning and Death; and Puzzles and Paradoxes. In each part, we have attempted to provide a balance between historically important works and articles that are more contemporary. Although one always has to be somewhat selective in deciding which pieces to include in an introductory anthology, we sought to include as many classic discussions of these topics as space permitted. Included at the end of each selection are study questions (except for the selections in the seventh part of the book); we believe that these will be of considerable help to students.

There are a number of ways in which this textbook is unique among introductory anthologies. First, we have included a very large selection of excellent essays so that you can focus your course as you see fit. Second, the last section of our textbook includes brief discussions of a number of puzzles and paradoxes, many of which have played a central role in major developments throughout the history of analytical philosophy. These short discussions create an excellent opportunity for further classroom exploration. And third, we have included some supplemental material that we think you will find both helpful and highly innovative.

The Instructor's Manual

The ideas and information that you will find on the following pages are, of course, merely suggestions. But we hope that the information that we have included will lighten some of the more time-consuming and administrative tasks of teaching so that you can focus the bulk of your efforts on the students.

First, we have included some **Sample Syllabi**, or, perhaps, templates for some syllabi. (You'll have to flesh them out as you see fit.) Given the number of articles that are included in the anthology, the possibilities for syllabi are practically endless, but we have provided four. The first syllabus is organized around a "greatest hits" theme, the second has a historical emphasis, the third a contemporary emphasis, and the fourth is more suited for a course that introduces students to philosophy through film.

Second, we have included **Summaries** of all the material in the textbook (except for the puzzles and paradoxes at the end of the book). These are meant to aid you in deciding which articles to include on your syllabus. If you are unsure whether an article would fit into the structure of your course, take a look at our summary of it to help you decide. Of course, the summaries are not intended as rigorous explanations of every point in a particular article. We have attempted to provide you with the substance of each article without getting into the details.

Finally, we have drafted a few **Sample Exam Questions** for each selection in the anthology (again, with the exception of the puzzles and paradoxes). These take the form of short-answer or essay questions (as opposed to multiple-choice, for example) and could be used for exams, quizzes, or even homework assignments. We hope that these questions will ease the often daunting task of writing exams. (The asterisked numbers next to these questions indicate questions featured on the student website [see below].)

The Instructor's Manual does not exhaust the supplementary material, however. We have also launched a student website and blog that we hope will be useful to your students.

The Student Website

The student website is a place where students can get information about the textbook and about philosophy in general. The URL for the student website is www.oup.com/us/perry9e. Have a look!

Currently available are some short practice multiple-choice exams and discussion questions that students can use to test their reading comprehension of the articles in the book. Students will be able to access both the Logical Toolkit and our suggestions for writing papers well (both of which appear in Part I of the book) through the website. We will also include various links to other areas of the Web that have useful resources for students of philosophy (for instance, the *Stanford Encyclopedia of Philosophy* at <http://plato.stanford.edu> and the *Philosophy Talk* radio show at <http://www.philosophytalk.org>).

We are constantly thinking of ways that we can improve the textbook and its supplementary materials, and we love to hear from instructors who use the book in a classroom setting. Feel free to contact us with any suggestions. Happy teaching!

NOTE: *What follows are some sample syllabi, or perhaps templates for syllabi, to aid in selecting pieces that fit together well. Each syllabus can be easily modified for either the quarter or the semester system by omitting or adding pieces. These are merely suggestions, and you should modify them as appropriate for your particular circumstance.*

Sample Syllabus 1

Greatest Hits

Plato, “Apology: Defence of Socrates”

Begin with a discussion of philosophy in general, its roots in Ancient Greece, and the importance it places on examining our lives and the world around us.

Saint Thomas Aquinas, “The Existence of God”

Move on to discuss philosophy of religion, beginning with Aquinas’s historically important arguments for the existence of God. If time allows, discuss some current developments of the cosmological argument as put forward by philosophers such as William Lane Craig.

David Hume, “Dialogues Concerning Natural Religion”

Everyone, theist and atheist alike, should come to terms with the problem of evil. Hume’s piece is one of the classic discussions of this problem. Hume also touches on other issues, such as the designer argument, along the way.

René Descartes, “Meditations on First Philosophy”

Move from philosophy of religion to epistemology, reading the most important work in this field. Raise the problem of skepticism, especially as formulated in Meditations 1 and 2, and explain Descartes’s goal of building a solid foundation for knowledge. Examine how Descartes responds to the skeptical arguments.

Keith DeRose and Ted A. Warfield, Excerpt from *Skepticism* Reader

DeRose and Warfield provide a general framework for skeptical challenges and cover a few contemporary approaches to those challenges. Of particular value is their discussion of the “Ah, Com On!” objection that many students articulate when first hearing of skeptical hypotheses.

David M. Armstrong, “The Nature of Mind”

In the twentieth century, much interesting work about the nature of the mind was focused on developing an account that is consistent with materialism. Armstrong’s account is one of the most prominent of this sort.

Frank Jackson, “What Mary Didn’t Know”

In this paper, Frank Jackson invokes a creative and well-known thought experiment to argue that physicalism must be false.

Daniel Dennett, “Where Am I?”

This is an extremely interesting piece that touches on both the mind-body problem and the metaphysics of identity through time.

John Perry, “A Dialogue on Personal Identity and Immortality”

Again, an entertaining way to discuss some of the issues about personal identity is by reading this dialogue by John Perry.

John Stuart Mill, “Utilitarianism”

Having touched on philosophy of religion, epistemology, and metaphysics, finally move on to ethics. Compare different ethical theories by first talking about utilitarianism.

Immanuel Kant, “Groundwork of the Metaphysic of Morals”

Then discuss Kantianism.

Aristotle, “Nicomachean Ethics”

And finally discuss virtue ethics as it’s derived from Aristotle.

Peter Singer, “Famine, Affluence, and Morality”

Show your students what applied ethics is all about by reading this provocative piece by Peter Singer. Guaranteed to start interesting classroom discussions.

John Rawls, “A Theory of Justice”

An excerpt from Rawls’ classic, *The Theory of Justice*.

Judith Jarvis Thomson, “A Defense of Abortion”

End with Thomson’s classic article.

Sample Syllabus 2

Historical Emphasis

Plato, “Apology: Defense of Socrates”

There’s plenty to discuss in this classic Platonic dialogue.

David Hume, “Dialogues Concerning Natural Religion”

For a classic discussion of the problem of evil and other issues surrounding the existence of God, Hume is an excellent resource.

René Descartes, “Meditations on First Philosophy”

You could really spend an entire term discussing Descartes’s *Meditations* and the issues they raise.

David Hume, “An Enquiry Concerning Human Understanding”

The traditional problem of induction is an extremely rich starting point for an interesting discussion about the methods of science and the foundations of our own knowledge.

Plato’s “Allegory of the Cave”

This classic analogy reveals much about Plato’s thinking. It also pairs nicely with the selection from Zhuangzi, which might be thought of raising similar concerns as Plato about the sensible world, but which takes those concerns in a very different direction.

Mozi, “Excerpts from Mozi”

Mozi develops perhaps the world’s first consequentialist moral theory. In particular, Mozi focuses on the idea of being impartial, and how doing so can improve the world.

Albert Camus, “The Myth of Sisyphus”

Conclude with an existentialist classic.

Sample Syllabus 3

Contemporary Emphasis

Richard Taylor, “The Meaning of Human Existence”

Some interesting recent work addresses the old question of the meaning of life, and what better way to get an Introduction to Philosophy class started than by discussing some of it. This piece by Richard Taylor is a great place to start.

Susan Wolf, “The Meanings of Lives”

Another interesting related piece is this one by Susan Wolf, in which she continues the discussion of many of the issues Taylor raises.

David M. Armstrong, “The Nature of Mind”

In the twentieth century, much interesting work about the nature of the mind was focused on developing an account that is consistent with materialism. Armstrong’s account is one of the most prominent of this sort.

A. M. Turing, “Computing Machinery and Intelligence”

The advent of contemporary computers has raised many important philosophical questions about the relationship between the human mind and a machine like a computer. As one of the foundational figures in computer science, Turing has some interesting things to say about these questions.

John R. Searle, “Minds, Brains, and Programs”

Searle argues that the analogy between humans and computers can only go so far in this classic article on the philosophy of mind.

Roderick M. Chisholm, “Human Freedom and the Self”

Chisholm introduces what he takes to be the problem of human freedom, showing how human freedom appears to conflict with both determinism and indeterminism. He then develops a novel solution, one that requires some fairly radical claims about the nature of the human person.

John Martin Fischer, “Responsiveness and Moral Responsibility”

Fischer offers an account of moral responsibility that provides an alternative point of view to that of Roderick Chisholm. Fischer specifically argues that moral responsibility does not require alternative possibilities and concludes that it is therefore compatible with causal determinism.

John Rawls, “A Theory of Justice”

As one of the most important political philosophers of the twentieth century, John Rawls provides many interesting ideas about the nature of justice.

Robert Nozick, “Justice and Entitlement”

To continue the discussion about just societies, turn to Robert Nozick and learn about his entitlement theory.

Annette Baier, “The Need for More Than Justice”

Consider Baier’s feminist critique of John Rawls’s account of justice. Use this to transition from more theoretically oriented political and social philosophy to applied work.

Debra Satz, “Markets in Women’s Reproductive Labor”

End with a discussion about some contemporary issues in the study of gender and race. Here, Satz discusses the nature of women’s reproductive labor.

Kwame Anthony Appiah, “Racisms”

And here Appiah discusses different varieties of racism.

Sample Syllabus 4

Film-Oriented

Reading 1 Film 1	Thomas Nagel, “The Absurd” <i>The Unbearable Lightness of Being</i>
Reading 2 Film 2	René Descartes, “Meditations on First Philosophy” <i>The Matrix</i>
Reading 3 Film 3	John Perry, “A Dialogue on Personal Identity and Immortality” <i>Total Recall</i>
Reading 4 Film 4	Derek Parfit, “Personal Identity” <i>Memento</i>
Reading 5 Film 5	John Martin Fischer, “Responsiveness and Moral Responsibility” <i>Minority Report</i>
Reading 6 Film 6	Kwame Anthony Appiah, “Racisms” <i>Crash</i>
Reading 7 Film 7	Albert Camus, “The Myth of Sisyphus” <i>Crimes and Misdemeanors</i>

Logical Toolkit

Logic and Philosophy

Philosophy is very much about asking questions. Does God exist? What can we know? What keeps us the same through time and change? Is the mind distinct from the body? How do we know what is right and what is wrong? Because not everyone is going to agree on the correct answers to these questions, it is extremely important to give reasons why you think one answer is better than another. In giving reasons why you believe (or why others should believe) a particular answer, you are doing logic, even though you might not recognize it as such. Logic is just a way of articulating more clearly the reasoning that we ordinarily do when we tell someone why we believe something.

If we are to be persuaded that your position is correct, we need to have some way of assessing the reasons that you give for believing your position. For instance, we need to know whether your reasons really do lend support to the position. This is where learning a bit of logical apparatus can come in quite handy. So let's introduce some terminology.

Arguments

We'll start with the basic idea of an argument. As we use the term in philosophy, an argument is not just a verbal dispute about some matter. Rather, it is a way of articulating reasons. Or, to be more precise:

An *argument* is a series of statements where the last statement supposedly *follows from* or *is supported by* the first statements. The last statement is called the *conclusion*, and the first statements are called the *premises*.

Here's a relatively simple example:

1. Everyone who lives in Los Angeles lives in California.
2. Alvin lives in Los Angeles.
3. Therefore, Alvin lives in California.

Suppose we were trying to convince you that our friend Alvin lives in California. (Again, we probably wouldn't normally give you an argument to convince you of this, but this is a simple example just to get the idea of an argument under our belts.) We might give you the following reasons for believing that Alvin lives in California. First, we know that Alvin lives in Los Angeles. And second, we know that Los Angeles is in California, so anyone who lives in Los Angeles automatically lives in California. These two reasons are represented by premises 1 and 2, and they are meant to support the conclusion, which is number 3. Arguments in the articles that you read for class will most often not appear in this numbered form, but they can all be reconstructed in this form so that the reasoning is easy to see.

In this example, if you were to accept the two premises, you would have to accept the conclusion. So our argument is, in a certain sense, a *good* argument. But there are different ways that an argument can be good.

Validity

The first way an argument can be good is if its premises actually do support its conclusion. Recall that our definition of an argument is a series of statements in which the conclusion *supposedly* follows from or is supported by the premises. Well, there are some arguments with conclusions that actually do follow from the premises, and there are some arguments with conclusions that don't actually follow from the premises, even though they *supposedly* do. The first type of arguments are *valid* arguments, and the second type are *invalid* arguments. Or, a bit more carefully:

An argument is *valid* if its conclusion follows from its premises.

Or, more carefully still:

An argument is *valid* if it satisfies the following condition: If its premises were true, then its conclusion would *have to be* true.

The argument we gave previously is an example of a valid argument because if premises 1 and 2 were true, then 3 would have to be true. But the following is an example of an *invalid* argument:

1. Everyone who lives in Los Angeles lives in California.
2. Alvin lives in California.
3. Therefore, Alvin lives in Los Angeles.

If we were to put forth this argument while trying to convince you that our friend Alvin lives in Los Angeles, you shouldn't be convinced. Why not? Simply because the reasons that we gave for believing that Alvin lives in Los Angeles don't actually support that conclusion. For in this case, premise 1 could be true (it actually is true), and premise 2 could be true, but the conclusion might still be false (Alvin could live in San Francisco, for instance). Thus this is an *invalid* argument. The conclusion doesn't actually follow from the premises. It's not the case that if its premises were true, then its conclusion would have to be true.

In philosophy, as in life, we're mostly interested in putting forth valid arguments. At the very least, our conclusions must really follow from our premises. But although validity is a good first step, it's not the only way that an argument can be good.

Soundness

If we succeed in putting forth a valid argument, that's a good start. But we want more from our arguments. We also want our premises to actually be true. Recall that validity was about the *relationship* between premises and conclusion: If the premises were true, then the conclusion would have to be true. But sometimes that's a big "if." That is, sometimes we're not sure whether the premises are actually true. That's the next thing we care about. If our argument is valid and its premises are also true, then the argument is *sound*. More precisely:

An argument is *sound* if it is valid and has all true premises.

Or, more precisely still:

An argument is *sound* if it satisfies the following two conditions:

1. It is valid.
2. All of its premises are true.

Let us give another example to understand soundness better. Consider the following argument:

1. Abortion is the killing of an innocent person.
2. Killing innocent people is morally objectionable.
3. Therefore, abortion is morally objectionable.

This is a much more interesting argument than the one we gave about our friend Alvin. Indeed, it is likely to stir emotions. But we're not going to discuss the moral rightness or wrongness of abortion—we're just using this argument as an example so that we can better understand logic. Now, there are at least two ways that an argument can be good, so whenever you are confronted with an argument such as this, you should always ask yourselves two questions: First, is it valid? Second, is it sound?

We'll save you the suspense: This argument is indeed valid. Remember what that means, though. It doesn't mean that abortion is morally objectionable. All it means is that the premises of this argument really do support the conclusion of the argument. Or, in other words, *if* the premises were true, then the conclusion would have to be true. Whether this argument is valid is not a matter of controversy. What *is* a matter of controversy, however, is whether this argument is sound. That is, is it a valid argument with premises that are actually true? This is where opinions differ. For our purposes, it's enough to realize that if the premises of this argument actually are true, then the argument is sound (because it's also valid), and if the premises of this argument actually are false, then the argument is unsound (even though it's still valid).

Why do we care about putting forth sound arguments? Well, if you present someone with a valid argument and you can successfully argue that the premises of your argument are true, then the other person *must* accept the conclusion as well, on pain of irrationality. Because valid arguments are such that their conclusions really do follow from their premises, one cannot accept their premises without also accepting their conclusions. So if you are giving us

your reasons for, say, your belief in God, and you present us with a valid argument with premises with which we agree, then we must agree that God exists. Logic can be a very powerful tool.

Proving Validity

Sometimes, as in the cases above, it is easy to tell whether an argument is valid. Other times, it is quite difficult—as soon as an argument gets a bit complicated or controversial, we can't just “see” whether it is valid or not. It would therefore be nice if we had some more objective and rigorous methods for *proving* whether an argument is valid. Over the millennia, and especially in the last couple of centuries, philosophers, logicians, and mathematicians have made great strides toward this goal. The techniques developed by such thinkers can be extremely technical and complex; it often takes an entire semester to get a full understanding of the most basic logical systems. What follows then are some helpful, though rudimentary, tools for trying to prove and disprove validity.

Proving Validity: The Famous Forms Method

In determining whether an argument is valid or not, it is often best to abstract away from the argument's content and instead focus on the argument's form. Consider again the argument from above:

1. Everyone who lives in Los Angeles lives in California.
2. Alvin lives in Los Angeles.
3. Therefore, Alvin lives in California.

This argument is plainly valid, but its validity has nothing to do with Los Angeles, California, or Alvin in particular. There's nothing special about those things that make the argument valid. Instead, the argument is valid because it has a form. If we abstract away from the particulars, the argument's form is something like this:

1. Everyone who lives in *L* lives in *C*.
2. *A* lives in *L*.
3. Therefore, *A* lives in *C*.

With the case at hand, “*L*” stands for “Los Angeles”, “*C*” stands for “California”, and “*A*” stands for “Alvin”. But notice that *regardless* of what “*L*”, “*C*”, and “*A*” stand for, the argument will always be valid—we can “plug in” just about anything for these letters and we'll still end up with a valid argument. For instance, suppose instead that we let “*L*” stand for “Libya”, “*C*” stand for “Costa Rica”, and “*A*” stand for “Aphrodite”. The resulting argument would then look like this:

1. Everyone who lives in Libya lives in Costa Rica.
2. Aphrodite lives in Libya.

3. Therefore, Aphrodite lives in Costa Rica.

This argument, just like the argument we started with, is obviously valid. It's not *sound* since not all of its premises are true, but it is definitely valid: *if* the premises *were* true, the conclusion would *have to be* true.

So in determining whether an argument is valid, we ought to abstract away from the argument's content and instead only investigate the argument's form. Doing so can often help make otherwise complicated arguments manageable. Consider the following argument:

1. If the standard measures of a country's economic prosperity—GDP, unemployment, inflation, etc.—are not adequate, then economists have to invoke quantities, such as happiness, that cannot be measured in any reliable way.
2. If economists have to invoke quantities, such as happiness, that cannot be measured in any reliable way, then economics does not qualify as a truly scientific field of study.
3. Therefore, if the standard measures of a country's economic prosperity—GDP, unemployment, inflation, etc.—are not adequate, then economics does not qualify as a truly scientific field of study.

On the surface, this argument may look fairly complex, and so determining whether it is valid or not might be difficult. But if we abstract away from the content and look merely at its form, the argument is actually quite simple. It looks like this:

1. If *A*, then *B*.
2. If *B*, then *C*.
3. Therefore, if *A*, then *C*.

In this case, "*A*", "*B*", and "*C*" each stand for pretty long sentences:¹ "*A*" stands for "The standard measures of a country's economic prosperity—GDP, unemployment, inflation, etc.—are not adequate"; "*B*" stands for "Economists have to invoke quantities, such as happiness, that cannot be measured in any reliable way"; and "*C*" stands for "Economics does not qualify as a truly scientific field of study". It is these long sentences that make the original argument sound so daunting. But once we strip this away and look just at the argument's form, we see that it is simple and, more important, obviously valid. (Whether it is sound is another issue, of course.)

So our first rudimentary method for proving that an argument is valid is to show that the argument has a *form* that is obviously (or famously) valid. (Hence the name "the famous forms method.") What are the "famous" forms? Here are some examples, each of which is famous enough to get a name:

Modus Ponens

1. If **A**, then **B**.

2. **A.**
3. Therefore, **B.**

Modus Tollens

1. If **A**, then **B.**
2. Not **B.**
3. Therefore, not **A.**

Hypothetical Syllogism

1. If **A**, then **B.**
2. If **B**, then **C.**
3. Therefore, if **A**, then **C.**

Disjunctive Syllogism

- | | |
|---------------------------------|---------------------------------|
| 1. Either A or B. | 1. Either A or B. |
| 2. Not A. or: | 2. Not B. |
| 3. Therefore, B. | 3. Therefore, A. |

Constructive Dilemma

1. Either **A** or **B.**
2. If **A**, then **C.**
3. If **B**, then **D.**
4. Therefore, either **C** or **D.**

Conditional Proof

1. Assume: **A.**
2. Derive: **B.**
3. Therefore, if **A**, then **B.**

Reductio ad Absurdum

1. Assume: **A.**
2. Derive: **B** and not **B.**
3. Therefore, not **A.**

There are many more famous forms, far too many to list here. But the important point is this: if you come across an argument and it has one of these forms, or any form that is obviously valid, then you know the argument itself is valid.²

*Note to the reader: The rest of the section is more difficult.
Feel free to skip ahead to the next section.*

There is a serious limitation of the famous forms method, though: some arguments don't have a famous form and yet are still valid. That's one of the reasons the famous forms method is so rudimentary. But we can extend the famous forms method by seeing certain arguments as involving *combinations* of famous forms, in some sense. Consider this argument:

1. If Oscar is hungry, then Oscar is in a bad mood.
2. If Oscar is in a bad mood, then Oscar is frowning.
3. Oscar isn't frowning.
4. Therefore, Oscar isn't hungry.

This argument has the following form:

1. If A , then B .
2. If B , then C .
3. Not C .
4. Therefore, not A .

Although this argument is valid, it doesn't fit any of the famous forms mentioned above. (One way to see this is to note that it has four lines and the only form mentioned above that has four lines is constructive dilemma, which definitely doesn't fit this argument.) But we can nonetheless "break up" the argument into a combination of famous forms to prove that it is valid. Just consider premises (1) and (2) for the moment:

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

We know from hypothetical syllogism that "If A , then C " follows from these two premises. That is, we know if premises (1) and (2) are true, "If A , then C " *must be* true. As logicians sometimes put it, we can *derive* "If A , then C " from (1) and (2) using hypothetical syllogism. But now look again at premise (3):

3. Not C .

If we take "If A , then C " and combine it with this third premise—"Not C "—we know from *modus tollens* that "Not A " follows. That is, from "If A , then C " and premise (3), we can *derive* "Not A " using *modus tollens*. And that's very fortunate since "Not A " is exactly the conclusion of the argument!

We can visualize our steps like this, where the parenthetical note on the side tells us where we “got” each line from:

1. If A , then B . (Premise of the argument)
2. If B , then C . (Premise of the argument)

Sub-conclusion: If A , then C . (From (1), (2), and hypothetical syllogism)

3. Not C . (Premise of the argument)
4. Therefore, not A . (From the sub-conclusion, (3), and *modus tollens*)

So although the argument involving Oscar doesn’t fit any of our famous forms, we can use those famous forms to reason step by step to the conclusion, thereby proving that the conclusion follows from the premises.

Let’s try one more example like this where we have to “combine” the various famous forms:

1. Either it will rain tonight or it will snow tonight.
2. It won’t rain tonight.
3. If it snows tonight, then tomorrow’s class will be canceled.
4. Therefore, tomorrow’s class will be canceled.

Before reading ahead, we encourage you to try prove that this argument is valid on your own by combining the various famous forms . . . and now assuming you’ve tried it, here’s how we might proceed. First, we abstract away from the argument’s content and focus on its form, which is this:

1. Either A or B .
2. Not A .
3. If B , then C .
4. Therefore, C .

Again, this argument doesn’t fit any of our famous forms, but we can “combine” them to prove that this argument is valid. Look at the first two premises:

1. Either A or B .
2. Not A .

We know from disjunctive syllogism that “ B ” follows from (1) and (2). But now look at premise (3):

3. If B , then C .

If we combine “ B ” with “If B , then C ”, we know from *modus ponens* that “ C ” follows, which is exactly our conclusion! That is, from premises (1) and (2), we can derive “ B ” using

disjunctive syllogism. Then, from “*B*” and premise (3) we can derive our conclusion, “*C*”, using *modus ponens*. If we were to visualize our steps, it might look like this:

1. Either *A* or *B*. (Premise)
2. Not *A*. (Premise)

Sub-conclusion: *B*. (From (1), (2), and disjunctive syllogism)

3. If *B*, then *C*. (Premise)
4. Therefore, *C*. (From sub-conclusion, (3), and *modus ponens*.)

Now even this extended version of the famous forms method has significant limitations, and a full course in logic explores exactly what those limitations are, ways of getting around them, and other methods for providing validity. But for many of the arguments you’ll come across in this course and in everyday life, this extended version of the famous forms method is sufficient.

Proving Invalidity

The last section covered methods for proving that an argument is valid, but what about proving that an argument is invalid? Just as with proving validity, we can only examine the most rudimentary methods here. Nonetheless, even basic methods are quite helpful. We’ll look at two such methods: the countermodel method and the parity argument method (or well-known truth to falsehood method).

The Countermodel Method

If an argument is *valid*, that means if the premises were true, the conclusion would have to be true. So if the argument is *invalid*, it must be possible for the premises to be true while the conclusion is false. In the countermodel method, the goal is to simply construct such a possibility. Consider one of the arguments from above:

1. Everyone who lives in Los Angeles lives in California.
2. Alvin lives in California.
3. Therefore, Alvin lives in Los Angeles.

We noted that this argument is invalid: even if the premises are true, that doesn’t *guarantee* that the conclusion is true. For instance, what if Alvin lives in San Francisco, not Los Angeles? In that case, both (1) and (2) will be true and yet (3) will be false. That possibility shows that the argument is invalid.

What we have just come up with is a *countermodel*. Surely it is *logically possible* for Alvin to live in San Francisco—there is no outright *contradiction* in saying he does. But in that possibility, the premises are true while the conclusion is false. And because validity is such a strong

notion, even just one logical possibility where the premises are true and the conclusion is false is enough to show that the argument is invalid. More carefully:

A *countermodel* to an argument is a logical possibility where all of the premises of the argument are true while the conclusion is false.

Two points about countermodels. First, as has already been implied, if there is a countermodel to an argument, then the argument is invalid: it proves that truth of the premises does not *guarantee* the truth of the conclusion. Second, a countermodel need only be a *logical possibility*—a case that does not involve a contradiction. It doesn't have to be a *realistic* or even *scientific* possibility. The imagined scenario can involve wizards and dragons, giant mountains made of pure gold, time travel, and whatever else you like. The only requirement is that it not involve a *logical contradiction*: it can't involve one and the same sentence being both true and false at the exact same time. For instance, with the above argument, if we knew Alvin, we may think it *unlikely* or *unrealistic* that he lives in San Francisco. We may even know for a fact that he doesn't live there. But that doesn't matter. As long as it is *logically possible* for him to live in San Francisco—and it undoubtedly is—that's enough for a countermodel.

Let's try another case. See if you can come up with a countermodel which shows that the following argument is invalid:

1. Everyone who is a philosophy major is a deep thinker.
2. Lucy is not a philosophy major.
3. Therefore, Lucy is not a deep thinker.

There are literally an infinite number of countermodels to this argument, but here is just one: suppose that every philosophy major is a deep thinker, just as the first premise says, but that every physics major is also a deep thinker. And suppose Lucy is a physics major and not a philosophy major. Surely, this scenario is a logical possibility—it doesn't involve a contradiction. But in such a case, (1) and (2) are both true while (3) is false. Hence, we've found a countermodel, meaning the argument has been proven to be invalid.

The countermodel method is a relatively easy and straightforward method for proving invalidity, but it has severe limitations, especially when arguments get complex. For instance, consider this somewhat complicated argument:

1. Every epistemologist is a philosopher.
2. Some, but not all, epistemologists are reliabilists.
3. Every philosopher who is a reliabilist is an externalist.
4. Therefore, some epistemologists are not externalists.

This argument is indeed invalid, but constructing a countermodel can be difficult. First, if you don't know what an epistemologist, reliabilist, or externalist is, trying to imagine a scenario where the premises are true but the conclusion is false might sound like an impossible task. Second, because the argument is somewhat complicated, it can be hard to

keep everything “in your head” at once. So while the countermodel method, as described here, is a good method for simple arguments with familiar terms, it’s not so good for more complicated arguments. This is where the next method for proving invalidity can help.

The Parity Argument Method (or Well-Known Truth to Falsehood Method)

*Note to the reader: This sub-section is more difficult.
Feel free to skip ahead to the next section.*

Suppose you come across an argument where you *know* all of the premises are true but you also *know* that the conclusion is false. If you know those things, you also *know* that the argument is invalid. Why? Well, recall the definition of a valid argument: if the premises were true, the conclusion *would have* to be true. So if you know that the premises are true and yet the conclusion is false, then it obviously isn’t valid.

Now when we are trying to figure out whether an argument is valid, we typically don’t know whether the premises and conclusion are true or false. (If we did, we probably wouldn’t be examining the argument to begin with!) But what we can do instead is try to *compare* the argument in question with an argument where we know the premises are true but the conclusion is false. If we can find a legitimate comparison, then we know by parity of reason that both arguments are invalid.

Let’s work through the last argument we looked at to see how this works:

1. Every epistemologist is a philosopher.
2. Some, but not all, epistemologists are reliabilists.
3. Every philosopher who is a reliabilist is an externalist.
4. Therefore, some epistemologists are not externalists.

As we noted, trying to come up with a countermodel can be difficult, partly because we might not know some of these terms and partly because there are too many moving parts. The parity argument method starts, just as the famous forms method does, by abstracting away from the content and looking at the argument’s *form*, which is something like this:

1. Every *E* is a *P*.
2. Some, but not all, *Es* are *Rs*.
3. Every *P* who is an *R* is an *X*.
4. Therefore, some *Es* are not *Xs*.

By looking at the argument’s form, we get around the problem of possibly not knowing what the particular words mean. That’s the easy step. The harder step is trying to come up with an argument that has this same form, but where the premises are well known to be true while the conclusion is well known to be false. Consider this argument:

1. Every woman is a human.

2. Some, but not all, women are tall.
3. Every human who is tall is a mammal.
4. Therefore, some women are not mammals.

This argument has the exact same form as the argument we started with. All we've done is replace some of the terms: "*E*" now stands for "woman"; "*P*" stands for "human"; "*R*" stands for "tall"; and "*X*" stands for "mammal".

Here is what's so important about this argument though: the premises are *well-known truths* while the conclusion is a *well-known falsehood*. So we *know* that this argument is invalid. And since this argument has the exact same form as our original argument, the original argument must also be invalid.³

Let's do another example. Try this one on your own before reading ahead. (The steps are: (i) abstract away from the content and identify the argument's form, and (ii) come up with an argument that has the same form but where the premises are well known to be true and the conclusion is well known to be false.)

1. If Lin wins the next game, he'll advance to the finals assuming Crystal is eliminated.
2. Crystal will be eliminated if Bobby moves his knight to square H3.
3. Bobby won't move his knight to square H3.
4. Therefore, even if Lin wins the next game, he won't advance to the finals.

Again, we can use the parity argument method to show that this complicated argument is invalid. We start by noting that the argument has the following form:

1. If *L*, then *A* assuming *C*.
2. *C* if *B*.
3. Not *B*.
4. Therefore, even if *L*, not *A*.

We can then substitute in different sentences for "*L*", "*A*", "*C*", and "*B*" to get an argument where the premises are well known to be true while the conclusion is false. For instance, consider this (admittedly silly) argument:

1. If every mammal dies today, then Barack Obama will die today assuming he is a mammal.
2. Barack Obama is a mammal if he is a horse.
3. Barack Obama is not a horse.
4. Therefore, even if every mammal dies today, Barack Obama will not die.

This argument is so ridiculous, it's laughable. But it has the exact same logical form as the argument we started with. All we've done is replace "*L*" with "Every mammal dies today"; "*A*" with "Barack Obama will die today"; "*C*" with "Barack Obama is a mammal"; and "*B*" with "Barack Obama is a horse". Moreover, each of the premises is patently true, and yet the

conclusion is manifestly false: contrary to what the conclusion states, if every mammal dies today, Barack Obama will certainly die too. This shows that, even if the premises of the original argument are true, that doesn't guarantee that its conclusion is true. The original argument is therefore invalid.

So the parity argument method can help us overcome some of the limitations of the countermodel method. Still, this method has severe limitations of its own. For one, it requires a good amount of creativity. For another, when arguments get *extremely* complex, it is simply impractical to use this method. A full course in logic explores more powerful and systematic methods for proving invalidity.

Non-deductive Arguments

Not all invalid arguments are bad arguments. Why? Well, sometimes the premises of an argument are only meant to make the conclusion *likely* or *reasonable*, not necessarily *guaranteed*. For instance, suppose your car refuses to start one morning. You do a bit of investigative work and become convinced that the battery is dead. You reason as follows:

1. My car wouldn't start this morning.
2. I accidentally left the lights in my car on all of last night.
3. Last time my car was in the shop, the mechanic told me the battery would need to be replaced soon.
4. The car has plenty of gas in it—I just filled it up yesterday.
5. Therefore, the car battery is dead.

This argument looks compelling. But notice that it is *invalid* in our technical sense: even if the premises are true, the conclusion doesn't *have to be* true. For instance, what if the battery is fine, but there's a short in the electrical system? Or what if a friend is playing a prank on you, having disconnected your perfectly good battery? Or what if the battery is in working order, but aliens have abducted your engine and replaced it with a very realistic wax copy? As far-fetched as these scenarios might be, they still present *logical possibilities* where the premises are true while the conclusion is false. And that's enough to show that the argument is invalid.

Just because this argument is *invalid*, though, does that mean it's a *bad* argument? That it is to be dismissed? Of course not! The argument is perfectly respectable and worth taking seriously. That's because, in this case, you are only trying to figure out what is *most likely* the case, not what is absolutely *guaranteed* to be the case. Sure, it is logically possible that your car has a short, or that a friend is playing a prank on you, or that aliens have stolen your engine. But how likely are these logical possibilities given your evidence? Not very. The premises make it far more likely that your battery is dead, and insofar as that's all you were trying to show, the argument looks like a good one.

Arguments like this one point out that validity and soundness, as defined in the previous sections, are not the "end all-be all" of argumentation. We routinely employ arguments that

are not valid (and hence not sound) but still compelling. To help capture this idea, let's introduce a few more definitions. First, the definition of a "non-deductive argument":

*A non-deductive argument is an argument where the premises supposedly make the conclusion highly probable, although not guaranteed.*⁴

Notice that this definition, like our definition of an argument, only says that the premises *supposedly* make the conclusion highly probable. Sometimes the premises really do make the conclusion likely or reasonable, and sometimes they don't. If the premises of a non-deductive argument really do make the conclusion likely or reasonable, we'll call it a "strong" argument. More carefully, we'll introduce a second definition:

*A non-deductive argument is strong if it satisfies the following condition: if its premises were true, then its conclusion would be highly probable.*⁵

Just as an argument can be *valid* and have false premises, so an argument can be *strong* and have false premises. All that is required is that *if* its premises were true, the conclusion would likely be true as well. (If it helps, you can think of it this way: although such arguments are technically invalid, the connection between the premises and conclusion is still "strong.")

Our third and final definition is that of a cogent argument. A cogent argument is a strong one with all true premises. More precisely:

A non-deductive argument is cogent if it is strong and has all true premises.

If an argument is cogent, then you know the argument has made the conclusion highly probable. Of course, that doesn't necessarily mean the conclusion is true—that's what *sound* arguments establish. But at least the author has offered premises that are both true and relevant to the conclusion.

It turns out that there are many different types of non-deductive arguments, and determining whether any given non-deductive argument is strong (and hence cogent) will often require first determining what type it falls into. In what follows, we'll briefly review two kinds of non-deductive arguments that play a prominent role in philosophy.

Arguments from Analogy

Arguments from analogy are everywhere, but perhaps the best examples come from the medical sciences. For instance, suppose scientists have developed a promising new drug but need to know if it will have any adverse effect on humans. Instead of just giving doses to random people, which is risky and irresponsible, these scientists decide to first experiment on non-human animals, like pigs. They then note that the drug had no adverse effect on the pigs and, on that basis, infer that the drug will likely have no adverse effect on humans either. That is, they reason like this:

1. Drug *D* had no adverse effect on pigs.
2. Pigs are relevantly similar to humans—the skin, flesh, and organs of pigs are quite similar to those of humans.
3. Therefore, drug *D* will have no adverse effect on humans.

This is an argument from an analogy. At a general level, with arguments from analogy, we infer that certain individuals (probably) have a certain property since *relevantly similar* individuals have that property. That is, the general form of an argument from analogy looks like this:

1. Individual(s) *X* has property *P*.
2. Individual(s) *X* is relevantly similar to individual(s) *Y*.
3. Therefore, individual(s) *Y* has property *P*.

In the case above, *X* is the population of pigs the researchers administered the drugs to, *P* is the property of having no adverse effects from the drug, and *Y* is the population of humans.

When is an argument from analogy a *strong* argument? In short, the more relevant similarities that *X* and *Y* have, and the deeper these similarities, the better. For instance, as noted, the skin, flesh, and organs of pigs are quite similar to those of humans, and these similarities seem especially important for medical purposes. That makes the researchers' argument look strong. If they instead had tested the drug on termites and, on that basis, concluded that the drug would have no adverse effect on humans, the argument would not be strong. The relevant similarities between termites and humans are so few and superficial that however the drug affected the termites would likely tell us next to nothing about how it will affect humans.

The insistent reader might ask: But *exactly* how many relevant similarities must there be between *X* and *Y* for an argument from analogy to be strong? There is no definite answer. In some cases, just one or two relevant similarities will be enough, especially if those similarities are quite deep; in other cases, we will want more than that. This can make arguments from analogy a bit messy. Whereas deductive arguments are black and white—either an argument is valid or it isn't—non-deductive arguments, like arguments from analogy, often have some “gray area.” That's something to keep in mind when evaluating arguments from analogy (and most non-deductive arguments).

Last, it's important to stress that the similarities between *X* and *Y* must be *relevant*. Not just any old similarity will do. For instance, imagine someone offering an argument from analogy like this:

1. Abraham Lincoln was a good president.
2. Abraham Lincoln is relevantly similar to my uncle—my uncle, like Lincoln, is tall and likes to wear top hats.
3. Therefore, my uncle would also be a good president.

This argument, of course, is ridiculous. It may be that Abraham Lincoln and your uncle have some similarities—being tall and enjoying a good top hat—but they aren't *relevant* to the issue at hand. Whether or not someone is a good president has nothing to do with whether they are tall or frequently wear top hats. That's why the argument is so absurd. So in determining whether an argument from analogy is strong, it is important to check that the stated similarities are genuine and relevant.

Inference to the Best Explanation

Inferences to the best explanation are sometimes called *abductive* arguments, and they are another kind of non-deductive argument. For instance, suppose you and your friends are out walking deep in the woods when you come across a sizeable crater. Upon inspection, you find some puzzling items in the crater—there are mangled pieces of metal and even radioactive debris! You then recall that about a week ago, a nuclear satellite went missing in the area, although its remains haven't been found yet. You naturally conclude that the crater you are looking at is where the satellite crashed into the Earth.⁶ That is, you reason as follows:

1. There is a large crater out in the woods.
2. There are pieces of mangled metal as well as radioactive debris in the crater.
3. A nuclear satellite went missing in the area a week ago.
4. The crash site of the missing satellite has not been found.
5. Therefore, the satellite crashed into the Earth here.

Although invalid in our technical sense, this argument still looks promising. In this case, the conclusion is supported by the premises because the conclusion would *explain* at least some of the premises: if the conclusion were true, that would be a good explanation for premises (1) and (2). And although premises (3) and (4) aren't explained by the conclusion, they nevertheless make the conclusion look like an even better explanation of (1) and (2). An argument with this structure, where the conclusion purports to explain some of the premises, is called an *inference to the best explanation* or *abductive* argument.

When does an inference to the best explanation constitute a *strong* argument? Most basically, it's when the conclusion really is the *best* explanation of the relevant premises. What makes one explanation better than another? That's a vexed issue, but philosophers are largely in agreement over three criteria for evaluating proposed explanations: (i) *explanatory power*; (ii) *fit with background knowledge*; and (iii) *simplicity*. Let's briefly talk through each one.

First, an explanation's *explanatory power* has to do, at least partly, with how *likely* it makes the premises it supposedly explains. For instance, in the satellite case, the conclusion is meant to explain premises (1) and (2), and it seems to do a pretty good job of that. Why? Well, *if* the conclusion is true, premises (1) and (2) would be very likely. That is, if we assume that the nuclear satellite crashed at that spot, it would very likely create a crater filled with mangled metal and radioactive debris. Neither premise is *guaranteed* to be true if we assume the

conclusion, but each is still very probable. And that's why the conclusion—that the satellite crashed at that location—has a good amount of explanatory power.

Second, an explanation's *fit with background knowledge* has to do with how well the explanation fits with everything else we know about the world. For instance, suppose one of your friends comes to a very different conclusion about the crater: according to your friend, the crater was made by an alien spacecraft, one that is fueled by nuclear energy. This explanation might very well have a decent amount of explanatory power with respect to premises (1) and (2): it would explain what caused the crater as well as why there are mangled pieces of metal and radioactive debris. So why are we hesitant to consider this alternative explanation? It's because we simply don't know whether aliens exist and have the technology to visit our planet. By comparison, we know that nuclear satellites exist. We even *know* that one went missing in the area recently and has yet to be found—that's what premises (3) and (4) assert. That means your conclusion fits with our background knowledge much better than your friend's conclusion. So even if both explanations have the same amount of explanatory power, we have strong reason to prefer your conclusion over your friend's.

Third, an explanation's *simplicity* has to do with how many entities the explanation posits (or basic principles it employs). Suppose another friend of yours comes to yet another conclusion: the crater was caused by a small meteor, not the satellite; however, the satellite still flew over the crater, and as it did, pieces of metal and radioactive debris fell off the satellite and just happened to land in the crater. This explanation has a good amount of explanatory power with respect to premises (1) and (2). It also fits well with our background knowledge—we know that sometimes meteors fall out of the sky and that satellites can drop debris as they fall to the Earth. Still, your conclusion that the crater was caused by the satellite looks more plausible than your friend's conclusion that it was caused by the meteor and then filled with debris as the satellite passed overhead. That's partly because your conclusion is simpler than your friend's. Whereas your conclusion explains premises (1) and (2) with a single event—the satellite crashing—your friend's conclusion explains (1) and (2) with multiple events—the meteor crashing and then the satellite flying overhead, dropping materials as it went. Your friend's conclusion is a bit more complicated. So even if your friend's conclusion has a good amount of explanatory power and fits well with our background knowledge, it lacks the simplicity that your conclusion has. Absent any other evidence, that gives us reason to prefer your conclusion over your friend's.

This third criterion is the most controversial among the three, especially in the philosophy of science. It's obvious why we would prefer to *work with* simpler explanations, especially when it comes to the sciences: Why use complicated formulas and theories to explain something when simpler formulas and theories will do just as well? Who wouldn't want to save themselves the additional effort? That is, there are strong *practical* reasons to prefer simpler explanations. But just because one explanation is easier to work with than another, does that really mean it is more likely to be *true*? Why should we expect the truth to be (relatively) simple? This is an important debate in the philosophy of science but we will not engage in it here. Suffice it to say that, as a matter of fact, we do prefer simpler explanations to complicated ones, regardless of whether that is for purely practical reasons or not.

So explanations are usually judged according to (at least) these three criteria: *explanatory power*, *fit with background knowledge*, and *simplicity* (although this last one is controversial). If the proposed explanation is better than every other explanation in regard to these three criteria, then most philosophers would agree that it is indeed the best explanation. Unfortunately though, things are usually not so simple. Oftentimes, one explanation will do quite well with regard to one of these criteria but not so well with regard to the others. What should we do if one explanation is extremely simple and has a lot of explanatory power, but doesn't fit with our background knowledge terribly well? Should we prefer it to an explanation that is a bit more complicated and doesn't have as much explanatory power, but fits perfectly with our background knowledge? These are hard cases and there is no algorithm for determining the right answer. This is again where some "gray area" can appear. Oftentimes, it will be unclear whether a proposed explanation really is the best explanation, and so it is unclear whether the argument is a strong one.

Common Fallacies

A fallacy is a mistake in reasoning. Fallacies are commonly divided into two types: *formal* fallacies, where the mistake has to do with the *form* of the reasoning, and *informal* fallacies, where the mistake has to do with the *content* of the reasoning. The difference is important. If an argument commits a *formal* fallacy, then it is not a *valid* argument (as defined above). In contrast, if an argument commits an *informal* fallacy, then it may very well be *valid* (perhaps even sound!), but it is still problematic. We'll briefly cover two common formal fallacies and seven common informal fallacies.

Formal Fallacy: Affirming the Consequent

The following argument commits the fallacy of affirming the consequent:

1. If Amelia can vote in the United States, then Amelia is 18 years old.
2. Amelia is 18 years old.
3. Therefore, Amelia can vote in the United States.

The first premise of this argument is a *conditional*—that is, it is an "if . . . then" statement. The "if" part of a conditional is called the *antecedent*, and the "then" part of a conditional is called the *consequent*. Notice that premise (2) asserts the truth of the consequent of the conditional in premise (1), and then the argument concludes that the antecedent is therefore true. This is why this is called affirming the consequent, and it is an invalid form of reasoning. (Feel free to use either the countermodel method or the parity argument method to verify this.) Any argument that takes this form—a conditional, the consequent affirmed, and then the antecedent as conclusion—is invalid.

Why do people fall for the fallacy of affirming the consequent? It's partly because the fallacy looks similar to a perfectly valid form of reasoning already covered, namely, that of *modus ponens*. With *modus ponens*, we affirm the *antecedent* and infer the *consequent*, which is obviously a

valid inference. With the fallacy of affirming the consequent, we do the reverse: we affirm the *consequent* and infer the *antecedent*, which is an invalid inference. You might say that the fallacy of affirming the consequent is the “evil twin” of *modus ponens*. This is an example of how great care must be taken when examining an argument’s form so as to make sure we don’t confuse valid forms of reasoning with invalid ones or vice versa.

Formal Fallacy: Denying the Antecedent

Knowing what you now know about conditionals, you can probably guess what this fallacy looks like. Here’s one example:

1. If Amelia can vote in the United States, then she is 18 years old.
2. Amelia cannot vote in the United States.
3. Therefore, Amelia is not 18 years old.

Again, we have a conditional in the first premise, but in this case the second premise is a denial of the antecedent. The argument then concludes that the consequent must be false as well. But this too is a fallacious form of reasoning. (Feel free to use either the countermodel method or the parity argument method to verify this as well.) And again, any argument that takes this form—a conditional, the antecedent denied, and then the consequent denied as a conclusion—is invalid.

If affirming the consequent is the evil twin of *modus ponens*, then denying the antecedent is the evil twin of *modus tollens*. With *modus tollens*, we deny the *consequent* and infer the denial of the *antecedent*, which is a valid inference. With denying the antecedent, we do the reverse: we deny the antecedent and infer the denial of the consequent, which is invalid. This is another example of why we must be so careful in determining whether an argument is valid or not.

Informal Fallacy: Circular Reasoning

Suppose Alexis and Baht are having a debate over God’s existence, Alexis being a believer in God and Baht being a disbeliever in God. Baht asks Alexis how she knows God exists, to which Alexis responds, “Because, He (God) does!” One way of interpreting her argument is like this:

1. God exists.
2. Therefore, God exists.

Clearly, this is an unpersuasive argument. But note that, technically speaking, it may very well be a *sound* argument. It is certainly valid in our technical sense: if the premise were true, the conclusion would have to be true—after all, the premise and the conclusion are the same thing! That means if the premise in fact turns out to be true, the argument is technically sound.

The reason this argument is so unpersuasive, then, is not because it is invalid (or even unsound), but rather because it is *circular*. An argument is circular if its conclusion appears somewhere within its premises. The reason why no one should be persuaded by a circular argument is that one would have to *already* accept the conclusion of the argument *before* one accepted the premises. This gets things backward. Those who already accept the conclusion—like Alexis—will not need the argument to be persuaded, and those who do not already accept the conclusion—like Baht—have been given no reason to accept the premise.

Now typically, when an argument is circular it is not this obvious. More commonly, authors will use synonyms or complicated language to hide the argument's circularity. That's another reason it is so important to go through an argument carefully and slowly, separating out the premises from conclusion and asking questions along the way.

Informal Fallacy: Begging the Question

The fallacy of begging the question is similar to the fallacy of circular reasoning, although it is a bit subtler. Suppose Alexis offers another argument for God's existence:

1. The Bible says that God exists.
2. Everything the Bible says is true.
3. Therefore, God exists.

Again, this may very well be a sound argument. At the very least, it is certainly valid in our technical sense, meaning there is nothing wrong with the argument's form. But it's also unpersuasive because, presumably, one would need to accept its conclusion before one accepted premise (2). That is, Alexis likely believes that everything the Bible says is true because she believes God exists and "wrote" the Bible. This is the fallacy of *begging the question*. More precisely, an argument begs the question if one or more of its premises relies for its truth on the truth of the conclusion.

Notice how the fallacy of begging the question is slightly different than the fallacy of circular reasoning. In a circular argument, the conclusion *appears as a premise* in the argument. In contrast, when an argument begs the question, the conclusion doesn't necessarily appear as a premise, but rather is used implicitly to *support* a premise. This difference makes the fallacy of begging the question more difficult to spot: in order to know whether an argument begs the question, we have to know what is meant to support the premises, which isn't always obvious. The general rule of thumb is this, though: if the only people who would find a certain premise of the argument plausible are those people who already accept the conclusion, the argument begs the question.

Informal Fallacy: Ad Hominem Fallacy

Ad hominem roughly means "against the person." An author commits an ad hominem fallacy when, instead of engaging with the issue at hand, the author attacks the individual(s)

on the other side of the dispute. For instance, suppose Alexis and Baht are continuing their debate over God's existence, with Alexis offering a much better argument this time:

1. The universe had a beginning.
2. If the universe had a beginning, then it was brought about by God.
3. Therefore, the universe was brought about by God.

Baht is unconvinced and offers this rebuttal: "Of course you find this argument convincing—you are so desperate to believe in God that you'll accept any argument in favor of God's existence."

Baht's rebuttal commits the *ad hominem* fallacy. His comments don't engage Alexis's argument whatsoever—they don't raise doubts about a premise, question the argument's validity, point out an assumption the argument makes, or anything of the type. At best, his comments discredit the *author*, not the *argument*. But even if the author is unreliable or has suspicious motives, that doesn't necessarily mean the author's argument is problematic.

Of course, sometimes it is perfectly appropriate to question the author's credibility or character, especially if the author is presented as an expert or an authority. Nevertheless, it's important to remember that to discredit the author is not to discredit the argument.

Informal Fallacy: Ad Hoc Fallacy

An author commits the *ad hoc* fallacy when they offer a premise or rationale that doesn't have much independent support, but is crafted just to get around an objection or avoid a certain issue. Suppose Alexis and Baht have moved on a bit in the discussion and are now debating whether God created morality, at least assuming God exists. Alexis offers the following argument:

1. God created everything that exists.
2. Moral truth exists.
3. Therefore, God created moral truth.

Baht then raises a good objection: "Wait a second," he says. "Even if God exists, God didn't create himself—that would be nonsense. So God couldn't have created everything. How do we know morality isn't another exception?" Alexis says, "Ah, fair point. Let me revise my argument," and offers this:

- 1.* God created everything that exists except himself.
- 2.* Moral truth exists (and is distinct from God).
- 3.* Therefore, God created moral truth.

This revised argument, particularly premise (1*), is *ad hoc*. Baht's objection raises a deep issue: it seems that there are some things that it would be nonsensical for God to create, such as himself. What if it is also nonsensical for God to create moral truth? Surely that is a

question worth pondering. Alexis's revised argument doesn't deal with this deep issue. Instead, premise (1*) has been crafted to get around Baht's particular counterexample, thereby skirting the question. A better response from Alexis would be to explore whether moral truth is the kind of thing that would be sensible for God to create.

There are two points worth making about the ad hoc fallacy. First, if a premise or position is ad hoc, that doesn't necessarily mean the premise or position is *incorrect*. Alexis's revised argument may very well be sound! The problem with ad hoc premises and positions is that they are *undermotivated*. They have been crafted just to save the argument or view, not to get at the truth. Second, whether a premise or position is ad hoc is a bit of a judgment call and will often depend on the context. There's no precise way to determine whether an author has committed the ad hoc fallacy. Some philosophers have even thought that the ad hoc fallacy was too subjective to be of any value. Still, we think it is a useful fallacy to know about, at least as a way of articulating what you may find problematic about certain premises or positions. When a premise or position is ad hoc, it is a signal that we should evaluate it more closely.

Hasty Generalization

An author commits the hasty generalization fallacy when the author infers a sweeping principle or claim without a sufficient number of cases to support that principle or claim. For instance, suppose Alexis and Baht are debating whether atheism is on the rise, and Baht offers the following argument:

1. Both my friend and my cousin used to believe in God, but now they are atheists.
2. I don't know anyone personally who used to be an atheist but now believes in God.
3. Therefore, belief in God is declining overall.

Baht's argument involves a hasty generalization. His conclusion, that belief in God is declining overall, may very well be true, but the evidence he has offered is far too weak to establish this broad claim. Citing two cases where someone no longer believes in God, as well as the fact that he doesn't know anyone personally who has come to believe in God, is not nearly enough. He has jumped to his conclusion much too quickly.

While the hasty generalization fallacy is fairly obvious, it is all too easy to fall for. Hearing one or two personal stories is often much more moving and memorable than being told about in-depth studies, surveys, or statistical analyses. That's why public figures so often use such stories in their arguments. But as powerful as these stories may be, it's important to remember that as *evidence*, they often fall short of establishing the general claim the author is trying to establish.

Informal Fallacy: Tu Quoque Fallacy

Tu quoque roughly means “you too.” An author commits the *tu quoque* fallacy when, instead of answering a pressing objection to their own argument or view, they point out that their opponent’s argument or view suffers from a similar objection. For instance, suppose Alexis and Baht are at it again, this time debating whether God’s existence is compatible with free will. Baht argues that the two are not compatible, giving the following argument:

1. If God exists, then every one of our actions is the result of God’s divine decrees.
2. If every one of our actions is the result of God’s divine decrees, then we never freely perform any action.
3. Therefore, if God exists, then we never freely perform any action.

Now imagine Alexis responds by arguing that a similar problem for free will appears if God doesn’t exist. She argues as follows:

1. If God doesn’t exist, then every one of our actions is the result of the distant past and the laws of physics.
2. If every one of our actions is the result of the distant past and the laws of physics, then we never freely perform any action.
3. Therefore, if God doesn’t exist, then we never freely perform any action.

By offering this argument, Alexis is trying to show Baht that his worldview also rules out free will. That may very well be true, but it’s not especially helpful in evaluating Baht’s original argument. Does Alexis’s argument show that Baht’s argument has a false premise? Or that it is invalid? No, not at all. At best, it shows that Baht’s worldview faces a similar problem as her own worldview does. That doesn’t help us determine whether Baht’s argument is sound.

Now there are occasions where a *tu quoque* response is perfectly appropriate, namely, when weighing the pros and cons of various positions. For instance, if Baht’s overall goal is to show that atheistic worldviews are friendlier to free will than theistic ones, Alexis’s response is germane: her response is meant to show that atheism is no friendlier to free will than theism. But if Baht is just trying to show that God’s existence rules out free will, Alexis’s *tu quoque* response isn’t helpful.

Informal Fallacy: Strawman Fallacy

An author commits the strawman fallacy when they misrepresent their opponent’s position, argument, or objection, making it weaker than it really is. Suppose our good friends, Alexis and Baht, are wrapping up their discussion and a few other people are now listening in. Baht tries winning over the crowd by offering the following: “Alexis here believes there is a God. That is, she believes there’s some bearded old man in the sky, watching our every move, hoping we make a mistake because that means he gets to punish us with eternal hellfire.”

Baht has committed the strawman fallacy. Alexis almost certainly doesn’t think of God as an old man in the sky who enjoys punishing us when we do wrong. Her view is likely more

sophisticated and more plausible than that. In trying to win over the crowd, Baht has misrepresented Alexis's view, setting up a version of it that is far easier to knock down or dismiss.

A good remedy for the strawman fallacy is the principle of *charity*: when framing an author's position, argument, or objection, we should do so as charitably as possible. Put the author's position, argument, or objection in the strongest possible way. Don't attribute anything terribly implausible to the author unless the author explicitly affirms it. And even when the author does explicitly affirm something that sounds implausible, look for a more plausible interpretation that still fits with everything else the author has said. Of course, there are limits to the principle of charity. Sometimes authors really do say implausible things. Moreover, what may sound implausible to one person may not sound so implausible to another. Still, when evaluating an author's position, argument, or objection, start by finding the strongest version of it. Doing so will help us get at the truth.

Necessary and Sufficient Conditions

So much for arguments. Another important logical concept is that of necessary and sufficient conditions. The best way to get a handle on these concepts is through an example. So consider the following statement:

If you are a sophomore, then you are an undergraduate.

This statement is saying that being a sophomore is *sufficient* for being an undergraduate. In other words, *all* you need to be an undergraduate is to be a sophomore. (But that's not to say that's the *only* way to be an undergraduate.) In general, a statement of the form:

If X, then Y

is a statement that *X* is a sufficient condition for *Y*. Now consider the following statement:

If you can vote in the United States, then you are at least 18 years old.

This statement is saying that being at least 18 years old is *necessary* for being able to vote in the United States. In other words, one of the requirements for being able to vote in the United States is that you must be at least 18 years old. (But that's not to say that that is the *only* requirement.) In general, a statement of the form:

If X, then Y

is a statement that *Y* is a necessary condition for *X*. Occasionally you will come across a statement that purports to give both necessary and sufficient conditions for something. For example:

You have a sister if and only if you have a female sibling.

This statement says the same thing as the following two statements combined:

If you have a sister, then you have a female sibling.
If you have a female sibling, then you have a sister.

Or, in the language of necessary and sufficient conditions:

Having a sister is necessary and sufficient for having a female sibling.

Philosophers are often interested in the necessary and sufficient conditions for some interesting concept, such as knowledge. An interesting philosophical question is: What are the necessary and sufficient conditions for the claim that you have knowledge about some fact? Certainly it is necessary that what you think you know must actually be true for you to know it. But is that also sufficient? Probably not, as you may believe something is true even though you don't have any good reason to believe it, and so on.

A Priori and A Posteriori

It will be useful to have a few more pieces of philosophical terminology at our disposal. First, philosophers often distinguish between a priori and a posteriori. These are Latin terms that are especially useful in describing the way in which we are able to come to know certain propositions. Propositions that can be known a priori are those that can be known completely independent of experience. They are those propositions that we can know, so to speak, "from the armchair." For example, our knowledge that all triangles have three sides is a piece of a priori knowledge. There's no need to go around the world looking for triangles and counting up their sides to conclude that all triangles have three sides. On the other hand, propositions that can be known a posteriori are those that require experience of the world to come to know. For example, your knowledge that it is raining outside right now is a posteriori knowledge. To determine whether it is raining, you need to open your eyes and look at the world. No amount of armchair speculation will help.

Necessary and Contingent

Another distinction that comes in handy in philosophy is one between *necessary* and *contingent* truths. A necessary truth is a proposition that is true and could not have been false, whereas a contingent truth is a proposition that is true but might have been false. Most of the true propositions we ordinarily come across are contingent propositions. For instance, the fact that you are reading this right now is a contingent truth. You could very well have decided to do something else with your time. Even the fact that you exist is a contingent truth. Had your parents not met when they did, you could very well have never been born. In fact, we are so surrounded by contingent truths that it's difficult to think of an uncontroversial necessary truth. An example would be the fact that all triangles have three sides. No matter how the world could have been, triangles would always have had three sides—that statement

could not have been false. Of course, we could have used the word “triangle” to talk about four-sided figures, but that’s not to say that triangles could have been four-sided figures. The concept of a triangle is so intimately connected up with the concept of three-sidedness that it’s impossible to have one without the other. Another example is the fact that all bachelors are unmarried. This is a necessary truth because no matter how the world could have been, bachelors would have always been unmarried.

Although these terms are most often used to talk about true and false propositions, they are also sometimes used to distinguish between necessary and contingent *existence*. You and I exist only contingently—that is, we might not have existed. God, on many interpretations, is supposed to exist necessarily—that is, God could not have *not* existed.

NOTES

1. We have broken convention here. In the previous example, we had uppercase letters stand for individuals like Alvin, Los Angeles, and California, whereas in this example, we are using uppercase letters to stand for sentences. Traditionally, lowercase letters are instead used to stand for individuals while uppercase letters are exclusively used to stand for sentences, predicates, or categories. In developing a precise and robust logic, it is crucial to adhere to some such convention, but given our goals here, we won’t be so exact.

2. The astute reader will notice that the uppercase letters have been bolded here, whereas in the examples above and below they are not. That is because the bolded letters are technically “meta-linguistic variables” whereas the non-bolded letters are “atomic sentence letters.” The difference is roughly this: whereas atomic sentence letters are meant to stand for simple sentences in English (or any natural language), meta-linguistic variables are meant to stand for atomic sentence letters as well as certain combinations of atomic sentence letters. For example, while *A* can stand for “Oscar is hungry” or “It will rain tonight”, etc., **A** can stand for *A* or *B* or *C*, as well as “If *A*, then *B*” or “*B* or *C*”, etc. This distinction allows our “famous forms” method to be much more versatile. (Technically, something similar must be done with the line numbers as well.) You can start to see how logic gets complicated quickly!

3. Technically this only proves that the argument is *formally* invalid. There are funny cases where an argument is formally invalid although still valid, such as arguments with a necessary truth as the conclusion. (This raises difficulties for how to best define validity.) We’ll ignore this complication here.

4. Such arguments are often called “inductive” arguments. But sometimes inductive arguments pick out a narrower class of arguments, namely, those focused on enumerative induction and the like. To avoid possible confusion, we’ll stick with “non-deductive” to make it clear we are referring to a very broad class of arguments.

5. Two points for the interested reader. First, just how probable do the premises have to be to make the conclusion? Is a probability of .95 enough? What about .85? The answer to this question will often depend on the context. For instance, the probability required in a

scientific context, say, will often be higher than the probability required in an everyday affair. Second, technically, a strong non-deductive argument must meet the requirement of “total evidence,” where that means no additional evidence available at the time would change the degree of support. The requirement of “total evidence” is a bit complicated and raises all sorts of issues, so we will put it aside here.

6. This case is inspired by the incident of the Russian satellite, *Kosmos* 954, that disappeared in 1977, spreading radioactive debris over northern Canada.

Writing Philosophy Papers

Objective

However one nudges particular issues in philosophy, there is nonetheless a common thread: philosophy involves the construction and critical analysis of *arguments*. (For further discussion, see “Logical Toolkit.”) Learning how to and refining your ability to construct and analyze arguments are (arguably, at least) the ultimate objectives of any philosophy course. It is, therefore, the purpose of writing a philosophy paper, because your professor will use your papers to judge how well you are meeting the objectives of the course. Showing that you have comprehended and are able to evaluate the material will require you to present your ideas and arguments in a *clear, explicit, and organized* fashion. You should be able to do this by adhering to the following guidelines.

Give Reasons

An argument can be understood as *reasons* that support a *conclusion*. Your number one objective in writing a philosophy paper should typically be to *give reasons* that support the overall point you want to make. However, although your paper should have an overall point, or conclusion, you may want to make several secondary, related points in the body of your paper. These points should also be backed by reasons. The idea here is that whenever you make a claim, you must give reasons that tell your reader why they ought to accept your claim. This amounts to *explaining why* you have made the claim you have made. (A very frequent comment made by professors on students’ papers is some variation on “Why?” “Give reasons,” “Support your claims,” or “Explain.”)

Answer the Question

Make sure that you are clear about what the question is on which your paper will focus. All too often students answer what they *think* the question is, or discuss what they *think* the topic is, without ever addressing the real problem. Addressing the topic and answering the question requires you to *understand* the topic you have chosen and to *think carefully* about what you will need to say to clearly explain and evaluate the problem.

Organize Your Paper

Your paper should be structured in such a way that your argument proceeds in an orderly fashion. Defending a claim requires giving reasons in support of that claim, but you can do this in a more or less orderly fashion. Among others, two mistakes can lead to a muddled argument: